Hewlett Packard
Enterprise

HPE 3PAR OS 3.3.1 GA/EGA/MU1 Release Notes

Abstract

This document describes the features and issues included in HPE 3PAR OS 3.3.1 GA/EGA/MU1 and is intended for use by Hewlett Packard Enterprise customers, partners and field representatives.

Part Number: QL226-99683a
Published: November 2017
Edition: 2

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Contents

Upgrade Considerations

The HPE 3PAR OS can be upgraded concurrently with I/O activity on the attached hosts, provided certain conditions are met. For more information on planning for online upgrades, refer to the latest version of the *HPE 3PAR Operating System Upgrade Pre-Planning Guide*. To obtain a copy of this documentation, go to the **Hewlett Packard Enterprise Information Library**.

OS upgrade prerequisite: Upgrade Tool version U008 or later must be staged prior to the HPE 3PAR OS upgrade to 3.3.1 MU1.

The Upgrade Tools are 3PAR OS upgrade enabling patches that do not affect array operation outside of the upgrade process. These tools are intended to improve the online or offline upgrade experience by performing preparatory steps to ensure the StoreServ is in a known state, including pre-checks, post-checks and other validations.

CAUTION:

Mandatory Patch Required for Use of File Persona with 3.3.1 MU1

In order to use File Persona with 3.3.1 MU1, install the mandatory 3.3.1 MU1 P07 patch after upgrading to 3.3.1 MU1. This patch contains important content to ensure stable operation of and compatibility for File Persona with MU1. If this patch is not installed:

1. Enabling file services for the first time will be prohibited. A message indicates that the patch needs to be installed.
2. Management requests may return unexpected results or fail unexpectedly. If File Persona has been enabled and the system has been upgraded to 3.3.1 MU1, do not attempt to modify the configuration of the system before installing the required patch.

Supported Platforms

For information regarding the supported HPE 3PAR StoreServ Storage systems, see the HPE Single Point of Connectivity Knowledge (SPOCK) website:

http://www.hpe.com/storage/spock

Notes

WARNING:

3PAR deduplication and compression are resource intensive operations, and as loads increase to these volumes, File Persona volume performance can decrease significantly. The load applied to volumes with these services enabled may need to be controlled in order to manage the impact to other volumes specifically volumes used by File Persona feature set as part of a File Provisioning Group.

Components

Table 1: Components and Versions

Component	Version
Maintenance Update	3.3.1.215
Patches	None
CLI Server	3.3.1.215
CLI Client	3.3.1.215
System Manager	3.3.1.215
Kernel	3.3.1.215
TPD Kernel Code	3.3.1.215
CIM Server	3.3.1.215
WSAPI Server	3.3.1.215
Console Menu	3.3.1.215
Event Manager	3.3.1.215
Internal Test Tools	3.3.1.215
LD Check Tools	3.3.1.215
Network Controller	3.3.1.215
Node Disk Scrubber	3.3.1.215
PD Scrubber	3.3.1.215
Per-Node Server	3.3.1.215
Persistent Repository	3.3.1.215
Powerfail Tools	3.3.1.215
Preserved Data Tools	3.3.1.215
Process Monitor	3.3.1.215
Software Updater	3.3.1.215
TOC Server	3.3.1.215

Table Continued

Component	Version
VV Check Tools	3.3.1.215
Upgrade Check Scripts	170330.U004 (3.3.1.215)
File Persona	1.3.0.74-20170309
SNMP Agent	1.10.0
SSH	6.0p1-4+deb7u5
VASA Provider	3.0.12
Firmware Database	3.3.1.215
Drive Firmware	3.3.1.215
UEFI BIOS	05.02.54
MCU Firmware (OKI)	4.8.60
MCU Firmware (STM)	5.3.17
Cage Firmware (DC1)	4.44
Cage Firmware (DC2)	2.64
Cage Firmware (DC3)	08
Cage Firmware (DC4)	2.64
Cage Firmware (DCN1)	4082
Cage Firmware (DCN2)	4082
Cage Firmware (DCS1)	4082
Cage Firmware (DCS2)	4082
Cage Firmware (DCS5)	2.78
Cage Firmware (DCS6)	2.78
Cage Firmware (DCS7)	4082
Cage Firmware (DCS8)	4082
QLogic QLA4052C HBA Firmware	03.00.01.77
QLogic QLE8242 CNA Firmware	04.15.27

Table Continued

Component	Version
QLogic 260x HBA FC Firmware	174.03.70
QLogic 27xx/268x HBA FC Firmware	174.03.70
QLogic 83xx HBA FCoE Firmware	08.01.05
QLogic 8300 HBA iSCSI Firmware	05.07.35
Emulex LP11002 HBA Firmware	02.82.x10
Emulex LPe12002 HBA Firmware	02.10.x02
Emulex LPe12004 HBA Firmware	02.10.x02
Emulex LPe16002 HBA Firmware	11.1.220.6
Emulex LPe16004 HBA Firmware	11.1.220.6
3PAR FC044X HBA Firmware	200A8
LSI 9201-16e HBA Firmware	17.11.03
LSI 9205-8e HBA Firmware	17.11.03
LSI 9300-8e HBA Firmware	10.00.08

What’s New in the OS

New and enhanced features include:

3PAR OS 3.3.1

- Inline Compression—Inline for optimal efficiency
- Data Packing—Combines data reduction and flash efficiency technologies to maintain peak capacity efficiency over time
- Adaptive data reduction—New support for inline compression and data packing designed to reduce the data footprint
- Adaptive Sparing 2.0
- Express Layout Enhancements—Express Layout is now supported for all drives, and not just solid-state drives (SSDs)
- Self Identifying Drives—3PAR systems can now automatically recognize a newly introduced drive without needing a software patch
- More Raw Capacity—Support for more raw capacity. Twice the SSD raw capacity supported compared to HPE 3PAR OS 3.2.2
- Loop topology connection mode for direct connection to 16 Gbps FC 3PAR StoreServ target
- Larger Volume Sizes—Full and thin provisioning virtual volume maximum sizes increased to 64 TiB

- `setcpg` growth and warning limits are no longer capped at 1 PiB
- New TDVV format—Enhanced deduplication and reporting
- Write Cache behavior options during single node operational states—New options to turn on write back cache to improve performance.
- Default RAID type is 6 for all drive types
- IPv6 now supports default gateways

3PAR File Persona

- NTFS Security Mode and cross protocol locking for seamless group file sharing—SMB and NFS
- Static and Dynamic User Mapping for mapping AD and LDAP users for cross protocol access
- File Lock Enterprise Mode to meet corporate governance requirements
- Larger File Provisioning Group size of 64 TiB with up to 250 million files for simpler scaling of large data sets
- Online File System Check to complement inherent file system integrity
- 3PAR Web Service API to automate File Persona management
- Enhancements to the Object Access API to support file copy and partial file access
- Support for Sophos antivirus scan engine
- Antivirus bulk quarantine support
- Inclusion of share folder ACLs in the VFS configuration backup/restore process
- Support for FTP/FTPS shares
- Internationalization of user names, share names, and File Store names
- Thin Persistence support for File Provisioning Groups
- Growth of File Provisioning Groups by growing the underlying volumes (rather than adding additional volumes)
- Incremental improvements to file random IO performance

SmartSAN 2.0

- 3PAR StoreServ Management Console (SSMC) 3.1 Integration
- 3PAR Federation Zoning
- Expanded ecosystem and diagnostics

3DC Peer Persistence—Now supports a tertiary passive site in addition to the two existing active sites.

Remote Copy—Async streaming supported using RCIP over 10 GbE ports

IMPORTANT:

Remote Copy Async Streaming does not support Compressed volumes.

VMware Virtual Volumes (VVols)

- Now support 3PAR Remote Copy replication for 1:1 mapping of virtual maps to storage volumes
- Support for iSCSI

Combo Cards Supported on 3PAR 8000 Systems

- 16 Gb FC/10 GbE NIC combo HBA
- 10 GbE iSCSI/10 GbE NIC combo HBA

DC PCM Support—New 48 VDC power cooling module (PCM) to offer DC power on 3PAR 8000 systems

Enhanced serviceability—Actionable alerts that contain spare part numbers of failed components

Alert messages are now internationalized and can be displayed in Japanese or simplified Chinese via the Service Processor or StoreServ Management Console (SSMC).

Direct Attach Cable (DAC) Support

The HPE 3PAR StoreServ Storage System DAC qualification matrix was expanded to accommodate new Active DAC cables including AP818A, AP820A, new passive cables QK701A and QK702A, and new HPE BladeSystem cables 487655-B21, 537963-B21 and 487658-B21. These new supported DAC cables are all HPE qualified/ supported with 3PAR. See the complete listing of 3PAR DAC cables supported in the *3PAR Platforms and Required DAC OS Support* table.

NOTE:

- The term “direct” refers to the direct attach of the cable to the SFP+ housing, instead of attaching to a SFP+ module that plugs into the SFP+ housing.
- DAC cable support for 3PAR StoreServ 8000 and 20000 storage platforms requires OS version 3PAR OS 3.2.2 MU3 and later.

Table 2: 3PAR Platforms and Required DAC OS Support

		3PAR StoreServ Platforms and Required DAC OS Support				
DAC Description	**DAC Part #**	**7000**	**10000**	**8000**	**9000 and 20000**	**Speed/ Protocols Supported**
HPE 3COM (H3C)						
HPE X240 10G SFP+ to SFP+ 0.65m DAC	JD095C	3.1.3 or later	3.1.3 or later	Not supported	Not supported	10GbE, iSCSI, FCoE, File*, RCIP
HPE X240 10G SFP+ to SFP+ 1.2m DAC Cable	JD096C	3.1.3 or later	3.1.3 or later	3.2.2 MU3	3.2.2 MU3	10GbE, iSCSI, FCoE, File*, RCIP

Table Continued

		3PAR StoreServ Platforms and Required DAC OS Support				
DAC Description	**DAC Part #**	**7000**	**10000**	**8000**	**9000 and 20000**	**Speed/ Protocols Supported**
HPE X240 10G SFP+ to SFP+ 3m DAC Cable	JD097C	3.1.3 or later	3.1.3 or later	3.2.2 MU3	3.2.2 MU3	10GbE, iSCSI, FCoE, File*, RCIP
HPE X240 10G SFP+ to SFP+ 5m DAC	JG081C	3.1.3 or later	3.1.3 or later	3.2.2 MU3	3.2.2 MU3	10GbE, iSCSI, FCoE, File*, RCIP
HPE x240 QSFP+ 4x10G SFP+ 1m DAC Cable	JG329A	3.2.2 MU3	3.2.2 MU3	3.2.2 MU3	3.2.2 MU3	10GbE, iSCSI, FCoE, File*, RCIP
HPE X240 10G SFP+ to SFP+ 7m DAC	JC784C	3.1.3 or later	3.1.3 or later	3.3.1 or later	3.3.1 or later	10GbE, iSCSI, FCoE, File*, RCIP
HPE x240 QSFP+ 4x10G SFP+ 3m DAC Cable	JG330A	3.2.2 MU3	3.2.2 MU3	3.2.2 MU3	3.2.2 MU3	10GbE, iSCSI, FCoE, File*, RCIP
HPE x240 QSFP+ 4x10G SFP+ 5m DAC Cable	JG331A	3.2.2 MU3	3.2.2 MU3	3.2.2 MU3	3.2.2 MU3	10GbE, iSCSI, FCoE, File*, RCIP
HPE Procurve						
HPE 10-GbE SFP+ 1m DAC	J9281B	3.1.3 or later	3.1.3 or later	3.2.2 MU3	3.2.2 MU3	10GbE, iSCSI, FCoE, File*, RCIP
HPE 10-GbE SFP+ 3m DAC	J9283B	3.1.3 or later	3.1.3 or later	3.2.2 MU3	3.2.2 MU3	10GbE, iSCSI, FCoE, File*, RCIP
HPE X242 10G SFP+ to SFP+ 7m DAC	J9285B	3.1.3 or later	3.1.3 or later	3.3.1 or later	3.3.1 or later	10GbE, iSCSI, FCoE, File*, RCIP
HPE StoreFabric						
HPE C-series 3m Passive Copper SFP+ Cable	K2Q21A	3.1.3 or later	3.1.3 or later	3.2.2 MU3	3.2.2 MU3	10GbE, iSCSI, FCoE, File*, RCIP
HPE C-series 5m Passive Copper SFP+ Cable	K2Q22A	3.1.3 or later	3.1.3 or later	3.2.2 MU3	3.2.2 MU3	10GbE, iSCSI, FCoE, File*, RCIP

Table Continued

		3PAR StoreServ Platforms and Required DAC OS Support				
DAC Description	**DAC Part #**	**7000**	**10000**	**8000**	**9000 and 20000**	**Speed/ Protocols Supported**
HPE C-series 7m Passive Copper SFP+ Cable	QK701A	3.1.3 or later	3.1.3 or later	3.3.1 or later	3.3.1 or later	10GbE, iSCSI, FCoE, File*, RCIP
HPE C-series 10m Passive Copper SFP+ Cable	QK702A	3.1.3 or later	3.1.3 or later	3.3.1 or later	3.3.1 or later	10GbE, iSCSI, FCoE, File*, RCIP
HPE 1m B-series Active Copper SFP+ Cable	AP818A	3.1.3 or later	3.1.3 or later	3.3.1 or later	3.3.1 or later	10GbE, iSCSI, FCoE, File*, RCIP
HPE 3m B-series Active Copper SFP+ Cable	AP819A	3.2.2 MU4	3.2.2 MU4	3.2.2 MU4	3.2.2 MU4	10GbE, iSCSI, FCoE, File*, RCIP
HPE 5m B-series Active Copper SFP+ Cable	AP820A	3.1.3 or later	3.1.3 or later	3.3.1 or later	3.3.1 or later	10GbE, iSCSI, FCoE, File*, RCIP
HPE Blade System						
HPE BladeSystem c-Class 10 GbE SFP+ to SFP+ 3m Direct Attach Copper Cable	487655-B21	3.1.3 or later	3.1.3 or later	3.3.1 or later	3.3.1 or later	10GbE, iSCSI, FCoE, File*, RCIP
HPE BladeSystem c-Class 10 GbE SFP+ to SFP+ 5m Direct Attach Copper Cable	537963-B21	3.1.3 or later	3.1.3 or later	3.3.1 or later	3.3.1 or later	10GbE, iSCSI, FCoE, File*, RCIP
HPE BladeSystem c-Class 10 GbE SFP+ to SFP+ 7m Direct Attach Copper Cable	487658-B21	3.1.3 or later	3.1.3 or later	3.3.1 or later	3.3.1 or later	10GbE, iSCSI, FCoE, File*, RCIP

Notes:

- DAC cable support for HPE 3PAR StoreServ 8000 and 20000 platforms requires HPE 3PAR OS version 3.2.2. MU3 and later.
- All protocols are supported only with HPE 3PAR OS 3.2.2 MU3 and later.
- File* protocol is supported only with HPE 3PAR OS 3.2.2 and later.

Modifications to the HPE 3PAR OS

The following issues have been addressed in this release.

Issue IDs: 106328

Issue summary: Upgrade checks are too aggressive when performing an offline upgrade, preventing an upgrade when it should proceed.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.1.3, 3.2.1, 3.2.2

Issue description: The `checkupgrade` command is used to determine the system readiness to perform an upgrade. Offline upgrades have fewer restrictions because host I/O interruption is a given. The `checkupgrade` command was using online criteria for performing the checks despite an offline upgrade being performed, blocking the upgrade from proceeding when it should have been allowed to proceed.

Symptoms: An offline upgrade may not proceed due to a check that is only applicable for online upgrades being executed.

Conditions of occurrence: When using SPOCC to compete an offline HPE 3PAR OS upgrade.

Impact: Medium

Customer circumvention: None

Customer recovery steps: Resolve the condition that resulted in the upgrade check failure before attempting the upgrade again.

Issue IDs: 126114

Issue summary: Certain data backup solutions cannot access the secondary array in Remote Copy Peer Persistence configurations.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.1.3, 3.2.1, 3.2.2

Issue description: Allows data backup solutions, such as VADP (VMware vStorage API for Data Protection), to access data from the secondary site in Remote Copy Peer Persistence configurations. With the HPE 3PAR OS, the backup solution must use the Generic (non-ALUA) host persona when presenting volumes in a Remote Copy Peer Persistence group to the backup application.

Symptoms: Data backup solutions cannot read data from a Remote Copy secondary array.

Conditions of occurrence: Volumes in Remote Copy Peer Persistence groups on the secondary array when the backup solution tries to access the data on those volumes.

Impact: Medium

Table Continued

Customer circumvention: Set up the data backup solution to access the Remote Copy Peer Persistence primary system.

Customer recovery steps: Use primary system instead of the secondary system for backup operations.

Issue IDs: 141617

Issue summary: Unified Extensible Firmware Interface (UEFI) restart failure alert delivery can be delayed for an indefinite amount of time if an EEPROM read encounters a transient failure.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.2

Issue description: Transient read problems of a controller node's EEPROM data can postpone the delivery of restart failure alerts indefinitely. Because a reread of the data is based on a restart of the system manager process, the delivery of the alerts can be suppressed. This can cause what appears to be a stale alert to be posted at some later time.

