Hewlett Packard Enterprise

HPE 3PAR OS 3.3.1 GA/EGA/MU1/MU2/MU3 Release Notes

Abstract

This document describes the features and issues included in HPE 3PAR OS 3.3.1 GA/EGA/MU1/MU2/MU3 and is intended for use by Hewlett Packard Enterprise customers, partners and field representatives.

Part Number: QL226-99935a
Published: October 2018
Edition: 2

Notices

Acknowledgments

Contents

HPE 3PAR OS 3.3.1 GA Release Notes.....6

Upgrade Considerations.....6
Supported Platforms.....6
Notes.....6
Components.....6
HPE 3PAR OS 3.3.1 GA Release Notes.....9
What's New in the OS.....9
Modifications to the HPE 3PAR OS.....13
Known Issues with the OS.....28
HPE 3PAR 3.3.1 File Persona GA Release Notes.....40
Modifications to File Persona.....40
Known Issues with File Persona.....42
HPE 3PAR 3.3.1 CLI GA Release Notes.....47
Installation Notes for the CLI.....47
Supported Operating Systems.....48
What's New in the CLI.....48
Modifications to the CLI.....52
HPE 3PAR 3.3.1 CIM API GA Release Notes.....58
What's New with the CIM API and SNMP Software.....58
Modifications to the 3PAR CIM API.....59
HPE 3PAR 3.3.1 WSAPI GA Release Notes.....60
What's New with the Web Services API Software.....60
Modifications to the 3PAR Web Services API.....61

HPE 3PAR OS 3.3.1 EGA Release Notes.....63

Online Upgrade Considerations.....63
Affected components.....63
Modifications.....63
Verification.....70

HPE 3PAR OS 3.3.1 MU1 Release Notes.....73

Upgrade Considerations.....73
Supported Platforms.....73
Notes.....73
HPE 3PAR OS 3.3.1 MU1 Release Notes.....73
What's New in the OS.....74
Modifications to the HPE 3PAR OS.....74
Known Issues with the OS.....85
Modifications to File Persona.....90
HPE 3PAR OS 3.3.1 CLI Release Notes.....90
What's New in the CLI.....90
Modifications to the CLI.....92
HPE 3PAR OS 3.3.1 MU1 CIM API Release Notes.....93
Modifications to the 3PAR CIM API.....94
HPE 3PAR WSAPI 3.3.1 MU1 Release Notes.....94
What's New with the Web Services API Software.....94
Modifications to the 3PAR Web Services API.....95

HPE 3PAR OS 3.3.1 EMU1 Release Notes..... 96
Upgrade Considerations.....96
Supported Platforms.....96
Components..... 97
Modifications to the OS..... 98

HPE 3PAR OS 3.3.1 MU2 Release Notes.....102
Update Considerations.....102
Supported Platforms.....102
Notes..... 102
HPE 3PAR 3.3.1 MU2 Release Notes.....102
What’s New in the VASA/VVol.....102
Modifications to the HPE 3PAR OS.....103
Known Issues with the OS.....114
HPE 3PAR 3.3.1 MU2 File Persona Release Notes.....118
What’s New in File Persona.....118
Modifications to File Persona.....118
HPE 3PAR 3.3.1 MU2 CLI Release Notes..... 127
What’s New in the CLI..... 128
HPE 3PAR 3.3.1 MU2 CIM API Release Notes..... 128
What's New in the CIM API..... 128
HPE 3PAR 3.3.1 MU2 Web Services API Release Notes.....128
What's New with the Web Services API Software 128

HPE 3PAR OS 3.3.1 MU3 Release Notes.....130
Upgrade Considerations.....130
Supported Platforms.....130
Notes..... 130
HPE 3PAR 3.3.1 Operating System MU3 Release Notes..... 130
What’s New in the OS.....130
Modifications to the HPE 3PAR OS..... 131
Known Issues with the OS.....185
HPE 3PAR 3.3.1 File Persona MU3 Release Notes.....193
What’s New in File Persona.....193
Modifications to File Persona.....194
Known Issues with File Persona.....214
HPE 3PAR 3.3.1 CLI MU3 Release Notes..... 226
Supported Operating Systems.....226
What’s New in the CLI..... 226
Modifications to the CLI..... 229
HPE 3PAR 3.3.1 CIM API MU3 Release Notes..... 238
What's New with the CIM API and SNMP Software 238
Modifications to the 3PAR CIM API..... 239
HPE 3PAR 3.3.1 Web Services API MU3 Release Notes.....243
What's New with the Web Services API Software 243
Modifications to the 3PAR Web Services API.....244
Component Versions 245
Drive Firmware..... 248

Support and other resources..... 252
Accessing Hewlett Packard Enterprise Support.....252

Accessing updates..... 252
Websites.....253
Customer self repair.....253
Remote support.....253
Documentation feedback.....253

Warranty and regulatory information.....254
Warranty information..... 254

HPE 3PAR OS 3.3.1 GA Release Notes

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 Planning Guide*. To obtain a copy of this documentation, go to the **Hewlett Packard Enterprise Information Library**.

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

Table Continued

Component	Version
Internal Test Tools	3.3.1.215
LD Check Tools	3.3.1.215
Network Controller	3.3.1.215
Controller 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
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

Table Continued

Component	Version
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

HPE 3PAR OS 3.3.1 GA Release Notes

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 controller 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 Adapters Supported on 3PAR 8000 Systems

- 16 Gb FC and 10 GbE NIC four-port combo HBA
- 10 Gb iSCSI and 10 GbE NIC four-port combo HBA

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

Table Continued

		3PAR StoreServ Platforms and Required DAC OS Support				
DAC Description	**DAC Part #**	**7000**	**10000**	**8000**	**9000 and 20000**	**Speed/ Protocols Supported**
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
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						

Table Continued

		3PAR StoreServ Platforms and Required DAC OS Support				
DAC Description	**DAC Part #**	**7000**	**10000**	**8000**	**9000 and 20000**	**Speed/ Protocols Supported**
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

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: 141238

Issue summary: Unexpected controller node restart due to a duplicate ID.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2

Issue description:

Internal system IDs may be reused when the same ID is already in use causing an unexpected controller node restart.

Symptoms: Controller nodes restart unexpectedly.

Conditions of occurrence: Normal array operations.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

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

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

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.

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.

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

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

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 ID: 165016

Issue summary: The host sees path loss and multipath events during a rolling upgrade.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1

Issue description: The host experiences a brief loss of path to 3PAR array during a rolling upgrade. The host plogi request gets dropped by the 3PAR array.

Symptoms: The host sees `rejecting I/O to offline device messages` in /var/log/messages.

Conditions of occurrence: Rolling upgrade.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: The lost paths are supposed to be automatically re-established by host a few seconds later.

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.

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: 179732

Issue summary: An unexpected controller node restart may occur when dirty cache pages are not cleared during snapshot removal.

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

Affected software versions: 3.2.2 GA-EMU4

Issue description: When a volume or snapshot is removed or offline, its dirty cache pages are not cleared. These pages then hold the CPU which may eventually cause the array to become unresponsive or an unexpected controller node restart.

Symptoms: Controller node restarts unexpectedly or the array becomes unresponsive.

Conditions of occurrence: Volumes are removed, closed or offline.

Impact: High

Customer circumvention: None

Customer recovery steps: None

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

Customer circumvention: None

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

Issue IDs: 180613

Issue summary: System Manager does not restart.

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

Affected software versions: 3.2.2.MU3, 3.2.2.MU4

Issue description: After an unexpected array restart the system manager does not restart automatically and the controller nodes do not integrate into the cluster.

Symptoms:

Table of contents (TOC) quorum not reached.

System Manager does not restart automatically and waiting for manual intervention.

Conditions of occurrence:

An unexpected array restart.

Massive burst of TOC updates resulting in out of memory space.

Impact: High

Customer circumvention: None

Customer recovery steps: None

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.

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 controller nodes.

Symptoms: The array will restart three times.

Conditions of occurrence:

1. A controller node goes down.

2. The array unexpectedly restarts while the controller 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

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

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: 191212, 215059

Issue summary: System manager restart occasionally may lead to unexpected termination of system manager.

Affected platforms: All StoreServ

Affected software versions:

3.2.2.MU2, 3.2.2.MU3

Issue description: When restarting the system manager on an array using persistent ports, the system manager may terminate unexpectedly.

Symptoms: After any operation that restarts the system manager, the system manager continues to restart unexpectedly.

Conditions of occurrence:

Array using persistent ports.

Any operation changing partner-ports' mode or failover/failback status followed by a restart of the system manager.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue IDs: 203495/201975

Issue summary: Defrag IO logs is not well handled in controller 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 controller 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 controller node restart or cluster down after a node down.

Conditions of occurrence: Controller 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:

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

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.

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.

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 controller 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

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.

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

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

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: 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.

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.

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.

Issue ID: 221709

Issue summary: A 16G Remote Copy (RCFC) link on an array running 3.3.1 GA/EGA or 3.3.1 MU1 and connected to an array running 3.2.2 (GA/EGA or any MUs) may not come up after a controller node with the link reboots. This can happen when an array is going through an online upgrade from 3.2.2 to 3.3.1 or after the array has been upgraded to 3.3.1. This issue is corrected in 3PAR OS 3.3.1 EMU1.

Affected platforms: All StoreServ

Affected software versions: 3.3.1.GA, 3.3.1.MU1, 3.3.1 EGA

Issue description: A 16G Remote Copy (RCFC) link on an array running 3.3.1 GA/EGA or 3.3.1 MU1 and connected to an array running 3.2.2 (GA/EGA or any MUs) may not come up after a controller node with the link reboots. This can happen when an array is going through an online upgrade from 3.2.2 to 3.3.1 or after the array has been upgraded to 3.3.1. This issue may also cause an online upgrade of an array from 3.2.2 to 3.3.1 GA/EGA/MU1 to fail because of the error "Target <target-name> does not have active remote copy links on multiple controller nodes."

Symptoms:

The Remote Copy link information from the CLI command showrcopy will show status "Down" for one or more RCFC links.

An online upgrade of an array from 3PAR OS 3.2.2 to 3.3.1 GA/EGA/MU1 may fail with the error "Target `<target-name>` does not have active remote copy links on multiple nodes" if the other array in the Remote Copy configuration is running 3PAR OS 3.2.2 (GA or any of the MUs).

Conditions of occurrence: The issue occurs if all of the following conditions are met.

Impact: High

Customer circumvention: When doing Online Upgrade with 16Gb RCFC config from 3PAR OS 3.2.2 to 3PAR OS 3.3.1GA/EGA/MU1 on multiple arrays in a Remote Copy configuration, apply the 3PAR OS upgrade to the array with highest system serial number first and then the next highest serial number etc. Note, this issue is fixed in 3PAR OS 3.3.1 EMU1, and 3PAR OS upgrades to 3PAR OS 3.3.1 EMU1 will not encounter this issue.

Customer recovery steps: When this issue occurs, the corresponding Remote Copy links on both arrays will be marked as "Down". To recover, reset the RCFC port with the higher WWN (which can be seen using the "showrctransport" CLI command. Resetting the port can be done using the "controlport rst" CLI command or its SSMC equivalent.

Issue ID:223358

Issue summary: Under certain conditions `sdmetack` may not get launched to check snapshots.

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

Affected software versions: 3.3.1.GA, 3.3.1.EGA, 3.3.1.MU1

Issue description: After a power fail event or a cluster outage event all volumes in an `sd_meta_corrupt` state need to run `sdmetack`. On rare occasions a race condition exists such that the list of volumes needed check is created before all the snapshots for compressed volumes come on line. This skips adding these snapshots to the list. When `sdmetack` kicks off these omitted snapshots will be missed.

Symptoms: Should `sdmetack` be required to run and completes; if there are snapshots left in the `sd_meta_corrupt` state you have hit this issue.

Conditions of occurrence: A power failure or other event where `sdmetack` needs to run.

Impact: Low

Customer circumvention: Other than not using compressed volumes, none.

Customer recovery steps: If the above symptom is observed manual running of `sdmetack` will be required.

HPE 3PAR 3.3.1 File Persona GA Release Notes

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.)
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.

Table Continued

Issue ID	Summary	Description
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.
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.

Table Continued

Issue ID	Summary	Description	Corrective Action
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.
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.

Table Continued

Issue ID	Summary	Description	Corrective Action
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.
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.

Table Continued

Issue ID	Summary	Description	Corrective Action
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.
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.

HPE 3PAR 3.3.1 CLI GA Release Notes

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
creategroupsv	New -addtoset, -match option
creategroupvvcopy	New -comprand -deupcompression options
createsched	importvv now allowed, command limit 1023 bytes
createsralertcrit	Additional space categories, New %_average condition comparisons; Added SYSSPACE type
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

Table Continued

Command	Description
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
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

Table Continued

Command	Description
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
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.

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 controller 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 controller 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 controller 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 controller node time and `hwclock` differ, then `checkhealth` should log a corresponding error.

Symptoms: There is a time difference between the controller 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 controller node time.

Impact: Low

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`

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

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)

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

HPE 3PAR 3.3.1 CIM API GA Release Notes

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 controller 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.

Impact: Medium

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

Customer recovery steps: None

HPE 3PAR 3.3.1 WSAPI GA Release Notes

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.

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

Online 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 Planning Guide*. For more information regarding the required order for upgrade and installation of software components, see the *HPE 3PAR OS 3.3.1 EGA Upgrade Instructions*. To obtain a copy of this documentation, go to the Hewlett Packard Enterprise Information Library.

⚠ **WARNING:** 3PAR Remote Copy asynchronous streaming configurations do not support compression. Do not use the asynchronous streaming replication mode with compressed volumes.

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.

Supported Platforms

This HPE 3PAR OS release supports HPE 3PAR StoreServ Storage. For more information, see the HPE Single Point of Connectivity Knowledge (SPOCK) website:

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

The minimum Service Processor version that supports HPE 3PAR OS 3.3.1 EGA is Service Processor (SP) 5.0.0.0 + latest SP patch.

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)

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 controller node will not boot after clean shutdown when second controller 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 two-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 two-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: Do not create compressed volumes over 4TB.

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 and Patch 02.

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**

Issue ID: 221709

Issue summary: A 16G Remote Copy (RCFC) link on an array running 3.3.1 GA/EGA or 3.3.1 MU1 and connected to an array running 3.2.2 (GA/EGA or any MUs) may not come up after a controller node with the link reboots. This can happen when an array is going through an online upgrade from 3.2.2 to 3.3.1 or after the array has been upgraded to 3.3.1. This issue is corrected in 3PAR OS 3.3.1 EMU1.

Affected platforms: All StoreServ

Affected software versions: 3.3.1.GA, 3.3.1.MU1, 3.3.1 EGA

Issue description: A 16G Remote Copy (RCFC) link on an array running 3.3.1 GA/EGA or 3.3.1 MU1 and connected to an array running 3.2.2 (GA/EGA or any MUs) may not come up after a controller node with the link reboots. This can happen when an array is going through an online upgrade from 3.2.2 to 3.3.1 or after the array has been upgraded to 3.3.1. This issue may also cause an online upgrade of an array from 3.2.2 to 3.3.1 GA/EGA/MU1 to fail because of the error "Target <target-name> does not have active remote copy links on multiple controller nodes."

Symptoms:

The Remote Copy link information from the CLI command showrcopy will show status "Down" for one or more RCFC links.

An online upgrade of an array from 3PAR OS 3.2.2 to 3.3.1 GA/EGA/MU1 may fail with the error "Target `<target-name>` does not have active remote copy links on multiple controller nodes" if the other array in the Remote Copy configuration is running 3PAR OS 3.2.2 (GA or any of the MUs).

Conditions of occurrence: The issue occurs if all of the following conditions are met.

Impact: High

Customer circumvention: When doing Online Upgrade with 16Gb RCFC config from 3PAR OS 3.2.2 to 3PAR OS 3.3.1GA/EGA/MU1 on multiple arrays in a Remote Copy configuration, apply the 3PAR OS upgrade to the array with highest system serial number first and then the next highest serial number etc. Note, this issue is fixed in 3PAR OS 3.3.1 EMU1, and 3PAR OS upgrades to 3PAR OS 3.3.1 EMU1 will not encounter this issue.

Customer recovery steps: When this issue occurs, the corresponding Remote Copy links on both arrays will be marked as "Down". To recover, reset the RCFC port with the higher WWN (which can be seen using the "showrctransport" CLI command. Resetting the port can be done using the "controlport rst" CLI command or its SSMC equivalent.

Issue ID:223358

Issue summary: Under certain conditions `sdmetack` may not get launched to check snapshots.

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

Affected software versions: 3.3.1.GA, 3.3.1.EGA, 3.3.1.MU1

Issue description: After a power fail event or a cluster outage event all volumes in an `sd_meta_corrupt` state need to run `sdmetack`. On rare occasions a race condition exists such that the list of volumes needed check is created before all the snapshots for compressed volumes come on line. This skips adding these snapshots to the list. When `sdmetack` kicks off these omitted snapshots will be missed.

Symptoms: Should `sdmetack` be required to run and completes; if there are snapshots left in the `sd_meta_corrupt` state you have hit this issue.

Conditions of occurrence: A power failure or other event where `sdmetack` needs to run.

Impact: Low

Customer circumvention: Other than not using compressed volumes, none.

Customer recovery steps: If the above symptom is observed manual running of `sdmetack` will be required.

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

HPE 3PAR OS 3.3.1 MU1 Release Notes

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 Planning Guide*. To obtain a copy of this documentation, go to the **Hewlett Packard Enterprise Information Library**.

OS upgrade prerequisite: The latest Upgrade Tool 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 Using 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 P19 patch if you have already upgraded 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.

IMPORTANT: When File Persona is enabled/configured, upgrade from 3.3.1 MU1 to 3.3.1 EMU1 is not supported if P07, P08, or P19 have been installed on 3.3.1 MU1. If File Persona has not been configured and is not in use, then upgrade is supported even with P07, P08 or P19 installed.

Customers who have configured File Persona and are running 3.3.1 MU1 + P07, P08 or P19 should continue to apply all recommended patches to 3.3.1 MU1, but must wait for a future HPE 3PAR OS version beyond 3.3.1 EMU1 to become available in order to upgrade.

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.

HPE 3PAR OS 3.3.1 MU1 Release Notes

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

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

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 controller 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

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.

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.

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 GA

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 ID: 202380

Issue summary: The array unexpectedly restarts when using compressed Read Only (RO) snapshots.

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

Affected software versions: 3.3.1 GA

Issue description:

The array unexpectedly restarts when multiple compressed RO snapshots exist and when compressed Read Only (RO) snapshots are removed.

Symptoms: The array or a single controller node unexpectedly restarts.

CLI commands become unresponsive.

Attempts to remove a VV are repeatedly unsuccessful.

Conditions of occurrence: Presence of compressed volumes with multiple read-only snapshots.

Impact: High

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

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

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, 3.3.1 EGA

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

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 controller 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.

Issue ID: 221709

Issue summary: 16G Remote Copy (RCFC) link(s) can become "down" after Online Upgrade from 3PAR OS 3.2.2 to 3PAR OS 3.3.1. This issue is corrected in 3PAR OS 3.3.1 EMU1.

Affected platforms: All StoreServ

Affected software versions: 3.2.2.GA and all MUs, 3.3.1.GA, 3.3.1.MU1

Issue description: When doing 3PAR OS Online Upgrade from any version of 3PAR OS 3.2.2 to 3.3.1GA/EGA or 3PAR OS 3.3.1.MU1 in a 16G RCFC configuration, then RCFC link(s) may unexpectedly become "down".

Symptoms: The Link information from the CLI command showrcopy will show status "down" for RCFC link(s).

Conditions of occurrence: Remote Copy configuration with 16Gb FC links. 3PAR OS Online Upgrade from 3PAR OS 3.2.2 and its MUs.

Impact: High

Customer circumvention: When doing Online Upgrade with 16Gb RCFC config from 3PAR OS 3.2.2 to 3PAR OS 3.3.1 always start (apply the 3PAR OS upgrade) to the array with highest system serial number.

Customer recovery steps: Reset the RCFC port(s) that have "down" status on the array with higher serial number. Note, this issue is fixed in 3PAR OS 3.3.1 EMU1, and 3PAR OS upgrades to 3PAR OS 3.3.1 EMU1 will not encounter this issue.

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**
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**

Table Continued

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 controller 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

214861

Issue summary: A virtual volume (VV) cannot grow and may become unavailable.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1 MU1

Issue description: A virtual volume (VV) cannot grow and may become unavailable if the set size (ssz) of the common provisioning group (CPG) is less than the number of drives of that drive type available in the CPG.

Symptoms: VVs within a CPG are unable to grow.

Conditions of occurrence: The set size of the CPG is equal to or greater than the number of drives of that drive type present in the CPG.

Impact: High

Customer circumvention: Consider the number of PDs that match the CPG specification (for example, -ha, -p -devtype). The maximum set size for the CPG must be no more than the number of available PDs, minus the number of PDs for fault tolerance, where the fault tolerance is determined by the RAID level.

Customer recovery steps: Configure the CPG so that the set size is less than the number of PDs in the CPG and minus the number of PDs required for the RAID level fault tolerance.

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.

Issue ID: 218553

Issue summary: The System Manager restarts unexpectedly during virtual volume conversions when compression garbage collector is running on that virtual volume.

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

Affected software versions: 3.3.1 GA, 3.3.1 MU1

Issue description: There is a race condition between the conversion and compression garbage collection. This collision can lead to the System Manager restart.

Symptoms: System Manager restarts unexpectedly.

Conditions of occurrence: Using **`tunevv, updatevv, importvv, promotevv, createvvcopy`** on a compressed volume.

Impact: Low

Customer circumvention: Avoid using the CLI commands **`tunevv, updatevv, importvv, promotevv, createvvcopy`** on a compressed volumes.

Customer recovery steps: None.

Issue ID: 221709

Issue summary: A 16G Remote Copy (RCFC) link on an array running 3.3.1 GA/EGA or 3.3.1 MU1 and connected to an array running 3.2.2 (GA/EGA or any MUs) may not come up after a controller node with the link reboots. This can happen when an array is going through an online upgrade from 3.2.2 to 3.3.1 or after the array has been upgraded to 3.3.1. This issue is corrected in 3PAR OS 3.3.1 EMU1.

Affected platforms: All StoreServ

Affected software versions: 3.3.1.GA, 3.3.1.MU1, 3.3.1 EGA

Issue description: A 16G Remote Copy (RCFC) link on an array running 3.3.1 GA/EGA or 3.3.1 MU1 and connected to an array running 3.2.2 (GA/EGA or any MUs) may not come up after a controller node with the link reboots. This can happen when an array is going through an online upgrade from 3.2.2 to 3.3.1 or after the array has been upgraded to 3.3.1. This issue may also cause an online upgrade of an array from 3.2.2 to 3.3.1 GA/EGA/MU1 to fail because of the error "Target <target-name> does not have active remote copy links on multiple controller nodes."

Symptoms:

The Remote Copy link information from the CLI command showrcopy will show status "Down" for one or more RCFC links.

An online upgrade of an array from 3PAR OS 3.2.2 to 3.3.1 GA/EGA/MU1 may fail with the error "Target `<target-name>` does not have active remote copy links on multiple controller nodes" if the other array in the Remote Copy configuration is running 3PAR OS 3.2.2 (GA or any of the MUs).

Conditions of occurrence: The issue occurs if all of the following conditions are met.

Impact: High

Customer circumvention: When doing Online Upgrade with 16Gb RCFC config from 3PAR OS 3.2.2 to 3PAR OS 3.3.1GA/EGA/MU1 on multiple arrays in a Remote Copy configuration, apply the 3PAR OS upgrade to the array with highest system serial number first and then the next highest serial number etc. Note, this issue is fixed in 3PAR OS 3.3.1 EMU1, and 3PAR OS upgrades to 3PAR OS 3.3.1 EMU1 will not encounter this issue.

Customer recovery steps: When this issue occurs, the corresponding Remote Copy links on both arrays will be marked as "Down". To recover, reset the RCFC port with the higher WWN (which can be seen using the "showrctransport" CLI command. Resetting the port can be done using the "controlport rst" CLI command or its SSMC equivalent.

Issue ID:223358

Issue summary: Under certain conditions `sdmetack` may not get launched to check snapshots.

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

Affected software versions: 3.3.1.GA, 3.3.1.EGA, 3.3.1.MU1

Issue description: After a power fail event or a cluster outage event all volumes in an `sd_meta_corrupt` state need to run `sdmetack`. On rare occasions a race condition exists such that the list of volumes needed check is created before all the snapshots for compressed volumes come on line. This skips adding these snapshots to the list. When `sdmetack` kicks off these omitted snapshots will be missed.

Symptoms: Should `sdmetack` be required to run and completes; if there are snapshots left in the `sd_meta_corrupt` state you have hit this issue.

Conditions of occurrence: A power failure or other event where `sdmetack` needs to run.

Impact: Low

Customer circumvention: Other than not using compressed volumes, none.

Customer recovery steps: If the above symptom is observed manual running of `sdmetack` will be required.

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.

HPE 3PAR OS 3.3.1 CLI Release Notes

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
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.

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.

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

HPE 3PAR OS 3.3.1 MU1 CIM API Release Notes

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 does not complete 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 does not complete 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

HPE 3PAR WSAPI 3.3.1 MU1 Release Notes

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.

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

HPE 3PAR OS 3.3.1 EMU1 Release Notes

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 Planning Guide*. To obtain a copy of this documentation, go to the **Hewlett Packard Enterprise Information Library**.

OS upgrade prerequisite:

The latest Upgrade Tool must be staged prior to the HPE 3PAR OS upgrade to 3.3.1 EMU1.

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, postchecks and other validations.

CAUTION: Mandatory Patch Required for Use of File Persona with HPE 3PAR OS 3.3.1 EMU1.

In order to use File Persona with 3.3.1 EMU1, install the mandatory HPE 3PAR OS 3.3.1 EMU1 P19 patch after upgrading to 3.3.1 EMU1. This patch contains important content to ensure stable operation of and compatibility for File Persona with HPE 3PAR OS 3.3.1 EMU1. 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 HPE 3PAR OS 3.3.1 EMU1, do not attempt to modify the configuration of the system before installing the required patch.

IMPORTANT: When File Persona is enabled/configured, upgrade from 3.3.1 MU1 to 3.3.1 EMU1 is unsupported if P07, P08, or P19 have been installed on 3.3.1 MU1. If File Persona has not been configured and is not in use, then upgrade is supported even with P07, P08 or P19 installed.

Customers who have configured File Persona and are running 3.3.1 MU1 + P07, P08 or P19 should continue to apply all recommended patches to 3.3.1 MU1, but must wait for a future HPE 3PAR OS version beyond 3.3.1 EMU1 to become available in order to upgrade.

CAUTION: It is highly recommended that the array has all available and applicable patches applied before beginning the upgrade to 3.3.1 EMU1.

Supported Platforms

For details of supported HPE 3PAR StoreServ Storage, see the Single Point of Connectivity Knowledge (SPOCK) website at **http://www.hpe.com/storage/spock**.

Components

Component	Version
CLI Server	3.3.1.269 (MU1)
CLI Client	3.3.1.269
System Manager	3.3.1.315 (P18)
Kernel	3.3.1.269 (MU1)
TPD Kernel Code	3.3.1.315 (P18)
TPD Kernel Patch	3.3.1.315 (P18)
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)
VV Check Tools	3.3.1.315 (P18)
Upgrade Check Scripts	171005.U008
File Persona	1.3.0.74-20170309 (MU1)
SNMP Agent	1.10.0
SSH	6.0p1-4+deb7u5
VASA Provider	3.0.14 (MU1)
Firmware Database	3.3.1.276 (P09)
Drive Firmware	3.3.1.276 (P09)
UEFI BIOS	05.02.54 (MU1)

Table Continued

Component	Version
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
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	11.1.220.10
Emulex LPe16004 HBA Firmware	11.1.220.10
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

Modifications to the OS

HPE 3PAR OS 3.3.1 EMU1 combines all of the modifications and features provided by HPE 3PAR OS 3.3.1 Patch 09, Patch 11 and Patch 18.

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 09	Patch 09 provides support for new second source drives and drive FW updates.	None	**HPE 3PAR OS 3.3.1 MU1 Patch 09 Release Notes**
Patch 11	Patch 11 improves SSMC connectivity when LDAP is used.	None	**HPE 3PAR OS 3.3.1 MU1 Patch 11 Release Notes**
Patch 18	Patch 18 adds quality improvements including OS upgrade and controller node down recovery.	Obsoletes P14 and P17	**HPE 3PAR OS 3.3.1 MU1 Patch 18 Release Notes**

3PAR OS 3.3.1 EMU1 also includes the following modifications:

Issue ID: 214315

Issue summary: In some environments, the gFC driver might deliver a false positive detection of IO resource shortage. This is resolved.

Affected platforms: All StoreServ

Affected software versions: 3.3.1.GA, 3.3.1.EGA, 3.3.1.MU1

Issue description: A piece of code was added to simulate IO resource shortage. The simulation code might get triggered, leading to a false positive resource shortage detection.

Symptoms: The target port types show as `free` from the CLI command **`showport`**.

Conditions of occurrence: Occurs with 16 Gb FC adapters.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 215674

Issue summary: 3PAR 16Gb array ports may auto-negotiate to switches at 8Gb instead of 16Gb.

Affected platforms: All StoreServ

Affected software versions: 3.3.1.GA, 3.3.1.MU1

Issue description: 16 GB HBA LPE16002/LPE16004 HBA adapters in the StoreServ may negotiate to 8GB if TTS (Transmitter Training Signal) via FEC (Forward Error Correction) is disabled on the switch port.

Symptoms: 16 GB HBA connecting at 8GB to the Fibre Channel Switch.

Conditions of occurrence:TTS is disabled on the switch port.

Impact: Medium

Customer circumvention:

Change all 16GB FC ports on the array to use TTS with the command below:

```
portcfgfec --enable -tts <port>
```

NOTE: This command will cause the port to reset and must be performed only if partner ports are healthy.

Customer recovery steps:

Change all 16GB FC ports on the array to use TTS with the command below:

```
portcfgfec --enable -tts <port>
```

NOTE: This command will cause the port to reset and must be performed only if partner ports are healthy.

Use the **`portcfgfec --show <port>`** command to confirm `16G FEC via TTS Configured:` states `ON` after the `--enable`.

For example:

```
brocade:admin> portcfgfec --show 12
Port: 12
FEC Capable: YES
10G/16G FEC Configured: ON
16G FEC via TTS Configured: OFF
FEC State: Active
```

Issue ID:223358

Issue summary: Under certain conditions `sdmetack` may not get launched to check snapshots.