Symptoms: UEFI restart failure alerts are not reported in a timely manner if a transient read failure is encountered, despite a controller node having been unable to restart previously.

Conditions of occurrence: A transient read failure can delay the posting of a UEFI restart failure alert indefinitely.

Impact: Low

Customer circumvention: None

Customer recovery steps: None

Issue IDs: 142277

Issue summary: `removecert` removed certificates for both `ekm-server` and `ekm-client` when just a individual `ekm` service was specified.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1 GA & All MUs

Issue description: This issue has been corrected. A `removecert` command will now only remove a certificate of the specified `ekm` service.

Symptoms: `removecert` for `ekm-client` or `ekm-server` would remove certificates for both `ekm-client` and `ekm-server`.

Conditions of occurrence: Having an `ekm_client` and `ekm_server` certificate installed and removing a single one.

Impact: High

Table Continued

Customer circumvention: None

Customer recovery steps: Re-import the removed certificates.

Issue IDs: 144868

Issue summary: Controller nodes with full internal boot drives cause `sysmgr` to not start if controller nodes are restarted in that state.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.1.2, 3.1.3, 3.2.1, 3.2.2

Issue description: A full internal boot drive file system on a controller node will cause `sysmgr` and other system services to not start.

Symptoms: While starting an online upgrade, system manager does not restart.

Conditions of occurrence: The root file system of a node drive has run out of space.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue IDs: 146146

Issue summary: An unhelpful message is displayed when an attempt to add more File Persona (FP) nodes to a system with FP installed in some nodes but not in a running state.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1

Issue description: Addition of more File Persona (FP) nodes requires FP to be running on nodes which have it already configured. The error message displayed when FP on those nodes is in a shutoff state was unhelpful and provided no guidance as to the reason for this. The error message produced when adding new nodes to an existing FP cluster which are not running has been updated to: "File Persona is installed on nodes x,y but not running. To configure additional nodes run the command: `startfs -enable`."

Symptoms: `startfs` used to add new nodes to the File Persona configuration yields the message "File Persona must be running to allow additional nodes to be configured."

Conditions of occurrence: File Persona is installed but not running and an attempt is made to add FP on more nodes.

Impact: Low

Table Continued

Customer circumvention: Check that FP is running on all nodes it has previously been installed onto before attempting to install FP on more nodes. Run the command `showfs` to display the FP status.

Customer recovery steps:

1. Run `showfs` to determine that FP nodes are not in a running state.
2. Run `startfs -enable` to start any nodes which are currently not running.

Issue IDs: 146489, 146490

Issue summary: Change to SSH ciphers to align with industry best practices for security and network integrity.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: SSH clients used prior to 3.3.1.GA

Issue description: SSH Client update may be necessary! SSH Ciphers have a been changed; only the following ciphers groups are now supported.

Supported Ciphers

- KexAlgorithms: `diffie-hellman-group-exchange-sha256`
- Ciphers: `chacha20-poly1305@openssh.com`, `aes256gcm@openssh.com`, `aes128-gcm@openssh.com`, `aes256-ctr`,`aes192-ctr`, and `aes128-ctr`.
- MACs: `hmac-sha2-512-etm@openssh.com`, `hmac-sha2-256-etm@openssh.com`, `hmac-ripemd160-etm@openssh.com`, `umac-128-etm@openssh.com`, `hmac-sha2-512`, `hmac-sha2-256`, `hmac-ripemd160`, and `umac-128@openssh.com`.

Previously supported Ciphers

- KexAlgorithms: `curve25519-sha256@libssh.org`, `ecdh-sha2-nistp256`, `ecdh-sha2-nistp384`, `ecdh-sha2-nistp521`, `diffie-hellman-group-exchange-sha256`, `diffie-hellman-group-exchange-sha1`, `diffie-hellman-group14-sha1`, `diffie-hellman-group1-sha1`
- Ciphers: `aes192-ctr`, `aes256-ctr`, `aes128-ctr`, `aes192-cbc`, `aes256-cbc`, `aes128-cbc`, `3des-cbc`
- MACs: `hmac-sha1` and `hmac-sha1-96`

Customers using the OpenBSD SSH client can examine their supported ciphers to determine compatibility by examining man 5 `ssh_config`. There must be at least 1 Cipher in common in each three Cipher groups for the client to be compatible with HPE 3PAR OS.

Symptoms: SSH access to the array may be impacted when using clients which were used with prior versions of HPE 3PAR OS.

Conditions of occurrence: Updating to 3.3.1GA or later and attempting to use an older SSH cypher.

Table Continued

Impact: High

Customer circumvention: None

Customer recovery steps: SSH Client update or configuration.

Issue IDs: 146805

Issue summary: In a Remote Copy configuration, when a full sync on the primary array is stopped before it completes and a promotion happens on secondary array, subsequent resync could cause data inconsistency. This issue only applies to periodic group.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: Detected in 3.2.1 and 3.2.2; fixed in 3.3.1

Issue description: When full sync on primary array is stopped before it completes, a promotion occurs on secondary array to overwrite the base volume. As a result of the promotion, data between primary and secondary became inconsistent. A subsequent resync continues from the point where the previous full sync left off leading to miscompare. This issue only applies to periodic group.

Symptoms: There is data inconsistency on the remote copy target volumes.

Conditions of occurrence:

1. Full sync on primary is stopped before it completes.
2. A promotion automatically occurs on the secondary array to overwrite the base volume.
3. A subsequent resync is started on primary array.

Impact: Low

Customer circumvention: To prevent getting this issue, make sure arrays do not run out of space within the CPG. You can set the snapshot space allocation warning and user space allocation warning using the `setvv` command.

Customer recovery steps: Do another full sync to recover.

Issue IDs: 146991

Issue summary: CPG alerts in `showcpg` output may not automatically clear.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.1.3, 3.2.1, 3.2.2

Issue description: Prior to 3.3.1, the CPG Alerts fields in `showcpg` output may indicate an alert is set after the underlying condition has been resolved.

Symptoms: Response from CLI `showcpg -alert` may indicate a W/F/L alert is set ('Y') after the associated condition and alert have been cleared.

Table Continued

Conditions of occurrence:

- A CPG Grow operation which triggers a Warning, Fail or Limit alert.
- The condition which caused the alert is resolved.
- The corresponding alert (W/F/L) indicator to remain set ('Y') after the associated condition was resolved.

Impact: Low

Customer circumvention: Issue is resolved in 3.3.1

Customer recovery steps: The user can correct the display by issuing a redundant `setcpg` command to the affected CPG. For example, if the current CPG occupancy percentage warning is 50%, then issuing a CLI `setcpg -aw 50` to the affected CPG will clear the condition.

Issue IDs: 153893

Issue summary: `movetodomain` may cause the system manager to restart (recursive thread stack overflow).

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1, 3.2.2

Issue description: Using `movetodomain` with a very complex web of related VVs, LDs, CPGs, sets, RC groups and hosts may be unsuccessful. Recursion is no longer used to discover the complete list of objects that have to be moved to the new domain.

Symptoms: `movetodomain` may not succeed on complex web of objects, and you may receive the following message: "`Eagle IPC transport error: EA_PROCESS_DOWN --Message canceled because of process down.`"

Conditions of occurrence: Using the CLI command `movetodomain` to operate on a large number of objects that are related.

Impact: High

Customer circumvention: Plan ahead and set up virtual domains before creating several hundred hosts, VVs, CPGs, sets, and RC groups.

Customer recovery steps: None

Issue IDs: 156155

Issue summary: Array becomes unresponsive if the system manager restarts while region moves are in progress.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.2

Table Continued

Issue description: In extreme cases where multiple very large conversions are happening at once when the system manager restarts, then processing a lot of mirroring regions causes the system manager to become unresponsive.

Symptoms: Longer system manager restart times when system manager restarts in the middle of region movement on very large VVs.

Conditions of occurrence: The system manager is restarted while moving regions on large VVs. System manager has to restart.

Impact: Low

Customer circumvention: None

Customer recovery steps: Wait for system manager to complete its restart.

Issue IDs: 158195

Issue summary: User is unable to remove a Virtual Volume using `removevv`.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.2

Issue description: A scenario was created where the admin space was marked to be dropped and not able to be removed. Once this happened, the `removevv` command refused to remove the VV it thought was in the middle of having its admin space dropped.

Symptoms: A VV cannot be removed and returns the message: "Cannot remove volume as the entire snapshot tree is being removed."

Conditions of occurrence: An unexpected system manager or controller node restart when removing an entire VV tree using admin drop (normal removes don't use this).

Impact: Low

Customer circumvention: Do not perform controller node reboots while running `removevv`. Avoid operations known to restart the system manger while running `removevv`, such as installing a patch that contains the system manager component.

Customer recovery steps: None

Issue IDs: 159520

Issue summary: A VV block can occur every second when a large number of VV conversions are in progress, which can lead to host I/O stalling.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.1.2, 3.1.3, 3.2.1, 3.2.2

Table Continued

Issue description: A condition exists on the array that is preventing the VV blocking mechanism to work as designed while converting multiple VVs. This generally leads to the VV conversion failing.

Symptoms: Host I/O appears to be stalled while VV conversions are in progress.

Conditions of occurrence: Something prevents blocks attempting to convert more than 30 VVs simultaneously.

Impact: Low

Customer circumvention: Don't convert more than 30 VVs at once.

Customer recovery steps: None

Issue IDs: 160406

Issue summary: Host I/O stalls after attempting volume removal.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: All versions since 3PAR OS 3.2.1 MU3 Patch 38

Issue description: System resources attempt to access the same internal system locks multiple times with different requests in between the duplicate lock requests that results in a deadlock which results in the array's inability to share data.

Symptoms: The array becomes unresponsive and requires restart.

Conditions of occurrence: It is a timing issue. Theoretically, issuing a `freespace` command at the same time as removing a VV which had data on it could cause the issue. Because it's a timing issue, the probability to encounter the issue is low.

Impact: High

Customer circumvention: Do not run `freespace` while there is a volume removal in process.

Customer recovery steps: None

Issue IDs: 169491

Issue summary: `srdatac` log file grows too large because the system does not rotate the log file.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.2 MU2

Issue description: When the `sradatac` log file has no limit on the log file size, which leads to excessive use of space on the node disk for this log file.

Symptoms: Excess space on the node disk being used by the `srdatac` log file.

Table Continued

Conditions of occurrence: Excessive writing to `srdatac` log file when System Reporter is experiencing startup issues.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue IDs: 178014

Issue summary: Adaptive Optimization (AO) does not complete data region moves because a memory buffer cannot be allocated..

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.1.2, 3.1.3, 3.2.1, 3.2.2

Issue description: Inability to allocate a memory buffer in one individual LD can cause 64 LDs to fail region statistic collection, resulting in inability to run Adaptive Optimization accurately against a significant number of LDs.

Symptoms: AO does not move data between tiers as expected.

Conditions of occurrence: The only indication that the buffer allocation will adversely affect AO is seen in the /var/log/tpd/aomover log file: "Error in getstatldrg ... LD XYZ region stats not active".

Impact: Low

Customer circumvention: None

Customer recovery steps: Use customer circumvention steps.

Issue IDs: 180117

Issue summary: Reduced RAID protection after recovery from replaced or unavailable drive.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1 MU5 P53, 3.3.1

Issue description: When a drive will be replaced, the RAID system relocates data away from that drive in order to preserve the desired RAID protection. After the drive has been replaced, the RAID system will migrate back to the new drive to maintain the balanced I/O load. In certain circumstances, it is possible that the RAID protection will be degraded as a result of the migration back.

Symptoms: Reduced RAID availability seen in `showld -d`.

Conditions of occurrence: An unavailable or replaced drive that contains user data.

Impact: Medium

Table Continued

Customer circumvention: None

Customer recovery steps: Manually move the affected data regions to spares, which will pick the best RAID level available.

Issue IDs: 181090

Issue summary: In rare cases it was possible for any System Reporter (SR) cli command (or SSMC SR report) with the `-compareby` option to return an incomplete set of results.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1, 3.2.2

Issue description: System Reporter requests with the `-compareby` option always included a defined number of objects for which to return data. Because of an error in the query logic, it was possible for a reduced number of objects to be included in the final results.

Symptoms: System Reporter (SR) CLI command (or SSMC SR report) with the `-compareby` option return an incomplete set of results.

Conditions of occurrence: Run SR where the range of time specified (`-btsecs` and `-etsecs`) for SR spans the internal SR database files. The user cannot easily determine if the SR DB files are spanned.

Impact: Low

Customer circumvention: In order to completely avoid the problem it is necessary to avoid using the `-compareby` functionality. The likelihood of encountering the problem of a reduced data set can be greatly reduced by requesting data in smaller time windows (`-btsecs` to `-etsecs`), and making use of more granular data (hourly or daily) as appropriate for longer time windows.

Customer recovery steps: None

Issue IDs: 183278

Issue summary: Event log is flooded with internal connection messages.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.2

Issue description: An "infinite" loop in `srdatac` causes it to send CLI commands continuously, which causes an event for each iteration.

Symptoms: An excessive number of events, about one every second, similar to: "Debug Informational CLI server process event sw_cli User logged in Id:516 User:3parsvc Level:super Addr:127.0.0.1 (client local) app:CLI"

Conditions of occurrence: Occurs when a single controller node which is not the System Reporter owner node is restarted.

Table Continued

Impact: Low

Customer circumvention:

Re-starting the System Reporter processes can temporarily stop the flood of events:

```
cli stopsr -f
cli startsr -f
```

Customer recovery steps: None

Issue IDs: 184670

Issue summary: On four and eight node systems, an unexpected array restart closely following an unexpected controller node down can prohibit cluster integration.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1, 3.2.2

Issue description: First, there is a single controller node outage event. Following this event, during node rejoin, there is another unexpected event, such as a power loss. When the array is restarting, another controller node experiences a resource contention it can't handle because of the dual unexpected event. This small timing window and sequence of events has been resolved. This can only occur on systems with four or more nodes.

Symptoms: The array will restart three times.

Conditions of occurrence:

1. A controller node goes down.

2. The array unexpectedly restarts while the node in step #1 was coming back online.

3. When the entire array restarts from #2, another controller node, not the same controller node in step #1 is not able to completely recover due to resource contention. When this specific scenario occurs, the array restarts three times to clear the conditions to come back online.

Impact: High

Customer circumvention: None

Customer recovery steps: None

Issue IDs: 185414

Issue summary: `showcage -d` lacked an enclosure overall state field.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.2

Table Continued

Issue description: Because `showcage -d` was lacking an enclosure overall state field, the enclosure status obtained through other software, like SSMC, would not have an equivalent counterpart in `showcage` cli. Conditions like a missing IO card connection or an outdated firmware would cause SSMC to show a "degraded" enclosure overall state, while in `showcage -d` there will be no equivalent 'degraded' state.

Symptoms: SSMC displays a "degraded" overall status for the enclosure but there's no equivalent "degraded'" status in `showcage -d`.

Conditions of occurrence: Having an enclosure that has a missing I/O card connection or an outdated firmware.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue IDs: 189474

Issue summary: Unbalanced performance with a disproportionate mixture of merge cache buckets for 100k and 150k SSDs.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1 MU5

Issue description: On a storage array with both SSD 100 and SSD 150 drives, where there are a lot more of one drive type than another, hosts may see much larger I/O latencies for I/O targeted to the smaller population of drives.

Symptoms: Long I/O latencies for the host only when using the smaller pool of SSD.

Conditions of occurrence: A large number of SSD 100/SSD 150 and a small number of the other. There is also a significant IOPs host load.

Impact: Medium

Customer circumvention: Install the drive types in a balanced setup, or do not mix drive types.

Customer recovery steps: Until the system is balanced, relocate data away from the drive type with fewer drives.

Issue IDs: 191018

Issue summary: Physical VV copy takes a long time copying to a VV that is a much larger size.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.1.2, 3.1.3, 3.2.1, 3.2.2

Table Continued

Issue description: In order to finish a VV copy to a larger destination VV, the difference in size needs to be zeroed in order to ensure that the volumes are equal. This zeroing can add significant time. The issue is improved by adding logic to detect that the destination VV is completely empty and therefore does not need to have any zero writes applied.

Symptoms: Physical copy takes longer than expected.

Conditions of occurrence: Physical copy from a source VV to another VV of significantly larger size.

Impact: Low

Customer circumvention: A faster option can be to size the destination VV the same as the source VV then, after the copy is complete, grow the destination VV to its desired final size.

Customer recovery steps: None

Issue IDs: 203495/201975

Issue summary: Defrag IO logs is not well handled in node down recovery.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1.215

Issue description: When defrag is IO running and a node down happens, the logs for defrag IO are not handled. When another IO comes to the same offset after recovery, it will cause another node down due to the unhandled log. The result is the recovery node will reboot or the cluster down.

Symptoms: Unexpected node restart or cluster down after a node down.

Conditions of occurrence: Node down happens during defrag IO and logs from defrag are left over.

Impact: Medium

Customer circumvention: Install P01.

Customer recovery steps: After one more node down, it will be automatically recovered.

Known Issues with the OS

Issue IDs: 94331

Issue summary: The Management Console Volume Raw Space pie chart on the Physical Disks Summary tab incorrectly displays value on StoreServ with Adaptive Optimization software active.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions:

Table Continued

Issue description: The Volume Raw Space pie chart on the Physical Disks Summary tab incorrectly displays value for the selected device type on a StoreServ with Adaptive Optimization software active. This is due to the Management Console just adding up the virtual size of the virtual volume initially created from a Common Provisioning Group with the selected device type. With Adaptive Optimization software active, some of the virtual volume's regions might have been moved to another tier, and this needs to be taken into account when calculating the raw space for this pie chart.

Symptoms: The Management Console Volume Raw Space pie chart on the Physical Disks Summary tab incorrectly displays value.

Conditions of occurrence: Occurs when Adaptive Optimization is active.

Impact: Low

Customer circumvention: None

Customer recovery steps: None

Issue IDs: 112187

Issue summary: The `startfs` commands does not complete and time outs without configuring the File Persona cluster.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.2.GA-3.2.2.MU4, 3.3.1.GA

Issue description: In rare circumstances, `startfs n:sp n:sp...` may not complete after displaying the message "Executing `createfsvm fs_cpg`." This will be accompanied by an alert indicating that the `createfsvm` task has failed.

Symptoms: The `startfs` command hangs does not complete the tasks to create the File Persona configuration on one or more node does not complete.

Conditions of occurrence: Normal operation

Impact: Medium

Customer circumvention: The `startfs` command should be rerun after the previous invocation of the `startfs` command, including the tasks started by it, and any configuration created is automatically rolled back.

Customer recovery steps: Rerun the `startfs` command after the rollback recovery is complete.

Issue IDs: 131710

Issue summary: SR commands can return errors.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.1.1.*, 3.1.2.*, 3.2.1.*, 3.2.2.GA-3.2.2.MU4, 3.3.1.GA

Table Continued

Issue description: SR command can return a message if it internally requires large amounts of data.

Symptoms: SR commands return an "EA_PROCESS down" message.

Conditions of occurrence: Send an SR command that reads large amounts of data internally.

Impact: Medium

Customer circumvention: Do not use SR commands if seen.

Customer recovery steps: None. The system automatically recovers.

Issue IDs: 133562

Issue summary: iSCSI IO latency spikes

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1.GA - 3.2.1.MU5, 3.2.2.GA - 3.2.2.MU2

Issue description: iSCSI IO latency spikes as the IO requests and transfers would stall for up to 30 seconds before getting a response.

Symptoms: IO requests and transfers would stall for up to 30 seconds before getting a response.

Conditions of occurrence: The driver was using an interrupt mask that would cause an interrupt to be missed causing the IO delay by up to 30 seconds, depending on the next NOP_In/Out occurrence..

Impact: Low

Customer circumvention: Work around can be applied for reducing the heartbeat_interval to 1 to cause the iSCSI NOP_IN to occur every second:

```
tcli -e "kvar set -n iscsi_heartbeat_misses -v 120"
tcli -e "kvar set -n iscsi_heartbeat_interval -v 1"
```

Customer recovery steps: The system would recover from the IO pause on its own within the heartbeat time interval which is 30 seconds by default.

Issue IDs: 160232

Issue summary: Volumes with TPGID in range 3 to 256 are not allowed to join RC group.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.1.3 MU3, 3.2.2 MU3 - 3.2.2 MU4, 3.3.1

Issue description: When volumes are migrated from other arrays using Online Import Utility (OIU), it is possible for its TPGID to be in the range 3 to 256. When we try to add these volumes to Remote Copy group, it will produce the message "`tpgid <tpgid vlaue> does not match with group <group name>'s tpgid <257/258>`". Volumes with TPGID 0, 1 or 2 do not have this issue.

Table Continued

Symptoms: Volumes cannot be added to Remote Copy group.

Conditions of occurrence: Adding volume with TPGID in the range 3 to 256 to an RC group.

Impact: Medium

Customer circumvention: None

Customer recovery steps: Change the TPGID of the volume to 1 or 2 using command `setvv -settpgid <1/2> <vvname>`. After changing the TPGID, it can be added to RC group.

Issue IDs: 165063

Issue summary: Online conversions, online copy, online promote, `updatevv`, and imports have long I/O stall times on 20000 systems.

Affected platforms: StoreServ 20000

Affected software versions: 3.2.2.GA - 3.2.2.MU4, 3.3.1.GA

Issue description: Online conversions, online copy, online promote, `updatevv`, and imports have long I/O stall times on StoreServ 20000 systems due to internal structure invalidation.

Symptoms: Long I/O stall times during online conversions, online copy, online promote, `updatevv`, and imports.

Conditions of occurrence:

- Have a StoreServ 20000 system
- Start an online conversions, online copy, online promote, `updatevv`, or import
- See a long I/O stall time

Impact: High

Customer circumvention: Avoid online conversions, online copy, online promote, `updatevv`, and imports on StoreServ 20000 systems.

Customer recovery steps: The hosts will time out. Use standard recovery for host timeouts.

Issue IDs: 187897

Issue summary: Disk enclosures report a power control module (PCM) inlet temperature sensor reporting a "non_critical/under_warning" falsely implying that the inlet temperature is too cold.

Affected platforms: StoreServ 7000, StoreServ 8000

Affected software versions: 3.3.1

Issue description: Array logging event/alert: "non_critical/under_warning" for drive cage FW enclosure PCM0 or PCM1 inlet sensor.

Table Continued

Symptoms: Array logging event/alert: "non_critical/under_warning" for drive cage FW enclosure PCM0 or PCM1 inlet sensor.

Conditions of occurrence: Drive cage FW 406a or prior and cold data centers (< 10 degrees Celsius)

- System running drive cage FW version 406a on cage models DCN1, DCS1, DCS2, DCN2, DCS7, DCS8.
- Inlet temperature low enough to confuse drive cage FW into interpreting PCM0/1 inlet temp as below low temp threshold.

Impact: High

Customer circumvention: Ignore event. The event/alert is misleading, but low temperature threshold violations do not trigger any array recovery behavior that would cascade into an outage or data loss.

Customer recovery steps: None

Issue IDs: 192368

Issue summary: `cachesvr` process memory consumption may cause other processes to stop.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1.GA - 3.2.1.MU3, 3.2.2.GA - 3.2.2.MU3

Issue description: Over time the `cachesvr` process on the cluster master node may exhaust free memory, causing other user processes to halt. When this occurs, the affected process will restart and may continue to halt until the `cachesvr` process is restarted. Once the `cachesvr` process is restarted, its memory utilization is reset and the problem will not occur for some time, based upon system configuration and management activities performed.

Symptoms: `cachesvr` process memory size grows over time and causes other process to halt with the message "Unable to allocate xxxxxxxx bytes."

Conditions of occurrence: The issue is most likely to occur on systems which have large configurations and which execute frequent array management interactions.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue IDs: 193758

Issue summary: Large number of `updatevv` operations could lead to rare and unexpected IO stalls.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 322GA-322MU4, 3.3.1

Table Continued

Issue description: A large number of `updatevv` operations could lead to rare and unexpected IO stalls.

Symptoms: IO stalls could be encountered on StoreServ which goes through frequent and large number of `updatevv` operations.

Conditions of occurrence: Frequent and intense `updatevv` operations on snapshot volumes.

Impact: Medium

Customer circumvention: Reduce the frequency of events leading to intense `updatevv` operations.

Customer recovery steps: None

Issue IDs: 193846

Issue summary: `tunesys` does not apply the `-fulldiskpct` or `-chunkpct` options to the intra-node phase when active-active PDs are present.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1.GA (all PDs)

Issue description: An issue has been found with `tunesys` when custom values for `-fulldiskpct` or `-chunkpct` are supplied to control the chunklet movement phase and LD re-layout phases of the intra-node tuning, respectively. This affects all drive types.

Symptoms: `-fulldiskpct` and `-chunkpct` are used to customize intra-node re-balancing. When these options are used, expected tunes are not generated.

Conditions of occurrence: `tunesys -fulldiskpct <value> -chunkpct <value>` does not generate expected intra-node tunes.

Impact: Low

Customer circumvention: None

Customer recovery steps: Run manual intra-node tunes in consultation with HPE support.

Issue IDs: 196124

Issue summary: The CLI command `startfs -enable` does not complete due to the number of `rsh` connections open exceeding the allowed limit.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1 GA

Issue description: A configuration with a large number of FPGs (>32 on an 8 node, >64 on a 4 node) causes the CPG to run out of space, the ensuing intentional deactivation of affected FPGs may cause subsequent `startfs enable` commands not to work.

Table Continued

Symptoms: The `startfs -enable` command failed with error " Failed to get bridge list: Could not run {/sbin/brctl show} on node0: node0: Connection refused."

Conditions of occurrence: A large number of FPGs > 32 on 8 node, > 64 on 4 node; the CPG containing the FPGs is full and File Persona has shut down the FPGs; or `startfs -enable` is run.

Impact: High

Customer circumvention: Ensure the CPG which has the FPGs never runs out of space.

Customer recovery steps: None

Issue IDs: 196633

Issue summary: `setcpg` can default the RAID type of SD space to RAID 6.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1.GA

Issue description: An issue has been reported with the CLI `setcpg` command if no RAID type is explicitly defined in the new option list. In this case the existing RAID type will be removed from the list of stored options, and the CPG will silently inherit the system default of RAID 6. This applies to all `devtypes` (SSD,FC,NL).

Symptoms: After `setcpg` is used to update the CPG creation options customers may experience any or all of the following:

- VV Creation or growth failures
- Snapspace growth failures resulting in stale snapshots

Conditions of occurrence: This will only happen on systems where it is not possible to create RAID 6 `setsize 8` sets with cage availability (e.g. where RAID 5 or RAID 1 was configured previously).

Impact: Medium

Customer circumvention: Always explicitly specify ALL options when `setcpg` is used from the CLI. (This issue does not affect changing the CPG settings via the SSMC.)

Customer recovery steps: Use `setcpg` to refresh the CPG creation options to include all relevant parameters; in particular this should include the RAID type, set size, device type and availability.

Issue IDs: 196758

Issue summary: The `tunevv` command may unexpectedly not work or change a volume to the default RAID 6 `setsize 8` if the target CPG has an undefined RAID type.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1.GA

Table Continued

Issue description: An issue has been reported with the `tunevv` command where, if the target CPG has no RAID type defined, the tune may either not work or change the volume to RAID 6 `setsize 8` unexpectedly. (Note that the `tunesys` command will warn the user and will not rebalance any volumes where any associated CPG does not have a defined RAID type. This check is missing from the `tunevv` command.)

Symptoms: If the target CPG has no defined RAID type the following may occur:

- The tune may fail if the system does not have resource to create tune destination LDs with RAID 6 `setsize 8`, cage availability.
- The tune will succeed but will modify the volume to be the new system default RAID type of RAID 6.

Conditions of occurrence: This may occur if the target CPG of the tune has no configured RAID type.

Impact: Medium

Customer circumvention: Make sure that the target CPG of all tunes have a specified RAID type.

Customer recovery steps: Use `setcpg` to refresh the CPG creation options to include all relevant parameters; in particular this should include the RAID type, set size, device type and availability.

Issue IDs: 199218

Issue summary: Imports and `updatevv` have long host I/O stall times.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1.GA

Issue description: Imports or `updatevv` with a large list of VVs will have long I/O stall times.

Symptoms: Long host I/O stall time.

Conditions of occurrence:

- Start an import or `updatevv` with a large list of VVs
- Long host I/O stall time

Impact: High

Customer circumvention: Avoid using imports or `updatevv` with a large list of VVs.

Customer recovery steps: The hosts will time out. Use standard recovery for host timeouts.

Issue IDs: 199904/168180

Issue summary: StoreServ controller node unexpectedly restarts while handling IO.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Table Continued

Affected software versions: 3.3.1

Issue description: StoreServ controller node(s) unexpectedly restarts while handling host IO.

Symptoms: Restart of StoreServ controller node.

Conditions of occurrence: This is a corner case situation with blockless region moves happening. Region moves could be due to tuning, conversions.

Impact: Medium

Customer circumvention: Disable blockless region move with help from HPE support.

Customer recovery steps: StoreServ self recovery as in the case of any situation needing a controller node restart.

Issue IDs: 200606

Issue summary: `showvv -s` can display negative numbers for Used size for compressed volumes.

Affected platforms: StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1

Issue description: The `showvv -s` command, used to show space information, can sometimes display a negative value for the one of the used size columns (Snp, Usr, Total) for compressed volumes.

Symptoms: An obviously incorrect and negative value in one or more of the used size columns for a compressed volume.

Conditions of occurrence: This is a transient and infrequent occurrence when running `showvv -s` on compressed volumes.

Impact: Low

Customer circumvention: The `HostWr` column will display an accurate value for the amount of data written to the volume.

Customer recovery steps: The condition will resolve itself as more data is written.

Issue IDs: 201016

Issue summary: Mix of Pentium and Cluster memory buffers in request to Harrier2 when Remote Copy Asynchronous Streaming and Compression are enabled on the same volume.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1

Issue description: Target cluster will allocate Pentium memory resources by default. Under high resource usage these may be swapped with Cluster memory in the Remote Copy layer. Remote Copy may then send copy requests to the Harrier2 with a mixture of PM and CM leading to the assertion.

Table Continued

Symptoms: Node panic from Assertion point.

Conditions of occurrence: Compression volume added to Remote Copy Asynchronous group. Running low on resources while running IO to the volume can lead to bug occurring on target cluster.

Impact: High

Customer circumvention: Do not have compression volumes in Remote Copy Asynchronous groups.

Customer recovery steps: None.

Issue IDs: 201039

Issue summary: Performance of existing File Persona workloads may decrease more than expected when adding block workloads leveraging deduplication and compression.

Affected platforms: StoreServ 7000c, StoreServ 8000, StoreServ 20000

Affected software versions: 3.2.2, 3.3.1

Issue description: Deduplication and compression are resource intensive operations, and as the IO load to volumes with these services increases, the performance of other volumes that may or may not be using these services can decrease significantly. This impact can include both internal volumes used by the File Persona feature set as part of a File Provisioning Group and volumes consumed by external hosts.

Symptoms: Symptoms: Lower than expected performance.

Conditions of occurrence: Introduction of block workloads leveraging deduplication and compression.

Impact: Medium

Customer circumvention: The load applied to volumes with deduplication and/or compression enabled may need to be controlled in order to manage the impact to other volumes. One way to control the impact from these services is via the use of the 3PAR Priority Optimization feature set. You can create and modify threshold limits including I/O per second, bandwidth and latency on the volumes leveraging deduplication and/or compression in order to reduce their impact on the performance of other volumes and services.

Customer recovery steps: Reduce the newly introduced workload and then implement the circumvention recommendations.

Issue IDs: 201182

Issue summary: Recovery of File Persona FPGs (File Provisioning Groups) with names longer than 12 characters may require additional time.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.2, 3.3.1

Issue description: In the event that a File Persona FPG needs to be checked during a recovery, long FPG names will require support personnel to perform additional actions, potentially prolonging any outage.

Table Continued

Symptoms: Attempts by support personnel to perform an online check of the FPG does not work due to a long name.

Conditions of occurrence: FPGs with names greater than 12 characters exist; an FPG recovery check (`fsck`) is required.

Impact: Medium

Customer circumvention: Limit FPG names to 12 characters.

Customer recovery steps: None

Issue IDs: 203126

Issue summary: Express layout with a minimal configured system must use restricted set sizes.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1.GA

Issue description: In order to provide RAID protection, the maximum set size of an LD must be restricted. Considering the number of PDs that match the LD specification (for example, `-ha`, `-p`, `-devtype`), the maximum set size for the LD must be no more than the number of PDs, less the fault tolerance.

Symptoms: A failed disk immediately leads to a degraded LD, and the RAID protection shown in the LD is not actually available.

Conditions of occurrence: An LD layout selecting PDs where the set size of the LD, plus the fault tolerance of the RAID type is less than the number of those PDs.

Impact: High

Customer circumvention: Ensure the set size is limited as described.

Customer recovery steps: Tune the LD onto a new LD that follows the limitation.

Issue IDs: 206190

Issue summary: When an HPE 3PAR Online Upgrade from a release prior to 3.3.1 GA or 3.3.1 EGA is performed while a Windows Cluster online migration is in progress, it can result in an unexpected restart of the array.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1 GA, 3.3.1 EGA

Issue description: Performing an HPE 3PAR Online Upgrade from a release prior to 3.3.1 GA or EGA while a Windows Cluster online migration is in progress can result in cyclic System Manager restarts and ultimately an unexpected array restart.

Table Continued

Symptoms: The Cluster Shared Volumes for the Windows Cluster will go offline.

The HPE 3PAR OS Online Upgrade does not complete.

Conditions of occurrence: Performing a Windows Cluster online migration.

Performing an HPE 3PAR OS Online Upgrade.

Impact: High

Customer circumvention: Allow Windows Cluster online migration to complete successfully before performing the HPE 3PAR OS Online Upgrade.

Customer recovery steps: Wait for the array to come back online, wait for Windows Cluster Shared Volumes to come back online, and then restart these applications.

By using StoreServ Management Console, resume the peer motion action. Allow the Windows Cluster online migration to complete successfully.

Once the migration is complete, perform the HPE 3PAR OS Online Upgrade.

Modifications to File Persona

Issues that have been addressed in this release.