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

Affected software versions: 3.3.1.GA, 3.3.1.EGA, 3.3.1.MU1

Issue description: After a power fail event or a cluster outage event all volumes in an `sd_meta_corrupt` state need to run `sdmetack`. On rare occasions a race condition exists such that the list of volumes needed check is created before all the snapshots for compressed volumes come on line. This skips adding these snapshots to the list. When `sdmetack` kicks off these omitted snapshots will be missed.

Symptoms: Should `sdmetack` be required to run and completes; if there are snapshots left in the `sd_meta_corrupt` state you have hit this issue.

Conditions of occurrence: A power failure or other event where `sdmetack` needs to run.

Impact: Low

Customer circumvention: Other than not using compressed volumes, none.

Customer recovery steps: If the above symptom is observed manual running of `sdmetack` will be required.

HPE 3PAR OS 3.3.1 MU2 Release Notes

Update Considerations

The HPE 3PAR OS can be updated concurrently with I/O activity on the attached hosts, provided certain conditions are met. For more information on planning for online updates, 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**.

NOTE: Supported upgrade paths may be found on the **HPE Single Point of Connectivity Knowledge (SPOCK)** website.

OS update prerequisite: The latest Upgrade Tool must be staged prior to the HPE 3PAR OS upgrade to 3.3.1 MU2. The Upgrade Tools are 3PAR OS update enabling patches that do not affect array operation outside of the update process. These tools are intended to improve the online or offline update experience by performing preparatory steps to ensure the StoreServ is in a known state, including pre-checks, post-checks and other validations.

IMPORTANT: If upgrading from an earlier version of 3.3.1 to 3.3.1 MU2, see the **HPE 3PAR OS and Service Processor Software Update Guide (HPE 3PAR OS 3.3.1 HPE 3PAR Service Processor 5.x)** for instructions on updating your specific software.

CAUTION: It is highly recommended that the array has all available and applicable patches applied before beginning the update to 3.3.1 MU2.

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

This product includes code licensed under the GNU General Public License, the GNU Lesser General Public License, and/or certain other open source licenses. A complete machine-readable copy of the source code corresponding to such code is available upon request. This offer is valid to anyone in receipt of this information and shall expire three years following the date of the final distribution of this product version by Hewlett Packard Enterprise Company. To obtain such source code, send a check or money order in the amount of US $10.00 to:

Hewlett Packard Enterprise Company

Attn: General Counsel

3000 Hanover Street

Palo Alto, CA 94304 USA

Please specify the product and version for which you are requesting source code.

HPE 3PAR 3.3.1 MU2 Release Notes

What's New in the VASA/VVol

New and enhanced features include:

- The VASA Provider no longer allows clients to connect using TLS/SSL methods other than TLSv1.2.
- In strict TLSv1.2 mode VASA/VVol supports the following cipher suites:

```
ECDHE-RSA-AES256-GCM-SHA384      ECDHE-RSA-AES256-SHA384
DH-DSS-AES256-GCM-SHA384         DHE-DSS-AES256-GCM-SHA384
DH-RSA-AES256-GCM-SHA384         DHE-RSA-AES256-SHA384
DHE-RSA-AES256-SHA256            DHE-DSS-AES256-SHA256
DH-RSA-AES256-SHA256             DH-DSS-AES256-SHA256
ECDH-RSA-AES256-GCM-SHA384       ECDH-RSA-AES256-SHA384
AES256-GCM-SHA384                AES256-SHA256
ECDHE-RSA-AES128-GCM-SHA256      ECDHE-RSA-AES128-SHA256
DH-DSS=AES128-GCM-SHA256         DHE-DSS-AES128-GCM-SHA256
DH-RSA-AES128-GCM-SHA256         DHE-RSA-AES128-GCM-SHA256
DHE-RSA-AES128-SHA256            DHE-DSS-AES128-SHA256
DH-RSA-AES128-SHA256             DH-DSS-AES128-SHA256
ECDH-RSA-AES128-GCM-SHA256       ECDH-RSA-AES128-SHA256
AES128-GCM-SHA256                AES128-SHA256
```

Modifications to the HPE 3PAR OS

The following issues have been addressed in this release.

Issue ID: 182665

Issue summary: Physical Disks (PD) may lose both paths nearly simultaneously.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 MU4, 3.2.2 MU6, 3.3.1 GA-MU1

Issue description: PD may lose both paths with a check condition of 06/29, which can lead to host I/O timeouts.

Symptoms: Both paths are lost on PDs.

Conditions of occurrence: Normal operations.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 187217

Issue summary: In rare situations, a single controller node may unexpectedly restart during internal region moves during controller node reintegration or **`tunevv`**.

Affected platforms: All StoreServ

Affected software versions: 3.1.2 MU3, 3.2.1, 3.2.2, 3.3.1 GA, 3.3.1 MU1

Issue description: During controller node integration into the cluster or when **`tunevv`** is running, a single controller node may unexpectedly restart if the if an inconsistency in the metadata for the tuned volume is encountered.

Symptoms: A controller node unexpectedly restarts while running **`tunevv`**.

Conditions of occurrence: The CLI command **`tunevv`** is running or a controller node is attempting to join the cluster.

Impact: Medium

Customer circumvention: Avoid running **`tunevv`**.

Customer recovery steps: None. The array recovers itself.

Issue ID: 190961

Issue summary: The array unexpectedly restarts when internal operations are performed on virtual volumes.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU1

Issue description: The array unexpectedly restarts when VV close operations, which occur, for example, when **`checkvv -offline`** is running, are being performed on VVs while the System Manager transfers mastership to another controller node.

Symptoms: The array unexpectedly restarts when a VV is transitioning to an offline state.

Conditions of occurrence: The controller node running the System Manager is either intentionally rebooted or unexpectedly restarts while processing VV close operations.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 193779

Issue summary: Unexpected controller node restarts occur on arrays with four or more controller nodes and SSD drives.

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

Affected software versions: 3.3.1 GA - MU1

Issue description: Corrects the situation where a high I/O load to SSD based volumes on HPE 3PAR StoreServ systems, with four or more controller nodes, may experience unexpected controller node restarts due to the formation of a multi-node deadlock within the array.

Symptoms: Unexpected controller node restart when SSD drives are heavily utilized.

Conditions of occurrence: On arrays with four or more controller nodes and SSD drive types with high I/O load to volumes using SSD drives.

Impact: Medium

Customer circumvention: Avoid unbalanced array configurations and overloading the array.

Customer recovery steps: None.

Issue ID: 196834

Issue summary: A controller node restart during snapshot creation can lead to metadata inconsistency for the virtual volume family.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1 GA, 3.3.1 MU1

Issue description: A controller node restart while creating a snapshot may result in metadata inconsistencies if the snapshot is not fully defined when the controller node restarts.

Symptoms: VV family going into a metadata inconsistent state.

Conditions of occurrence: Concurrence of the snapshot creation and controller node restart.

Impact: Medium

Customer circumvention: Do not restart controller nodes during snapshot creation.

Customer recovery steps: Recover the VV by running **`checkvv`**.

Issue ID: 197461

Issue summary: An unexpected array restart occurs when converting TDVV2 to TDVV3.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1 GA, 3.3.1 MU1

Issue description: An unexpected array restart occurs when converting TDVV2 to TDVV3. This can cause long block times on large virtual volumes.

Symptoms: An unexpected array restart happens during VV conversions or online copy.

Conditions of occurrence: Use online conversions or online copy.

Impact: High

Customer circumvention: Refrain from using the online conversion or online copy features.

Customer recovery steps: None. The array restarts automatically.

Issue ID: 201081

Issue summary: Unexpected controller node restart when using compression.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA, 3.3.1 MU1

Issue description: When using compressed volumes, memory management refers to an incorrect cache page which in turn causes an unexpected controller node restart.

Symptoms: Single controller node restarts unexpectedly.

Conditions of occurrence: Compressed virtual volume.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 202908

Issue summary: A false thermal event may cause an unnecessary shutdown of an array.

Affected platforms: StoreServ 7000, StoreServ 8000

Affected software versions: 3.2.1, 3.2.2, 3.3.1

Issue description: A false thermal event may stimulate customers to proactively shutdown an array unnecessarily.

Symptoms: An `Cluster thermal shutdown` alert is observed.

Conditions of occurrence: Normal operation.

Impact: High

Customer circumvention: Ignore thermal event log messages with the signature, `Status change Critical Cluster thermal shutdown hw_node: x Node y, due to high temperature conditions, the storage system is being shutdown.`

Customer recovery steps: None. This is a false alert.

Issue ID: 204754

Issue summary: A single controller node restarts unexpectedly.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU1

Issue description: A single controller node unexpectedly restarts due to code concurrency. This may be exacerbated by large Remote Copy configurations.

Symptoms: Controller node unexpectedly restarts.

Conditions of occurrence: Normal operation.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 206128

Issue summary: Unsupported Peer Motion zoning leads to an unmanageable array.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1

Issue description: If Peer Motion zoning between source and destination arrays does not follow the recommended 1:1 zoning, and/or the number of peer paths between the arrays exceeds 2, the System Manager becomes nonfunctional.

Symptoms: The array becomes unmanageable.

Conditions of occurrence: Peer Motion zoning between source and destination arrays does not follow the recommended zoning.

Impact: Medium

Customer circumvention: Follow the recommended 1:1 Peer Motion zoning.

Customer recovery steps: None.

Issue ID: 208018

Issue summary: Applying the incorrect license changes the existing W19/WWNBASE ID.

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

Affected software versions: 3.2.1, 3.2.2, 3.2.2 MU6, 3.3.1 MU1

Issue description: Applying the incorrect license changes the existing W19/WWNBASE ID. This causes the WWPNs of the HBAs to change and a controller node, if rebooted, no longer joins the cluster.

Symptoms: System W19/WWNBASE ID gets changed.

WWN base of the **`rcopy`** ports and Host port WWNs will change on reset or on reconfiguration.

If a controller node is rebooted or replaced, the W19 serial number will prevent it from joining the currently running cluster.

Conditions of occurrence: Applying an incorrect license with a mismatched System W19/WWNBASE ID or Serial Number.

Impact: High

Customer circumvention: Validate the license key matches the W19 (system ID) before installing a new license.

Customer recovery steps: None.

Issue ID: 211084

Issue summary: Controller nodes restart unexpectedly upon modifying switch configuration/zoning.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description: Controller nodes restart unexpectedly with the message `Fatal exception` when switch configuration/zoning is modified.

Symptoms: Controller nodes restart unexpectedly when zoning or configuration changes are invoked.

Conditions of occurrence: Switch port ID is changed due to either switch configuration or switch port zoning.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 214240

Issue summary: The secondary array in an Asynchronous RC configuration becomes unresponsive.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description: The primary array periodically attempts to take coordinated snapshots (CSS) for asynchronous Remote Copy groups on the secondary array.

While the secondary array is busy deleting a backlog of snapshots, it is unable to service the primary request and the request times out.

In response the primary retries the CSS on 15 second intervals, adding snapshot requests to the growing work queue of the secondary array.

When the secondary array is eventually able to process the snapshot backlog, all the waiting requests are completed in rapid succession. This results in a many snapshots being created in a few minutes.

Symptoms: The secondary array may be unresponsive or slow to respond to management commands.

Many snapshots in `removing` or `removing_retry` state.

Conditions of occurrence: Remote copy asynchronous replication with TDVV volumes.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 214448

Issue summary: Degraded performance of snapshot removal when multiple snapshots are present or removed.

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

Affected software versions: 3.2.2, 3.3.1 GA, 3.3.1 MU1

Issue description: When snapshots are removed, they will no longer scan the VV family for additional snapshots that may need to be removed.

Symptoms: Slow snapshot removal.

Conditions of occurrence: Removing large number of snapshots within same deduplication family.

Impact: Medium

Customer circumvention: Remove a smaller number of snapshots and allow snapshot removal to complete before initiating further snapshot removal.

Customer recovery steps: None.

Issue ID: 215793

Issue summary: **`mkvg`** commands do not complete on a volume greater than 2TB in size.

Affected platforms: All StoreServ

Affected software versions: All

Issue description: **`mkvg`** commands do not complete on a VV greater than 2TB in size on AIX versions 6.1/7.1.

Symptoms: **`AIX mkvg`** commands do not complete.

Conditions of occurrence: Presenting a VV larger than 2 TB to AIX hosts running 6.1/7.1 with VIOS configured to create virtual SCSI disks.

Impact: High

Customer circumvention: Use VVs less than 2TB in size on AIX.

Customer recovery steps: None.

Issue ID: 218032

Issue summary: Host temporarily loses access to VV imported using Peer Motion with TDVV.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1

Issue description: Peer Motion to deduplication provisioned volumes might cause host clusters to temporarily lose access to the volumes. This occurs in rare cases when deduplication garbage collection happens to run towards the end of migration.

Symptoms: Host clusters lose access to the volumes being migrated.

Conditions of occurrence: Importing TDVV volumes using Peer Motion.

Impact: Low

Customer circumvention: Import to non-deduplication volumes.

Customer recovery steps: None.

Issue ID: 219819

Issue summary: The **`dryrun`** option for the compression estimator does not complete successfully.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU1

Issue description: The **`dryrun`** option for the compression estimator does not complete successfully.

Symptoms: The **`dryrun`** compression estimator task does not complete.

Conditions of occurrence: Using the compression estimator **`dryrun`** option.

Impact: Low

Customer circumvention: Limit the number of virtual volumes (less than 30) when using the compression estimator **`dryrun`** option.

Customer recovery steps: Rerun the estimator with a limited number of virtual volumes.

Issue ID: 219998

Issue summary: 3PARInfo tool does not show VV name of exported volume.

Affected platforms: All StoreServ

Affected software versions: All

Issue description: 3PARInfo tool may not display the VV name of exported volumes when there is an issue collecting the VV information on the array.

Symptoms: VV name is not populated in 3PARInfo tool data.

Conditions of occurrence: Normal operation.

Impact: Low

Customer circumvention: None

Customer recovery steps: Reissue 3PARInfo requests.

Issue ID: 221514

Issue summary: Unexpected controller node restarts occur when using compressed volumes.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA, 3.3.1 MU1

Issue description: Unexpected controller node restarts occur when using compressed volumes due to a deadlock condition on the array.

Symptoms: Single controller node restarts.

Conditions of occurrence: Using compressed volumes.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 221985

Issue summary: Controller nodes unexpectedly restart while attempting to integrate into the cluster.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 9000

Affected software versions: 3.2.2 GA-MU6, 3.3.1 GA-MU1

Issue description: Controller nodes unexpectedly restart while attempting to integrate into the cluster which can lead to an unexpected restart of the entire array.

Symptoms: Controller nodes restart or the array unexpectedly restarts.

Conditions of occurrence: A planned or unplanned controller node restart.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 222974

Issue summary: In a rare condition, host IO may stall if an error condition is present on a SAS cage.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 9000, StoreServ 20000, StoreServ 20000 R2

Affected software versions: 3.3.1

Issue description: In a rare condition, host IO may stall if an error condition is present on a SAS cage while trying to collect diagnostic information.

Symptoms: Host I/O stalls.

Conditions of occurrence: Normal operation.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 224727

Issue summary: When removing a volume or snapshot the System Manager may become unresponsive.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA, 3.3.1 MU1

Issue description: When removing a volume or snapshot the System Manager may become unresponsive. In this state, array management may become unresponsive.

Symptoms: Array management becomes unresponsive.

Conditions of occurrence: Most likely to occur while removing volumes from arrays with large cache sizes.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 227824

Issue summary: Synchronous mode Remote Copy groups will not start if the volumes were created using Peer Motion or the Online Import Utility.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU1

Issue description: Synchronous mode Remote Copy groups will not start if the volumes were created using Peer Motion or the Online Import Utility.

Symptoms: Synchronous RC groups do not start.

Conditions of occurrence: Starting RC groups containing imported volumes.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 228606

Issue summary: SSD speed mismatch message appears while running **`tunesys`**.

Platforms affected: All StoreServ

Affected software versions: 3.3.1 GA-MU1

Issue description: SSD speed mismatch message, `Mismatch CPGminspeed = 100, LDminspeed = 150`, appears in the task log while running **`tunesys`**.

Symptoms: **`tunesys`** task produces the message `Mismatch CPGminspeed = 100, LDminspeed = 150`.

Conditions of occurrence: CPG contains both SSD 100 and SSD 150 drives.

Impact: Medium

Customer circumvention: Use the **`tunesys -nold`** option for tunes.

Customer recovery steps: None.

Issue ID: 229075

Issue summary: A compressed virtual volume experiences SD metadata inconsistencies.

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

Affected software versions: 3.3.1 GA-MU1

Issue description: A compressed virtual volume may experience SD metadata inconsistencies while I/O is occurring to the volume and compression garbage collection is also running.

Symptoms: The CLI command **`showvv`** on a compressed VV shows `sd_metadata_inconsistent`.

Conditions of occurrence: Use of compressed volumes.

Impact: High

Customer circumvention: None.

Customer recovery steps: The `sd_meta_inconsistent` state must be cleared by running **`checkvv -fixsd`**.

Issue ID: 230334

Issue summary: Controller nodes become unresponsive while removing or updating volumes.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description: Controller nodes become unresponsive while removing or updating volumes which leads to the cluster manager removing the controller node from the cluster.

Symptoms: Controller nodes or the entire array unexpectedly restart.

Conditions of occurrence: A large amount of I/O occurring on the same virtual volume (VV) family and a removal or update command is run on that VV family.

Impact: High

Customer circumvention: Reduce the I/O load before performing removals or updates of VVs within the VV family.

Customer recovery steps: None.

Issue ID: 232878

Issue summary: **`setvvolsc -remove`** does not complete successfully when the Remote Copy licenses are not installed.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU1

Issue description: When the VVol Storage Container is not empty, and being removed using the **`setvvolsc -remove`** command, the command will not complete successfully if the Remote Copy license is not available.

Symptoms: Removal of a VVol Storage Container using **`setvvolsc -remove`**, displays the error message `This system is not licensed for Remote Copy`.

Conditions of occurrence: The storage container is not empty (has existing VVols) when **`setvvolsc -remove`** is attempted.

Remote Copy is not licensed on the array.

Impact: High

Customer circumvention: Use vSphere to remove all VVol-based VMs from the VVol storage container data store before using **`setvvolsc -remove`**.

Customer recovery steps: Use vSphere to remove all VVol-based VMs from the VVol storage container data store, before using `setvvolsc -remove`.

Patches Included in This Release

HPE 3PAR OS 3.3.1 MU2 combines all of the modifications and features provided by HPE 3PAR OS 3.3.1 GA, EGA, MU1, EMU1, plus 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 3.3.1 MU1 P15	Provides support for drive firmware updates.	OS-3.3.1.269-P09	**https://support.hpe.com/hpsc/doc/public/display?docId=a00027067en_us**
HPE 3PAR OS 3.3.1 MU1 Patch 19	Required patch to support File Persona version 1.4.2 with 3.3.1 MU1.	OS-3.3.1.269-P08	**https://support.hpe.com/hpsc/doc/public/display?docId=a00026783en_us**
HPE 3PAR OS 3.3.1 Patch 21	Quality improvements to SD metadata, compression, deduplication and others.	OS-3.3.1.269-P18	**https://support.hpe.com/hpsc/doc/public/display?docId=a00040475en_us**
HPE 3PAR OS 3.3.1 MU1 Patch 24	Corrects an issue with alert processing via the SP.	None.	**https://support.hpe.com/hpsc/doc/public/display?docId=a00040750en_us**
HPE 3PAR OS 3.3.1 Patch 25	Provides several critical quality improvements.	OS-3.3.1.269-P24	**https://support.hpe.com/hpsc/doc/public/display?docId=a00043628en_us**

Known Issues with the OS

Issue ID: 185740

Issue summary: During recovery from an unexpected array restart, controller nodes will go through additional recovery sequence and virtual volumes may remain unstarted.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU1

Issue description: While unsuccessfully attempting to read metadata for an IO in progress, if the array experiences a power failure or unexpected restart, the array goes through a recovery operation. After the recovery operation, the VVs that unsuccessfully attempted the read operation on the metadata will remain in the `not_started,internal_consistency_error` state.

Symptoms: VVols remain in `not_started,internal_consistency_error` state.

Conditions of occurrence: Unexpected array restart occurs when there is an unsuccessful metadata read operation in progress.

Impact: High

Customer circumvention: None.

Customer recovery steps: Run **`checkvv -offline -y <vvname>`**.

Issue ID: 209003

Issue summary: A common provisioning group (CPG) has space, but virtual volume growth is unsuccessful.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU1

Issue description: During controller node integration, a single virtual volume performs slowly. Host applications may time out.

Symptoms: A virtual volume does not grow, and remains in this state.

Conditions of occurrence: Thinly provisioned volume is expanding within its virtual space.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 211785

214861

Issue summary: A virtual volume (VV) cannot grow and may become unavailable.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1 MU1

Issue description: A virtual volume (VV) cannot grow and may become unavailable if the set size (ssz) of the common provisioning group (CPG) is less than the number of drives of that drive type available in the CPG.

Symptoms: VVs within a CPG are unable to grow.

Conditions of occurrence: The set size of the CPG is equal to or greater than the number of drives of that drive type present in the CPG.

Impact: High

Customer circumvention: Consider the number of PDs that match the CPG specification (for example, -ha, -p -devtype). The maximum set size for the CPG must be no more than the number of available PDs, minus the number of PDs for fault tolerance, where the fault tolerance is determined by the RAID level.

Customer recovery steps: Configure the CPG so that the set size is less than the number of PDs in the CPG and minus the number of PDs required for the RAID level fault tolerance.

Issue ID: 218553

Issue summary: The System Manager restarts unexpectedly during virtual volume conversions when compression garbage collector is running on that virtual volume.

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

Affected software versions: 3.3.1 GA, 3.3.1 MU1

Issue description: There is a race condition between the conversion and compression garbage collection. This collision can lead to the System Manager restart.

Symptoms: System Manager restarts unexpectedly.

Conditions of occurrence: Using **`tunevv, updatevv, importvv, promotevv, createvvcopy`** on a compressed volume.

Impact: Low

Customer circumvention: Avoid using the CLI commands **`tunevv, updatevv, importvv, promotevv, createvvcopy`** on a compressed volume.

Customer recovery steps: None.

Issue ID: 219941

Issue summary: Running **`updatevv`** may result in the volume going offline at the host.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 MU2 - MU6, 3.3.1 GA - MU2

Issue description: Running **`updatevv`** results in a volume going offline temporarily. This event may exceed the host's timeout and retry settings, causing the volume to go offline at the host.

Symptoms: Volume is temporarily unavailable to the host.

Conditions of occurrence: Running **`updatevv`** without the **`-removeandrecreate`** option.

Impact: High

Customer circumvention: Use the **`-removeandrecreate`** option with **`updatevv`**.

Customer recovery steps: None.

Issue ID: 225658

Issue summary: Lightweight Directory Access Protocol (LDAP) may disconnect if authorization parameters or a user name is incorrectly supplied.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 MU2, 3.2.2 MU4, 3.3.1 GA, 3.3.1 MU1, 3.3.1 MU2

Issue description: Lightweight Directory Access Protocol may disconnect if authorization parameters or a user name is incorrectly supplied which requires the user to login again.

Symptoms: User is disconnected and must login again.

Conditions of occurrence: Attempting to connect to the array using LDAP.

Impact: Medium

Customer circumvention: Define the LDAP authorization parameter Kerberos-realm and ensure that the user name does not start with the "\" character.

Customer recovery steps: Redefine authorization parameters to include the Kerberos-realm.

Issue ID: 228712

Issue summary: During controller node up processing, host I/O may stall on a single virtual volume.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA-MU1

Issue description: During controller node up processing, host I/O may stall on a single virtual volume.

Symptoms: Sluggish host I/O when a controller node is joining the cluster.

Conditions of occurrence: A controller node has been rebooted, and is rejoining the cluster.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 230407

Issue summary: The array unexpectedly restarts if Flash Cache simulation is enabled during upgrade or if the System Manager restarts.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA-MU1

Issue description: If Flash Cache simulation is enabled either during upgrade, or if System Manager restarts, controller nodes or the entire array may unexpectedly restart.

Symptoms: The array unexpectedly restarts and the CLI command **`showflashcache`** reports that the Mode is equal to SIM.

Conditions of occurrence: Flash cache simulation is enabled and the System Manager is restarted, a controller node is rebooted or unexpectedly restarts, or an HPE 3PAR OS upgrade is performed.

Impact: High

Customer circumvention: Disable Flash Cache simulation.

Customer recovery steps: Avoid running the Flash Cache simulator for extended periods of time, and not while attempting controller node service operations or OS upgrades.

Issue ID: 231482

Issue summary: When a controller node reboots or restarts, it does not join the cluster.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU1

Issue description: Under heavy workloads, a controller node that is either rebooted intentionally or unexpectedly restarts does not rejoin the cluster.

Symptoms: Controller nodes fail to join the cluster after a controller node reboot or restart.

Conditions of occurrence: A controller node attempts to join the cluster while the array is experiencing high workload.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Temporarily reduce the workload on the array.

HPE 3PAR 3.3.1 MU2 File Persona Release Notes

What's New in File Persona

Remote Copy auto failover for FPGs

Activates FPGs automatically on the secondary system. Occurs if the primary system fails when adding FPGs to a Remote Copy Group that uses the AutoFailover policy.

Remote Copy manual failover/failback for FPGs

Simplifies processes associated with adding FPGs to a Remote Copy Group, failing over the Remote Copy Group, and failing back the Remote Copy Group (SSMC 3.3.1 and later).

File lock compliance mode

Increases security with File Lock Compliance to meet regulations defined by U.S. Securities and Exchange Commission rule 17a-4.

Authentication improvements

Includes the following improvements:

- **LDAP performance**
- **Redundant LDAP providers**

  Specify multiple LDAP servers. Ensure resilience if a single-server failure occurs.
- **Local user mapping**

  Create user mappings between Active Directory users and Local users.
- **Minimum UID/GID lowered from 1000 to 100**

  Integrate simply with Linux environments that include user accounts in the 100 to 1000 range and require access to files presented by File Persona.

Major version on-disk upgrade

Uses the latest File Persona features with FPGs originally created on software versions earlier than 3.2.2 MU2.

SMB v1 protocol control

Allows the administrator to configure communication paths using SMB v1.

The SMB protocol facilitates communication paths between a client and File Persona and between File Persona and Active Directory (AD).

On a new configuration, SMB v1 defaults to disabled. The administrator can enable SMB v1 for each of the paths after determining that the clients or AD still require it.

With an upgrade, SMB v1 remains enabled for backward compatibility. The administrator can disable SMB v1 for each of the paths after confirming that no requirement exists for clients or AD.

Network diagnostics

Adds commands to perform `ping` and `traceroute` requests from the perspective of the File Persona instance. Simplifies diagnosis of network configuration issues during setup.

Modifications to File Persona

HPE 3PAR OS 3.3.1 MU2 addresses the following issues:

Issue ID: 72021

Issue summary: Corrects an issue where an alert indicating a temporary failure is received, while other CLI commands are failing repeatedly.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: If there is an issue with the operation of the node listed as `Active` by the `showfs` command, the other nodes may report one or more alerts indicating a temporary failure condition. Once the node listed as Active is healthy again, this issue will be automatically resolved.

Symptoms: Alert indicating a temporary failure is retrieved while other CLI commands are failing repeatedly.

Conditions of occurrence: The node listed as Active in the `showfs` command is in an abnormal state.

Impact: Medium

Customer circumvention: None

Customer recovery steps: Use the "`stopfs <node>`" command to stop the active node that is in an abnormal state and allow one of the other nodes to become Active.

Issue ID: 91629

Issue summary: Node for File Services restarts after upgrade to 3.2.2 MU3

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: An application that sets and cancels Directory Change Notification many times for each file during open/read/write/close file access pattern can trigger inefficient memory usage by the file services SMB server. Over time this memory usage can cause the file services to restart, which migrates the FPGs to their alternate node.

Symptoms: After several weeks of running, the file services for a node restart, causing the FPGs to be migrated to their alternate node.

Conditions of occurrence: Custom SMB application that uses an unusual pattern of Directory Change Notifications (set/cancel/set/…) while doing high I/O loads.

Impact: High

Customer circumvention: None

Customer recovery steps: Cluster automatically fails the file systems over to HA node. Customer must migrate the FPGs back to the original node.

Issue IDs: 94268

Issue summary: Corrects an issue when snapshot operations fail and the snapshot component is not functional.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: The issue occurs when snapshot reclamation operations such as snapshot creation, deletion or listing are running. An error message or exception "`Cannot get actor reference. Actor system is down`" could be seen.

Symptoms: Snapshot operations failing with an internal exception stating "`Cannot get actor reference. Actor system is terminated`".

Conditions of occurrence:

1. The file snapshot functionality is sensitive to system load, and could produce unexpected results under heavy snapshot operations
2. File services for a node could become unresponsive when reclamation operations are running.

Impact: High

Customer circumvention: Avoid running reclamation operations during peak hours. When reclamation is running, avoid running other snapshot operations.

Customer recovery steps: Restart the File Persona file system services using the `stopfs` and `startfs` CLI commands.

Issue ID: 96032/99297

Issue summary: When using the Open Files functionality in Microsoft Management Console (MMC), "`Error 6: The handle is invalid`" is often returned.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: Customers monitoring open file count with MMC will not receive "`Error 6`" on shares with frequent open/close/delete operations.

Symptoms: In-accurate open file count and "`Error 6: The handle is invalid`"

Conditions of occurrence: Attempting to use the "`Open Files`" functionality of the 'Shared Folders' plugin in MMC. Due to the active nature of their file system, file handles are closed in the time period between when the MMC client asks our server for the file list and when that list is returned. If this occurs "Error 6: The handle is invalid" is returned to the MMC an no file list is displayed.

Impact: High

Customer circumvention: Reduce the frequency of create/delete cycles on MMC monitored shares.

Customer recovery steps: Reducing the frequency of MMC polling, or reducing the open/close/delete frequency are the only actions (without the patched code) to avoid the "Error 6" issue.

Issue ID: 97041

Issue summary: A failover request can be unsuccessful when an SMB connection request comes in after a failover request has been made and SMB is still closing existing connections.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: Request for a manual failover is unsuccessful.

Symptoms: Requested file system failover is unsuccessful, and file system remains presented for original node.

Conditions of occurrence: SMB clients requesting new connections after a failover request was made, but not yet completed.

Impact: Medium

Customer circumvention: Do not allow new connections to the SMB server while failover is in progress.

Customer recovery steps: Retry of the failover after the initial failover will succeed unless new clients continue to try to make new connections.

Issue ID: 97354

Issue summary: Under some rare conditions, when a directory has a large number of sub-directories or files, `create` or `rename` operations in that directory may result in the disappearance of some files.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: Using `create` or `rename` operations on directories whose index was removed can cause some files in that directory to no longer be visible.

Symptoms: Some files in the affected directory will not be visible.

Conditions of occurrence:

- When offline FSCK is run on an FPG and when FSCK ends up detaching the directory index when correcting the directory entry in one of the directory pages where the directory has more than one 8K page.
- When a Snapshot Purge operation directory index is closed prematurely.

Impact: High

Customer circumvention: Perform one of the following actions:

- Upgrade to 3.2.2 MU4 with P85.
- Upgrade to 3.3.1 MU1 with P07.

Customer recovery steps:

1. Upgrade to one of the releases mentioned above.
2. Have Support run FSCK.

Issue ID: 97551

Issue summary: Offline FSCK reconnected lost+found directory names could not be renamed if the FPG had taken a snapshot.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: Each snapshot is associated with an epoch. Each file which is visible in a snapshot has an epoch range. Using the epoch range, the system is able to determine which file belongs to which snapshot, along with the changes that exist in that file which are different in different snapshots. When these files were lost FSCK brought them back as LOST + FOUND entries, but it did not update their epochs (different for each snapshot taken). This made them visible across all snapshots, causing the file rename operation to not succeed.

Symptoms: Rename operation does not complete on lost+found files.

Conditions of occurrence: When directory entries become inconsistent, FSCK tries to bring them back with the help of disk data as lost and found entries in the same folder where they were originally previous to the inconsistency. If they cannot be linked to their parent directory, they are placed in the Lost+Found folder in the root directory .

In the case of snapshots, directory entries with the same name can be part of different snapshots based on their birth and death epochs. If the directory entries have the same epoch, then all files will be visible to all snapshots, and the effect of different directory entries visible to different snapshots will no longer be possible. Hence renaming will not succeed.

Impact: Medium

Customer circumvention: N/A

Customer recovery steps: If the epochs have already been updated and lost+found files have been generated, the latest changes will not be able to bring back the correct epochs for the lost files/dentries. Otherwise with new changes, the lost files should be recovered with correct epochs.

Issue ID: 97565

Issue summary: Correct an issue where NFS Share is inaccessible after failover/ failback.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: The issue occurs after upgrade to 3.2.2 MU4 P51, during failover File Persona node has been powered off. After this operation, NFS Shares are not accessible. This causes /etc/export entry for the NFS shares to vanish. This issue occurs when the NFS Share path has been removed without removing the share. NFS Manageability component stops re-exporting NFS shares if it encounters any issue during the re-export of NFS Shares.

Symptoms: Inaccessible NFS share after failover/ failback operation.

Conditions of occurrence: Upgrade to version 3.2.2 MU4 P51 followed by Failover/Failback causes an NFS Share access issue. If any of the NFS share directory is deleted without removing the NFS share, re-export of NFS shares fails for that NFS share and it will not proceed to add other NFS shares. Due to this, not all NFS shares may be exported.

Impact: Medium

Customer circumvention: Do not remove any directory exported over NFS without removing the NFS Share.

Customer recovery steps: Recreate the removed directory and failover and failback fpg. A fix has been provided in NFS manageability component to continue processing other NFS export entries by skipping exports with non-existing directories.

Issue ID: 97762

Issue summary: Windows 10 client backup to a File Persona SMB share does not complete and returns the message "The sector size of the physical disk on which the virtual disk resides is not supported."

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: The Windows 10 version of Windows backup utility displays the following message:

```
Cannot create a file when that file already exists.

Details: The sector size of the physical disk on which the virtual disk resides
is not supported.
```

Symptoms:

- Windows 10 backup to File Persona SMB share is unsuccessful.
- Mount of ISO by Windows 10 from a File Persona SMB share is unsuccessful.

Conditions of occurrence: Using Windows 10 client to backup to a File Persona Share, or mount and ISO file from a File Persona share.

Impact: High

Customer circumvention: Using Windows client version earlier than Windows 10 works successfully.

Customer recovery steps: None

Issue ID: 98767

Issue summary: Delete operation is unsuccessful and no error is returned when the directory in question contains an Alternate Data Stream (ADS).

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: Customer cannot remove a directory that appears to be empty after renaming a file to which an Alternate Data Stream was attached.

Symptoms: Deletion of a directory is unsuccessful without an error message.

Conditions of occurrence: When you rename a file that has an Alternate Data Stream to an existing file name which also has an Alternate Data Stream, delete the target file and its associated ADS, then attempt to remove the empty directory.

Impact: Medium

Customer circumvention: Before renaming a file, verify that the destination file name does not already exist.

Customer recovery steps: None

Issue ID: 98778

Issue summary: In certain rare case scenarios, where case-insensitive lookups for files are involved over a CIFS client, a file services failover can be observed. It leads to data unavailability for the period of failover. This situation can be caused by multiple file operations such as `stat`, `create`, and so on, in parallel on the client under a heavy load.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: The customer will not experience data unavailability due to case-insensitive lookups over CIFS, even under heavy load, with the fix in place. This issue is observed when in-memory inconsistency in the dentry cache triggered by parallel modifications of the entry cache results in accessing illegal memory addresses. This leads to a file services failover.

Issue summary: In certain rare case scenarios, where case-insensitive lookups for files are involved over a CIFS client, a file services failover can be observed. It leads to data unavailability for the period of failover. This situation can be caused by multiple file operations such as `stat`, `create`, and so on, in parallel on the client under a heavy load.

Symptoms: FPGs in a degraded state due to being activated on their backup node and an alert indicating a failure of file services on the primary node for the FPGs.

Conditions of occurrence:

1. When files accessed over CIFS client trigger case-insensitive lookups in the file persona file services software.
2. Multiple parallel operations on the same FPG to create/delete/modify directory entries.
3. FPG node under heavy load triggers shrinking of the entry cache.

Impact: High

Customer circumvention: The issue can be avoided if the file names are unique irrespective of the letter case.

Customer recovery steps: None. This issue leads to temporary unavailability of file services data only for the duration of file services failover. The file services failover is automatic and does not need customer intervention.

Issue ID: 99998

Issue summary: Archive-bit on File Persona SMB-share not set when Microsoft Word modifies a file.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: After creating a file with Microsoft Word, the file does not have the archive bit set.

Symptoms: Backup utilities that rely on the archive bit, like Windows Backup Utility, will not back up files that were created by Microsoft Word or similarly behaving applications.

Conditions of occurrence: The DOS archive bit is not set on a files created by any application, like Microsoft Word, that keeps a temporary copy of a file while it is being modified, and then renames the file to the final name when the file is saved.

Impact: Medium

Customer circumvention: None

Customer recovery steps: Set the archive bit manually using the Powershell command line.

Issue ID: 100166

Issue summary: SMB service self-restart causes momentary interruption in SMB share access.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: When the SMB service restarts, access to the SMB share is interrupted for less than a minute while the service comes backup.

Symptoms:

- New share mapping cannot complete.
- I/O on existing mapped shares cannot complete

Conditions of occurrence: System under high authentication loads can, in rare circumstances, encounter this issue.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None. System self-heals.

Issue ID: 101057

Issue summary: When SMB shares are created under fstore level on File Persona, the permissions (ACLS) inherited should be converted to explicit to match Windows Server behavior. As the default ACL at the root of a share is server specific, and in File Persona it has inherited ACES (from the parent), the user should be cautious when modifying ACLs at the root of the share from a Windows Client.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU1 P07

Issue description: When fixed, the customer will not see a warning pop-up message regarding inherited permissions when modifying permissions on an existing directory or file that is a child object of a directory that was shared after being populated with files and directories.

Symptoms: Warning pop-up message regarding inherited permissions when modifying permissions on an existing directory or file.

Conditions of occurrence: Depending on the Windows version and the default Windows configuration, when using some Windows tools to modify the ACL at the root of the share, the Windows client might also request the server to modify/delete some of the inherited aces on the share folder and its children (if children exist). PLEASE NOTE that this behavior is different from a Windows Server, where the ACL at the root of a share does not have inherited aces.

Impact: Low

Customer circumvention: None

Customer recovery steps:

Different windows versions and Windows Explorer GUI versions will display pop-ups with warnings and/or options to avoid this by converting the inherited aces on the folder to explicit aces. For example, in Windows 2008R2, to allow permissions to be added to a File Persona SMB share folder from Windows without losing the existing permissions perform the following workaround:

1. Right-click the share folder.
2. Select **Properties > Advanced > Disable Inheritance > Convert inherited permissions into explicit permissions on this object**.
3. Click **OK**.
4. Add the user(s) or group (s) in the security tab.

An alternative is to use the File Persona 3PAR CLI:

```
setfshare smb - acl + | <permlist>
```

The specified aces in `<permlist>` will be pre-pended to the other existing ACEs in the share folder for ACL without affection the attributes of the other ACE and without affecting the ACLS of the children directories.

```
setfshare smb -acl <permlist>
```

The specified ACL will be applied (replacing the existing ACL) to the share folder, but any existing children directories will keep their existing ACL. Child directories created after the share folder ACL is modified will inherit from the new share folder ACL.

HPE 3PAR 3.3.1 MU2 CLI Release Notes

IMPORTANT: Ensure that any applications that use CLI, CIM, WSAPI, or VASA/VVol components are TLS v1.2 compliant. Non-compliant host applications may stop communicating with the array if TLS1.2 strict enforcement is selected.

What's New in the CLI

Changed Commands

Command	Description
`setcim`	New **`-pol`** options **`tls_strict`** and **`no_tls_strict`**
`setwsapi`	New **`-pol`** options **`tls_strict`** and **`no_tls_strict`**
`showcim`	New **`tls_strict`** and **`no_tls_strict`** policies in **`showcim -pol`**
`showwsapi`	New Policy field in **`showwsapi -d`**

HPE 3PAR 3.3.1 MU2 CIM API Release Notes

IMPORTANT: Ensure that any applications that use CLI, CIM, WSAPI, or VASA/VVol components are TLS v1.2 compliant. Non-compliant host applications may stop communicating with the array if TLS1.2 strict enforcement is selected.

What's New in the CIM API

A new CLI **`setcim`** command policy named **`tls_strict`** requires HTTPS connections to the CIM API to use only TLS 1.2 and only with the following set of secure ciphers.

- DHE-RSA-AES256-GCM-SHA384
- DHE-RSA-AES128-GCM-SHA256
- ECDHE-RSA-AES256-GCM-SHA384 (new)
- ECDHE-RSA-AES256-SHA384 (new)
- ECDHE-RSA-AES256-SHA (new)

The new policy `no_tls_strict`, which is the default, supports TLS 1.2 with the above ciphers, and TLS 1.1 and 1.0 with the following cipher in addition to those which were previously supported.

ECDHE-RSA-AES256-SHA (new)

Indications sent from the CIM server over HTTPS connections will respect the TLS policy setting.

HPE 3PAR 3.3.1 MU2 Web Services API Release Notes

IMPORTANT: Ensure that any applications that use CLI, CIM, WSAPI, or VASA/VVol components are TLS v1.2 compliant. Non-compliant host applications may stop communicating with the array if TLS1.2 strict enforcement is selected.

What's New with the Web Services API Software

A new CLI **`setwsapi`** command policy named **`tls_strict`** requires HTTPS connections to the WSAPI to use only TLS 1.2 and only with the following set of secure ciphers.

- DHE-RSA-AES256-GCM-SHA384
- DHE-RSA-AES128-GCM-SHA256
- ECDHE-RSA-AES256-GCM-SHA384 (new)
- ECDHE-RSA-AES256-SHA384 (new)
- ECDHE-RSA-AES256-SHA (new)

The new policy `no_tls_strict` supports TLS 1.2 with the above ciphers, and TLS 1.1 and 1.0 with the following cipher in addition to those which were previously supported.

ECDHE-RSA-AES256-SHA (new)

The default policy is `tls_strict`.

HPE 3PAR OS 3.3.1 MU3 Release Notes

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**.

IMPORTANT: If upgrading from an earlier version of 3.3.1 to 3.3.1 MU3, see the **HPE 3PAR OS and Service Processor Software Update Guide (HPE 3PAR OS 3.3.1 HPE 3PAR Service Processor 5.x)** for instructions on updating your specific software.

OS upgrade prerequisite: The latest Upgrade Tool must be staged prior to the HPE 3PAR OS upgrade to 3.3.1 MU3. 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: It is highly recommended that the array has all available and applicable patches applied before beginning the upgrade to 3.3.1 MU3.

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

This product includes code licensed under the GNU General Public License, the GNU Lesser General Public License, and/or certain other open source licenses. A complete machine-readable copy of the source code corresponding to such code is available upon request. This offer is valid to anyone in receipt of this information and shall expire three years following the date of the final distribution of this product version by Hewlett Packard Enterprise Company. To obtain such source code, send a check or money order in the amount of US $10.00 to:

Hewlett Packard Enterprise Company

Attn: General Counsel

3000 Hanover Street

Palo Alto, CA 94304 USA

Please specify the product and version for which you are requesting source code.

HPE 3PAR 3.3.1 Operating System MU3 Release Notes

What's New in the OS

New and enhanced features include:

3PAR OS 3.3.1 MU3

- Faster BIOS updates through enhanced logic to update only modified blocks of the image.
- Faster Drive Enclosure upgrades: Enhancements to drive enclosure management features.
- Parallelized firmware update logic for faster drive enclosure firmware upgrades.
- Faster Node Reboots: Enhanced logic to perform targeted checks and avoid unwanted VV checks during startup.
- Improved metadata traversing algorithms to speedup VV checks during startup.
- Faster **`updatevv`**: Enhanced logic to perform operations in parallel when used for large VV sets to reduce time for completion and IO block period.
- Enhanced **`tunevv`**, online physical copy, and online **`promotesv`** commands with new **`-pri`** option to set priority for the tasks
- Audit logs now include activity for privileged user account (root) to ensure audit logs include all actions for all local and remote users.
- Enhanced RAID-6 layout to improve performance for all set size including most used 6+2 set size.
- Updates to enhance HPE 3PAR OS security.
- The number of event log files is now configurable through the setsys EventLogNum{number} command with a range of 1 to 30 event log files (default 10). The size of each individual event log file is now configurable through **`setsys EventLogSize {number}`** with a range of 1Mib up to a maximum of 10Mib (default 4MiB).
- FIPS-2 compliance is now supported on management interfaces, configurable with the **`controlsecurity`** command. The supported Management Interfaces are CIM, CLI, EKM used for Data at Rest Encryption, LDAP Authentication, SNMP, Syslog, SSH, WSAPI, and VASA. ISCSI CHAP is not supported.

Modifications to the HPE 3PAR OS

The following issues have been addressed in this release.

Issue ID: 140776

Issue summary: When an SFP on an Emulex LPe 16004 is removed or inserted, the information in `showport -sfp` is not updated.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1 GA - MU2

Issue description: When an SFP on an Emulex LPe 16004 is removed or inserted, the information in `showport -sfp` is not updated.

Symptoms: **`showport -sfp -d`** returns stale SFP information if the SFP was removed, swapped, or inserted.

Conditions of occurrence: Insert SFP or remove/replace SFP in Emulex LPe 16004 HBA.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Restart the controller node or remove/insert the FibreChannel cable.

Issue ID: 152317

Issue summary: Controller node unexpectedly restarts when an inter-node packet send process takes more time to complete.

Affected platforms: All StoreServ

Affected software versions:3.2.2 and 3.3.1

Issue description: If sending an inter-node packet takes a long time to complete, there is a possibility of hitting race condition in packet cleanup path that can resulting in a controller node unexpectedly restarting.

Symptoms: An unexpected controller node restart.

Conditions of occurrence: Heavy I/O or remove copy workload resulting in long inter-node packet processing.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 155807

Issue summary: A controller node restart occurs during Port Persistence failback if a loss of a SCSI initiator happens at the same time on the same vport that is disappearing.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1 GA, 3.3.1 MU1

Issue description: Prevents an unexpected controller node restart, when using Persistent Ports, if an initiator disappears from the V-port while a port failback is in progress.

Symptoms: A controller node restart due to kernel assertion.

Conditions of occurrence: A Port Persistence failback at the same time a SCSI initiator is being lost on the network by the vport that is being failed back.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None. The controller node will restart itself.

Issue ID: 159670

Issue summary: The CLI command **`removehost`** will not complete successfully when there is an active VLUN associated with a host. The host descriptors are removed when they should not have been.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU2

Issue description: If descriptors are setup during **`createhost`**, they are incorrectly cleared during a **`removehost`** command when it times out.

Symptoms: **`sethost -desc <hostname>`** descriptors are cleared after a failed **`removehost`** command. They should not be cleared if the **`removehost`** can not proceed.

Conditions of occurrence: Any condition that prevents **`removehost`** from completing successfully, such as having VLUNs attached.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Recreate descriptors for host (or hosts).

Issue ID: 163864

Issue summary: Audit user enhancements enable customer to get the output of **`netstat –avntp`** and **`iptables –L`**.

Affected platforms: All StoreServ

Affected software versions: 3.1.3 GA. 3.3.1 EGA, 3.3.1.MU1

Issue description: This enhancement allows the customer to run **`iptables -L`** and **`netstat -avntp`** in the audit user environment.

Symptoms: You cannot run **`iptable`** or **`netstat`** from within the audit user environment.

Conditions of occurrence: Not applicable.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 164373

Issue summary: Added sufficient validation to avoid creating unmapped logical disks while creating full volume with zero detect option.

Affected platforms: All StoreServ

Affected software versions: All versions

Issue description: Sufficient validations are added to avoid manually removing the unmapped logical disks.

Symptoms: Unmapped logical disks while creating full virtual volumes.

`showld` displays the logical disks.

`checkhealth` shows the unmapped logical disks.

Conditions of occurrence: Creating Full volume with invalid polices.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Remove the unmapped logical disks using **`removeld`**.

Issue ID: 169481

Issue summary: Admitting a TPVV to a Remote Copy Group incorrectly shows an erroneous failure message when HPE 3PAR OS 3.3.1 is communicating with an older version.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA

Issue description: When trying to admit a noncompressed volume to a Remote Copy group using the **`-createvv`** option, it states that the target system does not support compression.

Symptoms: Admitting a noncompressed volume to Remote Copy group when using **`-createvv`** option states that the target system does not support compression.

Conditions of occurrence: CPVV has been converted to a TPVV.

Attempt to admit the volume to Remote Copy group while using the **`-createvv`** option.

Impact: Low

Customer circumvention: If this issue occurs, manually create and admit affected volumes on the target array.

Customer recovery steps: None.

Issue ID: 185349

Issue summary: **checkvv** reports DDCs consistent, allowing System Manager to bring these DDCs to normal state. This report occurs even when DDS remains in ICE state, if a manual check is executed.

Affected platforms: All StoreServ

Affected software versions: 3.2.1 MU1, 3.2.2, 3.3.1

Issue description: While running **checkvv** if without running a group check, DDCs might be found consistent and brought online (normal) even when DDS is in ICE state. DDCs are presented and accepting IO while the DDS cannot perform any deduplication operation which might result in unexpected behavior. **checkvv** should be aware of a DDS in ICE state, and avoid reporting DDCs as consistent.

Symptoms: **showvv** will show `DDS volume in internal_consistency_error`, and one or more DDCs in normal state.

Conditions of occurrence: DDS is marked ICE during pfail recovery, and a manual check is done on the DDC after the initial post recovery **checkvv**.

Impact: High

Customer circumvention: Do not run manual check against DDCs online. DDC should be forced offline so a group **checkvv** is run, allowing **checkvv** to correct inconsistencies in the deduplication group.

Customer recovery steps: DDS should be brought offline, and **checkvv** should do a group check to address inconsistencies in both the DDS and DDCs, bringing all volumes to normal state.

Issue ID: 188681

Issue summary: Metadata is not properly processed during controller node recovery with Asynchronous Streaming Remote Copy configurations on the primary array. This leads to an unexpected controller node restart.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description: Partial Cluster memory pointers are shared between two I/Os in the same sequence number.

One of the IO has been acknowledged, but the ticket log remains.

During controller node down recovery, the reference count for those shared cluster memory pointers are incremented for the acknowledged I/O, but not de-incremented.

Symptoms: Cluster memory pointer leaks.

Conditions of occurrence: Overlapping I/Os followed by controller node down in asynchronous streaming configuration.

Impact: Medium

Customer circumvention: Using sync or periodic modes instead of asynchronous streaming mode.

Customer recovery steps: None.

Issue ID: 191345

Issue summary: Single controller node restart due to array overflow in 8Gb FC driver.

Affected platforms: StoreServ 7000, StoreServ 8000, StoreServ 10000

Affected software versions: 3.2.1 MU3

Issue description: The 8Gb FC driver maintains a 512 element array to hold mapping between remote port indicator (RPI) provided by firmware and device identifier (DID). Due to a certain SAN condition, hosts/fabric keeps sending RSCN. In response, the driver/firmware sends LOGO to initiator DIDs. As part of LOGO, firmware continues generating and incrementing default RPIs. At some point, the driver receives RPI which is more than 512. This causes array overflow and controller node restart.

Symptoms: Host path loss.

Conditions of occurrence: Frequent login/logout events by hosts.

Numerous RSCNs notifications.

Number of initiators exceeding more than 512 on 8G FC driver.

Impact: Medium

Customer circumvention: Check setup, including fabric switch and hosts.

Customer recovery steps: Single controller node restart.

Issue ID: 192035

Issue summary: SSMC/IMC shows `cage0 DP2 offline` even though this data port is online.

Affected platforms: StoreServ 7000, StoreServ 8000

Affected software versions: All versions

Issue description: HPE 3PAR OS API to SSMC and other user Interfaces incorrectly report the drive enclosures (cages) external SAS connector states.

Symptoms: SAS connectors that are online show as offline in SSMC/IMC.

Conditions of occurrence: This issue is a user interface issue that is always visible.

Impact: Low

Customer circumvention: Use the 3PAR CLI to upgrade to MU3.

Customer recovery steps: None.

Issue ID: 193772

Issue summary: The System Manager application may restart if the **`setpd ldalloc on|off`** command is used on a two-node system and one controller node is down.

Affected platforms: All StoreServ

Affected software versions: 3.1.1 and later

Issue description: If the **`setpd ldalloc on/off`** command is used to disable or enable disks on a two-node system, and one controller node is not integrated into the cluster, a System Manger application restart may occur.

Symptoms: System Manager restart after issuing a **`setpd ldalloc`** command.

Conditions of occurrence: This restart occurs on any two-node system when a controller node is down and the **`setpd ldalloc`** command option is used to enable or disable disks.

This command can be used frequently as part of the drive removal process when physically reconfiguring systems.

Impact: Low

Customer circumvention: Do not issue the **`setpd ldalloc`** command on a two-node system if one controller node is not integrated into the cluster.

Customer recovery steps: A System Manager restart may cause tasks to time out. These tasks may need to be restarted.

Issue ID: 193779

Issue summary: Detection of a Harrier2 RPC send timeout as the first fatal link error reported in the cluster (note that sends timeouts commonly occur as side effects of other link errors or controller node down events, so ordering is critical for proper diagnosis).

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

Affected software versions: 3.2.2GA-3.3.1.MU1

Issue description: Queues which store traffic bound for two different destinations are connected and carry traffic between certain sources and destinations. This traffic forms a dependent network of queues and traffics is formed. If the head of line of each of at least six queues in such a network is such that it must wait for availability in a downstream queue which is full ,and mutual circular dependency exists within this network, no progress is possible.

- The multinode system deadlock arises even though software and hardware are functioning correctly, but together they cause the unintended consequence of deadlock.
- The ASIC sends timeout serves to break the deadlock. Once the send timeout is reported, the reporting ASIC link interface begins consuming all traffic queues to it, creating forward progress. One controller node is forced to restart by the cluster manager, but the other controller nodes are able to make progress. Unless a second deadlock cycle forms before the first controller node restarts, this is the only symptom. Thus the ASIC error actually prevents the cluster from remaining deadlocked.
- To prevent this multinode deadlock from forming, give one of the traffic flows another path through the controller node, preventing at least one queue in the cycle from containing traffic bound for more than one destination and guaranteeing that traffic will eventually flow.

Symptoms: The cluster-level symptom is an error, `tpd: CM: membership removed by mask <xxx>, sender <y>!`

Conditions of occurrence: Affects dual-ASIC systems with at least four controller nodes. Number of possible deadlock cycles increases exponentially with the number of controller nodes.

Reproduction seen only with significant traffic to SSD volumes. Odds appear to increase in arrays to which controller nodes have been added.

Impact: Medium

Customer circumvention: Maintain a balanced configuration and avoid overloading the array.

Customer recovery steps: If cluster outage with more than one controller node then host application I/O recovery would be required.

Issue ID: 195256

Issue summary: LUNs become unavailable due to excessive ODX token invalidation processing.

Affected platforms: All StoreServ

Affected software versions: All

Issue description: LUNs will remain accessible during ODX token invalidation processing.

Symptoms: LUNs continually return `BUSY` for write and SCSI-3 persistent reservation requests.

`sdt_token_update: Failed to remove token` appears frequently in the `showeventlog` output.

Conditions of occurrence: Use of ODX on many LUNs.

Impact: Medium

Customer circumvention: HPE support has a script that can periodically invalidate tokens to help prevent this issue from happening.

Customer recovery steps: None.

Issue ID: 195356

Issue summary: When using secure shell to access the CLI commands **`statrcvv`** and **`histrcvv`**, the CLI may abort with `out of memory` errors.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA

Issue description: When there are a large number of Remote Copy virtual volumes and using secure shell to access the CLI, the commands **`statrcvv`** and **`histrcvv`** abort with `out of memory` errors. The total amount of memory used is more than the 512 MiB limit for this large number of VVs.

Symptoms: Various CLI memory errors like `unable to alloc ####` bytes.

Conditions of occurrence: A large number of Remote Copy VVs.

Impact: Medium

Customer circumvention: Use the Remote CLI client instead of secure shell to access the commands.

Or use filtering optons **`-g`** or **`-t`** to limit the number of VVs.

Customer recovery steps: None.

Issue ID: 197465

Issue summary: A powerfail event may not recover correctly if defrag is running.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1 GA - MU1

Issue description: If there is a powerfail or other event that causes the cluster to go down while defrag is running, powerfail recovery may not handle defrag log entries correctly. The controller node handling the defrag log entry may produce an error message. Since this is during powerfail recovery, the cluster will go down. The same behavior will repeat until powerfail recovery errors on three attempts. On the fourth attempt System Manager will wipe powerfail recovery data from the controller nodes, causing the potential loss of data.

Symptoms: An unexpected error message during powerfail recovery.

Conditions of occurrence: Cluster-wide error messages or powerfail event while defrag is running.

Impact: High

Customer circumvention: Disable defrag.

Customer recovery steps: Run data integrity scans of their data.

Issue ID: 199872

Issue summary: If free space is left in only one cage out of 3 or more cages per controller node pair, you are unable to create a snapshot administration logical disk using **`-ha mag`**.

Affected platforms: All StoreServ

Affected software versions: 3.x.x.GA and 3.x.x

Issue description: If free space is left in only one cage out of three or more per controller node pair, you will be unable to create SA LD with **`-ha mag`** .

Symptoms: Unable to grow SA space.

Conditions of occurrence: If there are three or more cages attached per controller node pair, and all are full except one cage, then you will not able to create a snapshot administration logical disk using the **`createald`** command with **`-ha mag`**.

Impact: Medium

Customer circumvention: To avoid this issue, keep all cages in a balanced state.

Customer recovery steps: As a workaround, add **`-ssz 3`** for SA space in the common provisioning group parameter.

Issue ID: 199904

Issue summary: StoreServ controller node unexpectedly restarts while system is performing data movement activities such as **`compactcpg`**, **`tunevv`**, or **`tunesys`**.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description: StoreServ controller node unexpectedly restarts while system is performing data movement activities such as **`compactcpg`**, **`tunevv`**, or **`tunesys`**.

Symptoms: Restart of StoreServ controller node.

Conditions of occurrence: The system is performing data movement activities such as **`compactcpg`**, **`tunevv`**, or **`tunesys`**.

Impact: Medium

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

Customer recovery steps: StoreServ self-recovery as in any situation of controller node restart.

Issue ID: 200331

Issue summary:

Improved error detection and handling for the following drive models:

HCBF0600S5xeN010, HCBF1200S5xeN010, HCBF1200S5xeF010, HCBF1800S5xeN010

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description: We have developed a utility called Drive Health Assessment (DHA), which enables identification of drives that are at risk of degrading even before they show visible symptoms, to assess the status of every HCBF disk drive in the StoreServ, and to guide containment action. The utility operates by running a script on the array that retrieves certain log data from the drive system area. The script collects this log from all Cobra-F drives in the array. The data generated is phoned home for analysis. The analysis determines which drives are showing signs of degradation, and then proactive replacement begins.

Symptoms: Not applicable.

Conditions of occurrence: Not applicable.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 200537

Issue summary: Unique TPGIDs were not maintained upon dismissing and admitting Remote Copy volumes.

Affected platforms: All StoreServ

Affected software versions: 3.2.2.MU3, 3.2.2.MU4, 3.3.1.GA

Issue description: Current code version overwrites primary group ID with same virtual volume = 258 on original secondary array. Group ID on original primary array is not overwritten with 257, since the VV is “NEW” (tpgid=1). 258 is on both sides.

Symptoms: Both the arrays will show TPG id 258 for the primary and secondary vv.

Conditions of occurrence: A virtual volume which will be part of a secondary group will be dismissed and added to a new primary peer persistence Remote Copy group. Both of the arrays will show TPG ID258 for the Primary and secondary vv.

Impact: High

Customer circumvention:

1. Create PP group.
2. Perform switchover.
3. Dismiss virtual volume from the group.
4. Create a new group from original secondary array.
5. Remove/create new VV from original primary array.
6. Admit VV into this new group.

Both sides report TPG ID258.

Customer recovery steps: If a new volume is created, the TPG ID can be explicitly set using a CLI command to get the TPG ID to the proper value.

Issue ID: 200909

Issue summary: After a controller node starts up, file services are not started automatically on the controller node.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1 GA - MU2

Issue description: When a controller node starts up, enabling file services on the controller node requires all of the volumes in use by file services to be in a ready state. If it takes longer than expected for these volumes to reach a ready state, file services may be left in a shutoff state. If this occurs during an upgrade, a failure will be reported in the upgrade.

Symptoms: After a controller node starts up, the **`showfs`** command reports file services on the controller node to be in a shutoff state.

Conditions of occurrence: Reboot of a controller node while file services are enabled.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Use the **`startfs -enable <node>`** command for the controller node where file services were in the shutoff state. If an upgrade was in progress, after the controller node reports upgrading again, resume the upgrade process.