Issue ID	Summary	Description
67397	A request to stop file services on a node may result in them restarting.	Infrequently, a request to stop file services on a node may result in the services restarting instead of going to a stopped state.
68476	Cannot change only the VLAN tag of a node IP address.	The VLAN tag for a node IP address could not be changed without first moving the IP address to a different subnet temporarily.
76213	Antivirus scanning impacts read/write performance for small files.	Small file performance was significantly degraded when antivirus support was enabled.
76395	Password expiration policy changed for local users requires reset before effective.	Password expiration policy for local users has changed to "never expires." In previous releases, the default required passwords to be changed for local users after 30 days.
76846	Renaming a parent directory when child directory is open with directory change notification causes SMB users to be disconnected from node.	All SMB users could be temporarily disconnected from a node if a parent directory was renamed while a child directory was open with a directory change notification.

Table Continued

Issue ID	Summary	Description
77559	Local users and groups do not show up in Windows if Active Domain is missing in Provider Order.	Local users and groups could not be enumerated from a client if the system was joined to Active Directory, but Active Directory was not included in the provider stacking order.
78078	File Persona services become unavailable temporarily.	The management of File Persona services could periodically become unavailable for some time and then become available again on their own.
80075	Intermittent failure in scheduled snapshots/ snapshot reclamation.	The tracking of a snapshot space reclamation task would be interrupted and would require support assistance to recover.
80897	Share directory is not created when creating share using MMC.	Starting with 3.3.1 GA, to ensure proper behavior in conjunction with the cross protocol support added in the release, if a share is created through MMC, it is now expected that the user must: **1.** Go through explorer. **2.** Create the directory. **3.** Share the directory once it is created.
89743	When a File Provisioning Group (FPG) has a large number of objects, FPG performance may be decreased.	When an FPG object count approaches the 250,000,000 threshold, FPG performance may be decreased as the object count increases. With HPE 3PAR OS 3.3.1, the following system alert (message code 0x0720001) has been added when this threshold has been reached: "`FPG cc_fpg102 object count is approaching or has exceeded the maximum supported, 250000000. FPG performance may decrease as the object count increases.`"
92322	Only files and directories from the live view are included in the Files Used field displayed by the `showfpg -d` command.	The "Files" value in the `showfpg -d` output now includes snapshot versions of files and other internal metadata objects.
92967	SMB protocol access scenario leads to excessively high CPU usage.	Using a certain SMB protocol access scenario could lead to excessively high CPU usage (and lower performance.)

Table Continued

Issue ID	Summary	Description
93127	Filename wild carding from CMD "DOS" does not work correctly on Windows Server 2012 R2.	Looking for files using a wildcard pattern containing multiple '.' characters from a Windows Server 2012 client resulted in unexpected response.
94964	Snapshot plugin sometimes fails with cannot get actor reference, and actor system is terminated.	File Store snapshot creation would fail with the message "cannot get actor reference. Actor system is terminated", and a restart of file services on the impacted node was required to recover.
95776	The update record status is not handled properly after an unexpected restart of file services during the upgrade process.	Unexpected restart of file services during the upgrade process could leave the upgrade in a state where support intervention was required to complete the upgrade.

Known Issues with File Persona

Issue ID	Summary	Description	Corrective Action
74861	"Unknown error 528" error message on NFSv3 during `setfacl`. Unknown error 528 may be encountered when using Network File System (NFS) version 3 (NFSv3) to set file permissions using the `setfacl` utility or from access contention handling when accessing the file Access Control List (ACL).	Unknown error 528 may be encountered when using Network File System (NFS) version 3 (NFSv3) to set file permissions using the `setfacl` utility or from access contention handling when accessing the file Access Control List (ACL). This issue may occur in any NFSv3 implementation but is more likely to occur in a Lightweight Directory Access Protocol (LDAP) authenticated environment. Per NFSv3 specifications, clients should retry operations of this type, should the command fail. See section 4.5 in the NFSv3 specifications at: **https://www.ietf.org/rfc/rfc1813.txt**	To prevent this issue, user must either utilize a client that complies with the NFSv3 specification for retries, or do not use `setfacl` via script or utility that would allow multiple operations to occur in a short period of time. To recover from this issue, retry the failed operation. Several retries may be needed during periods of heavy `setfacl` call load.
75737	Setting Access Control Entries via a UID/GID that cannot be resolved will fail.	Setting access control entry via UID or GID fails if ID cannot be properly resolved to user or group name.	Make sure the UID and GUID are added to the name server before trying to use them on a file or directory.

Table Continued

Issue ID	Summary	Description	Corrective Action
75911	Metadata inconsistency reported on NFS I/O after failover event.	In some versions of NFS clients, on rare occasions while using V4 could result in file metadata inconsistencies during heavy I/O and failover.	Using the `noac` option during NFS mount would help address these situations of incorrect file attribute cache handling. But using the `noac` option will have a significant performance impact, and it is recommend to use it only for those applications which exhibit these issues.
77773	Avoiding name collisions when creating users and groups in AD.	When creating a user in AD, there are two name fields, one called "User logon name" and the other called "User logon name (pre-Windows 2000)."	To prevent possible name collisions and confusion with names stored in ACLs, the following is recommended: **1.** Make sure that neither of the two name fields is the same as the name of any other user or group in the domain. **2.** Set both of the two name fields to the same name when creating a user.
79212	Need better messaging (alert) when data is unavailable due to time sources being out of sync.	If the system is not configured for NTP before starting file services, and the system is joined to active directory, if the system time and active directory time are not in sync, some unexpected behaviors may occur.	It is important to configure NTP on the system before starting file services if you are planning to use Active Directory for authentication.
82177	Severe performance problems for file operations.	If files have UID values that cannot be mapped to a known user via one of the enabled authentication providers, accessing those files can result in higher than expected CPU utilization and lower performance.	Ensure that users can be mapped successfully to a name.
83268	Internal error: Mapping operation failed : 40,404	This condition happens when the “ToName” user or group has been configured with a UID/GID value of less than 1.	Ensure that UID/GID values less than 1 are not used in the "ToName" user or group.

Table Continued

Issue ID	Summary	Description	Corrective Action
83635	Creating SMB share on existing VFS using MMC, breaks share enumeration on the CLI.	Do not use Windows management tool MMC to create shares at the root of the VFS. Doing this will cause shares to stop enumerating.	To restore enumeration, remove the share using MMC.
83701	User can change permissions of C$ share, but eventually fails with error.	Do not use Windows management commands to add ACEs to c$ share. Attempting to change permissions at this top level will fail.	To get the permissions applied correctly, the command must be run at a lower level in the directory structure.
86217	Status of AD server in health is always 'Online'	The AD server connection health is not currently monitored.	The administrator of the Active Directory (AD) server can verify it is up and running. The cluster administrator can verify the AD host name is resolvable and pingable.
88762	Tight loop of HTTP requests or FTP requests creates large log on LDAP server.	When files are accessed frequently over FTP or Object Access API shares, there will be a high number of authentication requests to the LDAP server when using LDAP for authentication. If the log file is not managed on the LDAP server, then the file system of the LDAP server can be filled and cause the LDAP service to stop responding.	Make sure an appropriate log rotation policy is in place on the LDAP server when using it for authentication.

Table Continued

Issue ID	Summary	Description	Corrective Action
89456	Excessive I/O load during multiple Roaming user logoff may cause sync issue.	Excessive stress through creation of a huge I/O load across multiple roaming profile users (42 sessions) and then deletion followed by re-creation at the same time may have data sync issues observed for few of the files/ folders during Logoff. The error following error message is displayed: "Windows cannot copy file <Local Windows path> to location <Share path>. This error may be caused by network problems or insufficient security rights. DETAIL - Access is denied."	It is recommended to copy those files/folders specifically in such a scenario.
91456	Race condition during saves to SMB share using Notepad on nearly full FPG results in user data not being saved and no user error returned.	When using certain applications such as Notepad that do not honor indications of disk full during write requests (only during preallocation), and when writing to a nearly full FPG that consists of more than one VV, the application may indicate that data has been saved when in fact the disk was full.	Make sure to respond to the alerts indicating the FPG is 80% or 90% full and grow the FPG.
92080	Stopping Active management node immediately after cluster expansion can loose LDAP configuration.	After successfully starting file services on additional nodes and configuring networking for those newly added nodes, the existing LDAP configuration can take up to 10 minutes to get replicated to all the new nodes. If the currently active node (as shown by showfs) is stopped during this time, the LDAP configuration may be disabled.	If this occurs, the user will need to reconfigure the LDAP provider using `setfs` command. To avoid this issue, avoid stopping any node within 10 minutes of configuring additional nodes.
93279	Spurious `monitor.startprocess.ok` event reported.	Occasionally, an event with the identifier `monitor.startprocess.ok` may be reported unexpectedly.	This event can be safely ignored.

Table Continued

Issue ID	Summary	Description	Corrective Action
93701	Unable to use the same name for local user and local group.	Same name for local group and user is not supported with AD.	Use LDAP as the name provider.
94190	Manual intervention may be required to reestablish connectivity if AD server connectivity is interrupted.	If connectivity to the Active Directory server is interrupted, manual intervention may be required to reestablish connectivity.	Connectivity can be reestablished by issuing the `stopfs` command followed by the `startfs -enable` command for any impacted node. Alternatively, support can be engaged to accurately diagnose the issue and recover without restarting the entire file services for the node.
94267	All snapshots fail when Snapshot component is not functional. Cannot get actor reference, and actor system is terminated.	When all snapshot operations fail with “Snapshot component is not functional. Cannot get actor reference”, manual intervention may be required to reestablish snapshot capabilities.	Snapshot capabilities can be reestablished by issuing the `stopfs` command following by the `startfs -enable` command for any impacted node. Alternatively, support can be engaged to accurately diagnose the issue and recover without restarting the entire file services for the node.
96847	No snapshots listed even though snapshots exist..	When there is a significant load of snapshot related activity, for example, several snapshot creation / deletion / reclamation jobs are run in parallel, sometimes `showfsnap` command returns "No snapshots listed."	Re-trying the same operation after some time when the load eases will be listed accordingly. If a create/delete snapshot operation failed with error “Futures timed out,” internally the operation would have completed successfully, and can be validated using the `showfsnap` command.

Table Continued

Issue ID	Summary	Description	Corrective Action
97092	With AD configured after LDAP in auth stack and with unreachable LDAP, server may cause status to reported as Starting.	With LDAP configured before Active Directory in Auth stacking order, any AD user/ group lookup requests will go through the LDAP provider first before sending it Active Directory. If LDAP is down/not-reachable, any AD user/ group lookup requests becomes unresponsive, and the management interface and reporting of Starting state via `showfs` may be unresponsive.	If this occurs, checking and repairing the health of LDAP provider should restore the ability to manage the system.

Table Continued

Issue ID	Summary	Description	Corrective Action
97253	Executing multiple `showfsquota` commands can cause system to respond slowly or cause subsequent commands to fail.	When LDAP server is unavailable (LDAP is configured), executing the `showfsquota` CLI command multiple times might cause the system to respond very slowly or fail the execution of subsequent commands.	An admin should ensure that the LDAP server is up and running. Admin is notified through system alerts when the LDAP server has gone down.

97662	Unable to rediscover VTLs after node reboot.	If a node is rebooted, VTL tapes associated with NDMP backup may no longer be seen.	Perform the following steps to rediscover attached VTLs:
			1. Execute following command on the HPE 3PAR CLI:
			`showfsndmp -vtl vtldevices`
			It will list VTL device IPs similar to the following:
			`VtlDeviceIp`
			`1.1.1.1`
			`1.1.1.2`
			2. Execute following command by providing all above IPs separated by commas:
			`setfsndmp vtl +1.1.1.1,1.1.1.2`
			All VTLs will be rediscovered.

Installation Notes for the CLI

Deprecated Commands and Options

The deprecated options for the `cli`, `createuser`, and `setpassword` commands have been removed from the documentation.

Compatibility Changes in this Release

Remote CLI Client versions prior to 3.2.2 cannot connect to version 3.3.1 of the 3PAR OS without using the `-nosockssl` option.

NOTE:

The 3.3.1 Remote CLI Client is not backward compatible with 3.2.2 GA, 3.2.2 MU1, and releases prior to 3.2.1 MU5.

Compatibility changes in the next release

The following options will be removed:

`cli`: `-pwf`, `-user`, `-password`, and variable environment *TPDPWFILE*

`createuser`: `-e`

`setpassword`: `-save`, `-saveonly`, `-file`

Operating systems no longer supported:

- Red Hat Enterprise Linux 5 (RHEL 5)
- SUSE Linux Enterprise Server 10 (SLES 10)
- Ubuntu 12.04 LTS

Installation Directory

Default installation locations are new in 3PAR CLI 3.3.1:

- **Windows 32-bit:** `C:\Program Files\Hewlett Packard Enterprise\HPE 3PAR CLI`
- **Windows 64-bit:** `C:\Program Files (x86)\Hewlett Packard Enterprise\HPE 3PAR CLI`
- **UNIX and Linux:** `/opt/hpe_3par_cli`

In Windows, the Programs Menu has changed: Start->Programs->HPE 3PAR CLI->HPE 3PAR CLI <version>

Supported Operating Systems

For the list of supported operating systems, see the *3PAR CLI Remote Client* document on the SPOCK website at **SPOCK**.

Support for the following additional operating systems is provided in this release:

- Red Hat Enterprise Linux 6 Update 7 (RHEL 6.7)
- Red Hat Enterprise Linux 6 Update 8 (RHEL 6.8)

- Red Hat Enterprise Linux 7 Update 2 (RHEL 7.2)
- Red Hat Enterprise Linux 7 Update 3 (RHEL 7.3)
- SUSE Linux Enterprise Server 12 (SLES 12)
- Ubuntu 16.04 LTS
- Windows 10 Enterprise
- Windows Server 2016

What's New in the CLI

A Linux Control group has been added to restrict memory used by CLI and `tpdtcl` processes running on the array. This limitation under severe low memory situations will improve overall system stability. Under severe memory pressure, the performance of the Remote CLI may be hindered and potentially cause CLI sessions to terminate. These include tasks and other programs invoked indirectly by the CLI or `tpdtcl` server.

New Commands

- `removefsarchive`
- `setfsarchive`
- `showfsarchive`
- `srstatiscsi`
- `srstatiscsisession`
- `srstatvv`
- `srsysspace`
- `startfsarchive`
- `stopfsarchive`

Changed Commands

Command	Description
`checkhealth`	New `-d` option
`checkvv`	New `-compr_dryrun` option
`controlsr`	New subcommands `setperiod` and `setretention`
`createfpg`	Max size 64 TiB
`createfshare`	New subcommand `ftp`
`createfstore`	New mandatory `-secmode` option

Table Continued

Command	Description
createsv	New -addtoset option
createvv	Added three new policies for host DIF support; extended -f option to skip DIF policy change warning message; Compression changes
growfpg	Max size 64 TiB
histpd	New -devsvtime option
importvv	New -compr and -dedup compression options
locatecage	Support locate commands on HPE 3PAR StoreServ 8000 Storage system
removedomain	Added -pat option
removedomainset	Added -pat option
removefshare	New subcommand ftp
removehost	Added -pat option
removehostset	Added -pat option
removevvset	Added -pat option
setfpg	New -upgrade option
setfs	New subcommand usermap
setfsav	New -quar_file; SOPHOS added to -vendor
setfshare	New subcommand ftp
setfstore	New -secop_errsuppress and -secmode options
setrcopygroup	New policy mt_pp
setrcopytarget	New subcommand autotunelinks

Table Continued

Command	Description
setsralertcrit	Allows more changes, Merges SSD100 and SSD150 metrics
setsys	Added OverprovRatioLimit, OverprovRatioWarning, allowR5OnFCDrives,DisableCompr, AllowWrtbackUpgrade, and AllowWrtbackSingleNode
setvv	New policies: 3par_host_dif, std_host_dif, no_host_dif
showcpg	New -listcols and -showcols, output format changes
showfs	New -usermap option
showfsarchive	New -importfile, -export options and subcommand export
showfshare	New subcommand ftp
showfstore	Output changes
showhost	Output changes for -agent
showiscsisession	New -d option
showld	New -ck option
shownode	New -pci type "combo"
showportdev	New -d option for subcommand tzone, new subcommand uns
showsys	New -vvspace option
showtask	Limit increased to 2000
showuserconn	Output for -d lists memory
showvlun	New -pathsum columns
showvv	New showvv -pol output for host DIF settings; New compression output changes, changes to output of showvv -s and showvv -d
sr*	New -compareby option
srcpgspace	Compression output changes
srhistvlun	VVol filtering

Table Continued

Command	Description
`srrgiodensity`	Added `-totpct` option
`srstatvlun`	New `-vlun`, VVol filtering
`srvvspace`	VVol filtering. Compression output changes.
`statpd`	Added `-devsvtime` option
`tunesys`	New `-force`, `-slsz`, `-slth`, `-compactmb`, `-cleanwait`, `-maxnodetasks` and `-ss`

Modifications to the CLI

Issue IDs: 79971
Issue summary: `checkhealth` doesn't detect degraded SFPs in converged network adapters (CNAs). **Affected platforms:** StoreServ 10000 **Affected software versions:** 3.1.1 (MU2) **Issue description:** `checkhealth` doesn't detect degraded SFPs in converged network adapters (CNAs). **Symptoms:** None **Conditions of occurrence:** `checkhealth` doesn't detect degraded SFPs in converged network adapters (CNAs). **Impact:** Low **Customer circumvention:** None **Customer recovery steps:** None

Issue IDs: 126970
Issue summary: New controller nodes that are connected and not yet powered on or admitted may go unreported by `checkhealth`. These controller nodes may prevent a successful upgrade. **Affected platforms:** StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000 **Affected software versions:** 3.1.1 (MU2) **Issue description:** New controller nodes that are connected and not yet powered on or admitted may go unreported by `checkhealth`. These controller nodes may prevent a successful upgrade. **Symptoms:** Upgrade stalls.

Table Continued

Conditions of occurrence: A StoreServ with controller nodes not powered or not admitted to the cluster, but the cables are connected and the system is aware that something is plugged into those node slots.

Impact: Medium

Customer circumvention: Avoid leaving new controller nodes in a state where they are cabled, but not admitted.

Customer recovery steps: Power on affected nodes and run the CLI command `admithw`.

Issue IDs: 136799

Issue summary: `checkhealth` should detect phantom connections due to a stall on a socket read.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1 (MU3)

Issue description: The CLI `checkhealth` network should flag `tpdtcl` SSL sessions that do not finished authenticating within 5 minutes. These are presumed to be stalled

Symptoms: Login stalls with message, "Too many CLI connections."

Conditions of occurrence: CLI connection stall.

Impact: Medium

Customer circumvention: None

Customer recovery steps: Quit unresponsive CLI connection process.

Issue IDs: 138748

Issue summary: `checkhealth` does not provide a warning when the node time and `hwclock` (hardware clock) differ.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1 (MU2)

Issue description: If the node time and `hwclock` differ, then `checkhealth` should log a corresponding error.

Symptoms: There is a time difference between the node time and `hwclock`.

Conditions of occurrence: There are no specific conditions for this issue to appear except for a notable time difference (more than 60 seconds) between the hardware clock and the node time.

Impact: Low

Table Continued

Customer circumvention: None

Customer recovery steps: `hwclock --systohc` forces the current software clock's time to match the hardware clock.

Issue IDs: 146487

Issue summary: TLS v1.0 and 1.1 have been disabled to align with industry best practices for security and network integrity.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: All TLS client software

Issue description: TLS v1.0 and 1.1 have been disabled to align with industry best practices for security and network integrity.

Symptoms: TLS clients which are configured for older TLS versions may no longer connect to the 3PAR array after the array is updated to 3.3.1.

Conditions of occurrence: Update to 3.3.1GA.

Impact: High

Customer circumvention: None

Customer recovery steps: Update, or reconfigure, affected TLS clients to use TLS 1.2.

Issue IDs: 152319

Issue summary: CLI on HP-UX stalls when /home is NFS mounted and the NFS server is not available.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1, 3.2.2

Issue description: If /home is NFS mounted and NFS server is not available, Remote CLI client on HP-UX stalls.

Symptoms: Remote CLI client on HP-UX stalls.

Conditions of occurrence: /home is NFS mounted and NFS server is not available. Customer is trying to use the Remote CLI client. This issue is seen only on HP-UX.

Impact: High

Customer circumvention: Use SSH or 3.3.1 HPE 3PAR CLI Remote Client to connect the HPE StoreServ system. For a list of supported versions of each operating system, go to the Single Point of Connectivity Knowledge (SPOCK) for HPE Storage Products at **http://www.hpe.com/storage/spock**.

Customer recovery steps: This issue occurs because `ActiveTcl` is trying to access the /home/andreask directory, which most likely is not available in the customer setup. Creation of /home/andreask locally can mitigate this issue.

Issue IDs: 155314

Issue summary: Starting in 3.3.1, the HPE 3PAR CLI will have a new default certificate directory. This will cause previously accepted certificates to be ignored.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1

Issue description: Starting in 3.3.1, the HPE 3PAR CLI will have a new default certificate directory.

Old:

Linux, HP-UX, Solaris and AIX: $HOME/.hp3par

Windows: %USERPROFILE%\.hp3par

New:

Linux, HP-UX, Solaris and AIX: $HOME/.hpe3par

Windows: %USERPROFILE%\.hpe3par

If already using TPDCERTDIR environment variable or the `-certdir` option, no additional changes are needed.

Symptoms: When attempting to connect using the 3.3.1 HPE 3PAR CLI, the authenticity of the storage system cannot be established. Any applications that sit on top of the CLI may not be expecting this new message/dialog and may fail.

Conditions of occurrence: Use of the 3.3.1 HPE 3PAR CLI and not using the `TPDCERTDIR` environment variable or `-certdir` option.

Impact: High

Customer circumvention: Users of older HPE 3PAR CLI versions prior 3.3.1 will need to move/copy/link certificates located in the old directory to the new directory. A separate copy may be needed if using older versions of the CLI to communicate with older arrays with the same shared home directory. Copying the certificate files would be more convenient than accepting each existing certificate. As an alternative to copying the certificate files, the `TPDCERTDIR` environment variable or `-certdir` option can be used to point to the previous certificate directory being used.

Customer recovery steps: None

Issue IDs: 159572

Issue summary: CLI TLS Cipher Changes.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: All Prior to 3.3.1GA

Issue description: Cli TLS Cipher Changes:

Supported: `AES128-SHA`, `AES256-SHA`, `DHE-RSA-AES128-SHA`, `DHE-RSA-AES256-SHA`

Previously Supported: `DHE-RSA-AES256-GCM-SHA384`, `DHE-RSA-AES128-GCM-SHA256`

Table Continued

Symptoms: CLI clients which are configured for prior HPE 3PAR OS versions may no longer connect to the HPE 3PAR StoreServ Storage system after the array is updated to 3.3.1.

Conditions of occurrence: The HPE 3PAR array is running 3.3.1 or later and a non-supported cypher is used.

Impact: High

Customer circumvention: None

Customer recovery steps: If connectivity issues occur, reconfigure the clients to use currently supported cipher from the above list.

Issue IDs: 167576

Issue summary: Array unexpectedly reconfigures Remote Copy Fibre Channel (RCFC) ports to host mode when executing `admithw`.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.2

Issue description: `admithw` reconfigures all Fibre Channel ports, including RC ports, that are in a "free" state to host connection mode.

Symptoms: A possible loss of RC ports used during HPE 3PAR OS or hardware upgrade when `admithw` is executed.

Conditions of occurrence: Having RC in use, but temporary free or disconnected, during `admithw` execution.

Impact: High

Customer circumvention: Guarantee that before executing `admithw`, all FC ports, including RC ports, are properly connected and not showing as `free` in `showport`.

Customer recovery steps: Reconfigure any incorrectly configured RC port back to Remote Copy mode.

Issue IDs: 179378

Issue summary: Users with edit or higher permissions are able to use `updatevv` on virtual volumes in their domains.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: All versions before 3.3.1

Table Continued

Issue description: Previously, a super-user would have to issue the command `setuseracl <username> updatevv <virtual volume name>` to allow a non-super user to utilize the `updatevv` command. This process is no longer required given the user is granted edit or higher permissions for the domains to which the virtual volumes belong. The user can then use `updatevv` without requiring a super-user issue the `setuseracl` command.

Symptoms: When a non-super user, issues the command `updatevv <virtual volume name>` the user will get a "permission denied" message, given the command `setuseracl` was not issued for them.

Conditions of occurrence: The user does not have edit or higher permissions for the domain to which the virtual volume belongs.

Impact: Low

Customer circumvention: None

Customer recovery steps: None

Issue IDs: 184028

Issue summary: WSAPI audit trail support: `tpdtcl` needs to put original request IP/port info in the `eventlog` and `showuserconn`.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1, 3.2.2

Issue description: The event log now includes the remote IP and port of WSAPI sessions. This will also change the `showuserconn` output to include the port number *For example:* `100.100.100.100:port`. The port will also be included for CLI, SSMC, SSH and MC connections in both `eventlogs` and `showuserconn`.

Symptoms: WSAPI sessions always have an array local address of 127.0.0.1 or 127.127.0.1 to 127.127.0.8. Port info is missing for the IP addresses.

Conditions of occurrence: WSAPI connections always have local IP.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue IDs: 186303

Issue summary: `checkhealth` does not cover a DDS or `VVol VV internal_consistency_error` issue.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1 (MU3)

Table Continued

Issue description: `checkhealth` VV missing checks

Symptoms: `checkhealth` addresses internal consistency errors for system volumes.

Conditions of occurrence: `checkhealth` addresses internal consistency errors for system volumes.

Impact: Medium

Customer circumvention: `checkhealth` addresses internal consistency errors for system volumes.

Customer recovery steps: None

What's New with the CIM API and SNMP Software

New and enhanced features include:

- CIM API
  - Support for compression.
  - Disabled SSL zlib compression to address the "CRIME" vulnerability.
  - HTTPS is now enabled by default while HTTP is disabled by default. This is only true for new systems: firmware upgrades will not change the existing configuration.
  - A new "SparePartNumber" property was added to the Alert Indication class to indicate the customer-orderable replacement part number for faulty components.
- SNMP
  - The 3PAR MIB has been updated with a cpuStatsMIB that contains CPU statistics for each Node in a StoreServ array.
  - SNMP Alerts now contain fields for event tier and spare part information. The spare part information is shown if it is available for hardware tier alerts.

Modifications to the 3PAR CIM API

Issue IDs: 145085
Issue summary: A cimserver `IndicationSubscription` cannot be deleted. **Affected platforms:** StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000 **Affected software versions:** 3.2.1, 3.2.2 **Issue description:** `CIM_IndicationFilter` instances that exist only in the root/tpd but not interop namespace cannot be enumerated and deleted. **Symptoms:** The cimserver API will return a NOT FOUND error when attempting to delete a `CIM_IndicationSubscription`. **Conditions of occurrence:** `CIM_IndicationFilter` is created in root/tpd namespace only. **Impact:** Low **Customer circumvention:** None **Customer recovery steps:** Create the exact same `CIM_IndicationFilter` in interop namespace also.

Issue IDs: 161149

Issue summary: Volumes created with `CreateStorageVolumeFromStoragePoolWithTemplate` do not use the snapshot CPG specified by the storage setting.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.2

Issue description: The snapshot CPG specified by the `TPD_StorageSetting` template is not configured for volumes created with the CIM API call `CreateStorageVolumeFromStoragePoolWithTemplate`.

Symptoms: `CreateStorageVolumeFromStoragePoolWithTemplate` creates a storage volume without the snapshot CPG specified by the `SnapDSPName` property of the TPD_StorageSetting template instance.

Conditions of occurrence: Call the `CreateStoragevolumeFromStoragePoolWithTemplate` API function with a `TPD_StorageSetting` that has the property `SnapDSPName` specified with a valid CPG name.

Impact: Medium

Customer circumvention: None

Customer recovery steps: Stop and restart the cimserver by running the following CLI command: `setvv -snp_cpg <cpgName> <vvname>`

Issue IDs: 192537

Issue summary: Frequent polling of cage status by applications using the CIM API may cause invalid events indicating a cage interface card failure when none has occurred.

Affected platforms: StoreServ 7000, StoreServ 20000

Affected software versions: 3.2.2

Issue description: Customers with applications issuing frequent CIM API requests for controller nodes, drive cage, power supply, battery, or magazine information observe erroneous events that indicate an interface card failure.

Symptoms: The event log will contain events indicating the failure and recovery of Interface cards even though no failure has occurred:

2016-11-29 13:35:45 CET 0 Major Component state change hw_cage:4,hw_cage_ifc:0 Cage 4, Interface Card 0 Failed

2016-11-29 13:36:16 CET 0 Informational Component state change hw_cage:4,hw_cage_ifc:0 Cage 4, Interface Card 0 Normal

Conditions of occurrence: The CIM API (CIM server) is enabled as shown by the `showcim` CLI command. A customer application such as "CA Unified Manager v8.4" is polling the CIM API for controller node, drive cage, power supply, battery or magazine information.

Table Continued

Impact: Medium

Customer circumvention: Disable the CIM API with the `stopcim` command.

Customer recovery steps: None

What's New with the Web Services API Software

New and enhanced features include:

- Support for Compression
- Support for File Persona—Create/Update/Delete functions for VFSs, FPGs, file stores, file shares, quotas, snapshots, and directory permissions
- Improved API response time
- Audit trail for the Web Services API in the HPE 3PAR OS system event log
- Added a `uuid` field to volume set and host set objects
- Added `id`-based and `uuid` -based filtering for volume sets and host sets
- Added ability to query virtual copy objects, given a parent virtual volume
- Added ability to specify a volume set target during the creation of a virtual copy
- Added a list of patches installed on the system, accessible at URI …/api/v1/system
- Added detailed task message for single instance of GET tasks
- Returns `deviceName` as part of `portdevices` query
- Supports `hostDIF` volume policy
- Now supports the following System Parameters: `remoteSyslogSecurityHost`, `hostDIFTemplate`, `disableChunkletInitUNMAP`, `personaProfile`, `remoteCopyHostThrottling`, `AllowR5OnFCDrives`, and `AllowR5OnNLDrives`.
- Additions to Remote Copy functionality:
  - Pattern matching for queries of RC groups
  - Added an option (`allowRemoteCopyParent`) so promotion of a virtual copy can proceed even if the RW parent volume is currently in a Remote Copy volume group, if said group has not been started
  - Detailed information for remote copy links
- Additions to System Reporter (SR):
  - Added ability to query SR VLUN statistic data based on VLUN filters. The SR VLUN statistic data is limit to VLUNs that are matching the specified combination of `host`, `VV`, `LUN id` and `port`.
  - Added `privateSpaceMiB`, `sharedSpaceMiB`, `freeSpaceMiB`, and `totalSpaceMiB` fields to SR CPG space and CPG information.
  - Added `compression` and `hostWriteMiB` fields to SR volume space.
  - Added SR data for CPU
- Cluster Extension capabilities:

- Embedded 3PAR Cluster Extension storage failover logic in 3PAR OS with access by 3PAR Web Services API.
- Changed Cluster Extension Host software for Microsoft Windows to include Microsoft Windows Cluster integration logic only and to use 3PAR Web Services API to perform planned migration and disaster recovery for the Microsoft failover cluster integrated applications.

Modifications to the 3PAR Web Services API

Issue IDs: 160211
Issue summary: Intermittent `NON_EXISTENT_VOL` message reported by WSAPI after volume creation **Affected platforms:** StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000 **Affected software versions:** 3.2.2 **Issue description:** If a volume creation and volume query is done in quick successions via WSAPI, a message may be generated where WSAPI reports a `NON_EXISTENT_VOL` for the volume query request, even though the volume is successfully created. This has been resolved. **Symptoms:** If a volume creation and volume query is done in quick successions via WSAPI. **Conditions of occurrence:** WSAPI client issues a POST /volumes to create a volume and then GET /volumes/<new volume name> in quick succession. **Impact:** Low **Customer circumvention:** WSAPI client can wait a bit after a volume creation before issuing the GET request. **Customer recovery steps:** None. The volume is actually created.

Issue IDs: 160385
Issue summary: ZLIB compression is enabled in WSAPI and is a known vulnerability in TLS1.x. **Affected platforms:** StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000 **Affected software versions:** 3.1.3, 3.2.1, 3.2.2 **Issue description:** HTTP usage of ZLIB compression in TLS 1.x must be disabled to prevent exposure to the CRIME (Compression Ratio Information-leak Made Easy) security vulnerability. **Symptoms:** TLS compression was enabled for WSAPI HTTPS connection, which could be vulnerable to CRIME, see CVE-2012-4929 TLS/CRIME. **Conditions of occurrence:** WSAPI client communicates with WSAPI server over HTTPS (port 5989) with TLS compression enabled.

Table Continued

Impact: Low

Customer circumvention: WSAPI client can disable HTTPS TLS compression on its end.

Customer recovery steps: None

Issue IDs: 189113

Issue summary: WSAPI returns an error when System Reporter records exceed limit.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.2

Issue description: When System Reporter returns a large number of records, the error code returned by WSAPI is not clear and clients would not know how to fix the issue.

Symptoms: WSAPI request will return Error code 329 when System Reporter query results in a large number of records.

Conditions of occurrence: It can mostly occur while using `groupby`, and there are large number of objects on the system but not limited to this condition.

Impact: Medium

Customer circumvention: Reduce the scope of the request, such that the number of records are reduced.

Customer recovery steps: Retry the operation after reducing the scope of the request.

HPE 3PAR OS 3.3.1 EGA Release Notes

Purpose

The HPE 3PAR OS 3.3.1 EGA provides several quality improvements.

Guidance

All HPE 3PAR StoreServ Storage Systems running 3PAR OS 3.3.1 are susceptible to the issues corrected in this patch.

Prerequisites

- SP prerequisite: SP-5.0.0.0-22913
- OS prerequisites: OS-3.3.1.215

Modifications

The following issues are addressed in this release:

Issue IDs:159516
Issue summary: Reduced I/O block times for consistent imports **Affected platforms:** StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000 **Affected software versions:**3.2.2 MU4, 3.3.1 GA **Issue description:** Reduces host I/O stall times near the end of a Peer Motion migration where consistency groups are being used. **Symptoms:** Host may see longer I/O stall times of about 1 to 2 minutes near the end of migration. **Conditions of occurrence:** Using consistency groups for migration with large number of volumes or large sized volumes. **Impact:** High, Medium **Customer circumvention:** Avoid using consistency groups for migration as a workaround. **Customer recovery steps:** None.