Issue ID: 200972

Issue summary: Offloaded data transfer tokens are invalidated.

Affected platforms: All StoreServ

Affected software versions: All

Issue description: Offloaded data transfer tokens are invalidated due to token inconsistencies across array controller nodes after a controller node restart.

Symptoms: Messages in the scsid log indicating a token mismatch and clearing of all tokens for a given virtual volumes.

Conditions of occurrence: Use of Offloaded data transfer during a controller node restart.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 201664

Issue summary: If the master controller node restarts while there is an issue with the disks where reads succeed but writes do not, the system will end up going down and will require intervention to start again.

Affected platforms: All StoreServ

Affected software versions: 3.1.1 and later

Issue description: If the master controller node restarts while there is an issue with the disks where reads succeed but writes do not, the system will end up going down and will require intervention to start again.

Symptoms: System goes down after disk issues and cannot restart.

Conditions of occurrence: Physical disk can be read from but not write.

Impact: High

Customer circumvention: None.

Customer recovery steps: Use `ignoretocpfail` option.

Issue ID: 202630

Issue summary: When array unexpectedly restarts, the memory dump is incomplete and cannot be debugged.

Affected platforms: StoreServ 8000, 8440, or 8450 array with high buffer cache usage

Affected software versions: 3.2.2, 3.3.1 GA

Issue description: Data not necessary for debugging was included in the debugging memory dump. This data causes this dump to be too large to fit in the appropriate start drive partition. The dump is incomplete.

Symptoms: Memory dump due to 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: 202836

Issue summary: Logical unit numbers are inaccessible upon changing the MTU settings of an iSCSI port.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1 GA - MU1

Issue description: Upon changing the MTU settings of a tagged VLAN, the port's priority information is out of sync with the switch. The port becomes sluggish.

Symptoms: Unable to ping ports.

Unable to mount logical unit numbers.

Conditions of occurrence: Change the MTU settings of a VLAN.

Impact: Medium

Customer circumvention: Avoid changing MTU settings for a tagged VLAN.

Customer recovery steps: Reset the affected port using **`controlport rst <n:s:p>`**.

Issue ID: 203090

Issue summary: 3PAR OS Online upgrades to 3PAR OS 3.3.1 fail to complete successfully.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1 MU1 - MU2

Issue description: 3PAR OS online upgrade failures caused by hung FC HBA ports.

Symptoms: System Manager component of 3PAR OS can get stuck trying to upgrade, resulting in the upgrade operation being aborted.

Conditions of occurrence: 3PAR systems configured with 16Gb FC HBAs with one or more of the FC ports in hung state at the time the 3PAR OS upgrade is initiated.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 204754

Issue summary: Single controller node restarts unexpectedly due to a code concurrence issue. This is resolved.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA, 3.3.1 MU1

Issue description: Single controller node restarts unexpectedly in some environments due to a code concurrence issue.

Symptoms: A controller node suddenly restarts during normal operation.

Conditions of occurrence: High demand for VV page table metadata triggers high activity with VV hash operations.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Single controller node restart. Normally restarted controller node joins cluster automatically.

Issue ID: 204959

Issue summary: Controller node automatically resets occurs after auto-reset should have been disabled.

Affected platforms: All StoreServ

Affected software versions: 3.1.1, 3.2.1, 3.2.2, 3.3.1

Issue description: After a controller node is shut down with **shutdownnode**, it should remain down and not automatically reset. If the previous master was the one shut down, or System Manager restarted afterwards, the timer for the automatic reset would start again.

Symptoms: The master controller node restarts 45 minutes after shutting down.

Conditions of occurrence: Restart of System Manager.

Impact: Low

Customer circumvention: Do not assume that the controller node will stay shut down if it was the master controller node before.

Customer recovery steps: Shut the controller node down again.

Issue ID: 205259

Issue summary: System Manager log entries indicate that an mcall is running for a long time while holding the config lock.

Affected platforms: StoreServ 7000, StoreServ 8000

Affected software versions: 3.2.1, 3.2.2, 3.3.1

Issue description: Some 7000 and 8000 drive enclosure CLI commands take a long time to complete. These commands hold the System Manager config lock while running.

In 3.3.1 MU3, the drive enclosure management functionality is no longer part of the System Manager process. This functionality does not compete with System Manager for the same resources (mutexes, semaphores, and so forth). Even when a CLI operation takes a long time to complete, it does not affect the rest of the system operation.

Symptoms: The presence of System Manager log entries such as the following.

```
config_lock_tattler Thread <name> with mcall MCJBT_TTY(222) has been
outstanding for an extended period of time
```

Conditions of occurrence: Execution of drive enclosure CLI commands such as **ddump** take a long time to complete.

Impact: Low

Customer circumvention: Avoid the usage of drive enclosure CLI commands during normal operation. Only use CLI commands when necessary, such as when investigating another issue.

Customer recovery steps: None.

Issue ID: 205406

Issue summary: Remote copy disaster recovery operation failed leaving the Remote Copy groups in inconsistent state.

Affected platforms: All StoreServ

Affected software versions: 3.1.1, 3.2.1, 3.2.2, 3.3.1

Issue description: During Remote Copy disaster recovery operation, volume promotion will not succeed if region move is in progress. This failure will leave the Remote Copy groups in inconsistent state.

Symptoms: A Remote Copy disaster recovery operation is performed while a region move is in progress.

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

Impact: Medium

Customer circumvention: Do not perform disaster recovery operation while region move is in progress.

Customer recovery steps: Use **`setrcopygroup`** with appreciate options to bring the groups back into consistent state.

Issue ID: 205515

Issue summary: When moving back chunklets to a physical disk, the select spare algorithm selected duplicated chunklets causing the move to fail.

Affected platforms: All StoreServ platforms

Affected software versions: 3.2.1.MU5, 3.2.2-3.2.2.MU4, 3.3.1.GA, 3.3.1.MU1

Issue description: Moving back the selected duplicate chunklets causes original move to fail.

Symptoms: **`Movebackpd`** fails, causing `servicemag resume failed`.

Conditions of occurrence: Multiple failed physical disks that caused chunklets being moved around and then moved back.

Impact: Medium

Customer circumvention: Do not let the failed physical disks sit too long. Replace it as soon as possible. The more failed physical disks in the system at the same time, the greater the chance of this occurring.

Customer recovery steps:

1. Do not issue **`servicemag resume`** on those replacement drives.
2. Issue **`controlmag onloop`** for the mag of the replacement drive.
3. Run **`controlpd spinup <wwn>`** of the replacement drive.
4. Run **`controlpd clearerr <wwn>`** of the replacement drive.
5. Run **`setpd ldalloc on <pdid>`** of the replacement drive.
6. Run **`moverelocpd -f <pdid>`** of the replacement drive. You will see the `duplicated_dest` error on some of the chunklets. Issuing this command once will not cause the system to go into a loop like **`servicemag resume`** does.
7. Wait for `moverelocpd` to finish.
8. Once it finishes, run **`showpdch -from <pdid>`** to determine what chunklets are not relocated. The pdid here is the original pdid before the replacement.
9. Use the following script to do a manual chunklet relocation: **`showpdch -from <pdid> -nohdtot | while read pd ch rest; do echo $pd $ch; movech -f -perm -ovrd $pd:$ch; done&`**
10. All the chunklets with `duplicated_dest` error are moving in the background.

Issue ID: 206128

Issue summary: Peer Motion/Federation Destination array may not function as intended when the peer zoning between source and destination arrays does not follow the recommended zoning.

Affected platforms: All StoreServ

Affected software versions: 3.1.3, 3.2.1, 3.2.2, 3.3.1

Issue description: When the peer zoning between source and destination arrays in Peer Motion/Federation configuration does not follow the recommended 1:1 zoning, or when the number of peer paths between the arrays exceeds 2, stale SCSI objects are added to the kernel, even when migrations work. This action causes **`System Manager`** not to be functional.

Symptoms: Some of the array management CLIs may hang.

Conditions of occurrence: Peer zoning between source and destination arrays does not follow the recommended zoning.

Impact: Medium

Customer circumvention: Follow the recommended 1:1 peer zoning with just two paths between the arrays.

Customer recovery steps: When it is determined that stale peer SCSI objects have piled up in the kernel, restart the controller nodes to get rid of them.

Issue ID: 207547

Issue summary: Controller node error messages due to a Remote Copy ticket being in an invalid state.

Affected platforms: All StoreServ

Affected software versions: 3.3.1.GA

Issue description: Silent failure to read a volume in the reread path causes the Remote Copy ticket status to be released in an invalid state. This causes controller node panic and error messages.

Symptoms: Remote Copy reread does not succeed.

Conditions of occurrence: When a controller node goes down, Remote Copy rereads the data and sends it to a target array.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 208116

Issue summary: Metadata corruption during the controller node-down recovery while processing a delayed copy on write (DCOW) exception entry on the base data cache page.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU1

Issue description: When processing the DCOW exception entry on the data cache page of the base VV during the controller node-down recovery, the replicant controller node finds an inconsistent exception entry between RO snapshot and base VV.

Symptoms: Metadata inconsistency event.

VV into Internal Consistency Error.

Conditions of occurrence: A controller node restart.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Run **`checkvv`**.

Issue ID: 208350

Issue summary: Lightweight Directory Access Protocol (LDAP) connections may not successfully bind when using Simple Authentication and Security Layer as a binding mechanism. The process gets stuck in a tight loop and never exits.

Affected platforms: All StoreServ

Affected software versions: 3.1.3, 3.2.1, 3.2.2, 3.3.1

Issue description: Lightweight Directory Access Protocol connections do not successfully bind through Simple Authentication and Security Layer due to a Kerberos error. This error results in the binding getting stuck, and the `auth_helper` process never exits.

Symptoms: Multiple `auth_helper` processes stacking up, and utilizing near 100% CPU usage.

Conditions of occurrence: Lightweight Directory Access Protocol is configured with a binding type of Simple Authentication and Security Layer.

Impact: Medium

Customer circumvention: Set the binding type to simple.

Customer recovery steps: Contact HPE support.

Issue ID: 209286

Issue summary: System Manager becomes sluggish after running **compactcpg**.

Affected platforms: All StoreServ

Affected software versions: 3.1.1.xx, 3.2.2.xx, 3.3.1.xx

Issue description: Variable overflow caused infinite loop in function.

Symptoms: System Manager is sluggish.

Conditions of occurrence: If common provisioning group raid type is 6, row size 1 and have full vv with size greater than 6TB.

Impact: Medium

Customer circumvention: Do not to run **compactcpg** if common provisioning group raid type is 6, row size 1 and have a full vv with size greater than 6TB.

Customer recovery steps: Restart System Manager.

Issue ID: 211011

Issue summary: CLI show is sluggish during **updatevv**.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA, 3.2.2-3.2.2.MU4, 3.3.1 MU1

Issue description: In some environments, the **cachesvr** process may become sluggish during **updatevv** operations on virtual volume sets.

Symptoms: CLI show commands become sluggish during **updatevv** operations.

Conditions of occurrence: The issue is most likely to occur on systems on which issue updatevv operations to large vvsets.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Contact HPE Support.

Issue ID: 211227

Issue summary: After changing an IP address, NETC does not establish the new IP when a failover occurs.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 MU2 - MU4, 3.3.1 GA, 3.3.1 MU1

Issue description: A cluster loses its IP address when performing **`setnet startaddr`** after an IP has already been established, followed by a network master change. This renders the cluster inaccessible.

Symptoms: You are not able to connect remotely to the cluster.

The IP address gets set to 127.127.1.1.

Conditions of occurrence: Cluster has changed IP addresses since first initialization and a failover occurs after the first step happens.

Impact: High

Customer circumvention: Do not change IP addresses until a patch for this issue has been installed.

Customer recovery steps: Restarting NETC and NETCC or restarting the cluster will recover from this issue.

Issue ID: 211282 / 221630

Issue summary: Conversion of a thinly provisioned virtual volume without a snap common provisioning group to a fully provisioned virtual volume, will cause uncontrolled shutdown when the snap CPG is added back and snapshot or physical copy operations are used.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1

Issue description: Conversion of a thinly provisioned virtual volume without a snap common provisioning group to a fully provisioned virtual volume leaves internal data structures in an inconsistent state. Adding back the snap common provisioning group (through the **`setvv`** command) and performing operations that use snapshots, such as creation of normal read only or read/write snapshots or initiating a physical copy, will cause uncontrolled shutdown of all controller nodes in the cluster with the error message of `Fatal Exception`. After the shutdown, data structures return to a consistent state.

Symptoms: `Fatal Exception` on all controller nodes in `adm_get_blks` (or similarly patched function).

Conditions of occurrence: Conversion of thinly provisioned virtual volume without a snap common provisioning group to a fully provisioned virtual volume brings out this case.

Impact: High

Customer circumvention: Avoid converting a thinly provisioned virtual volume that has no snap common provisioning group to a fully provisioned virtual volume. This action results in inconsistent values in data structures. Those data structures are used when you add back the snap common provisioning group and try to do any snapshot-related actions on the fully provisioned virtual volume. When the snapshot-related actions use the inconsistent data structures, the controller node produces a fatal exception in the `adm_get_blks` function. Once the controller nodes restart, the data structures are made consistent.

Customer recovery steps: Once the controller nodes restart, the data structures are made consistent.

Issue ID: 211324

Issue summary: A fully provisioned File Provisioning Group may be unexpectedly deactivated when the backing CPG is nearly out of space.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU2

Issue description: When the CPG backing an FPG is nearly out of space, as seen in the Free column of the command `showspace -cpg <cpgName>`, the system tries to automatically deactivate all the thinly provisioned FPGs that are using that CPG. This action ensures that the FPG does not experience a write error due to lack of space for additional allocations to the backing VV(s). In this case, fully provisioned FPGs are also deactivated.

Symptoms: An alert indicates that the CPG is nearly out of space.

An alert indicates that the FPG has been deactivated.

Conditions of occurrence: CPG space is exhausted and fully provisioned FPGs are in use.

Impact: Medium

Customer circumvention: Ensure adequate CPG capacity when thin provisioning is used.

Customer recovery steps: Resolve the CPG space issue, then manually activate the FPGs again using the **`setfpg`** command.

Issue ID: 211936

Issue summary: Unexpected controller node restarts when the **`ddcscan`** encounters a virtual volume that has already been removed as a result of **`vvcopy`** or **`promotesv`** not completing successfully.

Affected platforms: All StoreServ

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

Issue description: Unexpected controller restarts when **`ddcscan`** encounters a virtual volume that has already been removed as a result of online **`vvcopy`** or **`promotesv`** not completing successfully.

Symptoms: Controller restart.

Conditions of occurrence: Online **`vvcopy`** or **`promotesv`** not completed successfully involving deduplication volume, and subsequent destination volumes being deleted.

Impact: Medium

Customer circumvention: Do not remove the destination volume manually after online **`vvcopy`** or **`promotesv`** does not complete successfully.

Customer recovery steps: None.

Issue ID: 212190

Issue summary: Loss of IO resources due to IO write error for large write.

Affected platforms: All StoreServ

Affected software versions: 3.3.1, 3.2.2

Issue description: When sending large (over 1 MB) IO across the controller nodes, IO will split into small chunks. Returning an IO error can cause loss of resources of the rest of chunks which have not been sent.

Symptoms: IO gets stalled. **`statcmp`** shows credit usage as a high number. Even when under no IO, this level does not go down.

Conditions of occurrence: Any error in the IO path and large write.

Impact: High

Customer circumvention: Do not send big IO.

Customer recovery steps: If you encounter the issue before reaching the maximum size of the credit limit, under light load, and still see resource pool are not return, restart the affected controller nodes, one at a time.

Issue ID: 212699

Issue summary: The system can go through uncontrolled shutdown when links are coming down.

Affected platforms: All StoreServ

Affected software versions: 3.3.1

Issue description: When the links are coming down, the fragment can be accessed by two or more threads. Due to the race condition in the code, the same fragment can be cleaned by two threads, which can result in system error messages.

Symptoms: System error messages during link down.

Conditions of occurrence: Any link removal condition.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Restart the system.

Issue ID: 212816

Issue summary: Administrative connections including StoreServ Management Console (SSMC) may be logged out of their session independent of the user actions when lightweight directory access protocol is configured. CLI or SSH connections may not be able to be established.

Affected platforms: All StoreServ

Affected software versions: 3.1.3 GA-MU3, 3.2.1 GA-MU5, 3.2.2 GA-MU4, 3.3.1 GA-MU1

Issue description: Due to a communications issue with the LDAP server, the SSMC will not successfully receive an authentication which requires the user to re-enter credentials. CLI or SSH sessions may not be established. CLI or SSH sessions that are already established are not interrupted. For SSMC users, the issue occurs during an active session.

Symptoms: The users have to log in again after an unsuccessful login attempt from CLI SSH. For SSMC, a popup window requests credentials again.

The SSMC administrator connection no longer establishes and the administrator must re-enter administrators credentials.

Conditions of occurrence: LDAP is configured on the array.

Administrative connections are used.

Impact: Medium

Customer circumvention: Using a local administrator account for the SSMC administrator user prevents the administrative account from encountering this issue. Other (active) user sessions will still encounter the issue.

Customer recovery steps: Users must log in again, and the Administrator must log in again to the SSMC and enter the appropriate credentials.

Issue ID: 213022

Issue summary: Online/Offline upgrade will be blocked from 3.2.1 MU3 to 3.3.1 MU2 (or later) if 3.2.1 MU3 P20 is not installed.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU2

Issue description: Online/Offline upgrade will be blocked from 3.2.1 MU3 to 3.3.1 MU2 (or later) if 3.2.1 MU3 P20 is not installed.

Symptoms: Online/Offline upgrade will be blocked from 3.2.1 MU3 to 3.3.1 MU2 (or later) if 3.2.1 MU3 P20 is not installed.

Conditions of occurrence: Attempting to upgrade to 3.3.1 MU2 from 3.2.1 MU3 without P20 installed.

Impact: High

Customer circumvention: Install 3.2.1 MU3 P20 and retry the upgrade.

Customer recovery steps: Install 3.2.1 MU3 P20 and retry the upgrade.

Issue ID: 213080

Issue summary: Remove the minimum overprovisioning limit.

Affected platforms: All StoreServ

Affected software versions: 3.3.1

Issue description: Remove the minimum overprovisioning limit of 1.5.

Symptoms: Not able to set the overprovisioning limit system parameter to something less than 1.5.

Conditions of occurrence: None.

Impact: Low

Customer circumvention: To understand the overprovisioning ratios and the limitations, read the overprovisioning sections in the CLI Admin Guide.

Customer recovery steps: To understand the overprovisioning ratios and the limitations, read the overprovisioning sections in the CLI Admin Guide.

Issue ID: 213124

Issue summary: Unexpected array restart due to out of order snapshot removal.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description: Snapshot removal does not succeed, followed by error messages.

Symptoms: Error messages on multiple controller nodes.

Conditions of occurrence: This issue is caused by an out of order snapshot removal situation. Snap removal produces an error, and child snap gets removed causing an out of order removal.

Impact: High

Customer circumvention: Remove snapshots individually rather than through a group command.

Customer recovery steps: Restart the controller node with the incomplete snapshot removal.

Issue ID: 213876

Issue summary: Deduplication group in `not_started` state due to metadata inconsistency.

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

Affected software versions: 3.3.1.GA, 3.3.1.MU1

Issue description: **`admck`** and **`ddcscan`**make a false assumption that if the TDVV3 DDS has a region count as zero, the exceptions will not have ext key.

Symptoms: DDS and TDVVs are in `not_started, needs_check` state.

Conditions of occurrence: System does the powerfail at the time when odd number of controller nodes are online.

Impact: Medium

Customer circumvention: None.

Customer recovery steps:**`checkvv`** brings the VVs to `normal` state.

Issue ID: 214117

Issue summary: Extend `MC_CHECK_CONSISTENCY tattler` timeout.

Affected platforms: All StoreServ

Affected software versions: 3.3.1.MU2

Issue description: Extend `MC_CHECK_CONSISTENCY tattler` timeout so it does not show up as an alert.

Symptoms: See a tattler alert about the `MC_CHECK_CONSISTENCY`.

Conditions of occurrence: See a tattler alert about the `MC_CHECK_CONSISTENCY`.

Impact: Low

Customer circumvention: Ignore it.

Customer recovery steps: Remove the alert.

Issue ID: 214392

Issue summary: Report Link Status completion error interrupt.

Affected platforms: StoreServ 20000

Affected software versions: 3.3.1

Issue description: Fiber Channel over Ethernet driver gets 2 RLS (report link status) completion error interrupt for mailbox command at the same time from the HBA firmware. During the first IOCB response handling, the driver releases the allocated `rls_buf` and marks it as NULL.

During the second interrupt handling, the driver again attempts to release the already released buffer. There is no check for NULL buffer.

This causes the controller node to restart.

Symptoms: Controller node restart.

Conditions of occurrence: This can occur during normal operation.

Impact: Medium

Customer circumvention: Controller node restart.

Customer recovery steps: Controller node restart.

Issue ID: 214447

Issue summary:

Port 1:2:1 loops on `LINK_LOGIN_WAIT` state. The Loop is down, link is not in a ready state, and **`portdb`** cannot complete for port 1:2:1. The message is `hba_state was REDO_PORT_DB`.

Affected platforms: StoreServ 7000 series

Affected software versions: 3.2.1, 3.2.2, 3.3.1

Issue description: The port gets Port Login requests from the initiator. In response, the driver keeps adding notify ack IOCBs to the request queue to send to firmware. After some time, the request queue becomes full, as IOCBs cannot be sent to FW successfully. The driver starts a one second periodic timer to check the queue depth. This timer keeps on looping, as the link is not ready, until the controller node runs out of memory due of excessive file cache usage, or due to excessive log messages.

Symptoms: Host path is lost.

Conditions of occurrence: Sluggish port.

Impact: Low

Customer circumvention: Port reset.

Customer recovery steps: Port reset.

Issue ID: 215046

Issue summary: Eliminate temp close step in **`updatevv`** .

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1 GA - MU2

Issue description: **`Updatevv`** has long block times if it is performed on a VV set with thousands of snapshots. This process uses a different internal approach that should be much faster and scale better when using this level of snapshots.

Symptoms: Long block times.

Conditions of occurrence: Perform **`updatevv`** on thousands of snapshots at the same time.

Impact: Low

Customer circumvention: Use smaller VV sets.

Customer recovery steps: Reconnect their host.

Issue ID: 215184

Issue summary: Cage reports `Failed, no ESI port`. The backend ports connect to the cage doing linkup/linkdown flipping frequently.

Affected platforms: StoreServ 10000

Affected software versions: 3.2.2, 3.3.1

Issue description: Issue is related to misbehaving drive which kept bringing the entire loop down.

Symptoms: The backend ports connected to the cage will have linkup/linkdown flipping frequently.

Conditions of occurrence: Class setup with FC backend. Port Login fails to a drive.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 215766

Issue summary: Prevent the array from becoming administratively sluggish when internal table backups are overextended.

Affected platforms: All StoreServ

Affected software versions: All

Issue description: Corrects an issue where the array’s internal table data backup mechanism is overextended which can lead to the array becoming administratively sluggish and could lead to data becoming unavailable.

Symptoms: CLI commands on the array become sluggish.

Event log entries are not logged.

Conditions of occurrence: The array's internal mechanism for backing up table data becomes overextended and slow to respond.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Restart of the cluster master controller node should restore system response in most cases.

Issue ID: 215892

Issue summary: Internal counter is unable to drain during controller node up process. This issue prevents the controller node from joining the cluster.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU1

Issue description: Counter is incremented too soon before the log page processes. As a result, the controller node fails to join the cluster because vv block does not succeed during controller node up process.

Symptoms: Controller node cannot join the cluster. Online upgrade cannot complete.

Conditions of occurrence: Online upgrade.

vv conversion/online vv copy.

Impact: Medium

Customer circumvention: Do not do online upgrade to 3.3.1.MU1 from a lower version.

Customer recovery steps: Reduce/stop I/O until the controller node joins.

Issue ID: 215894

Issue summary: Freeing the same log page on the remote L2 bmap location multiple times.

Affected platforms: All StoreServ

Affected software versions: 3.2.2.GA-MU4, 3.3.1 GA, 3.3.1 MU1

Issue description: Due to the incorrect interpretation of the return status of the remote bmap, the code attempts to free the same page multiple times leading to the metadata inconsistency or single controller node restart.

Symptoms: Inconsistent VV.

Conditions of occurrence: Freeing up the space from a thinly provisioned VV.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: During single controller node restart, the condition gets cleared automatically.

If reported metadata inconsistency on the VV, run **`checkvv`** manually to clear this condition.

Issue ID: 216556

Issue summary: Prevents controller node restarts when a defrag is attempted on a deduplicated page due to race condition between defrag and deduplication IO path for tdvv1 volumes.

Affected platforms: StoreServ 7000

Affected software versions: 3.2.1 MU1 - MU5, 3.2.2, 3.3.1 GA - MU1

Issue description: Prevents controller node restarts due to defrag being attempted on a deduplicated page due to race condition between defrag and deduplication IO path for tdvv1 volumes.

Symptoms: Controller restart.

Conditions of occurrence: Presence of tdvv1 volumes.

Impact: High

Customer circumvention: None.

Customer recovery steps: The system recovers automatically after controller node restart.

Issue ID: 216866

Issue summary: Syslog mirroring does not work for UDP servers.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU1

Issue description: In 3.3.1 MU1, HPE 3PAR added the functionality to allow up to 3 Syslog server to receive the same (mirrored) events. This new functionality did not support UDP Syslog servers. If multiple UDP Syslog servers are configured, only the first server will be connected.

Symptoms: Configuring Multiple UDP syslog servers will result in a connection to only one.

Conditions of occurrence: Configuring Multiple UDP syslog servers.

Impact: Low

Customer circumvention: Use multiple TCP or TLS syslog servers instead of UDP.

Customer recovery steps: None.

Issue ID: 217322

Issue summary: Leaked cache page under certain conditions.

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

Affected software versions: 3.3.1

Issue description: Under certain circumstances, systems using compressed or DECO volumes can run into an exhaustion of cache memory pages in processor memory. This occurs when defragmentation IO launched by the system causes the processor memory cache pages to be leaked.

Symptoms: Single controller node error message due to resource exhaustion of cache pages, or a leaked cache page being accessed for a host IO.

Conditions of occurrence: This is a race condition between defragmentation IO launched by the system and host IO.

Impact: High

Customer circumvention: Update the system to 3.3.1 MU3.

Customer recovery steps: None.

Issue ID: 217485, 230376

Issue summary: HPE Power Supplies sometimes report incorrect temperature values for sensors.

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

Affected software versions: 3.2.2 GA - MU6, 3.3.1 GA - MU2

Issue description: HPE Power Supplies sometimes report incorrect temperature values for its sensors. Those erroneous temperature values, if above the critical threshold defined by the drive enclosure firmware, will cause the drive enclosure to shutdown.

Symptoms: Enclosure shutdown after reporting critical temperature values for one or more power supplies.

Conditions of occurrence: All DCS5, DCS6, DCS9, and DCS10 drive enclosures are susceptible to this issue.

Impact: High

Customer circumvention: None.

Customer recovery steps: Restore the power to the enclosure when it happens. If this happens too often with a specific enclosure, contact HPE support to have the drive enclosure power supplies replaced.

Issue ID: 217532, 221700

Issue summary: If the Service Console is restarted in a situation where the array is experiencing a flood of events, local notification of alerts may be delayed.

Affected platforms: All StoreServ

Affected software versions: 3.3.1.GA-MU2

Issue description: In instances where there is a flood of events on the array, the service processor may miss notification on a set of those events. Notifications may also be delayed.

Symptoms: Delayed local notification of events or alerts from the service processor.

Conditions of occurrence: A high throughput of events being generated by the array.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Wait for the service processor to process the alerts.

Issue ID: 217854

Issue summary: An uncontrolled shutdown due to double decrements, after a compressed volume has already been detected with metadata inconsistency.

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

Affected software versions: 3.3.1 GA - MU2

Issue description: After a compressed volume reports metadata inconsistency, an uncontrolled shutdown will not happen.

Symptoms: A controller node restarts unexpectedly.

Conditions of occurrence: A compressed volume is running on the array.

A metadata inconsistency is detected on the compressed volume.

A read ahead request is issued, where the offset with the metadata inconsistency being the starting offset.

Impact: High

Customer circumvention: Install GA patch 20 or MU1 patch 21.

Customer recovery steps: No extra recovery is needed after the uncontrolled shutdown.

Issue ID: 218415

Issue summary: Compressed volumes may experience metadata inconsistency issue under some circumstances with certain data patterns.

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

Affected software versions: 3.3.1

Issue description: Under some circumstances with certain data patterns, compressed volumes, including those with deduplication location attribute, can encounter double allocation issues. The double allocation results in two different offsets using the same backing page, with the data of one of those offsets overwriting the data of the other. This can result in SD metadata inconsistency and decompression errors.

Symptoms: Volumes with Compression attribute will be reported in `sd_needs_fix` state. You may see an alert indicating that one or more volumes are in this state.

Conditions of occurrence: This is a small window when Compression Garbage Collection process and host IO are processing the same page that was previously uncompressed, and the new IO data becomes compressible.

Impact: High

Customer circumvention: None. The system may run out of space if Compression Garbage Collection is not used.

Customer recovery steps: Run **`checkvv -fixsd`** on affected volumes to make them normal again. You may also restore from other copies of data on or off the array.

Issue ID: 219773

Issue summary: Controller nodes go down when running the SD metadata checker.

Affected platforms: All StoreServ

Affected software versions: 3.3.1.GA-MU2

Issue description: `sdmetack` kernel dataless write path does not correctly handle translation table entries marked bad due to decompression error in partial page write.

Symptoms: Error message in `excp_get_new_tag`.

Conditions of occurrence: This requires another issue to get the virtual volume into the SD meta corrupt state. If any of the translation table entries are marked bad in the partial page write path, the `sdmetack` dataless write will cause an error message when it hits those entries.

Impact: High

Customer circumvention: Do not run `sdmetack` on affected virtual volumes without this patch.

Customer recovery steps: Move the `sdmetack` binary. This patch also includes improving the `disable_sdmetack mvar`, so that can be used instead for future issues.

Issue ID: 220440

Issue summary: Host does not log back in to 3PAR target port after the **`controlport`** reset command on target port.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU2

Issue description: When switch port is hard coded to 16G and 3PAR port is reset, I/O running on HP-UX Server hangs. Server sends ABTS (Aborts) to 3PAR and 3PAR does not respond.

Symptoms: Host I/O timeout. **`showport`** shows port type as `free` instead of `host`.

Conditions of occurrence: Switch port locked to 16G and 3PAR port reset.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Restart Server or reset Server HBA port(s).

Issue ID: 220731

Issue summary: Spare drive part number information is showing up as Unknown.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU1 - MU2

Issue description: The spares file /opt/tpd/spn/drive.cfg is updated when new drives are released.

Symptoms: If the part number file is outdated, alerts may display `UNKNOWN` in the Spare_PN field.

Conditions of occurrence: Systems having one of the following drive models:

SSKB0600S5xeN010

SSKB1200S5xeF010

SSKB1200S5xeN010

SSKB1800S5xeN010

HCFP0600S5xeN010

HCFP1200S5xeF010

HCFP1200S5xeN010

HCFP1800S5xeN010

Impact: Medium

Customer circumvention: Patch with new spares file to address issue is advised.

Customer recovery steps: Replace the drive.

Issue ID: 220823

Issue summary: Too many prints while running **`admck`** in certain scenario can fill up internal system space.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1

Issue description: Too many prints while running **`admck`** in certain scenario can fill up system space needing to manually free up the space.

Symptoms: **`defrag`** logs fill up the internal space resulting in system running out of space.

Conditions of occurrence: Presence of thin provisioned volumes.

Impact: Low

Customer circumvention: Create snapshots or run **`checkvv`** offline to throttle the prints, resulting in less internal storage space consumption.

Customer recovery steps: Manually free up space.

Issue ID: 221491

Issue summary: Too much input and output coming in causes VV access loss.

Affected platforms: All StoreServ

Affected software versions: 3.1.3, 3.2.1, 3.2.2, 3.3.1 GA - MU1

Issue description: Hitting one system internal resource limit causes a leak of another.

Symptoms: Input/output stalls. **`statcmp`** displays credit usage as a high number which does not go down when under no input/output.

Conditions of occurrence: The system is overused.

Impact: Medium

Customer circumvention: Monitor the resource usage and do not let it hit its maximum. **`statcmp`** can be used to monitor resources.

Customer recovery steps: If the system is under a light load and resources are not returned, restart the affected controller nodes one at a time.

Issue ID: 221541

Issue summary: Generate a new alert to differentiate customer installable patch upgrade and complete OS upgrade.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU2

Issue description: Patch updates may be installed by customers.

Symptoms: All updates generate the same alert indicating that an update is available and indicate that proper resolution is to contact the authorized service provider. However, patch updates are able to be installed directly by the customer.

Conditions of occurrence: A patch update is staged to the StoreServ.

Impact: Low

Customer circumvention: While customer self-install of patches is now recommended, it is still optional.

Customer recovery steps: None.

Issue ID: 221856

Issue summary: The final step of a snapshot removal is completed while the configuration lock is held in System Manager. The IOCTLs for this step are sent to each controller node sequentially, causing extra time spent in this stage and possible timeout of other processes in System Manager.

Affected platforms: All StoreServ

Affected software versions: 3.3.1.GA-MU2

Issue description: When snapshots are removed, the IOCTLs for the final stage will no longer be sent sequentially.

Symptoms: Slow snapshot removal.

Conditions of occurrence: Removing snapshot on systems with large cache.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Stop Remote Copy groups and allow the snapshot removal to complete prior to restarting remote copy. Do not create any snapshots and have minimal other items running on the array.

Issue ID: 222014

Issue summary: Events reporting correctable Control Cache DIMM errors identify the wrong DIMM.

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

Affected software versions: 3.2.2, 3.3.1

Issue description: When correctable DIMM errors cross a defined threshold, an event is generated so that the DIMM can be replaced. Under some circumstances, the wrong DIMM is identified in the event.

Symptoms: A DIMM is replaced, but the system continues to report correctable errors against the same DIMM number.

Conditions of occurrence: Correctable error threshold exceeded when the DIMM number is greater than CC_0.1.0. Data cache DIMMs are not affected.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 222387

Issue summary: **`System Manager`** displays the message `System is up and running`. Some physical disks still did not have a valid cage position assigned to them.

Starting with 3.3.1MU3, **`showsysmgr`** will print the following message when a PD without cage position is detected:

```
System has started, PD cage position refresh is pending.
```

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU2

Issue description: The **`System Manager`** might start and become operational while the physical disk cage position assignment is still ongoing. If the drives do not have a proper cage position, **`createcpg`**, among other commands, might report no space available in the array.

Symptoms: **`showpd`** shows drive without cage position, or with a question mark (?) besides the cage position field.

`createcpg` might issue the following message:

```
createcpg: Error: no available space for given
```

Conditions of occurrence: After wiping and servicing magazines, the physical disk might become online without a cage position for a short time. This issue is addressed in the next System Manager/enclmgmt refresh cycle (2 minutes).

Impact: Low

Customer circumvention: Wait for the physical disks to have the cage position assigned before executing command lines such as **`createcpg`**.

Customer recovery steps: None.

Issue ID: 222641

Issue summary: DNS servers do not get written out to `/etc/resolv.conf` when no good servers are accessible.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1 GA - MU2

Issue description: DNS servers do not get written out to `/etc/resolv.conf` when no good servers are available.

Symptoms: DNS resolution does not work.

Conditions of occurrence: None of the DNS servers are reachable by the InServ.

Impact: High

Customer circumvention: Ensure that at least one DNS server specified does work.

Customer recovery steps: Contact HPE support to add the DNS server manually.

Issue ID: 222817

Issue summary: RO snapshots of compressed volumes experience SD metadata inconsistency.

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

Affected software versions: 3.3.1 GA, 3.3.1 MU1, 3.3.1 MU2

Issue description: Under certain circumstances, RO snapshots of compressed volumes (compressed or compressed/deduplicated) can encounter sd metadata inconsistencies. This occurs due to a race condition between the compression garbage collector and deletion of non-leaf RO snapshots.

Symptoms: RO snapshots of compressed/deco volumes will be reported in `sd_needs_fix` state, or the customers may see an alert indicating that one or more volumes are in this state.

Conditions of occurrence: A race condition between the compression garbage collector and deletion of non-leaf RO snapshots.

Impact: High

Customer circumvention: None.

Customer recovery steps: Run **`checkvv -fixsd`** on affected volumes to make them normal again.

Issue ID: 222954

Issue summary: Lightweight Directory Access Protocol (LDAP) requires a CA certificate for connecting to domain controllers when using SSL/TLS with a simple binding configuration.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU2

Issue description: Lightweight Directory Access Protocol now requires a CA certificate to be imported to the array to prevent man-in-the-middle attacks. This requirement allows the connection to succeed for SSL/TLS using simple binding while protecting the system from malicious attacks.

Symptoms: You will not be able to authenticate against domain controllers using SSL/TLS when simple binding is defined.

Conditions of occurrence: SSL/TLS and simple binding must be defined in the authorization parameters or StartTLS must be defined.

Impact: Medium

Customer circumvention: Use the alternative binding configuration, Simple Authentication and Security Layer, if you do not wish to import a CA certificate. Because Simple Authentication and Security Layer uses Kerberos, man-in-the-middle attacks are not a concern when using this binding.

Customer recovery steps: Redefine the Lightweight Directory Access Protocol authorization parameters, specifically the binding type. When using simple binding, if you do not wish to import a CA certificate for Lightweight Directory Access Protocol, change the authorization parameters to use Simple Authentication and Security Layer binding. If TLS is used with Simple Authentication and Security Layer, then a CA certificate must still be imported. SSL is not valid with the Simple Authentication and Security Layer binding type.

Issue ID: 223254

Issue summary: Running **sdmetack** on compressed virtual volumes consumes CPU resources. In some cases, no input/output operations are possible.

Affected platforms: All StoreServ

Affected software versions: 3.3.1

Issue description: **sdmetack** dataless writes consume too many CPU cycles. This consumes the host IO.

Symptoms: Few input/output operations per second, poor performance while **sdmetack** is running.

Conditions of occurrence: Running **sdmetack** directly, or **checkvv -fixsd**.

Impact: High

Customer circumvention: Do not run **sdmetack** manually without this patch.

Customer recovery steps: Stop **sdmetack** processes and move the binary. After 3.3.1.MU2 is installed, **disable_sdmetack mvar** can be used.

Issue ID: 223365

Issue summary: Single controller node will not start after clean shutdown when the second controller node has a bad voltage regulator.

Affected platforms: StoreServ 7000

Affected software versions: 3.2.1 MU3 - MU5, 3.2.2, 3.3.1 GA - MU1

Issue description: After properly shutting down the system, if a power regulator issue prevents a controller node from starting, the system will hang because it is waiting for the missing controller node to start. This causes data unavailability and customer or support intervention.

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

On a 4 or 8 node system, after a proper shutdown, the array will state that it is waiting for a missing controller node, but will start after 10 minutes.

Conditions of occurrence: Shutdown array. Start array; all but one controller node starts.

Impact: High

Customer circumvention: None.

Customer recovery steps: Use the CLI command, **setsysmgr tocgen**.

Issue ID: 224547

Issue summary: Checks added to Target Port Group ID (TPGID) changes for Remote Copy volumes to minimize invalid values that result in data unavailable scenarios.

Affected platforms: All StoreServ

Affected software versions: 3.1.3 GA - 3.3.1 MU2

Issue description: The CLI command **`setvv -settpgid`** will validate the Target Port Group ID value specified for Remote Copy volumes to insure the primary and secondary systems have a different value.

Symptoms: The primary and secondary systems to not have correct Target Port Group ID values that results in the host unable to access (R/W) Remote Copy volumes.

Conditions of occurrence: The Target Port Group ID value on the primary and secondary system for all the volumes in a remote copy group are the same.

Impact: High

Customer circumvention: Do not change the Target Port Group ID of volumes in remote copy groups to values other than 257 or 258. The value on the primary and secondary systems must be different.

Customer recovery steps: Use the CLI **`setvv -settpgid`** command to set the Target Port Group ID value on volumes in a remote copy group to either 257 or 258 as appropriate. All the volumes within a Remote Copy group must have the same value and must be different than the value being used by the remote system.

Issue ID: 224575

Issue summary: Controller node restart caused by array overflow.

Affected platforms: StoreServ 20000

Affected software versions: 3.2.2, 3.3.1

Issue description: Controller node restart caused by array overflow.

Symptoms: Controller node restart.

Conditions of occurrence: Number of initiators exceeding 256.

Impact: High

Customer circumvention: Do not exceed a maximum of 256 initiators per target port.

Customer recovery steps: Restart the controller node.

Issue ID: 224637

Issue summary: Unwanted increase in memory usage due to not releasing the dynamically allocated memory.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU1, 3.3.1 GA

Issue description: An unwanted increase in memory usage occurs, due to the dynamically allocated memory not being set free through **`tpd_vv_alloc( ) interface`**.

Symptoms: The system is out of memory after a few months.

Conditions of occurrence: Compression VVs are deployed.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: A controller node restart clears the out of memory situation.

Issue ID: 224861

Issue summary: List of available initiators for target-driven peer zoning is being limited to 512.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA, 3.3.1 MU1, 3.3.1 MU2

Issue description: If there are more than 512 initiators registered in the fabric, the reported initiator list shows only 512 initiators. This list is used by SSMC to indicate the initiators available for target-driven peer zoning. All initiators would not be able to be used by SSMC for target-driven peer zoning.

Symptoms: Only 512 of the initiators are displayed in the initiator list.

Conditions of occurrence: If there are more than 512 initiators, only 512 are shown, rather than of all the initiators.

Impact: High

Customer circumvention: If there were more than 512 initiators, the customer would have to get the list of available initiators from the switch and then use 3PAR CLI commands to create/update the target-driven peer zones that include the initiators which were not being returned in the initiator list.

Customer recovery steps: None.

Issue ID: 225198

Issue summary: System Manager error in `vvol_priority.c:483` when domain-based QoS rule is in place.

Affected platforms: All StoreServ

Affected software versions: All

Issue description: An uninitialized variable used in processing QoS domain rules can lead to repeated System Manager shutdown or System Manager memory corruption.

Symptoms: Recursive System Manager shutdown.

Conditions of occurrence: There is an enabled QoS domain: rule.

Impact: High

Customer circumvention: Delete or clear QoS domain rules.

If the highest QoS rule ID is less than 4096, delete and recreate any QoS domain: rule that has a lower ID than the lowest vvset: rule. Both vvset: and domain: QoS rules can then be used and enabled.

Customer recovery steps: Contact HPE support.

Issue ID: 225657

Issue summary: Hash function leads to poor TOC performance when large fully provisioned volumes are used.

Affected platforms: All StoreServ

Affected software versions: All

Issue description: On systems that use a large total size of fully provisioned volumes, TOC performance during checkpoints degrade significantly. This degradation is due to individual markers describing the virtual volumes to LD mappings for a single virtual volumes, and being placed into the same bucket.

Symptoms: Periodic long delays when system state changes, potentially resulting in host timeouts.

Conditions of occurrence: Multiple large fully provisioned volumes must be present. Frequent snapshots or Remote Copy will cause the delay.

Impact: High

Customer circumvention: Move to thin provisioned volumes.

Customer recovery steps: Contact HPE Support.

Issue ID: 225480, 219612

Issue summary: Remote Copy, which requires a lot of CPU resources, times out.

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

Affected software versions: 3.3.1

Issue description: The **compr_gc** process uses many CPU resources during IO runs. This creates a lot of issues in case of working with Remote Copy, which also requires a lot of CPU resources.

Symptoms: **compr_gc** consumes CPUs resources, which blocks other tasks.

Conditions of occurrence: The periodic **compr_gc** process runs.

Impact: High

Customer circumvention: Replace the **compr GC** atomic operations with local variables and combine into one single atomic operation. In this way, we reduce the number of atomic operations in **compr gc** to a great extent and CPU utilization is reduced.

Customer recovery steps: None.

Issue ID: 225981

Issue summary: The message indicating a downgrade failed includes commands that you cannot run.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU2

Issue description: This will only happen during downgrades. The message indicating a downgrade failed includes CLI commands that you cannot run. Those commands have been replaced with usable commands.

Symptoms: Downgrades from 3.3.1.MU2 and later may fail due to unsupported RAID6 layouts.

Conditions of occurrence: Logical disks created with new RAID6 layouts (added in 3.3.1.MU2) exist, and the system is being downgraded to a release prior to 3.3.1.MU2.

Impact: Medium

Customer circumvention: Use the system parameter (in the downgrade error message) to indicate the layouts that were added in 3.3.1.MU2 are not to be used. Then tune the offending logical disks to use the older layouts.

Customer recovery steps: None.

Issue ID: 227577

Issue summary: Deduplication Stores (DDS) regions count might be incorrectly set based on Snapshot data logical disks (LD) node scan count. It must be set based on the Snapshot Admin (SA) logical disks node scan count.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA, 3.3.1 MU1, 3.3.1 MU2

Issue description: During DDS creation, InformOS scans the SD LD controller nodes. This is used to set the DDS regions number. The regions number is used to split the SA metadata space ownership between controller nodes based on regions where each region is owned by different controller nodes. Since this is SA metadata space, SAs should be used instead of SD LDs.

Symptoms: DDS regions might be set improperly resulting in degraded performance for DDS SA metadata Input/Output.

Conditions of occurrence: Every time DDS is created. In most cases this will not result in DDS regions being set improperly since SD LDs scan and SA LDs scan will result in the same number of controller nodes. It is incorrect to use the SD LD node scan.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: After DDS is created, there is no way to reset the DDS regions number. If it has been improperly set, InformOS will set DDS regions number to 0. This will be the DDS regions number until DDS is removed.

Issue ID: 228280

Issue summary: A controller node restarts when inconsistent metadata entries for compression volumes are encountered.

Affected platforms: StoreServ 8000, StoreServ 20000 R2

Affected software versions: 3.3.1 GA, 3.3.1 MU1, 3.3.1 MU2

Issue description: A controller node restarts when inconsistent meta data entries for compression volumes are encountered.

Symptoms: All or one controller restart.

Conditions of occurrence: Presence of compression volumes.

Impact: High

Customer circumvention: None.

Customer recovery steps: The array recovers automatically after controller node/array restart.

Issue ID: 228404

Issue summary: No logical disk ownership for flashcache for system with more than four controller nodes that has flashcache unit carry over from 3.2.1 code after the cluster had four different controller nodes restart in the past.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 GA - MU6, 3.3.1 GA - MU2

Issue description: No logical disk ownership for system with more than four controller nodes that created flashcache unit in 3.2.1. This will occur after four different controller nodes restart after the upgrade to 3.2.2 or any newer code base.

Symptoms: No logical disk ownership error for system with more than four controller nodes with flashcache logical disks.

Conditions of occurrence: System with more than four controller nodes that created flashcache unit in 3.2.1. This outage occurs after four different controller nodes restart after the upgrade to 3.2.2 or any newer code base.

Impact: High

Customer circumvention: Remove flash cache logical disk (fcld). and recreate them.

Customer recovery steps: None. Once cluster restart in 3.2.2 or newer code, the fcld will never hit this issue again.

Issue ID:228605

Issue summary: **`admithw`** reports HEWLETT-PACKARD:691970-001 as unsupported cables after an OS update.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description: When updating TPD from 3.2.2.x to 3.3.1.x, you will get some transient messages in **`admithw/checkhealth`** or even events/alerts in the eventlog regarding unqualified cables. These can be ignored as they are transient states when going from 3.2.2.x to 3.3.1.x. TPD version 3.2.2.x does not have the fix for unqualified cables that 3.3.1.x has.

Symptoms: When CLI Client is 3.2.2.476 and the rest of TPD is 3.3.1.337, you will notice unqualified cables messages in **`admithw`**/**`checkhealth`** or even events/alerts in the eventlog.

Conditions of occurrence:CLI Client version is mismatched with the rest of TPD versions.

Impact: Medium

Customer circumvention: Ignore the unqualified cables messages when updating TPD from 3.2.2.x to 3.3.1.x.

Customer recovery steps: Use **`showversion -b`** to check that CLI Server, CLI Client, System Manager, Kernel, TPD Kernel Code are all the same TPD version.

Issue ID: 229417

Issue summary: When a snapshot removal fails in a vv family for some reason, and another snap of the same family is tried next for removal, it could lead to out-of-order removal scenario resulting in controller node restart.

Affected platforms: All StoreServ

Affected software versions: 3.2.1 GA - MU5, 3.2.2 GA - MU6, 3.3.1 GA-3.3.1 MU1

Issue description: When a snapshot removal fails in a vv family for some reason, and another snap of the same family is tried next for removal it could lead to out-of-order removal scenario resulting in controller node restart.

Symptoms: Controller node restart.

Conditions of occurrence: Presence of thin provisioned volumes and two or more snapshots in the vv family.

Impact: Medium.

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 229744

Issue summary: TOC quorum loss when replacing disks without dismissing old ones.

Affected platforms: All StoreServ

Affected software versions: All

Issue description: If sufficient quantities of disks are replaced without dismissing the old ones, the system mistakenly acts as if the removed disks should still be present and treats it as if it has lost TOC quorum, forcing a system shutdown. It has been updated to recognize that the new disks are there in place of the old ones, and that the old ones should no longer be considered for quorum calculations.

Symptoms: After multiple disks are been replaced, removing an additional one causes the system to take a powerfail save.

The system is unable to start up after a powerfail save, indicating that only half of the expected TOC disks minus one are present, despite powercycling cages.

Conditions of occurrence: Multiple physical disks removed from the system and new ones put in their place, without using proper procedures. The **`servicemag resume`** command will ensure that the old physical disks are dismissed.

Impact: High

Customer circumvention: Always use **`servicemag`** to replace disks. If the old physical disks still show after it was believed that **`servicemag`** was used correctly, use **`dismisspd`** to dismiss the old physical disks.

Customer recovery steps: None.

Issue ID: 229824

Issue summary: When using LDAP, if you login with the format **`Domain\username`** and the kerberos-realm is not set, the authentication will fail.

Affected platforms: All StoreServ

Affected software versions: 3.2.1 MU2 - MU5, 3.2.2, 3.3.1 GA - MU2

Issue description: **`Auth_helper`** tries to find a valid kerberos realm, and append the given prefix to it. Because the kerberos realm is not set, it causes a segmentation fault when trying to parse the whole kerberos realm.

Symptoms: Denial of authentication.

`Auth_helper` coredumps found on active controller node

Conditions of occurrence: Kerberos-realm authorization parameter is not set.

Impact: Low

Customer circumvention: Define the kerberos-realm authorization parameter.

Customer recovery steps: None.

Issue ID: 229998

Issue summary: Can't run **`checkvv -y`** due to `does not need to be fixed` checkvv error.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU2

Issue description: Under certain powerfail recovery scenarios involving compression, a compressed VV can get into an internal state where it can not be restarted by checkvv -y.

Symptoms: **`showvv`** shows the compressed VV as `not_started,snapdata_invalid,sd_needs_fix`.

`Checkvv -y` fails to restart the VV reporting `<vv_name> does not need to be fixed`.

Conditions of occurrence: All controller nodes have errors, leading to powerfail recovery involving a compressed virtual volume.

Impact: High

Customer circumvention: None.

Customer recovery steps: Contact HPE support.

Issue ID: 230995

Issue summary: Legacy serial number was set to non-zero leading to serial number mismatch across the cluster controller nodes.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description: Legacy serial number was set to non-zero for the systems configured with 10 digit HP style serial number.

Symptoms: Legacy serial number is set to non-zero.

Serial number mismatch across the cluster controller nodes.

Controller nodes might not join the cluster on restart or during controller node replacement due to serial number mismatch.

Conditions of occurrence: Legacy serial number is set to non-zero. On restart or during controller node replacement the controller nodes might not join the cluster.

Impact: Medium

Customer circumvention: The legacy serial number should be set to zero on controller nodes configured with 10 digit HP style serial number.

Customer recovery steps: Reset the legacy serial number to zero if the system is configured with 10 digit HP style serial number.

Issue ID: 231311

Issue summary: Debian packages out of date with security vulnerabilities.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA, 3.3.1 MU1

Issue description: Debian packages out of date with security vulnerabilities.

Symptoms: Security scan reports vulnerabilities.

Conditions of occurrence:

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 231985, 231996

Issue summary: The VASA Provider allows connections from clients using TLS/SSL methods other than TLSv1.2.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 GA - MU6, 3.3.1 GA - MU1

Issue description: The VASA Provider negotiates with incoming client connections to determine the highest version TLS/SSL method mutually supported by the client and the VASA Provider. This means clients can connect to the VASA Provider using TLS/SSL methods other than TLSv1.2.

Symptoms: None.

Conditions of occurrence: None.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 232561

Issue summary: Failure during the flash cache LD removal occurs when removing the Adaptive Flash Cache while data movement operations are active.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1

Issue description: When removing the Adaptive Flash Cache while data movement operations are active, failure during the flash cache LD removal can occur. This failure will orphan the LDs and leave them in a bad ownership state. The bad state can lead to a controller node terminating unexpectedly.

Symptoms: Uncontrolled shutdown.

Orphaned fcachelds - unmapped LDs.

Conditions of occurrence: Removal of flashcache during data movement operations ((tunesys, updatevv, promotevv, etc...).

Impact: High

Customer circumvention: Avoid removing the flash cache during the following:

updatevv

tunevv (convertvv and converttpvv,) onlinecopy, onlinepromote/conversion, imports

promotevv -online

createvv -onlines

convert

region moval task --- <vvmap>

tunevv/ld/sys/node..., compactcpg

ao <region move>

rebalance ld (controller node comes back up <nup>)

relocation <l_refcnt , chunklet of ld relocating ...disk failure> , drive coming back for (playback or raidset state change)

writeback cache in single controller node

Customer recovery steps: Any orphaned LDs after the uncontrolled shutdown can be removed.

Issue ID: 232671

A race condition between controller node down recovery and new coming IO could could cause `data unavailable`.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU2

Issue description: If we get a single controller node error message, the replicant controller node could potentially hit the issue during controller node down recovery which could cause`data unavailable`.

Symptoms: Data is unavailable.

Conditions of occurrence: Online upgrade.

Impact: Medium

Customer circumvention: Reduce the number of VVs in the Remote Copy groups, and reduce the number of Remote Copy groups.

Customer recovery steps: Wait for the cluster to panic three times and it will recover.

Issue ID: 235136, 237643

Issue summary: Log entry reporting bogus IO abort time out.

Affected platforms: StoreServ 8000, StoreServ 9000, StoreServ 20000

Affected software versions: 3.3.1 MU1 - MU2

Issue description: Log entry reporting bogus IO abort time out.

Symptoms: FC port is in an unusable state, and controller node reboot times out.

Conditions of occurrence: Command stuck in driver layer in DATA_OUT phase exceeding abort timeout.

Impact: Low

Customer circumvention: None.

Customer recovery steps: Reboot the controller node.

Issue ID: 235189

Issue summary: Controller node restart because of lock order not maintained.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 MU4, 3.3.1 GA - MU2

Issue description: Controller node restart because of lock order not maintained.

Symptoms: Controller node restart.

Conditions of occurrence: Order of lock is not maintained.

Impact: Medium

Customer circumvention: Controller node restart.

Customer recovery steps: Controller node restart.

Issue ID: 235834

Issue summary: Certain drive models may lose writes which occurred prior to a self initiated reset.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 MU4 - MU6, 3.3.1 GA - MU2

Issue description: Certain drive models are marked as failed to avoid data inconsistency due to dropped writes when drives undergo self reset.

Symptoms: Can lead to data inconsistency.

Conditions of occurrence: Drive resets may occur due to issues within the drive firmware.

Impact: High

Customer circumvention: None.

Customer recovery steps: Use **`checkld`**, **`checkvv`**, or any host based consistency check if a physical disk has failed with a detailed state containing `Miscompare`.

Issue ID: 236118

Issue summary: Performing consistency checks during VV removal can cause unexpected single controller node restart.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU1

Issue description: When VVs are being removed at the same time that consistency checks are being performed, single controller node error messages can occur if the consistency check is run after the system has already removed a volume and before the volume entry is removed from the VV table that is used by the consistency check call.

This issue can occur when volumes are being converted to compressed volumes if a consistency check is issued when the volume that was converted is being removed. The consistency check can be called by either the Service Processor or any other process that issues a `checkhealth` consistency call. The Service Processor performs a consistency check on an hourly basis.

Symptoms: Single controller node restart with `Fatal exception` error.

Conditions of occurrence: Running **`checkhealth -svc -full consistency`** or **`checkhealth -svc -full`** while a VV is being removed.

Impact: High

Customer circumvention: The SP runs **`checkhealth -svc -full consistency`** as part of status collection. The SP can be stopped, but this means none of the other SP functions are active.

If the SP is stopped, avoid running **`checkhealth -svc -full consistency`** or **`checkhealth -svc -full`** as a task and interactively.

Customer recovery steps: The controller node that restarted will automatically rejoin the cluster.

Tasks that were running, like vv conversions, may need to be restarted.

Issue ID: 236286

Issue summary: During snapshot processing between controller nodes, a deadlock condition can be encountered. This can result in the host I/O being delayed and the array may become unresponsive.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1 MU1 - MU3

Issue description: The ioctl call to start the VV does not complete due to the collision with the `sdt GET_VVMAP` command.

Symptoms: Stuck ioctl

Conditions of occurrence: Concurrence of the scsi **`GET_VVMAP`** command and snapshot creation.

Impact: High

Customer circumvention: None.

Customer recovery steps: Controller node reboot.

Issue ID: 237098

Issue summary: VVol snapshotting takes around 8s leading to long stun times on VMware.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1 GA - MU1

Issue description: VVol snapshotting was taking 8s to complete. Each of the two CLI commands involved takes 4 seconds to respond. The long response time leads delay in VVOL snapshot completion.

Symptoms: Delay in VMware snapshot completion by VVols.

Conditions of occurrence: VMware snapshots are taken with VVols.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 237503

Issue summary: Used space on compressed VV increased and became same as reserved space when total host data (HostWr) became zero.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU2

Issue description: Usr-Used becomes the same as Usr-Rsvd after writing complete zeros on compressed VVs, which leads to no reclamation.

Symptoms: No reclamation happens, and out of space.

Conditions of occurrence: Lots of I/O operations on compressed VVs.

Impact: Low

Customer circumvention: Run **`tunevv`**.

Customer recovery steps: Run **`tunevv`**.

Issue ID: 237546

Issue summary: Controller node restarts when inconsistent meta data entries for compression volumes are encountered and a subsequent partial page write happens.

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

Affected software versions: 3.3.1 GA - MU2

Issue description: When a partial page write occurs after inconsistent meta data entries for compression volumes are encountered, the controller node reboots.

Symptoms: Controller node restart.

Conditions of occurrence: Presence of compression volumes.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 240493

Issue summary: System Manager memory leak.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU1

Issue description: System Manager calculate system capacity utility does not release memory (fail_lds). Restarting System Manager as a workaround can lead to other repercussions (DU) due to other side effects.

Symptoms: System Manager memory leak.

Conditions of occurrence: **`showsys -space`** will cause System Manager memory leak.

Impact: Medium

Customer circumvention: Do not run **`showsys`**.

Customer recovery steps: System Manager restarts.

Patches Included in This Release

HPE 3PAR OS 3.3.1 MU3 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 3.3.1 MU2 Patch 30	Improved error handling for certain drive models	None	**https://h20565.www2.hpe.com/hpsc/doc/public/display?docId=emr_na-a00050187en_us**
HPE 3PAR OS 3.3.1 MU2 Patch 32	Provides several critical quality improvements	None	**https://h20565.www2.hpe.com/hpsc/doc/public/display?docId=emr_na-a00051930en_us**
HPE 3PAR OS 3.3.1 MU2 Patch 34	Delivers quality improvements for StoreServ 9000 & 20000 cage FW	None	**https://h20565.www2.hpe.com/hpsc/doc/public/display?docId=emr_na-a00051850en_us**
HPE 3PAR OS 3.3.1 MU2 Patch 35	Provides support for drive FW updates	None	**https://support.hpe.com/hpsc/doc/public/display?docId=a00052396en_us**
HPE 3PAR OS 3.3.1 MU2 Patch 38	Provides quality improvements for consistency checks and VV removal	None	**https://h20565.www2.hpe.com/hpsc/doc/public/display?docId=emr_na-a00052440en_us**

Known Issues with the OS

Issue ID: 196915

Issue summary: Upgrade drive firmware on drives with the model name beginning with AREX.

Platforms affected: All StoreServ

Affected software versions: 3.2.1 MU5, 3.2.2. MU3 - MU6, 3.3.1 MU1 - MU3

Issue description: The AREX 3P03 and 3P07 firmware corrects a data inconsistency issue with previous firmware versions.

Symptoms: Data inconsistency was reported.

Conditions of occurrence: AREX0920S5xnFTRI, AREX1920S5xnFTRI, AREX3840S5xnFTRI drives with firmware prior to version 3P03.