Issue IDs:165063

Issue summary: Online conversions, online copy, online promote, **`updatevv`**, and imports have long I/O stall times.

Affected platforms: StoreServ 20000

Affected software versions:3.2.2 GA, 3.2.2 MU4, 3.3.1 GA

Issue description: Online conversions, online copy, online promote, **`updatevv`**, and imports have long I/O stall times due to internal structure invalidation.

Symptoms: Host may experience longer than normal service times at the end of migration.

Conditions of occurrence: Starting Online Imports, peer-motion imports or **`updatevv`**.

Impact: High

Customer circumvention: Avoid online conversions, online copy, online promote, **`updatevv`**, and imports on StoreServ 20000 systems.

Customer recovery steps: Use standard recovery for host timeouts.

Issue IDs:188463

Issue summary: Single node will not boot after clean shutdown when 2nd node has a bad voltage regulator.

Affected platforms: StoreServ 7000

Affected software versions:3.2.1 MU3, 3.2.1 MU5, 3.2.2 MU4, 3.3.1 GA

Issue description: After properly shutting down the system, if a power regulator failure prevents a controller node from booting, the system will not boot because it is waiting for the missing controller node to boot.

Symptoms: On a 2 node system, after a proper shutdown, the array does not boot while waiting for the other controller node to join the cluster.

Conditions of occurrence: When a 2 node array is shutdown and simultaneously encounters a power regulator failure.

Impact: High

Customer circumvention: None

Customer recovery steps: None

Issue IDs:199218

Issue summary: Imports and **`updatevv`** have long host I/O stall times.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions:3.3.1 GA

Issue description: Imports or **`updatevv`** with a large list of VVs will have long I/O stall times.

Symptoms: Longer than normal host service times on VLUNS.

Conditions of occurrence: Start an import or **`updatevv`** with multiple list of VVs, a VVset or consistency group.

Impact: High

Customer circumvention: Avoid using imports or **`updatevv`** with a large list of VVs.

Customer recovery steps: Use standard recovery for host timeouts.

Issue IDs:200023

Issue summary: The **`showpatch -hist`** command output shows the **Id** as NA.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions:3.2.2 MU4, 3.3.1 GA

Issue description: The **`showpatch -hist`** command output shows the **Id** as NA

Symptoms: The **`showpatch -hist`** command output shows the **Id** as NA

Conditions of occurrence: Running the CLI command **`showpatch -hist`**

Impact: Low

Customer circumvention: None

Customer recovery steps: None

Issue IDs:200464

Issue summary: The command **`updatevv -removeandrecreate`** skips the addition of some of the VVs within a virtual volume set. The resultant VVs are missing from virtual volume set.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.2.1 GA, 3.2.1 MUx, 3.2.2 GA, 3.2.2 MUx, 3.3.1 GA

Issue description: **`updatevv -removeandrecreate`**, may skip A VV while adding it in Virtual Volume Set (VVSet).

Symptoms:**`updatevv -removeandrecreate`** all snapshots may not be added back to the VVSET.

Conditions of occurrence: Using **`updatevv -removeandrecreate`**

Impact: High

Customer circumvention: Do not user **`updatevv -removeandrecreate`**.

Customer recovery steps: Create the snapshot manually in the VVSet.

Issue IDs:205041

Issue summary: When retention is applied, a scheduled task to create a snapshot is marked failed even though snapshot creation and removal are successful.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1 GA

Issue description: When scheduled task of createfsnap is created with a retention period, the creation of the snapshot and removal of the old snapshot is successful from PML, but CLI intermittently indicates a failure in task details.

Symptoms: Even though the snapshot creation and reclamation is successful, the task indicates that the operation has not completed successfully.

Conditions of occurrence: When system is serving a heavy load and the customer executes numerous snapshot tasks.

Impact: Medium

Customer circumvention: None

Customer recovery steps: No recovery steps are required since creation and removal of snapshots are successful.

Issue IDs:206194

Issue summary: When compressed or compressed deduplicated volume grows over 4TB, the VV master controller node may restart unexpectedly.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1 GA

Issue description: Unexpected controller node restart that may result in unexpected array restart

Symptoms: Master controller node restarts unexpectedly, subsequent master controller node may also restart unexpectedly, triggering a full array restart.

Conditions of occurrence: Use of compressed or compressed deduplicated volume larger than 4TB in size.

Impact: High

Customer circumvention: Install Patch 02 or 3.3.1-EGA.

Customer recovery steps: None

Issue IDs:206441

Issue summary: Unexpected array restarts in response to meta-data inconsistencies.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1 GA

Issue description: After removing all Thinly Deduplicated Virtual Volumes (TDVV) within a CPG, and a controller node reboot or system manager restart, the next TDVV creation may result in LDs being reused.

Symptoms: The array or controller node may not successfully restart.

Conditions of occurrence: A new TDVV is created in a new CPG, after all TDVV are removed from an existing CPG and the array, a controller node or System Manager is restarted.

Impact: High

Customer circumvention: After removing all TDVVs within a CPG do not immediately reboot or shutdown the array.

Customer recovery steps: None

Issue IDs:206840

Issue summary: Array unexpectedly restarts during Remote Copy operation when a read is requested from a disk during disk firmware upgrade.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000

Affected software versions: 3.3.1 GA

Issue description: During an online upgrade to 3.3.1, HDD/SSD firmware is upgraded. It is possible for two HDD/SSD to be involved in the firmware upgrade process, one is in logging mode while other one is in log playback mode.

Symptoms: Customer applications may abort if array unexpectedly restarts as data is temporarily unavailable.

Conditions of occurrence: Online upgrade with Remote Copy active.

Impact: High

Customer circumvention: Perform the online upgrade to 3.3.1-EGA

Customer recovery steps: None.

HPE 3PAR OS 3.3.1 EGA combines all of the modifications and features provided by HPE 3PAR OS 3.3.1 Patch 01.

Refer to the release notes documents for each patch for a full list of modifications, features and supported drives. To learn more about each patch, use the links provided to access the individual patch release notes.

3PAR OS 3.3.1 Patch	Description	Obsoletes	Links to Documentation
Patch 01	P01 provides several quality improvements.	None	**HPE 3PAR OS 3.3.1 Patch 01 Release Notes**
Patch 02	P02 provides several quality improvements.	None	**HPE 3PAR OS 3.3.1 Patch 02 Release Notes**

Affected components

Component	Version
CLI Client	3.3.1.228
System Manager	3.3.1.228 (P02)
TOC Server	3.3.1.228 (P02)
TPD Kernel Patch	3.3.1.228 (P02)

Verification

The installation of EGA can be verified from an interactive CLI session. Issue the CLI command `showversion -a -b` to verify that EGA is listed:

```
cli% showversion -a -b
Release version 3.3.1.215
Patches:  P01,P02

Component Name                   Version
CLI Server                       3.3.1.223 (P02)
CLI Client                       3.3.1.223
System Manager                   3.3.1.223 (P02)
Kernel                           3.3.1.215
TPD Kernel Code                  3.3.1.223 (P02)
TPD Kernel Patch                 3.3.1.223 (P02)
CIM Server                       3.3.1.215
WSAPI Server                     3.3.1.215
Console Menu                     3.3.1.215
Event Manager                    3.3.1.215
Internal Test Tools              3.3.1.215
LD Check Tools                   3.3.1.215
Network Controller               3.3.1.215
Node Disk Scrubber               3.3.1.215
PD Scrubber                      3.3.1.215
Per-Node Server                  3.3.1.215
Persistent Repository            3.3.1.215
Powerfail Tools                  3.3.1.215
Preserved Data Tools             3.3.1.215
Process Monitor                  3.3.1.215
Software Updater                 3.3.1.215
TOC Server                       3.3.1.223 (P02)
VV Check Tools                   3.3.1.217 (P01)
Upgrade Check Scripts            170517.U640 (3.3.1.226)
File Persona                     1.3.0.74-20170309
SNMP Agent                       1.10.0
SSH                              6.0p1-4+deb7u5
VASA Provider                    3.0.12
Firmware Database                3.3.1.217 (P01)
Drive Firmware                   3.3.1.215
UEFI BIOS                        05.02.54
MCU Firmware (OKI)               4.8.60
MCU Firmware (STM)               5.3.17
Cage Firmware (DC1)              4.44
Cage Firmware (DC2)              2.64
Cage Firmware (DC3)              08
Cage Firmware (DC4)              2.64
Cage Firmware (DCN1)             4082
Cage Firmware (DCN2)             4082
Cage Firmware (DCS1)             4082
Cage Firmware (DCS2)             4082
Cage Firmware (DCS5)             2.78
Cage Firmware (DCS6)             2.78
Cage Firmware (DCS7)             4082
Cage Firmware (DCS8)             4082
QLogic QLA4052C HBA Firmware     03.00.01.77
QLogic QLE8242 CNA Firmware      04.15.27
QLogic 260x HBA FC Firmware      174.03.70
QLogic 27xx/268x HBA FC Firmware 174.03.70
```

```
QLogic 83xx HBA FCoE Firmware    08.01.05
QLogic 8300 HBA iSCSI Firmware   05.07.35
Emulex LP11002 HBA Firmware      02.82.x10
Emulex LPe12002 HBA Firmware     02.10.x02
Emulex LPe12004 HBA Firmware     02.10.x02
Emulex LPe16002 HBA Firmware     11.1.220.6
Emulex LPe16004 HBA Firmware     11.1.220.6
3PAR FC044X HBA Firmware         200A8
LSI 9201-16e HBA Firmware        17.11.03
LSI 9205-8e HBA Firmware         17.11.03
LSI 9300-8e HBA Firmware         10.00.08
```

Upgrade Considerations

The HPE 3PAR OS can be upgraded concurrently with I/O activity on the attached hosts, provided certain conditions are met. For more information on planning for online upgrades, refer to the latest version of the *HPE 3PAR Operating System Upgrade Pre-Planning Guide*. To obtain a copy of this documentation, go to the **Hewlett Packard Enterprise Information Library**.

OS upgrade prerequisite: Upgrade Tool version U008 or later must be staged prior to the HPE 3PAR OS upgrade to 3.3.1 MU1.

The Upgrade Tools are 3PAR OS upgrade enabling patches that do not affect array operation outside of the upgrade process. These tools are intended to improve the online or offline upgrade experience by performing preparatory steps to ensure the StoreServ is in a known state, including pre-checks, post-checks and other validations.

CAUTION:

Mandatory Patch Required for Use of File Persona with 3.3.1 MU1

In order to use File Persona with 3.3.1 MU1, install the mandatory 3.3.1 MU1 P07 patch after upgrading to 3.3.1 MU1. This patch contains important content to ensure stable operation of and compatibility for File Persona with MU1. If this patch is not installed:

1. Enabling file services for the first time will be prohibited. A message indicates that the patch needs to be installed.
2. Management requests may return unexpected results or fail unexpectedly. If File Persona has been enabled and the system has been upgraded to 3.3.1 MU1, do not attempt to modify the configuration of the system before installing the required patch.

Supported Platforms

For information regarding the supported HPE 3PAR StoreServ Storage systems, see the HPE Single Point of Connectivity Knowledge (SPOCK) website:

http://www.hpe.com/storage/spock

Notes

WARNING:

3PAR deduplication and compression are resource intensive operations, and as loads increase to these volumes, File Persona volume performance can decrease significantly. The load applied to volumes with these services enabled may need to be controlled in order to manage the impact to other volumes specifically volumes used by File Persona feature set as part of a File Provisioning Group.

What's New in the OS

New and enhanced features include:

3PAR OS 3.3.1 MU1

- IPv6 support for Peer Persistence Quorum Witness
- Support replication of compressed volumes using Remote Copy asynchronous streaming (RCAS) mode of replication on platforms that support both compression and RCAS.
- A Drive Health Assessment (DHA) utility that enables identification of certain drive models that are at risk of becoming degraded before they show visible symptoms is transferred to HPE as part of normal data collection. Drive models that utilize this enhancement are HCBF0600S5xeN010, HCBF1200S5xeN010, HCBF1200S5xeF010, HCBF1800S5xeN010
- Allows combining the use of custom Role Based Access Control (RBAC) roles with Virtual Domains. Users may now be assigned custom roles as well as standard RBAC roles in individual Virtual Domains
- Added support for the Brocade 40G-QSFP-4SFP-C-501 DAC, Cisco QSFP-4X10G-AOC5M Active Optic, and Arista QSFP+ 4x10G SFP+ 3m DAC cables
- Updates to enhance HPE 3PAR OS security

Modifications to the HPE 3PAR OS

The following issues have been addressed in this release.

Issue ID: 152596
Issue summary: Encrypted systems may report alerts at startup that an encrypted system is not encrypted. **Affected platforms:** StoreServ 8000, StoreServ 9000, StoreServ 20000, StoreServ 20000 R2 **Affected software versions:** 3.2.2 GA - MU4, 3.3.1 GA, 3.3.1 EGA **Issue description:** A timing issue at startup caused encrypted systems to report an alert that controller node drives were encrypted but that the system was not encrypted. This happened because the system had not yet determined its own encryption status. **Symptoms:** Alerts indicated that the controller node drives were encrypted but that the system was not encrypted. These alerts typically were resolved within a few seconds. However alert monitoring tools were being triggered. **Conditions of occurrence:** Any system that supports and has encryption enabled. **Impact:** Low **Customer circumvention:** None. The alerts are automatically cleared after a few seconds. **Customer recovery steps:** None

Issue ID: 179894
Issue summary: Enhanced Smart Trip for disk models beginning with HVIPC helps identify drive errors earlier, and request disk replacement by notifying users to replace disks reporting errors. **Affected platforms:** All StoreServ

Table Continued

Affected software versions: 3.3.1 GA, 3.3.1 EGA, and all previous versions

Issue description: Disks exhibiting certain types of correctable errors will not be identified early for replacement.

Symptoms: HVIPC disk models report unusually high numbers of correctable errors, leading to eventual disk replacement.

Conditions of occurrence: On StoreServ 10000 with HVIPC drives installed, higher than normal correctable errors may be observed, leading to eventual disk replacement.

Impact: Medium

Customer circumvention: Perform maintenance when disks require replacement.

Customer recovery steps: None

Issue ID: 184101

Issue summary: Occasionally peer motion volume migration from 3par array to 3par array does not complete.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 GA – 3.2.2 MU4

Issue description: Peer motion import would return error string `Name -srctpg is too long, should be less than 5 characters Error: bad rv argument`. This was due to a misinterpretation of a unusual mode page.

Symptoms: Peer motion migrations would fail.

Conditions of occurrence: Edge case in data handling, when certain internal fields were set by source array describing the volume to be migrated.

Impact: Low

Customer circumvention: Convert the volumes to fully provisioned before migration.

Customer recovery steps: Retry Migration.

Issue ID: 193352

Issue summary: High volume of fixed events in the event log.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 GA, 3.2.2 MU1, 3.2.2 MU2, 3.2.2 MU3, 3.2.2 MU4

Issue description: High volume of fixed events, even though there is no problem in the StoreServ.

Symptoms: High Volume of events

Table Continued

Conditions of occurrence: Every thirty minutes, message will be flooded.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue ID: 196169

Issue summary: A high volume of events due to the PD health check in every 60 minutes for non SAS controller nodes will generate error event logs.

Affected platforms: StoreServ 10000

Affected software versions: 3.1.2.GA - MU5, 3.1.3 GA - MU3, 3.2.1.GA - MU5, 3.2.2.GA - MU4, 3.3.1 GA and EGA

Issue description: In Peer Motion configurations, a high volume of events were being logged due to the periodic Physical Disk (PD) health check.

Symptoms: High volumes of events.

Conditions of occurrence: StoreServ 10000 with peer motion configured.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue ID: 196653

Issue summary: Corrects an upgrade issue where an array unexpectedly restarts and the nodes do not join the cluster due to multiple drive failures.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 MU4, 3.3.1 GA, 3.3.1 EGA

Issue description: SSD drives with the 100 RPM designation have a chunklet failure threshold which is exceeded due to differences in failed chunklet calculations between HPE 3PAR OS versions.

Symptoms: During an OS upgrade, the array will unexpectedly restart and the controller nodes will not rejoin the cluster.

Conditions of occurrence: An HPE 3PAR OS upgrade is performed and the 100 RPM SSD drives chunklet failures exceed the threshold.

Impact: High

Table Continued

Customer circumvention: 100 RPM SSD drives chunklet failures should be within the threshold prior to performing an OS upgrade.

Customer recovery steps: None

Issue ID: 199964

Issue summary: Remote Copy Async Streaming with fibre channel links over a low bandwidth FCIP network may intermittently stop and restart when many Remote Copy volumes in one or more groups undergo initial simultaneous synchronization.