AREX0400S5xnNTRI drives with firmware prior to version 3P03.

AREX0480S5xnNTRI, AREX0920S5xnNTRI, AREX1920S5xnNTRI, AREX3840S5xnNTRI drives with firmware prior to version 3P07.

Impact: High

Customer circumvention: New firmware to address issue.

Customer recovery steps: Replace the drive.

Issue ID: 203642, 198145

Issue summary: A Harrier 2 `unexpected memory error` leads the system to stop responding for an hour.

Platforms affected: StoreServ 8000, StoreServ 9000, StoreServ 20000, StoreServ 20000 R2

Affected software versions: All

Issue description: An `unexpected memory error` is reported by Harrier 2 for one of the attached memory cores. This leads to panic, and the system stops responding for an hour until reset by System Manager. attempting to use Harrier 2 to DMA memory from CM to PM.

Symptoms: Panic followed by delay in panic processing.

`WARNING: XCBq2 engine not idle, can't skip XCB` printed to console during hang.

Conditions of occurrence: None.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Manually reset the hung controller node.

Issue ID: 204468

Issue summary: LDAP tasks fail to run when the `ldap-service-account` and `ldap-service-account` password parameters are not defined.

Platforms affected: All StoreServ

Affected software versions: 3.3.1 GA - MU3

Issue description: LDAP tasks will not run when the `ldap-service-account` and `ldap-service-account` password parameters are not defined. The service account must authenticate the tasks via LDAP.

Symptoms: LDAP tasks fail, and when displaying the task details with a message saying the user failed to authenticate.

Conditions of occurrence: The `ldap-service-account` and `ldap-service-account-password` are not defined in the LDAP authorization parameters.

Impact: Low

Customer circumvention: Define the `ldap-service-account` and `ldap-service-account-password` to be defined for an actual user existing in the LDAP server.

Customer recovery steps: None.

Issue ID: 211053

Issue summary: Controller node unexpectedly restarts during the cluster join when an inter-node link encounters hardware errors.

Platforms affected: StoreServ 7000

Affected software versions: 3.2.2, 3.3.1

Issue description: If hardware inter-node link errors occur during controller node joining the cluster, the join process can send packets repeatedly exhausting memory, resulting in controller node restart.

Symptoms: System experiencing hardware inter-node link errors causing unexpected controller node restarts.

Conditions of occurrence: System experiencing hardware inter-node link errors during node rejoining the cluster.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 213894

Issue summary: Fixing bug that results in VV block operations triggered by AO region moves involving multiple vv spaces getting stuck in kernel.

Platforms affected: All StoreServ

Affected software versions: 3.1.x, 3.2.x, 3.3.1 GA - MU3

Issue description: A VV block operation triggered by AO region moves got stuck in the kernel. This leads to System Manager becoming unresponsive. Subsequent restart of System Manager can lead to cluster down situations. This happens in a very narrow window where AO region moves of multiple SD spaces of the same volume happen in some combination of IOs to the volume, where some IOs happen in between two successive VV block calls to the same volume.

Symptoms: Examining the vv_hdr pointer reveals vv_nldio_wait_cmp to be 0x2 and there is some element in vv_ldio_wait_list. Further examining the vv_ldio_wait_list shows the first cmp to be a marker cmp cmp_flag = 0x100,cmp_type = CMP_PCI. This marker cmp is stuck in the queue.

Conditions of occurrence: Its a very narrow window in the code with block_less vv block operation. block_less operation only happen if some non-SA spaces are moved. This bug is a corner case that only happens if multiple spaces of the same vv are moved along with IOs.

Impact: High

Customer circumvention: None.

Customer recovery steps:

1. Restart the controller node.
2. It is likely the System Manager on other controller nodes may have an issue on failover. In that case, perform a powerfail wipe, and bring up the cluster again.

Issue ID: 215922

Issue summary: Host IO timeout with VV space allocation exceeds the allowed limit.

Platforms affected: All StoreServ

Affected software versions: 3.2.2, 3.3.1 GA - MU3

Issue description: Host IO timeout with VV space allocation exceeds the allowed limit.

Symptoms: `VV_ALLOC_FAIL` failure messages.

Host IO timeouts.

Conditions of occurrence: Rapid writes to VV from starting to end offset followed by unmap or rewrite.

Impact: Low

Customer circumvention: None.

Customer recovery steps: Host IOs expected to become normal once the system is able to get back the free space for VV.

Issue ID: 225413

Issue summary: Event Log string reports inconsistent cluster memory DIMM ID.

Platforms affected: StoreServ 8000, StoreServ 9000, StoreServ 20000

Affected software versions: 3.3.1 GA - MU3

Issue description: When reporting cluster memory errors that were self-corrected by hardware (correctable ECC), an event log entry indicates the wrong DIMM (specifically DIMM 0.0.0). There are multiple other strings in both the event logs and console logs that do indicate the correct DIMM number. The following console prints indicate the correct DIMM ID.

```
Nov 30 19:33:51 MXN71338CB-2 vmunix: [15053649.305074] debug info id 1,core 1,chaadr 0x6b7ae3a30, denal 0x26b7ae3a30,
errad
```

Symptoms: Logs indicate different DIMM IDs when reporting on correctable errors.

Conditions of occurrence: Cluster Memory DIMM is experiencing correctable ECC error on DIMM other than 0.0.0.

Impact: Low

Customer circumvention: Confirm actual DIMM ID by examining other events and console logs.

Customer recovery steps: None.

Issue ID: 228811

Issue summary: Panic under specific workloads (combined with updatevv operations) in tdvv3, when a cache page ends up being shared by both the base and snap vv, but the dedup is initiated for the snapshot.

Platforms affected: All StoreServ

Affected software versions: 3.3.1 GA - MU3

Issue description: When a previously written data pattern is written for the second time, the DDS triggers a conversion request to the original data block to convert its exception entry from DDC to DDS. In specific workloads (combined with updatevv operations), the original DDC entry belongs to a snapshot but the cache page belonging to this DDC block is shared with the base virtual volume. In this condition the dedupe handler wrongly processes the base vv and updates the base virtual volume's counters causing the panic.

Symptoms: Kernel panic with assertion, `Kernel panic[4]: tpd: Assertion point: file: hat_vol.c, line: 8326 Assertion failed: cnt >= 0`.

Kernel panic with stack trace, `BUG: unable to handle kernel NULL pointer dereference IP: [<ffffffffa00b72b5>] cmp_do_dln_del+0x58/0x85 [tpd]`.

Conditions of occurrence: VVs with dedupable data where snapshots are present and `updatevv` operations happen.

Impact: High

Customer circumvention: Stop **`updatevv`** operations.

Customer recovery steps: The system should recover once the controller nodes reboot.

Issue ID: 230419

Issue summary: Cluster restart due to the overwhelming mirror log entries of type of log page processing from dedup store.

Platforms affected: All StoreServ

Affected software versions: 3.2.2, 3.3.1 GA - MU3

Issue description: A controller node could potentially overwhelm the mirror log entries due to the excessive generation of the garbage from the dedup store and overburdening of that node in the cluster for handling VV mastership. It can reach to the point that no more mirror log entries can be allocated leading to the node restart. The backup VV master node also meets the same result during the recovery leading to the cluster restart.

Symptoms: Node restart followed by the cluster restart.

Conditions of occurrence: Concentration of VV mastership on a single controller node and huge garbage generation from the dedup store.

Impact: High

Customer circumvention: None.

Customer recovery steps: None, the array should be normal after the cluster restart.

Issue ID: 230877

Issue summary: With Adaptive Flashcache enabled, after creating a virtual volume and admitting it into Remote Copy group, array is not manageable anymore.

Platforms affected: All StoreServ

Affected software versions: 3.2.2 MU4, 3.3.1 MU2 - MU3

Issue description: This is a deadlock between multiple threads running during admitting or dismissing virtual volumes from remote copy group when Adaptive Flashcache is enabled on the array. Each thread waits for other thread to complete and thus System Manager stops responding.

Symptoms: Deadlock in System Manager process.

Conditions of occurrence: During normal operation, if Adaptive Flashcache is enabled, adding or removing virtual volumes from Remote Copy groups may result in a System Manger deadlock.

Impact: Medium

Customer circumvention: Avoid admitting and dismissing virtual volumes from Remote Copy group when Adaptive Flashcache is enabled.

Customer recovery steps: Remove Adaptive Flashcache and restart System Manager.

Issue ID: 231283

Issue summary: The**`tunesys`** command excludes logical disk for further analysis if it has bad disk creation pattern.

Platforms affected: All StoreServ

Affected software versions: 3.2.2, 3.3.1 GA - MU3

Issue description: The**`tunesys`** command excludes logical disk for further analysis if it has bad disk creation pattern.

Symptoms: The **`tunesys`** command will exclude from list with error.

Conditions of occurrence: The **`tunesys`** command is run while the system has a logical disk with a bad or empty pattern.

Impact: Medium

Customer circumvention: Avoid running **`tunesys`**.

Customer recovery steps: Run **`tuneld`** command manually on excluded logical disks.

Issue ID: 232101

Issue summary: Invalid iSCSI PDUs from the initiator can cause the host port to go offline.

Platforms affected: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1 GA - MU3

Issue description: Invalid iSCSI PDUs from the initiator can cause the host port to go offline.

Symptoms: Host ports go offline.

Too many firmware cores generated.

Conditions of occurrence: Initiator sends invalid iSCSI PDUs.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 233992

Issue summary: Coordinated snapshots created via the CLI do not expire as intended.

Platforms affected: StoreServ 8000, StoreServ 9000, StoreServ 20000, StoreServ 20000 R2

Affected software versions: 3.2.2 GA - MU6, 3.3.1 GA - MU3

Issue description: Coordinated snapshots created via the CLI do not expire as intended. Consequently, they remain on the secondary array until manually expired or deleted.

When coordinated snapshots are created by a scheduled job, if the snapshot's name is tagged with only the hour and minute of its creation, then unexpired/undeleted snapshots left over from one day will cause the creation of the corresponding snapshot for a following day to fail.

Symptoms: Coordinated snapshots created with the CLI do not expire.

New coordinated snapshots can collide with existing snapshots having the same name and not be created.

Conditions of occurrence: Coordinated snapshots created by the scheduler, using the CLI.

Impact: Low

Customer circumvention: Configure and setup a shell script on the secondary to manually set the expiration dates on the affected snapshots.

Customer recovery steps: Delete unneeded snapshots.

Issue ID: 242378

Issue summary: Data Inconsistencies may be seen during volume migration of non-ALUA hosts if proper migration procedure is not followed.

Platforms affected: All StoreServ

Affected software versions: 3.3.1 MU1 - MU3

Issue description: When migrating volumes which are exported to non-ALUA hosts, if unzoning the source array as required by the proper migration procedure is not followed, then Data Inconsistencies may be seen on the volumes being migrated. Note that failure to perform the unzoning operation is not a supported procedure. The migration tools clearly prompt with the message to perform unzone operation, perform the same before proceeding with the subsequent migration steps.

Symptoms: The hosts or the applications at the hosts may not function as expected due to Data Inconsistency on the volumes.

Conditions of occurrence: Non-ALUA hosts are migrated and the proper procedure of unzone operation is not followed.

Impact: Medium

Customer circumvention: Follow the proper migration procedure by following the prompts from migration tools for the unzone operation when migrating non-ALUA hosts.

Customer recovery steps: None.

Issue ID: 243445

Issue summary: Issuing a FPG recover operation targeted to a volume that is already associated with an activated FPG causes that FPG to become unavailable.

Platforms affected: StoreServ 7000, StoreServ 8000, StoreServ 20000, StoreServ 20000 R2

Affected software versions: All versions

Issue description: Issuing a FPG recover operation targeted to a volume that is already associated with an activated FPG will result in a message indicating the operation was not successful. If file services are in a Running state on all of the nodes in the system, the message is returned without impact to the activated FPG. If file services is in a Starting state on one of the nodes, this can result in the FPG becoming unavailable.

Symptoms: After attempting a recover operation on a volume, a message is returned indicating the operation was not successful, and the FPG using that volume is unavailable.

Conditions of occurrence: An FPG is in an activated state and a request is made to recover an FPG using one of the volumes of the activated FPG.

Impact: High

Customer circumvention: Do not attempt a recover operation on a volume that belongs to an FPG listed by the **`showfpg`** command. The recover operation is only to be used on volumes belonging to FPGs that were previously forgotten using the **`removefpg -forget`** option.

Customer recovery steps: Ensure file services are in a Running state on all of the nodes. Then attempt to deactivate and then activate the impacted FPG.

HPE 3PAR 3.3.1 File Persona MU3 Release Notes

What's New in File Persona

Remote Copy auto failover for FPGs

Activates FPGs automatically on the secondary system. Occurs if the primary system fails when adding FPGs to a Remote Copy Group that uses the AutoFailover policy.

Remote Copy manual failover/failback for FPGs

Simplifies processes associated with adding FPGs to a Remote Copy Group, failing over the Remote Copy Group, and failing back the Remote Copy Group (SSMC 3.3.1 and later).

File lock compliance mode

Increases security with File Lock Compliance to meet regulations defined by U.S. Securities and Exchange Commission rule 17a-4.

Authentication improvements

Includes the following improvements:

- **LDAP performance**
- **Redundant LDAP providers**

  Specify multiple LDAP servers. Ensure resilience if a single-server failure occurs.
- **Local user mapping**

  Create user mappings between Active Directory users and Local users.
- **Minimum UID/GID lowered from 1000 to 100**

  Integrate simply with Linux environments that include user accounts in the 100 to 1000 range and require access to files presented by File Persona.

Major version on-disk upgrade

Uses the latest File Persona features with FPGs originally created on software versions earlier than 3.2.2 MU2.

SMB v1 protocol control

Allows the administrator to configure communication paths using SMB v1.

The SMB protocol facilitates communication paths between a client and File Persona and between File Persona and Active Directory (AD).

On a new configuration, SMB v1 defaults to disabled. The administrator can enable SMB v1 for each of the paths after determining that the clients or AD still require it.

With an upgrade, SMB v1 remains enabled for backward compatibility. The administrator can disable SMB v1 for each of the paths after confirming that no requirement exists for clients or AD.

Network diagnostics

Adds commands to perform `ping` and `traceroute` requests from the perspective of the File Persona instance. Simplifies diagnosis of network configuration issues during setup.

Modifications to File Persona

Issue ID: 89753

Issue summary: No alert was raised when the FPG approached or exceeded the supported limits.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: A file system with 250 million or more files can start to become slower and less responsive. No alert is raised to warn the user that the file count could be the issue.

Symptoms: Gradual decreased performance of FPG I/O without any indication or alert as to why.

Conditions of occurrence: When an FPG has more than or is approaching 250 million files.

Impact: Low

Customer circumvention: Monitor the number of files in each FPG periodically using the `showfpg -d` command. View the details of the FPG in SSMC. If the file count is approaching 250 million files, create a new FPG to receive new writes.

Customer recovery steps: Migrate files to a new FPG and then remove files from the existing FPG until the file count is below the 250 million file limit.

Issue ID: 96608

Issue summary: Several changes are included to address temporary interruptions when accessing SMB shares.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: Various conditions are causing the SMB stack to restart (self-heal), and is resulting in `fcollect` support collections.

Symptoms: SMB share clients may notice momentary unresponsiveness in read or write access to files on the file share.

Conditions of occurrence: Accessing SMB shares.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 97625

Issue summary: Management of file services was temporarily unavailable. The `showfs` command reported a "Starting" state.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: Under conditions where authentication providers such as LDAP and Active Directory are not responding in a timely manner, management of file services could become temporarily unavailable.

Symptoms: The `showfs` command could display a "Starting" state, even though the data services are not actually restarting.

Conditions of occurrence: This could happen when a customer is bound to an LDAP for authentication. If the LDAP access is slow due to lot traffic or heavy network, then it negatively impacts the name resolution. Consequently, it impacts the file services manageability.

Impact: Medium

Customer circumvention: Ensure that any configured authentication provider is healthy and responsive.

Customer recovery steps: None.

Issue IDs: 97644

Issue summary: The system time used by the file services and the Active Directory was not synchronized, which resulted in unsuccessful file share creation and access.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: The system time for file services is not in the acceptable time threshold range (anything more or less than 5 minutes), in relation to Active Directory. Due to the time drift, the file services are no longer joined to the Active Directory domain. As a result, authentication and name resolution is unsuccessful for the file services. The problem persists until the system time for file services and Active Directory are synchronized.

Symptoms: All management operations or access to file shares that are configured to use the Active Directory authentication are unsuccessful. Creating a file share or accessing an existing one using MMC is unsuccessful too.

Conditions of occurrence:

The file services time drifts ahead from the Active Directory domain. As a result, file services are no longer joined to the Active Directory. This is because the NTP time is not set properly before the file services get started.

Impact: Medium

Customer circumvention: Use the same external time server (NTP) configuration as the Active Directory.

Customer recovery steps: To address the issue, upgrade to the latest HPE 3PAR OS 3.3.1 MU2. With the upgrade installed, a 3PAR alert is generated when the time between File Persona and the Active Directory drifts.

Issue ID: 98756

Issue summary: Share access was lost on one of the array controller nodes. The following message was displayed `Access Denied`.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: Various conditions can cause the SMB stack to restart (self-heal), and this can result in `fcollect` support collections.

Symptoms: SMB clients cannot access shares on a controller node even though the other controller nodes have it.

Conditions of occurrence: Clients accessing SMB shares on a controller node.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Perform a failover FPG to another controller node and then run the `stopfs <node>` followed by the `startfs -enable <node>` commands.

Issue ID: 98997

Issue summary: The `httpd` monitoring service is running on both controller nodes (4-nodes) after performing an FPG failover. This is happening even though the Object Access API file shares are not available on both controller nodes of the cluster.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: The `httpd` monitoring service is running on both controller nodes (4-nodes) after performing an FPG failover. Port 80 or Port 443 open on unexpected interfaces.

Symptoms: Port 80 or Port 443 open on unexpected interfaces.

Conditions of occurrence: An FPG failover to another controller node containing Object Access API file shares was unsuccessful.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 99290

Issue summary: A new Active Directory request was unsuccessful when a secure channel connection to the Active Directory was reset.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: When a secure channel connection to an Active Directory is reset, there are chances that the authentication requests could be unsuccessful.

Symptoms: Customer intermittently loses a secure channel connection to Active Directory. All new Active Directory requests do not complete successfully.

Conditions of occurrence: The Active Directory domain services get restarted while file services are in use.

Impact: High

Customer circumvention: Do not disrupt the network connection between the SMB client and the Active Directory server when an Active Directory request is executed.

Customer recovery steps: Reduce the frequency of open, close, and delete operations.

Issue ID: 99851

Issue summary: FPG became unavailable just before the NFS file share got full.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: FPG could become unavailable just before an NFS file share gets full. Even a retry could give the same result. Consequently, activating the FPG may no longer be possible without further intervention.

Symptoms: NFS share becomes unavailable during writing when the FPG is nearly full and snapshots are concurrently in use.

Conditions of occurrence: A snapshot was taken in an `fsfull` condition.

Impact: High

Customer circumvention: The issue can be avoided if the files are created and accessed with the same combination of cases of letters in the name.

Customer recovery steps: None.

Issue ID: 100011

Issue summary: Under high I/O load, the following message is displayed `The thread pool's task queue is full, limit: 75`.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: Sometimes management commands are unresponsive due to a high I/O load.

Symptoms: File Persona management commands become unavailable till the internal File Persona management service detects the issue and recovers from the situation automatically.

Conditions of occurrence: High I/O load.

Impact: Medium

Customer circumvention: Run fewer parallel instances of the `robocopy` command.

Customer recovery steps: None.

Issue ID: 101659

Issue summary: Higher than expected CPU load on active file services although no SMB file share was configured.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: Active Directory was not used in the setup, rather LDAP was the directory server. Hence, the File Persona authentication order does not have an entry for the Active Directory.

Symptoms: Higher than expected CPU load, as seen in the `statfs -cpu` command output.

Conditions of occurrence: The mechanism that searches for groups and users when using an LDAP provider is not efficient for large number of groups and users. Too many resources are consumed.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 101885

Issue summary: The offline FSCK utility took longer than expected to complete.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: An offline FSCK executed by support could take an unusual amount of time to complete when an FPG has a large number of files.

Symptoms: The offline FSCK utility was taking longer than expected to complete.

Conditions of occurrence: Offline FSCK was taking longer than expected to complete.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 102366

Issue summary: The user was unable to traverse the root of shares by default.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions 3.3.1 GA or later and earlier than 3.3.1 MU2 P36

Issue description: In the NTFS security mode, users were unable to traverse the root of the share by default. The default ACE for the root of the share was modified to allow for traversal by the `Everyone` user to address this issue.

Conditions of occurrence: When creating shares using default permissions.

Impact: Low

Customer circumvention: Create a share at the File Store root and modify the Share Folder ACL if this level of default access is not desired.

Customer recovery steps: None.

Issue ID: 102927

Issue summary: The `mkdir` operation at the client's side was unsuccessful. The following message was displayed `retry operation` when simultaneous cache clear management commands were issued.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions 3.3.1 GA or later and earlier than 3.3.1 MU2 P36

Issue description: The `mkdir` operation at the client's side could be unsuccessful. A message such as `retry operation` will be displayed when simultaneous cache clear management commands are issued. The `setfs auth -clearcache` command is issued at the same time as a `mkdir` request.

Symptoms: Directory creation is unsuccessful and is returning a retry message.

Conditions of occurrence: Executing the `setfs auth -clearcache` command simultaneously with the `mkdir` operation causes the `setfs auth -clearcache` command to be unsuccessful.

Impact: Medium

Customer circumvention: Do not issue the `setfs auth -clearcache` command during I/O operations.

Customer recovery steps: Redo the directory creation.

Issue ID: 103728

Issue summary: Unreachable file services system.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions 3.3.1 GA or later and earlier than 3.3.1 MU2 P36

Issue description: File services became temporarily unavailable on multiple controller nodes during an SMB share load test.

Symptoms: An unreachable system which cannot be pinged.

Conditions of occurrence: I/O load during an SMB share load test.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 103877

Issue summary: False NFS health check failure alerts were generated.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: False NFS health check alerts were generated. There was no impact on file access. There was no pattern to the occurrence of the alerts. Sometimes they occurred within a period of hours and sometimes over multiple days.

Symptoms: Alerts similar to the following example were generated:

```
Time     : 2017-08-17 17:03:53.40 JST
Node     : 0
Seq      : 66125
Class    : Alert
Severity : Major
Type     : NFS Share
Component: sw_fs_fstore_share_nfs:3702654806126534045:nfs-resource-health-check
Message  : File Services NFS Share:3702654806126534045:nfs-resource-health-check Failed
(FAILED)
```

Conditions of occurrence: Normal operation

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 104334

Issue summary: The `HPEidmapd` daemon got restarted unexpectedly when the `Authenticated^Users@NT^AUTHORITY` object was not found.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions 3.3.1 MU1 P07 or later and earlier than 3.3.1 MU2 P36

Issue description: The `Authenticated^Users@NT^AUTHORITY` object could not be found when the Name Service Switch interface was used. The `HPEidmapd` daemon could get restarted.

Symptoms: Alert indicating that the `HPEidmapd` service has restarted.

Conditions of occurrence:

- Name Service Switch configuration
- Query of the `Authenticated^Users@NT^AUTHORITY` object was not found.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 104469

Issue summary: An SMB directory folder was not created successfully.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions 3.3.1 GA or later and earlier than 3.3.1 MU2 P36

Issue description: An SMB directory folder was not created successfully when a share was mapped as a user from another domain.

Symptoms: An SMB directory folder was not created successfully.

Conditions of occurrence: This issue occurs when a client, who is a member of a domain, logs in as a non-admin user and tries to map a drive as a user from another domain with a primary group "Domain Users" and attempts to create folders and files.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 105799

Issue summary: Online FSCK restarted the file services on the controller node where the operation was active.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions 3.3.1 GA or later and earlier than 3.3.1 MU2 P36

Issue description: Online FSCK mapped the same part of the file in the memory multiples times, consequently straining the memory.

Symptoms: File services restarts on the controller node during online FSCK operation.

Conditions of occurrence: If a snapshot is taken on a file store and the files get modified at irregular intervals, then the snapshots are sparse.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 105838

Issue summary: When Microsoft Management Console (MMC) is used to display open files, MMC is displaying files for all FPGs owned by the File Persona node instead of files for a specific VFS.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: MMC is opening files for an entire FPG instead of the file share that is managed.

Symptoms: Inaccurate per file share open file count in MMC.

Conditions of occurrence: Microsoft Management Console (MMC) is used to display open files.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 106099

Issue summary: Microsoft Project could not open files. The following state for the `SMB Create` attribute was displayed `STATUS_SHARING_VIOLATION`.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: Microsoft Project sometimes does not open files properly. The following state for the `SMB Create` attribute is displayed `STATUS_SHARING_VIOLATION`

Symptoms: Microsoft Project is not opening files properly.

Conditions of occurrence:

- An SMB threads creates a file and enforces a particular shared lock mode
- An SMB thread tries to enforce compatible shared mode locks

Impact: Medium

Customer circumvention: None.

Customer recovery steps: If the file is accessed by the SMB protocol, then turn-off the share mode enforcement.

Issue ID: 106597

Issue summary: Failover was unsuccessful with the following message `Error handling umount notification before for host`. The Virtual File Server IP address was still deactivated.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: Rolling back and `umount` was unsuccessful.

Symptoms: During protocol deactivation, an unsuccessful `umount` did not initiate a roll back process and brought the protocol service down. This caused extended data unavailability for the impacted Virtual File Server.

Conditions of occurrence: Deactivation of an FPG under I/O load.

Impact: High

Customer circumvention: None.

Customer recovery steps: Issue a new request to deactivate the FPG until the request is successful, then reactivate the FPG.

Issue ID: 107555

Issue summary: Online FSCK was unsuccessful with the following message `Failed to add missing entries for tag`.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: Online FSCK could be unsuccessful when namespace connectivity is checked for link mismatch. Consequently, there could be filesystem activation issues.

Symptoms: FPG is not activated successfully after running Online FSCK

Conditions of occurrence: Running Online FSCK after deleting directories in between multiple snapshots.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 109148

Issue summary: The client IP list of an existing NFS file share could not be modified after File Persona upgrade.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: 3.3.1 MU1/EMU1 with P07, P08 , or P19

Issue description: After File Persona is upgraded, the command to modify the client IP list of pre-existing NFS file shares was unsuccessful with the message `NFS Modify Export failed with error: Modifying export path is not allowed`.

Symptoms: After File Persona is upgraded, the command to modify the client IP list of an existing NFS file shares is unsuccessful as shown below:

```
cli% setfshare nfs -clientip +<client IP> -fstore <file store> <vfs> <share
name>
NFS Modify Export failed with error: Modifying export path is not allowed.
```

Conditions of occurrence: Upgrade with NFS file shares configured.

Impact: Low

Customer circumvention: None.

Customer recovery steps: Delete and recreate the share.

Issue ID: 109538

Issue summary: Excessive SMB status calls were received although no SMB file shares were exported from the system.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: When excessive SMB status calls can be received even though no SMB file shares are exported from the system.

Symptoms: Performance of the controller node degrades due to excessive SMB status calls.

Conditions of occurrence: File services enabled.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 109769

Issue summary: Degradation in performance when copying large files (such as `.iso`) to an SMB file share.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions 3.3.1 GA or later and earlier than 3.3.1 MU2 P36

Issue description: Degradation in performance when copying large files (such as `.iso`) to an SMB file share.

Symptoms: Slow I/O, disk fragmentation.

Conditions of occurrence: Writing large files to an SMB share.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 110892

Issue summary: After upgrading from 1.4.2, the `srvsvc` container indicates a failed state and the system generates a log file.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions

Issue description: At startup, if FP was joined to a Windows domain, LSASS has to re-join and/or enumerate domains. LWIO calls NetLogon, which uses the customer DNS to resolve the Hostname to an IP. If the DNS has no entry for hostname, or has an incorrect reverse lookup pointer, the DNS response is unpredictable. In this case LWIO releases the allocated memory, resulting in double free memory.

Symptoms: HP-SMB cannot join the domain and the HP-SMB stack restarts.

Conditions of occurrence:

- Customer DNS does not have a reverse lookup zone.
- Hostname does not have a record pointer in the reverse lookup zone.
- Hostname has a record pointer in the reverse lookup zone but it is different from the Hostname (canonical name).

Impact: High

Customer circumvention:

Verify DNS entries for correct configuration.

Customer recovery steps: None.

Issue ID: 114353

Issue summary: After generating `lwsm` log files, File Services stop running.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: HPE 3PAR 3.3.1 MU1 and later.

Issue description: When attempting to join a domain, the name resolution using Active Directory return unexpected results and causes abnormal termination of HP-SMB stack. File Services (in particular SMB) stop running.

Symptoms: Customer encounters an abnormal termination from HP-SMB stack.

Conditions of occurrence:

1. Insufficient or incorrect Active Directory DNS configuration.
2. Joining domain with HP-SMB server.

Impact: Medium

Customer circumvention: Properly configure the Active Directory DNS.

Customer recovery steps: Upgrade the system to the latest patch.

Issue ID: 116735

Issue summary: An attempt to access a file or folder is unsuccessful.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: HPE 3PAR OS 3.3.1 MU1 plus limited availability versions; HPE 3PAR 3.3.1 MU2 plus limited availability versions.

Issue description: As part of authenticating a user and authorizing access to files and folders, the domain must resolve the user ID. An offline domain within the AD forest, coupled with a disabled rfc2307 mode, can prevent access to a file or folder.

Symptoms: An attempt to access a file or folder is unsuccessful.

Conditions of occurrence: 1) The system is joined to an Active Directory (AD) domain. 2) The AD domain exists in a forest with one or more one-way trusts. 3) At least one domain in the forest is offline. 4) rfc2307 mode is disabled.

Impact: Medium

Customer circumvention: Ensure that all domains in the forest are healthy. Use bi-directional trusts.

Customer recovery steps: Bring the unhealthy domain back online.

Issue ID: 116805

Issue summary: Access denied is returned when attempting to access files and subdirectories with the local Guest account.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: HPE 3PAR OS 3.3.1 MU1 plus File Persona limited availability versions; HPE 3PAR 3.3.1 MU2 plus File Persona limited availability versions.

Issue description: When the file store contains files or directories with ownership and/or access permissions for the local user Guest (Guest@LOCAL_CLUSTER), the issue is that when the Local user Guest account is enabled and Guest attempts to access those files or directories, “access denied” failures occur.

Symptoms: An attempt to access a file or folder is unsuccessful.

Symptom 1:

When the local Guest account is enabled:

When Local user Guest creates a directory or file from an SMB client, when the parent directory has create permissions for Guest, and a CREATOR_OWNER ACE exists, the directory/file will be created, but Guest will get “access denied” when trying to access the newly created file or directory that should be owned by Guest.

Symptom 2:

When the local Guest account is enabled:

If there are existing files owned by local user Guest in the file store on a previous release (e.g. 3.2.2-MU4), after upgrade to a later release with this problem (e.g. 3.3.1-MU1): Guest is be able to log on but is not able to access the directories/files previously owned and accessible by Guest, getting “access denied” errors.

Conditions of occurrence:

1) Local Guest account is enabled.

2) Files and/or directories exist that are owned by local user Guest or have ACLs that should make them accessible by Guest.

Impact: Medium

Customer circumvention:

Do not upgrade the on-disk version of the file system to 12.2 or greater if there are files with ownership or permissions by the Guest local user, as the permissions metadata may become corrupted. Wait until after upgrading to the latest release before upgrading the on-disk version.

Do not enable the local Guest account before upgrading to the latest release.

Customer recovery steps: Bring the unhealthy domain back online.

1. If the customer has not enabled the Guest account and/or has not created files owned by Guest, this is not an issue. No recovery is necessary.

2. If the customer does enable the Guest account but the customer has not upgraded the on-disk version prior to upgrading to the latest release, no recovery is necessary.

3. If the customer had files or directories owned or with permissions by Guest@LOCAL_CLUSTER, and the on-disk version was upgraded to 12.2 or greater before upgrading to the latest release, the metadata may have been corrupted, and access by the Guest user may continue to fail even after the release upgrade.

This can be fixed after upgrading to the latest release by performing the following steps:

a) Map a share to the file store from a Windows client as the domain Administrator.

b) Reset the ownership and/or ACLs of relevant files to Guest. For example:

c) Recursively change the ownership at the parent directory owned by Guest to be owned by the domain Administrator.

For example, in Windows 10, Windows Explorer Security Tab=>Advanced Security Settings=>Owner Change (Check “Replace all child object permission entries with inheritable permission entries from this object”)

d) Recursively change the ownership at the same parent directory back to Guest

---

Issue ID: 117510

Issue summary: With File Persona P07 and later, file system does not mount after fail over. System returns the message: `ftx_capsule_play_redo_on_page(): bs_access() of bfset 2 file-tag 261786531 failed(-1032) dirTag : 2`

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: HPE 3PAR OS 3.3.1 MU1 P7 and P8

Issue description:

Reloading the file extents without checking whether the metadata of file changed during the recovery (part of mount process).

Symptoms: With File Persona P07 and later, file system fails to mount after fail over and returns the message: `ftx_capsule_play_redo_on_page(): bs_access() of bfset 2 file-tag 261786531 failed(-1032) dirTag : 2`

System recovery of file system during becomes excessive.

Conditions of occurrence: Create a fragmented filesystem. Create large number of files so that bmt will have more extents. If the system is crashed and is getting mounted it has to first recovered for metadata consistency which will take long time.

Impact: Medium

Customer circumvention: None.

Customer recovery steps:

Run domain activate on the system and wait for completion of the recovery process. Remount the system.

Issue ID: 200909

Issue summary: File Services did not automatically start on the controller node even though the controller node had been started (during upgrade or otherwise).

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: When a controller node starts up, enabling file services on the controller node requires that all volumes in use by file services are in a ready state. If it takes longer than expected for these volumes to reach a ready state, file services may be left in a `shutoff` state.

Symptoms: Even though the controller node was started, the `showfs` command reported that file services on the controller node are in a `shutoff` state. The following was displayed:

```
cli% showfs
Node FSNode State     Active InCluster -Version- ---N:S:P--- BondMode MTU
   0 Yes    Upgrading No     No        -         0:4:2,0:4:1        -   -
   1 Yes    Shutoff   No     No        -         1:4:2,1:4:1        -   -
   2 No     Unknown   No     No        -         -                  -   -
   3 No     Unknown   No     No        -         -                  -   -
-------------------------------------------------------------------------------
   4 total
If this occurs during an upgrade, an error like the following may be reported:
Rebooting node 1.
Waiting for node to go down....
Waiting for node 1 to rejoin the cluster.....................................
Checking if the system is healthy enough to proceed...
Waiting for File Services update to complete................................
File Services update did not complete after 600 secs. (Unable to connect with
server.
```

Conditions of occurrence: Rebooting a controller node while file services are enabled.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Use the `startfs -enable <node>` command for the node when file services are in a `shutoff` state. If the upgrade is in progress, and the controller node is still reporting `Upgrading`, resume the upgrade process.

Issue ID: 211324

Issue summary: A fully provisioned FPG could unexpectedly deactivate when the backing CPG is almost out of space.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions

Issue description: When a CPG backing an FPG is nearly out of space (as can be seen in the `free` column of the `showspace -cpg <cpgName>` command output), the system tries to automatically deactivate all the thinly provisioned FPGs. These FPGs use the CPG to ensure that the FPG does not experience a write error due to lack of space for additional allocations for backing VVs. In this case, fully provisioned FPGs could get deactivated.

Symptoms: An alert like the following would indicate that the CPG is almost out of space:

```
Id          : 417
State       : New
Message Code: 0x027001d
Time        : 2018-05-18 14:42:59 MDT
Severity    : Critical
Type        : CPG growth failure non-admin
Message     : CPG FC_r6 SD and/or user space could not grow due to
unavailability of free space. CPG grow attempted using degraded availability
parameters (-ha mag) also failed.
An alert like the following would indicate that the FPG has been deactivated:
Id          : 418
State       : New
Message Code: 0x0720001
Catalog-Key : filesystem-event:filesystem.cmd.unmount
Time        : 2018-05-18 14:43:17 MDT
Severity    : Informational
Type        : File Provisioning Group
Message     : File Provisioning Group:<id>:<fpgName> Normal  (UNMOUNTED)
Details     : FPG Event: Unmounted FPG <fpgName> on host <nodeName>.
```

Conditions of occurrence: CPG space is exhausted but fully provisioned FPGs are still in use.

Impact: Medium

Customer circumvention: Whenever thin provisioning is used, it is critical that the administrator pays close attention to the remaining capacity of the CPGs. Respond to alerts by adding capacity to the CPG or otherwise free up space in the CPG. An early indicator that of a CPG that is dependent on a given storage class and is soon running out of capacity could be seen with the following type of alert:

```
Id          : 452
State       : New
Message Code: 0x0270010
Time        : 2018-05-18 14:26:49 MDT
Severity    : Major
Type        : FC raw space allocation 85% alert
Message     : Total FC raw space usage at 10453G (above 91% of total 11442G).
```

If this alert is resolved by adding space of the given storage class, this issue could be avoided

Customer recovery steps: Resolve the CPG space issue and then manually activate the FPG again using the `setfpg` command.

Issue ID: 211445

Issue summary: File Persona upgrade was unsuccessful because the File Persona node did not start after a StoreServ node upgrade.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions earlier than 3.3.1 MU2 P36

Issue description: The networking configuration for File Persona was in an unexpected state. This caused the upgrade to stall.

Symptoms: The `showfs` command shows a node in a `shutoff` state after an upgrade.

Conditions of occurrence: File services are enabled.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 228325

Issue summary: FPG growth in a Remote Copy Group in a failed over state was successful but the FPG was no longer present in the Remote Copy Group.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions

Issue description: When a Remote Copy Group is in a failed over state, FPG commands such as `growfpg` should not be issued. The FPG will successfully grow but the FPG might be taken out from the Remote Copy Group.

Symptoms: FPG growth was successful but the FPG was taken out from the Remote Copy Group. The following message is displayed:

```
Failed to admit new volumes to Remote Copy Group: Unable to restart RC Group
<rcGroupName> on target <rcTarget>: Error: Could not be started on target
<rcTarget>. Attempt to start group <rcGroupName> failed: Secondary group for
group <rcGroupName> on target system <rcTarget> doesn't match primary group.
```

Conditions of occurrence: Attempt to grow an FPG while an associated Remote Copy Group is in a failed over state.

Impact: Medium

Customer circumvention: Perform a recover operation on the Remote Copy Group before attempting to grow the FPG.

Customer recovery steps: Manually grow a VV in the secondary group to match the primary group and then restart the Remote Copy Group.

Issue ID: 230966

Issue summary: During a Remote Copy failover, if the FPG name was longer than 18 characters, the FPG could not be activated.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions

Issue description: The maximum length of VV name that could be used by File Persona is 23 characters. The VV name format generated when using the `creatfpg` or `growfpg` commands is `<fpgName>.N`. When creating an FPG, the name is limited to 21 characters. When an FPG is configured for Remote Copy, a suffix is typically added to the secondary volumes like "`.r`" or "`_sec`". This additional suffix could cause the VV name to exceed 23 characters, causing an unsuccessful activation of the FPG on the secondary system.

Symptoms: When an FPG was activated by using the `createfpg -recover` command on the secondary system, the following message was displayed:

```
Error: Unable to get minor number of device '/dev/tpddev/vvb/<vvid>': No such
file or directory
```

Conditions of occurrence: Attempting to recover an FPG, where the VV name is 24 characters or longer.

Impact: Medium

Customer circumvention: After setting up Remote Copy for an FPG, confirm that the VV names on the secondary system are 23 characters or shorter. If the names are longer than 23 characters, rename them to names that are shorter than 23 characters.

Customer recovery steps: Rename the affected volumes on the secondary system to 23 characters or shorter. Reissue the `createfpg -recover` command with the volumes for the impacted FPG.

Issue ID: 232375

Issue summary: FPG did not deactivate automatically if the backing CPG was almost out of space.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions

Issue description: When a CPG backing an FPG is almost out of space (as can be seen in the `free` column of the `showspace -cpg <cpgName>` command output), the system automatically deactivates all the thinly provisioned FPGs. The FPGs use the CPG to ensure that the FPGs do not experience a write error due to lack of space for additional allocations for backing VVs. If the FPG has active I/O, the deactivation request could be unsuccessful. The FPG could remain active even though there is little space left in the CPG. Continuing to write to an FPG in this state results in an alert such as the following:

```
FPG Event: FPG <fpgName> domain <domainID> has become unavailable. Reason: ADE
filesystem requested failover of FPG <fpgName>, segment 1, reason: disk I/O
error.
```

If the above alert gets displayed, contact HPE Support to run a filesystem check (FSCK) before the FPG can be reactivated.

Symptoms: An alert like the following could indicate that the CPG is nearly out of space:

```
Id          : 417
State       : New
Message Code: 0x027001d
Time        : 2018-05-18 14:42:59 MDT
Severity    : Critical
Type        : CPG growth failure non-admin
Message     : CPG FC_r6 SD and/or user space could not grow due to
unavailability of free space. CPG grow attempted using degraded availability
parameters (-ha mag) also failed.
If the deactivation of an FPG has also failed, an alert like the following
would be present:
Id          : 454
State       : New
Message Code: 0x00e000a
Time        : 2018-05-18 14:43:24 MDT
Severity    : Minor
Type        : Task failed
Message     : Task 26383 (type "background_command", name "setfpg_task") has
failed (Task Failed). Please see task status for details.
The failed task will have details like the following:
   Id Type               Name        Status Phase Step -------StartTime-------
------FinishTime------- -Priority- -User--
26383 background_command setfpg_task failed   ---  --- 2018-05-18 14:43:03 MDT
2018-05-18 14:43:24 MDT n/a        3parsvc

Detailed status:
2018-05-18 14:43:03 MDT Created     task.
2018-05-18 14:43:03 MDT Updated     Executing "setfpg_task" as 0:20955
2018-05-18 14:43:04 MDT Updated     Deactivating FPG <fpgName>
2018-05-18 14:43:24 MDT Error       Failed to deactivate <fpgName>: RPC error
umount error : filesystem mountpoint  directory is busy. Filesystem <fpgName>
host <nodeName>.
2018-05-18 14:43:24 MDT Error       FPG <fpgName> was not activated on host
```

```
<backupNodeName>.
2018-05-18 14:43:24 MDT Error       Task exited with status 1
2018-05-18 14:43:24 MDT Failed      Could not complete task.
```

Conditions of occurrence: The CPG space is exhausted but an FPG has active I/O, which prevents successful deactivation.

Impact: High

Customer circumvention: Pay close attention to the remaining capacity in the CPGs. Whenever thin provisioning is used, it is critical that the administrator pays close attention Respond to alerts by adding capacity to the CPG or free up space in the CPG. An alert of the following type will be displayed:

```
Id           : 452
State        : New
Message Code: 0x0270010
Time         : 2018-05-18 14:26:49 MDT
Severity     : Major
Type         : FC raw space allocation 85% alert
Message      : Total FC raw space usage at 10453G (above 91% of total 11442G).
```

This alert could be resolved by adding space of the given storage class.

Customer recovery steps: Halt all I/O to and from the FPG and manually deactivate the FPG with the `setfpg` command. Once the CPG space issue has been taken care of, manually activate the FPGs by using the `setfpg` command. If the `FPG Event: FPG <fpgName> domain <domainID> has become unavailable` alert has already been raised, contact HPE Support for assistance on how to run a filesystem check to get the FPG online back again.

Known Issues with File Persona

Issue ID: 109129

Issue summary: The system produces the following error when `domainSID` is not configured on the LDAP server: `Failed to obtain Domain SID from Ldap Server`

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions

Issue description: Unsuccessful bind when using POSIX to bind either a newly configured LDAP server or an LDAP server without `domainSID` configured.

Symptoms: `Setfs -ldap` command is not successful and returns the message: `Failed to obtain Domain SID from Ldap Server`.

Conditions of occurrence: Using POSIX to bind either a newly configured LDAP server or an LDAP server without `domainSID` configured.

Impact: Medium

Customer circumvention: Customers binding to an LDAP provider using POSIX schema must configure `domainSID` for the LDAP server (see, *HPE 3PAR File Persona User Guide*).

Customer recovery steps: None.

Issue ID: 110635

Issue summary: Online upgrade was unsuccessful with the following message `Failed to failover the filesystem(s).`

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions

Issue description: When the system is under an excessive load with many delayed ACKs, an online upgrade could be unsuccessful.

Symptoms: Online upgrade was unsuccessful.

Conditions of occurrence: Upgrade is performed when the system is under excessive load with lots of delayed ACKs.

Impact: Medium

Customer circumvention: Before performing an upgrade, use the `statcmp` or `srstatcmp` commands to ensure that there are no delayed ACKs reported. If there are delayed ACKs reported, reduce the load on the system before proceeding with an upgrade.

Customer recovery steps: None.

Issue ID: 110992

Issue summary: User mapping to a local provider cannot be configured for Active Directory users or groups with a UID or GID value greater than 1,000,000,000.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: 3.3.1 MU2 P36

Issue description: If there are pre-existing files with owners or groups with a UID or GID value greater than 1,000,000,000, the owners and groups of those files will not be mapped to a local user or group, Adding local users or groups with a UID or GID value greater than 1,000,000,000 is not supported.

Symptoms: If the `AD<=>Local` user mapping is configured for users/groups, files or directories with owners/groups with a UID/GID greater than 1,000,000,000 are no longer accessible.

Conditions of occurrence:

- There are existing files or directories with users or groups with a UID or GID greater than 1,000,000,000
- User mapping is enabled and configured for the `AD<=>Local` user mapping for the same users or groups

Impact: Medium

Customer circumvention: If the user or group UID or GID to be migrated is above the supported maximum value of 1,000,000,000 for the local provider, use LDAP mapping instead of local mapping.

Customer recovery steps: Remove the `AD<=>Local` user mapping rules for all affected users or groups.

Issue ID: 112689

Issue summary: Unauthorized NFS clients got access to an NFS share.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions

Issue description: NFS clients could encounter problems when mounting a nested directory from an export. The issue occurs when the export has a different client-access list when compared to the parent directory.

Symptoms: NFSv3 mount request succeeds for an export operation for a client that is restricted from accessing the export. In other words, the client is not part of the client-access list for the export.

Conditions of occurrence:

1. An export is configured on the parent directory with a client-access list as “*” and the nested export that gets created after having a restricted access to “10.x.x.22”
2. The nested export gets successfully mounted from "10.x.x.22" (restricted) till the NFSv3/NFSv4 mount of the parent export is attempted and data is accessed from the same client.
3. The issue occurs when the client-access list is specified in the subnetwork as well.

Impact: Medium

Customer circumvention: It is recommended to set a same client-access list for both the parent and child exports (nested) to avoid undesired NFSv3 access.

Customer recovery steps: Modify the parent and nested exports to have the same set of client-access lists.

Issue ID: 115337

Issue summary: The `setfs ad` command completes successfully, but the `showfs –ad` command indicates that the File Services nodes are not joined to Active Directory.

Affected platforms: StoreServ 7000, 8000, 9000, 20000, 20000 R2

Affected software versions: 3.3.1 MU2 P36

Issue description: An attempt to join a system using a previously working Active Directory provider appears to complete successfully. However, because the Active Directory provider no longer appears in the stacking order, the connection is unsuccessful. The system does not provide notification.

Symptoms: The `setfs ad` command completes successfully but the `showfs –ad` command indicates that the File Services nodes are not joined to Active Directory.

Conditions of occurrence: The system has been joined to Active Directory previously, but the Active Directory provider is not currently in the stacking order.

Impact: Medium

Customer circumvention: Use `showfs –auth` to verify that ActiveDirectory is listed in the stacking order before attempting to join the system to Active Directory using the `setfs ad` command.

Customer recovery steps:

Add the Active Directory provider back to the stacking order with the `setfs auth` command. Because Active Directory was previously joined, the command produces the following output:

```
 cli% showfs -ad
Domain Name  : <domainFQDN>
NetBIOS Name : <domainNetBIOS>
Forest       : <forestFQDN>
Status       : Online
```

Issue ID: 115360

Issue summary: The system returns a misleading message that `Another AD task is already running` after a successful request to join Active Directory.

Affected platforms: StoreServ 7000, 8000, 9000, 20000, 20000 R2

Affected software versions: 3.3.1 MU2 P36

Issue description: Because the active File Persona node is not the lowest numbered node, an attempt to join a system to Active Directory returns an unexpected message, even though the request completed successfully.

Symptoms: The system displays an unexpected message after an attempt to join the system to Active Directory.

Conditions of occurrence: The Active node for File Persona is not the lowest numbered node when attempting to join the system to Active Directory.

Impact: Low

Customer circumvention: None.

Customer recovery steps:

Ignore the message and use the `showfs -ad` command to confirm a successful join. The command returns a response similar to the following:

```
cli% showfs -ad
Domain Name  : <domainFQDN>
NetBIOS Name : <domainNetBIOS>
Forest       : <forestFQDN>
Status       : Online
```

Issue ID: 117638

Issue summary: The file services management service for a node indicates the node state as `starting` and does not proceed to a `running` state. This is a rare occurrence.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All

Issue description: In a rare scenario, the file services management service for a node continues to report the node state as `starting` and does not change to a `running` state. Active data services at the time of transition continue to run, but the system does not allow configuration changes to this node.

Symptoms: The `showfs` command shows a Starting state for one but not all of the nodes in the system.

Conditions of occurrence: The node has been up and running file services for more than a month without interruption.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: If a single node in the system reports a starting state continuously and for an extended period of time (greater than 30 minutes), stop and restart the node using the `stopfs <node>` command, followed by a `startfs -enable <node>` command.

Issue ID: 120330

Issue summary: Lost AD join after restart, including during upgrade or while switching domains.

Affected platforms: StoreServ 7000, 8000, 9000, 20000, 20000 R2

Affected software versions: 3.3.1 MU2 P39 and later

Issue description: Upgrading or switching domains can leave the HP-SMB in an unknown state. An attempt to join the nodes returns an incorrect message saying the nodes are already joined.

If HP-SMB can't connect to the customer Domain Controller at start-up (network disruption or unavailable DNS/DC), then domain/trust enumeration or refreshing a Kerberos Ticket with DC does not succeed. In this case, even if HP-SMB is joined to the domain, it cannot provide authentication/authorization services.

Symptoms: After upgrading, the system loses share access, and attempts to join the domain result in a message that nodes are already joined. A domain query of join status repots that the nodes are not joined to the domain.

Conditions of occurrence: Either DNS or the DC is not available at start-up and HP-SMB is already joined to a Domain.

Impact: Low

Customer circumvention: Make sure that the network infrastructure is healthy.

Customer recovery steps: None.

Issue ID: 120389

Issue summary: Cannot join new domain after leaving existing domain.

Affected platforms: StoreServ 7000, 8000, 9000, 20000, 20000 R2

Affected software versions: 3.3.1 MU2 P39 and later

Issue description: Joining Windows Domain (for new installation) or switching from Windows Domain A to Windows Domain B (for existing installation) or upgrade (File Persona was already joined), does not succeed. Either the new Windows DC FQDN is not registered in DNS or DNS does not have a reverse lookup pointer to the DC FQDN.

Symptoms: Leaving a domain and attempting to join a different domain is unsuccessful and returns a `Status 31` error. Repeated attempts to rejoin are also unsuccessful. When attempting to join a new Windows domain, or re-establish a Kerberos ticket time for a previously established join, File Persona sometimes cannot resolve the DC FQDN. Even if File Persona can resolve the request, HP-SMB sometimes cannot provide authentication.

Conditions of occurrence:

- The Active node for File Persona is not the lowest numbered node when attempting to join the system to Active Directory.
- DNS does not have a reverse pointer to DC FQDN.
- Previously, nodes were successfully joined to a domain.
- Leaving the joined domain and joining an entirely different domain after updating the DNS configuration.

Impact: Low

Customer circumvention: Make sure that the DNS contains correct forward and reverse entries to DC.

Customer recovery steps: Fix the DNS then retry Join Windows Domain operation.

Issue ID: 143701

Issue summary: All thinly provisioned FPGs deactivated when the backing CPG was almost out of space.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions

Issue description: When a CPG backing an FPG is almost out of space (as could be seen in the `free` column of the `showspace -cpg <cpgName>` command output), the system tries to automatically deactivate all thinly provisioned FPGs. The FPGs use a CPG to ensure that the FPGs do not experience a write error due to lack of space for additional allocations when backing the VVs. However, the alert that indicates that the FPG has been deactivated does not clearly show lack of space in the CPG.

Symptoms: An alert like the following would indicate that the CPG is nearly out of space:

```
Id          : 417
State       : New
Message Code: 0x027001d
Time        : 2018-05-18 14:42:59 MDT
Severity    : Critical
Type        : CPG growth failure non-admin
Message     : CPG FC_r6 SD and/or user space could not grow due to
unavailability of free space. CPG grow attempted using degraded availability
parameters (-ha mag) also failed.
              One or more alerts like the following would indicate that the
FPGs have been deactivated:
Id          : 418
State       : New
Message Code: 0x0720001
Catalog-Key : filesystem-event:filesystem.cmd.unmount
Time        : 2018-05-18 14:43:17 MDT
Severity    : Informational
Type        : File Provisioning Group
Message     : File Provisioning Group:<id>:<fpgName> Normal  (UNMOUNTED)
Details     : FPG Event: Unmounted FPG <fpgName> on host <nodeName>.
```

Conditions of occurrence: CPG space is exhausted but thinly provisioned FPGs are still in use.

Impact: Medium

Customer circumvention: Whenever thin provisioning is used, pay close attention to the remaining capacity in the CPGs. Respond to all alerts by adding capacity to the CPG or otherwise free up space in the CPG. An early indicator that CPGs dependent on a given storage class will soon run out of capacity could be seen with the following type of alert:

```
Id          : 452
State       : New
Message Code: 0x0270010
Time        : 2018-05-18 14:26:49 MDT
Severity    : Major
Type        : FC raw space allocation 85% alert
Message     : Total FC raw space usage at 10453G (above 91% of total 11442G).
```

This alert could be resolved by adding space of the given storage class.

Customer recovery steps: Resolve the CPG space issue and then manually activate the FPG again using the `setfpg` command.

Issue ID: 199108

Issue summary: When a patch (such as 3.3.1 MU2 P36) was unsuccessfully applied to the `tpd-fs` package, the correct status did not get reflected in the overall upgrade status.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All patches including the `tpd-fs` package

Issue description: When an update of the `tpd-fs` package is still in progress, but the update was unsuccessful, the overall status update could be reported as successful. The version of the file services on the controller nodes does not all match the version as reported by the `showfs` command.

Symptoms: Unsuccessful upgrade is reported in the verbose progress, but the final status of the upgrade is reported as successful. This could look something like the following:

```
[2017-01-23T19:57:59 HKT] Failed to upgrade node1fs.
[2017-01-23T19:57:59 HKT] current_version=1.2.2.6-20161017,
expected_version=1.2.3.2-20161117
[2017-01-23T19:58:00 HKT] Finalizing File Services. command=curl command to finalize
[2017-01-23T19:58:00 HKT] exit_status=0
[2017-01-23T19:58:02 HKT] File Services is healthy. (count=1)
[2017-01-23T19:58:02 HKT] File Services upgrade completed successfully.
```

After the upgrade, the output of the `showfs` command displays different versions across the set of nodes.

Conditions of occurrence: Application of a patch (such as 3.3.1 MU2 P36) which contains an update to the `tpd-fs` package.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: To correct the inconsistency in versions across the controller nodes, contact HPE Support.

Issue ID: 213022

Issue summary: Online or offline upgrade from 3.2.1 MU3 to 3.3.1 MU2 (or later) was unsuccessful.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: 3.3.1 MU2 or later

Issue description: Online or offline upgrade from 3.2.1 MU3 to 3.3.1 MU2 (or later) is unsuccessful if 3.2.1 MU3 P20 was not installed.

Symptoms: Online or offline upgrade from 3.2.1 MU3 to 3.3.1 MU2 (or later) is unsuccessful.

Conditions of occurrence: Attempting to upgrade from 3.2.1 MU3 to 3.3.1 MU2 without 3.2.1 MU3 P20 installed.

Impact: High

Customer circumvention: Install 3.2.1 MU3 P20 before attempting the upgrade.

Customer recovery steps: Install 3.2.1 MU3 P20 and retry the upgrade.

Issue ID: 228736

Issue summary: File services could not be started on a controller node.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions

Issue description: An internal service gets in a state where file services are not allowed on a controller node any longer.

Symptoms: When the `startfs –enable <node>` command is issued, output such as the following is seen in the task. File services remain unavailable.

```
2018-01-09 14:40:07 GMT Updated     Executing "startfsen_task -enable <node>" as 3:10968
…
2018-01-09 14:40:09 GMT Error       error: Failed to attach interface
2018-01-09 14:40:09 GMT Error       error: cannot fork child process: Cannot allocate
memory
```

Conditions of occurrence: Restarting file services after the controller node has been up and running for a long time.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Run the `shutdownnode check <node>` command to ensure the node can be rebooted safely. Restart the impacted controller node with the `shutdownnode reboot <node>` command. Alternatively, contact HPE Support for assistance in getting the service back in a correct state to avoid a controller node reboot.

Issue ID: 229706

Issue summary: The `growfpg` command was unsuccessful with the following message: `Error: command 'blockresize' requires <path>`.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions

Issue description: The `growfpg` command could be unsuccessful when an FPG grows.

Symptoms: When this issue occurs, the `growfpg` task does not complete successfully and returns a message like:

```
Unable to grow FPG <fpgName>. Error: Failed to refresh grown VV...error: command 'blockresize' requires <path>
option
```

Conditions of occurrence: Occurs when growing an FPG.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Issue the `growfpg -recover_storage` command and then retry the original `growfpg` command.

Issue ID: 233101

Issue summary: When a Remote Copy Group containing FPGs fails over, more than a permitted number of FPGs are listed on the target.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: 3.3.1 MU2 P36

Issue description: File Persona supports 16 FPGs per node pair. It is important to take this limit into account when configuring Remote Copy. If there are more FPGs present between the two sites and the FPGs failover, the FPGs cannot be activated. One additional FPG beyond the limit is recovered, but is left in a deactivated state.

Symptoms: After a Remote Copy failover, not all FPGs are listed. Only one FPG is listed in a deactivated state. In the following example, a 2-node system that supports 16 FPGs is shown:

```
cli% showfpg
                            ------(GiB)-------
FPG           --Mountpath--- --Size-- Available ActiveStates -DefaultCpg- --------VVs-------- State
Version
src_node1fs_0 /src_node1fs_0     0.00      0.00 DEACTIVATED  fs_cpg       src_node1fs_0.1_sec degraded 12.3
tgt_node0fs_0 /tgt_node0fs_0  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_0.1     normal   12.3
tgt_node0fs_1 /tgt_node0fs_1  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_1.1     normal   12.3
tgt_node0fs_2 /tgt_node0fs_2  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_2.1     normal   12.3
tgt_node0fs_3 /tgt_node0fs_3  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_3.1     normal   12.3
tgt_node0fs_4 /tgt_node0fs_4  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_4.1     normal   12.3
tgt_node0fs_5 /tgt_node0fs_5  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_5.1     normal   12.3
tgt_node0fs_6 /tgt_node0fs_6  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_6.1     normal   12.3
tgt_node0fs_7 /tgt_node0fs_7  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_7.1     normal   12.3
tgt_node1fs_0 /tgt_node1fs_0  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_0.1     normal   12.3
tgt_node1fs_1 /tgt_node1fs_1  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_1.1     normal   12.3
tgt_node1fs_2 /tgt_node1fs_2  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_2.1     normal   12.3
tgt_node1fs_3 /tgt_node1fs_3  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_3.1     normal   12.3
tgt_node1fs_4 /tgt_node1fs_4  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_4.1     normal   12.3
tgt_node1fs_5 /tgt_node1fs_5  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_5.1     normal   12.3
tgt_node1fs_6 /tgt_node1fs_6  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_6.1     normal   12.3
tgt_node1fs_7 /tgt_node1fs_7  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_7.1     normal   12.3
-------------------------------------------------------------------------------------------------------------
--
           17 total          16384.00  16373.12
```

Conditions of occurrence: Remote Copy configured with FPGs and more than the supported number of FPGs for the secondary site configured between the two sites.

Impact: Low

Customer circumvention: Do not configure more than the supported number of FPGs at the secondary site between the two sites when enabling Remote Copy.

Customer recovery steps: To revert to the supported number of FPGs, use the `removefpg -forget` command on the deactivated FPG. Failback the Remote Copy Group to the primary site. Configure fewer FPGs for Remote Copy or add controller nodes to support the aggregate set of FPGs between the sites.