Affected platforms: StoreServ 8000, StoreServ 9000 and StoreServ 20000

Affected software versions: 3.3.1 GA and previous versions

Issue description: Remote Copy Async Streaming or Remote Copy Periodic Async configurations with Remote Copy Fibre Channel (RCFC) links using Fibre Channel over IP (FCIP) with bandwidth less than 2Gbps may experience intermittent link restarts.

Symptoms: Remote Copy link restarts will be recorded in the event log. Time to synchronize the volumes may be extended.

Conditions of occurrence: This could occur during the initial synchronization of a large number of volumes simultaneously when RCFC link bandwidth is less than 2Gbps.

Impact: Medium

Customer circumvention: The following workarounds can be used to reduce the probability of this issue occurring.

1. For the short duration of the initial sync, provision high bandwidth for the links and reduce to the desired bandwidth after synchronization is complete.
2. Limit the number of volumes that synchronize concurrently based on the available bandwidth of the RCFC links.

Customer recovery steps: Restart the Remote Copy Group.

Issue ID: 200073

Issue summary: `sys:all_other` Quality of Service (QoS) rule overrides I/O throttling of virtual volumes even after moving it to a QoS defined vvset, until sysmgr is restarted

Affected platforms: All StoreServ

Affected software versions: 3.1.2 MU2, 3.1.3, 3.2.1

Issue description: A volume that is covered by the default QoS rule and then modified to be covered by a specific rule will be subject to both rules, instead of only the specific rule.

Symptoms: If the default rule has more strict limits than the specific rule, the volume will be subject to the more restrictive default.

Table Continued

Conditions of occurrence: A volume which is not part of a vvset with a QoS rule is subjected to the default rule. It then becomes part of a vvset with a QoS rule.

Impact: Medium

Customer circumvention: Disable the QoS default rule before creating a new volume, then add the specific QoS rule, and re-enable the QoS default rule.

Customer recovery steps: None

Issue ID: 200464

Issue summary: Corrects an issue where a Virtual Volume was not included in the vvset.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 GA to 3.2.2 MU4, 3.3.1 GA

Issue description: A Virtual Volume(s) is not added to the vvset.

Symptoms: After running the **`updatevv -removeandrecreate`** command on the vvset, a Virtual Volume(s) is missing from the output. Additionally, when the **`updatevv -removeandrecreate`** on and individual Virtual Volume(s) in vvset, it will not add the last Virtual Volume(s) in the vset.

Conditions of occurrence: This issue occurs if the vlunset is created from vvset and then vlunset is exported to hostset. The issue can be observed by running the command `updatevv -removeandrecreate` on vvset.

Impact: High

Customer circumvention: Do not create a VLUN set from vvset. Rather create an individual VLUNfor each Virtual Volume(s) in vvset and export individual VLUN to the host.

Customer recovery steps: Create new vvset.

Issue ID: 200537

Issue summary: Corrects an issue where peer volumes being replicated with Remote Copy and Peer Persistence may have the same Target Port Group ID (TPGID) assigned.

Affected platforms: All StoreServ

Affected software versions: 3.1.2 GA to 3.3.1 GA, 3.3.1 EGA

Issue description:

When a volume is dismissed and admitted to a new group after switchover, both the primary and secondary Remote Copy volume will have the same TPGID.

Symptoms: Volumes on both the primary and secondary side of the Remote Copy will be exported to the hosts, resulting in potential data unavailability.

Table Continued

Conditions of occurrence: A Virtual Volume (VV) is dismissed from a Peer Persistence configured Remote Copy Group and then added to a new group after a switchover.

Impact: High

Customer circumvention: Do not dismiss and readmit volumes to Remote Copy Groups after a switchover.

Customer recovery steps: None.

Issue ID: 201904

Issue summary: Improves defragmentation (defrag) for compression volumes.

Affected platforms: StoreServ 8000, StoreServ 9000, StoreServ 20000, StoreServ 20000 R2

Affected software versions: 3.3.1 MU1

Issue description:

Without defragmentation, compression volumes can become fragmented after a period of time.

For TPVV/TDVV, when the admck utility detects fragmentation, an auto defragment task will be triggered.

Symptoms: Fragmented space usage. More space is consumed than expected.

Conditions of occurrence: IO is fragmented for an extended period of time, or frequent write-same-zero operations are performed. Disk allocation is fragmented.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue IDs: 202473

Issue summary: Unexpected controller node restart due to a rare timing issue.

Affected platforms: StoreServ 8000, StoreServ 9000, StoreServ 20000, StoreServ 20000 R2

Affected software versions: 3.3.1 GA and EGA

Issue description: During normal cache management operations with compressed volumes, a rare timing event may lead to a double deallocation of a cache page.

Symptoms: Unexpected controller node restart.

Conditions of occurrence: Compressed volumes are running on the array.

Impact: High

Customer circumvention: None

Customer recovery steps: None

Issue ID: 202630

Issue summary: In the event of an unexpected controller node restart, diagnostic data may not be collected.

Affected platforms: StoreServ 8000

Affected software versions: 3.2.2 GA - MU4, 3.3.1 GA, 3.3.1 EGA

Issue description: Extraneous data was included in the diagnostic files, potentially causing them to be too large to fit in the allocated space resulting in an incomplete collection.

Symptoms: Diagnostic data collection following an unexpected controller node or array restart may be incomplete.

Conditions of occurrence: Unexpected controller node or array restart.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue ID: 204455

Issue summary: Host LUNS are not prevented from being exported on RCFC ports.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 EGA and all previous versions

Issue description: Host LUNS are not prevented from being exported on RCFC ports.

Symptoms: Inability to take snapshots on volumes exported on RCFC ports.

Conditions of occurrence: Host LUNS are exported on Remote Copy ports.

Impact: Medium

Customer circumvention: Do not have host visibility on Remote Copy Ports and do not export LUNS on these ports for host access.

Customer recovery steps: Remove the LUN exports currently defined on RCFC ports, offline the RCFC port, using the **`servicehost`** command to remove the lost host connection on that port, and restart the RCFC port.

Issue ID: 204706

Issue summary: A service alert indicating an internal error with the SQLite DB for System Reporter generated when first upgrading to software version 3.3.1.

Affected platforms: All StoreServ

Table Continued

Affected software versions: 3.3.1 GA and EGA

Issue description: When the System Reporter (SR) is upgraded circumstances on the array may allow a request to be issued to the new SR, before it is completely upgraded, resulting in the CLI Internal Error SQlite DB Mgs ID: 15001d being generated. The requests will succeed when retried after the SR upgrade process is complete.

Symptoms: After upgrading to 3.3.1 users may see the service alert: **`CLI Internal Error SQlite DB. . .`**

Conditions of occurrence: May occur after upgrade to 3.3.1 GA or EGA.

Impact: Low

Customer circumvention: The service alert **`CLI Internal Error SQlite DB...`** may be disregarded if observed when first upgrading to 3.3.1 GA .

Customer recovery steps: None

Issue ID: 205064

Issue summary: Adds support of the Brocade 40G-QSFP-4SFP-C-501 DAC cable.

Affected platforms: All StoreServ

Affected software versions: 3.1.3 GA - 3.1.3 MU3, 3.2.1 GA -3.2.1 MU5, 3.2.2 GA - 3.2.2 MU4, and 3.3.1 GA

Issue description: When the Brocade 40G-QSFP-4SFP-C-501 DAC cable is connected to a 10G port (iSCSI, FCoE, or NIC), the port indicates it is in a degraded state.

Symptoms: Degraded SFP message displays after running CLI command <cmd> showport -d -sfp</cmd>, and an Alert is generated.

Conditions of occurrence: Connection of Brocade 40G-QSFP-4SFP-C-501 DAC cable.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue ID: 205066

Issue summary: Adds support of the Cisco QSFP-4X10G-AOC5M Active Optic cable.

Affected platforms: All StoreServ

Affected software versions: 3.1.3GA through 3.1.3MU3, 3.2.1GA through 3.2.1MU5, 3.2.2GA through 3.2.2MU4, and 3.3.1GA

Table Continued

Issue description: When the Cisco QSFP-4X10G-AOC5M Active Optic cable is connected to a 10G port (iSCSI, FCoE, or NIC), the port indicates that it is in a degraded state.

Symptoms: Degraded SFP message displays after running CLI command **`showport -d -sfp`**, and an alert is generated.

Conditions of occurrence: Connection of Cisco QSFP-4X10G-AOC5M Active Optic cable.

Impact: Medium

Customer circumvention: Use the DAC cables recommended or supported by HPE.

Customer recovery steps: Replace the cable with the cable recommended or supported by HPE.

Issue ID: 205406

Issue summary: Remote Copy disaster recovery operartion did not complete, leaving the Remote Copy groups in an unexpected (inconsistent) state.

Affected platforms: All StoreServ

Affected software versions: 3.1.1 GA - 3.3.1 GA, 3.3.1 EGA

Issue description: During the Remote Copy disaster recovery operation, volume promotion will not complete if any region moves are in progress. This puts the Remote Copy groups in an unexpected state.

Symptoms: The CLI command **`showrcopy`** will indicate that the roles, in the group information, are not as expected. For example; one side of the RC configuration is the primary and the other side is primary-rev, or one side is in secondary and the other is secondary-rev.

Conditions of occurrence: Performing a Remote Copy disaster recovery operation while a region move is in progress.

Impact: Medium

Customer circumvention: Wait until all region moves are complete before performing Remote Copy disaster recovery.

Customer recovery steps: Use **`setcopygroup`** command with appropriate options to restore the Remote Copy groups to a normal state.

Issue ID: 206188

Issue summary: FC Multi-Queue feature was not enabled on 16GB FC HBA after an array update.

Affected platforms: StoreServ 8000, StoreServ 20000, StoreServ 20000 R2

Affected software versions: 3.3.1 GA and EGA

Issue description: 3PAR 3.3.1 OS upgrade from any version of 3.2.2 or 3.2.1 required additional node reboot after completion of OS upgrade before the Multi-Queue feature is enabled on the LPe16002 or LPe16004 16G FC ports.

Table Continued

Symptoms: 3PAR array performance may be less than expected.

Conditions of occurrence: Upgrading the HPE 3PAR OS from 3.2.2 or 3.2.1 to 3.3.1 GA or 3.3.1 EGA.

Impact: Medium

Customer circumvention: None

Customer recovery steps: Reboot each node once after the 3PAR 3.3.1GA OS upgrade is complete.

Issue ID: 207547

Issue summary: Remote Copy read failure results in unexpected controller node restart.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA, 3.3.1 EGA

Issue description: An internal timeout while reading a volume causes the Remote Copy ticket status to be in an invalid state leading to unexpected controller node reboots.

Symptoms: Remote Copy re-read timed out.

Conditions of occurrence: Any condition which can cause the Remote Copy read and re-read to fail. For instance, a multiple PD firmware upgrade where replication cannot read data from the disk within the timeout period.

Impact: High

Customer circumvention: Avoid situations which could potentially disrupt the Remote Copy read operations, like upgrading PD firmware without suspending Remote Copy groups.

Customer recovery steps: None.

Patches Included in This Release

HPE 3PAR OS 3.3.1 MU1 combines all of the modifications and features provided by HPE 3PAR OS 3.3.1 GA, EGA and the following patches.

NOTE:

To learn more about each patch, use the links provided to access the individual patch release notes.

Patch	Description	Obsoletes	Links to Documentation
HPE 3PAR OS 3.2.1 MU5 Patch 59	Provides support for drive FW updates and new drives.	OS-3.2.1.426-P55, OS-3.2.1.426-P58	**HPE 3PAR OS 3.2.1 MU5 Patch 59 Release Notes**

Table Continued

HPE 3PAR OS 3.2.1 MU5 Patch 71	Adds quality improvements including OS upgrade and node down recovery.	OS-3.2.1.426-P55	**HPE 3PAR OS 3.2.1 MU5 Patch 71 Release Notes**
HPE 3PAR OS 3.2.2 MU4 Patch 74	Patch 74 provides support for drive FW updates and new drives.	OS-3.2.2.612-P58, OS-3.2.2.612-P73	**HPE 3PAR OS 3.2.2 MU4 Patch 74 Release Notes**
HPE 3PAR OS 3.2.2 MU3 Patch 70	Patch 70 delivers several quality improvements.	OS-3.2.2.530-P47, OS-3.2.2.530-P55	**HPE 3PAR OS 3.2.2 MU3 Patch 70 Release Notes**
HPE 3PAR OS 3.2.2 MU4 Patch 80	Patch 80 provides several quality improvements.	OS-3.2.2.612-P76	**HPE 3PAR OS 3.2.2 MU4 Patch 80 Release Notes**
HPE 3PAR OS 3.2.2 MU4 Patch 84	Patch 84 provides several quality improvements.	OS-3.2.2.612-P76	**HPE 3PAR OS 3.2.2 MU4 Patch 84 Release Notes**
HPE 3PAR OS 3.3.1 GA/EGA Patch 04	Patch 04 provides improvements for slow disks and virtual volume management.	None	**HPE 3PAR OS 3.3.1 Patch 04 Release Notes**

Known Issues with the OS

Issue ID:181445

Issue summary: After an unexpected array restart, the normal consistency checks performed on Virtual Volumes may report as **not_started, needs_check**.

Affected platforms: All StoreServ

Affected software versions: 3.2.1 GA - 3.3.1 MU1

Issue description: Automatic**checkvv** at restart time corrects any metadata issues found, but does not start the VV. Manual intervention of running **checkvv** is required to have the volume start.

Symptoms: Virtual Volumes reporting status as `Not_started,needs_check`.

Conditions of occurrence: During the recovery from an unexpected array restart, the virtual volume **checkvv**.

Impact: Medium

Customer recovery steps: Manually run the **checkvv** command on the affected volumes.

Issue ID: 195256

Issue summary: Logical Unit Number (LUN) access lost due to excessive Offloaded Data Transfer token invalidations.

Affected platforms: All StoreServ

Affected software versions: All

Issue description: The 3PAR array is not cleaning up expired Offloaded Data Transfer (ODX) tokens in a timely manner, leaving open the possibility of getting flooded with token invalidation requests as writes come into the array hitting the same data area covered by previously populated ODX tokens. Excessive amounts of token invalidation requests require time to process, resulting in loss of access to a LUN.

Symptoms: LUN continuously returns back **`Busy`** as it tries to invalidate ODX tokens.

Conditions of occurrence: Heavy use of ODX across multiple LUNs.

Impact: Low

Customer circumvention: HPE support has developed a script that will periodically clean up expired ODX tokens. Contact HPE support about installing this script to avoid this problem.

Customer recovery steps: Access to the LUNs will be restored after the storm of token invalidation requests passes. Specific host actions may need to be taken to recover the LUN access on the host OS.

Issue ID:199872

Issue summary: An issue where a CPG with availability of magazine set is trying to grow using the **`-ha cage`** option.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 MU3, 3.3.2 MU4, 3.3.1 GA, 3.3.1 EGA, 3.3.1 MU1

Issue description: CPG with **`-ha mag`** option set trying to grow associated volumes with **`-ha cage`** and failing due to availability.

Symptoms: Error of `insufficient SA space` in CPG when trying to create a TPVV.

Alert with code `0x0270009` and type CPG growth failure will be seen when running showalert.

Conditions of occurrence: On a system with limited cage availability which has a CPG with **`-ha mag`** set may see this if trying to create a TPVV.

Impact: Low

Customer circumvention: Set **`setsize -saga as 3 '-ssz=3'`**.

Issue ID: 204959

Issue: If a system manager or controller node restart occurs, a previously halted controller node attempts to reboot and join the cluster.

Affected platforms: StoreServ 8000, Store Serv, 9000, StoreServ 20000, StoreServ 20000 R2

Affected software versions: All versions

Issue description: Normally, when a controller node goes down, it will be automatically reset once after 45 minutes to avoid unintentional controller node reboot issues. In the case that **`shutdownnode`** was used, this reset is disabled. However, if the System Manager is restarted or the master controller node is restarted (either due to an unexpected condition or manual action), the system disregards previous actions and starts a new 45 minute timer to reset any unbooted controller nodes.

Symptoms: Controller nodes that are intentionally halted are automatically restarted.

Conditions of occurrence: Controller nodes are halted or otherwise in a down state and the master controller node reboots or restarts, including shutdownnode of the master controller node, or the System Manager is restarted.

Impact: Low

Customer circumvention: If the master node was restarted or the System Manager restarted, anticipate that the system will attempt to reset any down controller nodes after 45 minutes even if the shutdown was intentional. Keep controller nodes powered off if they are intended to be kept down.

Customer recovery steps: Perform a controlled shut down of the controller node again and power it off until it is ready to be reintegrated into the cluster.

Issue ID: 211785

Issue summary: Virtual Volume allocation does not successfully complete if the setsize of the CPG is less than the number of healthy drives of that drive type present in the StoreServ.

Affected platforms: All StoreServ series

Affected software versions: 3.1.2, 3.1.3, 3.2.1, 3.2.2, 3.3.1

Issue description: Virtual Volume allocation does not successfully complete if the setsize of the CPG is less than the number of healthy drives of that drive type present in the StoreServ.

Symptoms: An alert is generated indicating that there is no space available to grow a volume in the CPG.

Conditions of occurrence: The setsize of the CPG is less than the number of healthy drives of that drive type present in the StoreServ.

Impact: Medium

Customer circumvention: Change the CPG set size and/or RAID.

Customer recovery steps: None

Issue ID: 213662
Issue summary: If the system contains only system volumes, and has cages with old firmware, the Service Processor or the **admithw** command might upgrade only a portion of the cages. **Affected platforms:** All StoreServ **Affected software versions:** 3.3.1 GA, 3.3.1 EGA, 3.3.1 MU1 **Issue description:** If the system contains only system volumes, and has cages with old firmware, the Service Processor or the **admithw** command might upgrade only a portion of the cages. This does not occur if there are customer volumes configured on the array. **Symptoms:** Alerts indicate `Interface Card Firmware Out of date`. The enclosure health shows `Degraded`The Service Processor reports `Cage not on current firmware` after it finishes the system upgrade. Check Health reports the same error. **Conditions of occurrence:** Cage firmware is not in the current state and **admithw** is performed. **Impact:** Low **Customer circumvention:** None **Customer recovery steps:** Re-run the action **Admit hardware** from the Service Processor until **checkhealth** reports no old cage firmware.

Modifications to File Persona

CAUTION:

A patch **must** be applied to the StoreServ array after upgrading to 3.3.1 MU1 before File Persona is used or modified. Do not perform file services related tasks or administrative operations until this patch is installed.

What's New in the CLI

New Commands

- `removecorequest`
- `setcorequest`
- `setfsaudit for File Access Auditing`
- `showcorequest`
- `showfsaudit File Access Auditing`

Changed Commands

Command	Description
`addsnmpmgr`	New `-notify` option
`createcert`	Add 4 syslog Services
`createfshare`	New `-audit` option
`importcert`	Add 4 syslog Services
`removecert`	Add 4 syslog Services
`removefsarchive`	subcommand auditlogs and `-fstore` now mandatory for archive operations, new `-importfile` option
`setfs`	New `nodeip` option, `-vlantag` is now optional
`setfsarchive`	`-fstore` now mandatory for admin operations, , new `-importfile`
`setfsav`	`KASPERSKY` now supported
`setfshare`	New `-audit` option
`setrcopygroup`	New `vvol` subcommand and `vvol -removetest`
`setsnmpmgr`	New `-notify` command
`setsys`	New parameter `ComplianceOfficerApproval`
`setuser`	New co role

Table Continued

Command	Description
`showcert`	Add 4 syslog Services
`showfsarchive`	subcommands `auditlogs` and `export`, new options `-importfile`, `-export`
`showrole`	new co role
`showwsapisession`	New `type` and `-filter`
`SR commands`	Add `percentile`, `per_group`, `per_time`, `only_compareby` to summary option

Modifications to the CLI

Issue ID: 163864
Issue summary: Enables additional commands in the audit user environment. **Affected platforms:** All StoreServ **Affected software versions:** 3.1.3 GA to 3.3.1 GA/EGA **Issue description:** This enhancement enables **`itables -L`** and **`netstat -avntp`** in the audit user environment. **Symptoms:** The **`itables -L`** and **`netstat -avntp`** were not supported in the audit user environment. **Conditions of occurrence:** Functionality was previously unsupported in the audit user environment. **Impact:** High **Customer circumvention:** None **Customer recovery steps:** None

Issue ID: 193846
Issue summary: Corrects a tuning issue where the **`tunesys`** process did not apply the **`-fulldiskpct`** or **`-chunkpct`** commands to the intra-node phase when active-active PDs are present. **Affected platforms:** All StoreServ **Affected software versions:** 3.2.2MU1 - 3.2.2 MU4 (SSD only), 3.3.1 GA, 3.3.1 EGA (all PD types) **Issue description:** A tuning issue was found with tunesys when custom values for **`-fulldiskpct`** or **`-chunkpct`** are supplied to control the chunklet movement phase and LD re-layout phases of the intra-node tuning respectively. In release 3.2.2.MU1 and later this only affects node-level re-balancing of SSDs. In release 3.3.1 this affected all disk types.

Table Continued

Symptoms: **`-fulldiskpct`** and **`-chunkpct`** are used to customize intra-node re-balancing. They are generally only used under direction from HPE support. When these options are used, expected tunes are not generated.

Conditions of occurrence: **`tunesys -fulldiskpct <value> -chunkpct <value> -`** does not generate expected intra-tunes.

Impact: Low

Customer circumvention: None

Customer recovery steps: Run manual intra-node tunes in consultation with HPE support.

Issue ID: 195084

Issue summary: Corrects an issue where the **`tunesys`** process terminated unexpectedly and generated the message **Error getting SD space from CPG**.

Affected platforms: All StoreServ

Affected software versions: 3.2.2.MU2+

Issue description: An incorrect calculation of the amount of space to allocate for the destination of a tune prevented the **`tunevv`** task from completing and generated the message **Error getting SD space from CPG**.

Symptoms: **`tunevv`** fails while migrating Virtual Volumes from one CPG to another with error **Error getting SD space from CPG**.

Conditions of occurrence: When running **`tunevv`** on Virtual Volume(s) with CPG params limiting to node pair without applied **`-nd param`**.

Impact: Low

Customer circumvention: None

Customer recovery steps: Use the **`setcpg`** command to set **`-p -nd <node(s)> param`** on affected cpg.

Issue ID: 196065

Issue summary: Corrects an issue where the **`tunesys`** process used an incorrect Virtual Volume(s) size.

Affected platforms: All StoreServ

Affected software versions: 3.1.1 and later

Issue description: Corrects an issue in the **`tunesys`** process where the total used size of the Virtual Volume(s) across all CPGs was used rather than only the space within the specified CPG.

Symptoms: Volumes were skipped by **`tunesys`** due to space issues when space was available.

Table Continued

Conditions of occurrence: Volume used space within the CPG less than the available space (but total size greater than the available space) and the tuning skipped.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Modifications to the 3PAR CIM API

Issue IDs: 181532

Issue summary: Enhance the StoreServ SNMP agent to generate unique notification traps for selected StoreServ alerts.

Affected platforms: All StoreServ

Affected software versions: all

Issue description: Prior to this change, the StoreServ's SNMP agent used a single notification message type to send all 3PAR alerts; all alerts shared the same SNMP trap OID.

With this enhancement, the customer may configure the 3PAR SNMP agent to generate notifications messages with unique OIDs for selected traps as defined by the 3PAR mib.

Symptoms: Customer software that depends upon the SNMP OID to identify the nature of a StoreServ trap will not work correctly.

Conditions of occurrence: The 3PAR SNMP Agent is used to process 3PAR system traps.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue IDs: 207552

Issue summary: cimserver sometimes deadlocks during patch installation causing event process to become unresponsive.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 GA - MU4

Issue description: cimserver sometimes deadlocks on exit during patch installation which caused delivery of alerts and events to other utilities, such as SSMC and WSAPI, to cease.

Symptoms: cimserver does not shutdown and restart, and does not process incoming reqests.

Alerts and events are not delivered to WSAPI and SSMC.

Conditions of occurrence: A patch is installed which restarts cimserver. For example, a patch that updates the cim api or the api libraries.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

What's New with the Web Services API Software

New and enhanced features include:

- Added `groupby` capability for all Versus Time and At Time System Reports.
- Added `compareby` capability for the following system reports: `cpgspacedata`, `volumespacedata`, `portstatistics`, `vlunstatistics`, and `physicaldiskstatistics`.
- Added max volume sizes as part of system query.
- Added iSCSI VLAN info as part of port query.

Modifications to the 3PAR Web Services API

Issue IDs: 209660
Issue summary: Get File services fails with internal server error when Active Directory is configured. **Affected platforms:** All StoreServ systems that support File Services **Affected software versions:** 3.3.1 GA and 3.3.1 EGA **Issue description:** WSAPI returns an `Internal Server Error` if it does not recognize the Active Directory status. **Symptoms:** If Active Directory is configured, GET on file services returns `Internal Server Error`. **Conditions of occurrence:** WSAPI client issues a GET /fileservices. **Impact:** Low **Customer circumvention:** None **Customer recovery steps:** None

Issue ID: 209785
Issue summary: WSAPI will return **Internal Server Error** if volume state was not recognized. **Affected platforms:** All StoreServ **Affected software versions:** 3.3.1 GA and 3.3.1 EGA **Issue description:** New properties were added to the Virtual Volume detailed state. WSAPI will return Internal Server Error when performing the get function on volumes. **Symptoms:** WSAPI will return Internal Server Error when performing the `GET` function on volumes.

Table Continued

Conditions of occurrence: Performing a GET on /v1/volumes and /v1/volumes/<vol_name> from WSPAI and any of the specified (/v1/volumes and /v1/volumes/<vol_name>) volumes is in one of the following states: **consistent, standby, sd_meta_inconsistent, sd_needs_fix or sd_meta_fixing**.

Impact: Low

Customer circumvention: None

Customer recovery steps: None

Components

Table 3: Components and Versions

Component	Version
Maintenance Update	3.3.1.269 (MU1)
Patches	P07
CLI Server	3.3.1.269 (MU1)
CLI Client	3.3.1.269
System Manager	3.3.1.269 (MU1)
Kernel	3.3.1.269 (MU1)
TPD Kernel Code	3.3.1.269 (MU1)
CIM Server	3.3.1.269 (MU1)
WSAPI Server	3.3.1.269 (MU1)
Console Menu	3.3.1.269 (MU1)
Event Manager	3.3.1.269 (MU1)
Internal Test Tools	3.3.1.269 (MU1)
LD Check Tools	3.3.1.269 (MU1)
Network Controller	3.3.1.269 (MU1)
Node Disk Scrubber	3.3.1.269 (MU1)
PD Scrubber	3.3.1.269 (MU1)
Per-Node Server	3.3.1.269 (MU1)
Persistent Repository	3.3.1.269 (MU1)
Powerfail Tools	3.3.1.269 (MU1)
Preserved Data Tools	3.3.1.269 (MU1)
Process Monitor	3.3.1.269 (MU1)
Software Updater	3.3.1.269 (MU1)
TOC Server	3.3.1.269 (MU1)

Table Continued

Component	Version
VV Check Tools	3.3.1.269 (MU1)
Upgrade Check Scripts	170811.U008
File Persona	1.4.0.78-20170713 (P07)
SNMP Agent	1.10.0
SSH	6.0p1-4+deb7u5
VASA Provider	3.0.14 (MU1)
Firmware Database	3.3.1.269 (MU1)
Drive Firmware	3.3.1.269 (MU1)
UEFI BIOS	05.02.54 (MU1)
MCU Firmware (OKI)	4.8.60 (MU1)
MCU Firmware (STM)	5.3.17 (MU1)
Cage Firmware (DC1)	4.44 (MU1)
Cage Firmware (DC2)	2.64 (MU1)
Cage Firmware (DC3)	08 (MU1)
Cage Firmware (DC4)	2.64 (MU1)
Cage Firmware (DCN1)	4082 (MU1)
Cage Firmware (DCN2)	4082 (MU1)
Cage Firmware (DCS1)	4082 (MU1)
Cage Firmware (DCS2)	4082 (MU1)
Cage Firmware (DCS5)	2.79 (MU1)
Cage Firmware (DCS6)	2.79 (MU1)
Cage Firmware (DCS7)	4082 (MU1)
Cage Firmware (DCS8)	4082 (MU1)
QLogic QLA4052C HBA Firmware	03.00.01.77 (MU1)
QLogic QLE8242 CNA Firmware	04.15.27

Table Continued

Component	Version
QLogic 260x HBA FC Firmware	174.03.70
QLogic 27xx/268x HBA FC Firmware	174.03.70
QLogic 83xx HBA FCoE Firmware	08.01.05
QLogic 8300 HBA iSCSI Firmware	05.07.35
Emulex LP11002 HBA Firmware	02.82.x10
Emulex LPe12002 HBA Firmware	02.10.x03
Emulex LPe12004 HBA Firmware	02.10.x03
Emulex LPe16002 HBA Firmware	11.1.220.9
Emulex LPe16004 HBA Firmware	11.1.220.9
3PAR FC044X HBA Firmware	200A8
LSI 9201-16e HBA Firmware	17.11.03
LSI 9205-8e HBA Firmware	17.11.03
LSI 9300-8e HBA Firmware	10.10.01

Drive Firmware

Drives and firmware versions

Model ID	Type	Capacity	V-Class	StoreServ 7000	StoreServ 8000	StoreServ 20000	StoreServ 9000	Firmware Version
HAKP2000S5xeN7.2	7.2K	2TB	Y	Y	Y	Y	N	3P03
HAKP4000S5xeN7.2	7.2K	4TB	Y	Y	Y	Y	N	3P03
HAKP6000S5xeN7.2	7.2K	6TB	Y	Y	Y	Y	N	3P03
HCBF0600S5xeN010	10K	600GB	Y	N	Y	Y	N	3P05
STHB0600S5xeN010	10K	600GB	Y	Y	Y	Y	N	3P02
HCBF1200S5xeF010	10K	1.2TB	N	Y	Y	Y	N	3P05
HCBF1200S5xeN010	10K	1.2TB	Y	Y	Y	Y	N	3P05
STHB1200S5xeF010	10K	1.2TB	N	Y	Y	Y	N	3P00
STHB1200S5xeN010	10K	1.2TB	Y	Y	Y	Y	N	3P02
HCBF1800S5xeN010	10K	1.8TB	Y	Y	Y	Y	N	3P05
STHB1800S5xeN010	10K	1.8TB	Y	Y	Y	Y	N	3P02
AREA0400S5xnNTRI	SSD	400GB	N	Y	Y	Y	Y	3P00

Table Continued

AREX0400S5xnNTRI	SSD	400GB	Y	Y	Y	Y	Y	3P02
DDYE0400S5xnNMRI	SSD	400GB	N	Y	Y	Y	Y	3P01
AREA0480S5xnNTRI	SSD	480GB	N	Y	N	N	N	3P00
DOPE0480S5xnNMRI	SSD	480GB	N	Y	Y	Y	N	3P08
AREA0920S5xnNTRI	SSD	920GB	N	Y	N	N	N	3P00
AREA1920S5xnNTRI	SSD	1.92TB	Y	Y	Y	Y	Y	3P00
AREX1920S5xnNTRI	SSD	1.92TB	Y	Y	Y	Y	Y	3P06
DDYE1920S5xnNMRI	SSD	1.92TB	N	Y	Y	Y	Y	3P01
DOPE1920S5xnNMRI	SSD	1.92TB	N	Y	Y	Y	N	3P08
AREA3840S5xnNTRI	SSD	3.84TB	N	Y	Y	Y	Y	3P00
AREX3840S5xnFTRI	SSD	3.84TB	N	Y	Y	Y	Y	3P02
AREX3840S5xnNTRI	SSD	3.84TB	N	Y	Y	Y	Y	3P06
DDYE3840S5xnNMRI	SSD	3.84TB	N	Y	Y	Y	Y	3P01

Table Continued

DOPM3840S5xnNMRI	SSD	3.84TB	N	Y	Y	Y	N	3P05
AREA7680S5xnFTRI	SSD	7.68TB	N	N	Y	Y	Y	3P01
AREA7680S5xnNTRI	SSD	7.68TB	N	N	Y	Y	Y	3P04
DDYM7680S5xnNMRI	SSD	7.68TB	N	N	Y	Y	Y	3P01
AREA15T4S5xnFTRI	SSD	15.3TB	N	N	Y	Y	Y	3P01
AREA15T4S5xnNTRI	SSD	15.3TB	N	N	Y	Y	Y	3P04

Websites

General websites

Hewlett Packard Enterprise Information Library

www.hpe.com/info/EIL

Single Point of Connectivity Knowledge (SPOCK) Storage compatibility matrix

www.hpe.com/storage/spock

Storage white papers and analyst reports

www.hpe.com/storage/whitepapers

For additional websites, see **Support and other resources**.

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:

  http://www.hpe.com/assistance

- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:

  http://www.hpe.com/support/hpesc

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates:

  Hewlett Packard Enterprise Support Center

  www.hpe.com/support/hpesc

  Hewlett Packard Enterprise Support Center: Software downloads

  www.hpe.com/support/downloads

  Software Depot

  www.hpe.com/support/softwaredepot

- To subscribe to eNewsletters and alerts:

  www.hpe.com/support/e-updates

- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:

www.hpe.com/support/AccessToSupportMaterials

IMPORTANT:

Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

http://www.hpe.com/support/selfrepair

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

Remote support and Proactive Care information

HPE Get Connected

www.hpe.com/services/getconnected

HPE Proactive Care services

www.hpe.com/services/proactivecare

HPE Proactive Care service: Supported products list

www.hpe.com/services/proactivecaresupportedproducts

HPE Proactive Care advanced service: Supported products list

www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central

www.hpe.com/services/proactivecarecentral

Proactive Care service activation

www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

To view the warranty for your product or to view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products* reference document, go to the Enterprise Safety and Compliance website:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional warranty information

HPE ProLiant and x86 Servers and Options

www.hpe.com/support/ProLiantServers-Warranties

HPE Enterprise Servers

www.hpe.com/support/EnterpriseServers-Warranties

HPE Storage Products

www.hpe.com/support/Storage-Warranties

HPE Networking Products

www.hpe.com/support/Networking-Warranties

Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

www.hpe.com/info/reach

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

www.hpe.com/info/ecodata

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

www.hpe.com/info/environment

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (**docsfeedback@hpe.com**). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.