Issue ID: 233575

Issue summary: An FPG AutoFailover was unsuccessful when the Compliance feature was enabled on the primary array but not enabled on the target or secondary array.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: 3.3.1 MU2 P36

Issue description: An FPG AutoFailover could be unsuccessful when the Compliance feature is enabled on the primary array but is not enabled on the target or secondary array. The issue occurs when a primary array that is compliance-enabled has a Compliance policy applied to one or more FPGs. During failover, the FPG that is compliance-enabled will not succeed as expected. As the FPGs are recovered serially, all the FPGs that come after the compliance-enabled FPG will not get enabled on the target array.

Symptoms: A `createfpg -recover` operation did not complete successfully and returned the following status in the task:

```
Could not activate FPG <fpgName> as it has compliance enabled object(s). To
activate this FPG, enable ComplianceOfficerApproval using "setsys" and try
again.
```

Conditions of occurrence: The Compliance feature is enabled on the primary array but not on the secondary array. Some but not all FPGs are compliance-enabled.

Impact: Medium

Customer circumvention: Ensure that the Compliance feature is enabled on both the source and target arrays.

Customer recovery steps: Enable FPGs that are not compliance-enabled manually using the `createfpg -recover set:<FPG-name>` command on the secondary array.

Issue ID: 233780

Issue summary: When failback after a failover was triggered using a pushbutton failover using SSMC, not all FPGs were activated successfully.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: 3.3.1 MU2 P36

Issue description: When multiple Remote Copy Groups are failed over or failed back in parallel, some of the FPGs might not be recoverable and activated successfully. They will be left in a deactivated state.

Symptoms: One or more FPGs could be in a deactivated state on the arrays to which the FPGs are intended to be activated. The error message when trying to manually activate such an FPG could look like the following:

```
Failed to activate <fpgName>: node <nodeName> has no access to the volumes
[[<vvId>]] used by filesystems [<fpgName>] (suggested action - check volume
access on the specified nodes)
```

Conditions of occurrence: Failover or failback of multiple Remote Copy Groups containing FPGs.

Impact: Medium

Customer circumvention: Failover or failback one Remote Copy Group at a time.

Customer recovery steps: Recover the missing FPGs manually using `removefpg -forget <FPG-name>` command. Follow the task to completion using `waittask -v <taskid>`. Recover the FPG by calling the `createfpg -recover set:<FPG-name>` command and following the task to completion.

Issue ID: 235843

Issue summary: The `checkvv -dedup_dryrun <vvName>` command was not permitted when the VV belonged to an activated FPG.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions 3.2.2 MU2 or later

Issue description: If data is being written on an FPG, do not run `checkvv` on the activated FPG. Even though the `-dedup_dryrun` command does not modify the volume, it is not allowed on an activated FPG.

Symptoms: The following message is displayed when attempting to run the `checkvv -dedup_dryrun <vvName>` command:

```
<fpgName>.<N> has an associated activated FPG. Cannot run checkvv on VV of an
activated FPG. Please deactivate FPG and re-run checkvv
```

Conditions of occurrence: Attempting to run a `dedup dryrun check` on a VV that is part of an activated FPG.

Impact: Medium

Customer circumvention: Select only one of the following options:

- Deactivate the FPG. Run the check and reactivate the FPG. This will make the FPG unavailable to clients.
- Create a Virtual Copy of the FPG volumes. Run the `checkvv` command on the volumes. Delete the Virtual Copy.

Customer recovery steps: None

Issue ID: 243445

Issue summary: An FPG recovery operation returns a message that the operation was not successful.

Affected platforms: StoreServ 7000c, 8000, 9000, 20000, 20000 R2

Affected software versions: All versions 3.2.2 MU2 or later.

Issue description: An FPG recovery operation targeted to a volume that is already associated with an activated FPG returns a message that the operation was not successful. There is no impact to the activated FPG when File Services are in a `running` state on all nodes in the system. When File Services is in a `starting` state on one of the nodes, the FPG can become unavailable.

Symptoms: After attempting a recover operation on a volume, the system returns a message that the operation was not successful, and that the FPG using that volume is unavailable.

Conditions of occurrence: This issue occurs when an FPG is in an activated state and the request to recover an FPG uses one of the volumes of the activated FPG.

Impact: High

Customer circumvention: Use the `showfpg` command to show whether the volume belongs to an FPG.

If the volume does belong to a listed FPG, do not attempt a recover operation. Use the recover operation only when volumes belonging to FPGs were forgotten using the `removefpg -forget` option.

Customer recovery steps: Make sure that File Services are in a `running` state on all of the nodes in the system.

Deactivate the FPG, if possible, and then reactivate it.

HPE 3PAR 3.3.1 CLI MU3 Release Notes

Supported Operating Systems

For releases following HPE 3PAR OS 3.3.1 MU3, Hewlett Packard Enterprise will no longer support the HPE 3PAR Command Line Interface on the following Operating Systems:

OS Family	Operating System
Linux	RHEL 5
	SLES 10
AIX	AIX 6.1
Microsoft Windows	Windows Server 2008 R2
Microsoft Windows Enterprise	Windows Enterprise 8

What's New in the CLI

- A new security feature was added to CLI, a banner for Remote CLI client that is displayed before prompting for user and password.

  Existing **`removesshbanner`**, **`setsshbanner`** and **`showsshbanner`** are deprecated and new commands **`removebanner`**, **`setbanner`** and **`showbanner`** are added.

  New options **`-ssh`**, **`-cli`** or **`-all`** are added to above three commands. The default is **`-ssh`** for the existing SSH banner.

  In order to help with scripting, a new CLI global option, **`-nobanner`**, was added to suppress the banner output to **`stderr`**.
- An option has been added to the CLI **`help`** command to display the license text associated with open source products bundled on the array. This information is accessible via the command **`help opensource`**.

New Commands

- **`controlsecurity`**
- **`removebanner`**
- **`setbanner`**
- **`showbanner`**
- **`statwsapi`**

Deprecated Commands

- **removesshbanner**
- **setsshbanner**
- **showsshbanner**

Changed Commands

Command	Description
controlport	Added **fs ping** and **traceroute** support
createflashcache	Updated help to show reduced flash cache sizes
createfsgroup	Update the GID range
createfshare	Moved audit options for NFS protocol to new option: **-audit {operation1:value1[,operation2:value2]...}** Added new audit option, **audit_acl**
createfsuser	Updated the range for UID
creategroupvvcopy	Add **-pri** to various region mover commands
createsralertcrit	Add **-defer** option Add bandwidth percentages fields to port type criteria Add **-groupby** option to most **sralertcrit** types
createvv	Min granularity for **-exp/-retain** is in minutes (M|m)
createvvcopy	Add **-pri** to various region mover commands
promotegroupsv	Add **-pri** to various region mover commands
promotesv	Add **-pri** to various region mover commands
removesshbanner	Deprecated, rename to **removebanner** and add options
removefpg	Additional note added related to compliance in Help text
removefsarchive	Additional note added related to compliance in Help text
removewsapisession	Add new option **-close_sse** to close the SSE connection channel
setsshbanner	Deprecated, rename to **setbanner** and add options
setcim	add new **-pol** setting **tls_strict**

Table Continued

Command	Description
setfpg	**-version** support added
setfs	LDAP resilience feature is added SMBv1 feature added
setfsaudit	For 3.3.1 MU2, **-size** and **-time** both are supported **tz** subcommand, local timezone support added for audit log files
setfshare	Moved audit options for NFS protocol to new option: **-audit {operation1:value1[,operation2:value2]...}** Added new audit option, **audit_acl**
setsralertcrit	Add **-defer** option Add bandwidth percentages fields to port type criteria Clarify that multiple criteria names would be specified separated by spaces Add **-groupby** option to most **sralertcrit** types
setsys	Add warnings for over provisioning Add new R6LayoutVersion to allow changing between RAID6 implementations Additional note added related to compliance in Help text **SingleLunHost** added
setuser	Removed CO info from notes as **browse->co** is no longer supported
setvv	Min granularity for -exp/-retain is in minutes (M|m)
setwsapi	Add the option **-evtstream** to enable/disable event stream state Add **-pol** option
showsshbanner	Deprecated, rename to **showbanner** and add options
showcim	add new **-pol** setting **tls_strict**
showfs	LDAP resilience feature added, **showfs -ldap** example updated Added info regarding supported versions new **nfs** to display global health
showfsaudit	Local timezone support added for audit log files

Table Continued

Command	Description
showsysmgr	New status to **showsysmgr** to indicate when a PD cage position refresh is pending
showwsapi	Add option **-evtstream**. The Event Stream State will be displayed in **-d** output Add policy info to **-d** output
tunesys	Added **-noinitialcompacts**
tunevv	Add **-pri** to various region mover commands
upgradecage	Added **-parallel** and **-status** options, removed deprecated **minlevel/maxlevel/model** options.

Modifications to the CLI

Issue ID: 127426

Issue summary: The CLI command **setfsquota** times out when updating quota for fstore created with **&** (ampersand) in fstore name.

Affected platforms: All StoreServ

Affected software versions: 3.1.3, 3.2.1, 3.2.2, 3.3.1, 3.3.1 MU1

Issue description: The CLI command **setfsquota** will no longer time out with the below error when updating quota for fstore created with **&** (ampersand) in fstore name.

```
root@MXN6232L1W-1 Fri Feb 09 05:12:24:~# setfsquota -fpg examplefpg -fstore
'examplefstore&' -scapacity 900 -hcapacity 1000 'examplevfs&'

Error: status=500,exception message=Error unmarshalling: Premature end of file.
```

Symptoms: The command to set quota for a fstore times out with error, `Error: status=500,exception message=Error unmarshalling: Premature end of file.`

Conditions of occurrence: A fstore was created with a name containing the **&** (ampersand) character.

Impact: Medium

Customer circumvention: Avoid creating a fstore with a name containing the **&** (ampersand) character in the affected software versions listed above.

Customer recovery steps: Use the **setfsquota -restore** command option.

Issue ID: 139804

Issue summary: Under some circumstances, it is not possible to start a new **tunesys** task for a time after canceling a previous instance.

Affected platforms: All StoreServ

Affected software versions: 3.1.1 and later

Issue description: The **tunesys** command starts a number of subordinate tasks while rebalancing a system. One of these tasks is **compactcpg** which is used to free space released by tuning. If **tunesys** is canceled while running **compactcpg**, the **tunesys** task moves to a canceled state. The exclusive locking mechanism used to prevent multiple instances is not released until the compact completes. A new instance cannot be started until the lock is released.

Symptoms: User is unable to start a new **tunesys** instance for a time after an existing instance is canceled. The following message is displayed:

```
A tunesys task is already running on this storage system.
```

Conditions of occurrence: **Tunesys** has been canceled. Attempting to restart, the message `A tunesys task is already running on this storage system` appears.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: A **tunesys** instance has been canceled, and a new instance will not start. The following message appears: `A tunesys task is already running on this storage system`, use the following command to find any running **compactcpg** instances:

```
showtask -active
```

Use the **canceltask** command to cancel any executing **compactcpg** commands.

Issue ID: 162449

Issue summary: Only one volume greater than 16TiB in size can be tuned at a time.

Affected platforms: All StoreServ

Affected software versions: 3.3.1, 3.3.1 MU1, 3.3.1 MU2

Issue description: A limitation was placed on the 3.3.1 through 3.3.1 MU2 releases which only allowed a single volume greater than 16TiB to be tuned at a time.

This limitation has been relaxed. Up to 5 large volumes can now be tuned at once.

Symptoms: In 3.3.1 through 3.3.1 MU2, if a second large volume tune was started, the following error message would be displayed: `A large volume tuning task is already running on this storage system.`

Conditions of occurrence: A second large volume tune is started.

Impact: Medium

Customer circumvention: In 3.3.1.MU3 and later, up to 5 large volumes may be tuned in parallel.

Customer recovery steps: No recovery is required for this issue.

Issue ID: 167884

Issue summary: **`Tunesys`** analysis can be inaccurate if a system is low on free space but space is available in some uncompacted CPGs.

Affected platforms: All StoreServ

Affected software versions: 3.1.1 and later

Issue description: If the system is low on free space but space is available in some uncompacted CPGs, the initial analysis phase of **`tunesys`** can give inaccurate results. The analysis phase needs a balanced amount of free space available on the system otherwise the projected balance figures will be inaccurate.

Symptoms: **`Tunesys`** analysis is inaccurate. Some tunes may be missed and the system may not be balanced completely after **`tunesys`** completes.

Conditions of occurrence: Uncompacted CPGs exist on the system.

Impact: Medium

Customer circumvention: Manually compact all relevant CPGs, then rerun **`tunesys`** with the **`-force`** option to ensure that all volumes are analyzed and rebalanced.

Customer recovery steps: None.

Issue ID: 179028

Issue summary: The error message produced is truncated and is missing diagnostic information when **`tunevv`** is used to tune a large volume in slices and there is not enough clean free space to tune a slice.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 and later

Issue description: When **`tunevv`** is used in sliced tuning mode to tune a large volume, a certain amount of clean free space is needed to perform the analysis. If there is not enough space to tune a slice, the error message is truncated.

Symptoms: **`Tunevv`** reports partial error messages when space is low on a system. The messages are missing important information on the configuration of the space needed (RAID type, set size, availability).

Conditions of occurrence: System is low on space.

Impact: Low

Customer circumvention: Two options are possible:

1. Manually compact CPGs or delete volumes to free up space before a tune. There should be a delay between compaction/volume removal and starting the tune to allow for chunklet cleaning.
2. Use a smaller slice threshold and size. The default is to tune volumes greater than 16TiB in size in slices of 2TiB. The range for these options are as follows:

   `-slth` - slice threshold - Minimum 128GiB, maximum 16TiB, default 16TiB in multiples of 128GiB.

   `-slsz` - slice size - Minimum 128GiB, maximum 2TiB, default 2TiB in multiples of 128GiB.

Customer recovery steps: None.

Issue ID: 192194

Issue summary: User authentication issue causing physical disk over allocation.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 MU3

Issue description: Knowledgeable users with service and browse roles set are able to run certain commands that could result in over-allocation of physical disk space.

These commands were not applicable to service and browse users and execution by these user roles caused the array to over-allocate physical disk space. This issue has now been corrected.

Symptoms: Commands could be executed by a user with service or browse role set.

These commands resulted in multiple unused logical disks being created until the following error was seen: `Error - unable to create new LD - error: Could not find enough available disk space.`

Conditions of occurrence: Run commands from a user with a service or browse role set.

Impact: High

Customer circumvention: Do not execute internal CLI commands from user accounts with service or browse roles.

Customer recovery steps:

If the unused LDs are associated with a CPG, run **`compactcpg <cpg_name>`**.

If the unused LDs are outside of any CPG, they must be deleted manually using **`removeld <ld_name>`**.

NOTE: In both cases, the space made available will not be available to the system until chunklet cleaning completes.

Issue ID: 198511

Issue summary: Unclear error messages returned from **stoprcopygroup -pat** command.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description:

The **-pat** option is used with **stoprcopygroup** on an array where no matching groups are Primary.

In this case the list of groups is empty, causing a generic help message to be returned.

CLI Change: If no groups were found, return `No matching groups` message.

The list of groups contains more than the supported number of group names.

The error message indicating the problem is returned along with the generic help message.

CLI Change: Clarify the error message to indicate the number of groups specified and remove the generic help text.

Symptoms: After issuing the **stoprcopygroup -pat** command under the above conditions, a command usage statement is returned.

Conditions of occurrence: Use of **stoprcopygroup -pat** on an array where no Primary remote copy groups exist.

Use of **stoprcopygroup -pat** where the number of groups in the resulting pattern matched set exceeds the limit (512) for a single request.

Impact: Low

Customer circumvention: Issue **stoprcopygroup -pat** on array with at least one Primary group matching the pattern.

Issue **stoprcopygroup -pat** such that the matching pattern result is less than or equal to 512 groups.

Customer recovery steps: None.

Issue ID: 203068

Issue summary: **Tunesys** generates minor alerts when started.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU2 and earlier

Issue description: When **tunesys** starts, it can generate a spurious Minor severity alert.

Symptoms: Running the **tunesys** command generates a Minor Severity alert of the following form:

```
    Notification  Minor     Command error   sw_cli
        Command: setexclusive query tunesys Error: Internal error: Could not
get entry for pr table cli_exclusive, Invalid key
```

Conditions of occurrence: Occurs occasionally when **tunesys** is started.

Will not occur on 3.3.1 MU3 and later.

Impact: Low

Customer circumvention: The alert can be safely ignored.

Customer recovery steps: None.

Issue ID: 204958

Issue summary: New security banner for CLI.

`setbanner/removebanner/showbanner` (new names) were extended for CLI banners and **`-nobanner`** added.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU3

Issue description:

Existing **`removesshbanner`**, **`setsshbanner`** and **`showsshbanner`** are deprecated and new commands **`removebanner`**, **`setbanner`** and **`showbanner`** are added.

New options **`-ssh`**, **`-cli`** or **`-all`** are added to above three commands. The default is **`-ssh`** for the existing SSH banner.

In order to help with scripting, a new CLI global option, **`-nobanner`**, was added to suppress the banner output to stderr.

Symptoms: None.

Conditions of occurrence: None.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 206642

Issue summary: **`checkhealth`** does not report degraded peer logical or physical disks.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 MU2, 3.3.1 GA - MU2

Issue description: The **`checkhealth pd`** and **`checkhealth ld`** CLI commands do not detect degraded physical disks and logical disks. The command has been updated so that these degraded objects will also be reported.

Symptoms: **`Checkhealth`** does not report any issues for degraded physical or logical disks. .ssr volumes are not removed.

Conditions of occurrence: This issue is caused by some degraded not started logical disks and degraded physical disks.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Upgrade to 3.3.1 MU3, so that **`checkhealth`** will find and report these disks. Check and remove degraded logical and physical disks. .ssr volumes start to be deleted.

Issue ID: 208555

Issue summary: **`Checkhealth`** now correctly detects System Report space approaching full.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU2

Issue description: The **`checkhealth`** command previously reported an issue with the System Reporter space if it was 90% full. With the release of 3.3.1, the behavior of SR space management changed such that 95% of SR space is used as a threshold rather than 75%. This threshold leads to an incorrect identification of an issue within **`checkhealth`**. The **`checkhealth`** command now only reports an issue if the SR space utilized is at 97%.

Symptoms: The **`checkhealth`** command reports an issue with the System Reporter space if it is 90% full.

Conditions of occurrence: Once System Reporter has collected significant data, **`checkhealth`** will always report that `/sr_mnt` is over capacity, even though System Reporter is actively managing its space.

Impact: High

Customer circumvention: None.

Customer recovery steps: Upgrade to MU3.

Issue ID: 209367

Issue summary: **`tunesys`** terminates with the error `Reason*: divide by zero`.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description: **`tunesys`** exits during the inter-node re-balancing phase. This is caused by a calculation error when incorrect usage information is used to monitor progress while waiting for chunklets to clean.

Symptoms: **`tunesys`** does not complete, and exits with an error.

Conditions of occurrence: Chunklet cleaning is in progress.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Run the **`tunesys`** command again.

Issue ID: 212298

Issue summary: Allow combination of **-rcp** and **-pri** options to the CLI commands **promotesv** and **promotegroupsv**.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1 GA - MU2

Issue description: The CLI commands **promotesv** and **promotegroupsv** can now be executed with both the **-rcp** and **-pri** options together, which allow promotion within a Remote Copy volume tree (-rcp) and specifying a task priority (-pri) in the same operation.

Symptoms: Attempting to promote a snapshot within a Remote Copy volume and assign a priority to the copy task.

Conditions of occurrence: Attempting to promote a snapshot within a Remote Copy volume and assign a priority to the copy task.

Impact: Low

Customer circumvention: The task priority can be set with **settask -pri <priority> <taskid>** after starting a **promotesv** task.

Customer recovery steps: The task priority can be set with **settask -pri <priority> <taskid>** after starting a **promotesv** task.

Issue ID: 213662

Issue summary: When upgrading the system to 3.3.1 GA, 3.3.1 MU1, or 3.3.1 MU2 if the system contains only system volumes, and has cages with the old firmware, the Service Processor might upgrade only a portion of cages. The same issue might occur when the `Admit hardware` action is started.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU2

Issue description: Alerts are generated with `Interface Card Firmware Out of date` message. The enclosure health shows `Degraded`.

Symptoms: The Service Processor reports `Cage not on current firmware` after it finishes the system upgrade. `Check health` reports the same error.

Conditions of occurrence: When the state of cage firmware is not refreshed in the system cache soon enough to reflect the current state, and the upgrade process reads the stale data and stops more cage upgrades.

Impact: Low

Customer circumvention: None.

Customer recovery steps: Start the action `Admit hardware` from the Service Processor several times until no `Cages not on current firmware` is reported.

Issue ID: 214104

Issue summary: **`setuser -h`** incorrectly states that `Browse role should be always in a specific domain in order to co-exist along with co role`. Based on case 208806, a user can have same role in all the domains where they hold accounts. To create any user with role co, always use the command **`createuser`**.

Affected platforms: All StoreServ

Affected software versions: 3.3.1

Issue description: A browse user cannot be elevated to CO role.

Symptoms: In older builds of 1.4, which has case 208806, **`setuser -h`** allows a user with browse role to be made CO by changing their domain to **`all`** and role to **`co`**.

Conditions of occurrence: Certain 3.3.1 MU1 builds have this issue. Later builds have the updated help text so that the user is not confused with the NOTE in **`setuser -h`** conflicting with the actual behavior.

Impact: Low

Customer circumvention: Do not use **`setuser`** command to make an existing browse user a CO user. Always use command **`createuser`** to create one.

Customer recovery steps: None.

Issue ID: 215044

Issue summary: The **`SingleLunHost`** system parameter has been added to limit Virtual Volume exports.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2 GA - MU5, 3.3.1 GA - MU2

Issue description:

Symptoms: Without the system parameter enabled, users can export a VV to the same host multiple times.

Conditions of occurrence: Creation of exports such that the same Virtual Volume is exported multiple times to a given host.

Impact: Medium

Customer circumvention: To avoid these exports, customers can also heed the warnings to avoid creating them.

Customer recovery steps: None.

Issue ID: 215420

Issue summary: Contrary to the `setsshbanner` documentation, there is no checking for the 4,095 byte limit.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2, 3.3.1 GA, 3.3.1 MU2

Issue description: There is no limit checking on the number of bytes that can be entered in `setsshbanner`. The documented limit is 4,095 bytes. The following error occurs: `Banner text exceeds 4095 bytes`.

Symptoms: There is no limit check for the number of bytes that can be entered for the `setsshbanner`.

Conditions of occurrence: Input for the SSH banner is more than 4,095 bytes.

Impact: Low

Customer circumvention: Limit input to 4,095 bytes as documented.

Customer recovery steps: None.

HPE 3PAR 3.3.1 CIM API MU3 Release Notes

What's New with the CIM API and SNMP Software

New and enhanced features include:

- The `IOServicetime` property has been added to the CIM API objects named `TPD_LUNStatisticalData` and `TPD_NodeStatisticalData`. `StartStatisticTime` and `StatisticTime` which define the time period over which statistics were gathered are now defined for both `TPD_LUNStatisticalData` and `TPD_NodeStatisticalData`. `IOServicetime` may be calculated from properties already present using the formula **`IOTimeCounter/TotalIOs`**, but the result would be in clock ticks instead of microseconds. The time for a clock tick varies with the StoreServ hardware.
- The SNMP `authenticationFailure` trap (OID .1.3.6.1.6.3.1.1.5.5) message now includes the IP address of the manager that caused the authentication failure.
- Both the CIM API and SNMP support FIPS 140-2 compliance using the OpenSSL FIPS 140-2 certified cryptographic module. FIPS 140-2 compliance is managed using the controlsecurity fips cli command.
- A new system parameter, **`SingleLunHost`**, has been added that limits exports of a given Virtual Volume (VV) to one per host. In addition, port exports and matched set exports are prohibited when this system parameter is enabled.

Modifications to the 3PAR CIM API

Issue ID: 192537

Issue summary: Frequent CIM API requests by partner software for StoreServ cage, magazine, and controller node manufacturing information are causing the system to erroneously report device failure events.

Affected platforms: StoreServ 7000, StoreServ 20000, StoreServ 20000R2

Affected software versions: 3.2.2.GA MU5, 3.3.1 GA - MU2

Issue description: Frequent CIM server API calls by third-party software that retrieve manufacturing information for cages, magazines, and controller nodes occur. These calls result in invalid error messages and events being logged because the internal interface is busy.

Symptoms: Alerts for system manager errors.

Invalid events for cage and magazine failures.

Conditions of occurrence: CA Unified Manager or similar tools are actively monitoring a StoreServ system.

Impact: High

Customer circumvention: Turn off the software that is issuing CIM API requests. For CA Unified Manager, disable monitoring.

Customer recovery steps: None.

Issue IDs: 198548

Issue summary: The **`cimserver`** process is causing System Manager to log nuisance debug messages in the syslog file.

Affected platforms: All StoreServ

Affected software versions: 3.2.1, 3.2.2.MU6, 3.3.1 GA - 3.3.1 MU2

Issue description: A debug message appears in the syslog log file many times. The message is filling up the log file and implies an error has occured when none has.

Symptoms: The following message appears many times in the /var/log/tpd/syslog file:

```
{MC_OBJDESC_GET } objdesc.c: pr_read_value_table() failed: 9
```

Conditions of occurrence: The cim server is enabled and running.

Impact: Low

Customer circumvention: Stop the cim server with the **`stopcim`** CLI command.

Customer recovery steps: None.

Issue ID: 207552

Issue summary: If CIM server deadlocks during patch installation, StoreServ alerts are not delivered.

Affected platforms: All StoreServ

Affected software versions: 3.2.2 GA - MU5, 3.3.1.GA

Issue description: CIM server occasionally deadlocks during its process shutdown. When this lockup occurs, `em_filter`, which delivers events to system processes, stalls while waiting for CIM server to process messages. Events are not forwarded as a result.

Symptoms: Processes like the Web Services API interface and StoreServ Management Console will not receive system event notifications.

Conditions of occurrence: A StoreServ patch is installed.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: Stop the CIM server process and restart it with the following commands:

```
stopcim -f
startcim
```

Issue ID: 208562

Issue summary: CIM server association between `TPD_NodeSystem` and `TPD_Fan` returns an incorrect result. Customer and/or partner software may report incorrect StoreServ controller node fan inventory information.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description: CIM server incorrectly fails to associate a fan with a cluster controller node.

Symptoms: Inventory software does not report a fan, or does not report which controller node a fan is connected to.

Conditions of occurrence: Software uses CIM API to gather inventory information on system fans.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 209992

Issue summary: Uninitialized scalar variable in CIM API could cause a program error.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 GA - MU2

Issue description: Uninitialized scalar variable in CIM API could cause a program error.

Symptoms: None.

Conditions of occurrence: None.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 213188

Issue summary: Update the CIM API internal interfaces to scale better when enumerating logical devices.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU2

Issue description: The CIM API does not take advantage of an internal interface that returns Logical Devices (LDs) in chunks, which allows CIM API to scale up to a large number of logical disks.

Symptoms: Improves performance.

Conditions of occurrence: None.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 213365

Issue summary: The CIM API erroneously deletes a target virtual volume set when an attempt to create a clone of another fails.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description: A call to **`CreateGroupReplicaFromElements`** to create a physical copy of a Source Storage Group (known as a vvset in 3par parlance) will delete the Target Storage Group if the number of elements in the two storage groups is not the same. Note, no volumes will be deleted.

Symptoms: The vvset represented by the Target Storage Group in the failed **`CreateGroupReplicaFromElements`** call will be deleted.

Conditions of occurrence: Run **`TPD_ReplicationServices.CreateGroupReplicaFromElements`** function for a clone with a Source Storage Group and a Target Storage Group specified. The source and target groups must have a different number of elements.

Impact: Low

Customer circumvention: For **`CreateGroupReplicaFromElements`** specify a source and target Storage Group with the same number of elements, or specify a target Storage Group with no elements.

Customer recovery steps:

Re-create the Storage Group that was deleted by re-running **`CreateGroupReplicaFromElements`** with the same arguments. The group will be created and the cloned volumes will be added to the group. The re-created group will now contain the volumes it originally contained.

Alternatively, run the CLI command **`creategroup`**.

Issue ID: 224614

Issue summary: The snmpagent unexpectedly restarts.

Affected platforms: All StoreServ

Affected software versions: 3.2.2, 3.3.1

Issue description: The snmpagent slowly leaks memory with each trap it delivers. Eventually it aborts and restarts when it runs out of memory.

Symptoms: The snmpagent unexpectedly restarts and snmpagent core files appear in /var/core/proc/saved.

Conditions of occurrence: snmp managers are configured. The snmpagent is delivering SNMP notifications, or traps, to these managers.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None needed. The snmpagent process automatically restarts.

Issue IDs: 227929

Issue summary: 3.3.1.MU1 ThreeParMIB does not work on RHEL with **`snmptranslate`** command.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU1, 3.3.1 MU2

Issue description: Users on the RHEL platform cannot import the ThreeParMIB.mib on RHEL 7.4 with the **`snmptranslate`** command.

Symptoms: Running the **`snmptranslate -m MIB:/usr/share/snmp/vendormibs/ThreeParMIB.mib .1.3.6.1.4.1.12925.1.2`** command results in the following error:

```
Bad parse of NOTIFICATION-TYPE: At line 2505 in /usr/share/snmp/vendormibs/
ThreePar-MIB.txt
```

Conditions of occurrence: As stated in symptoms.

Impact: Low

Customer circumvention: Edit the ThreeParMib.mib file and change the line that is causing the error:

```
2499         checkSnmpAlert NOTIFICATION-TYPE
2500                 OBJECTS     { component, details, nodeID, severity, timeOccurred,
id,
2501                               messageCode, state, serialNumber,
2502                               tier, sparePartNumber }
2503                 STATUS      current
2504                 DESCRIPTION "check snmp test alert"
2505                 ::= { storeServAlerts 4294967295 }   < Change 4294967295 to
4294967294
```

Customer recovery steps: None.

HPE 3PAR 3.3.1 Web Services API MU3 Release Notes

What's New with the Web Services API Software

New and enhanced features include:

- Added WSAPI support for Peer Persistence.
- Added WSAPI Support for events and alerts.
- Added WSAPI support for target driven zoning (Smart SAN).
- Added WSAPI support to filter volumes by provisioning type.
- Added support to add/remove a target to/from a remote copy group (**`admitrcopytarget/dismissrcopytarget`**).
- Added ability to modify the WWN of a virtual volume or of a virtual copy at creation time.
- Improved support for switch port inventory.
- Improved support for System Reporter options.
- Improved WSAPI connection security options.

Modifications to the 3PAR Web Services API

Issue IDs: 200657

Issue summary: WSAPI Spelling mistake.

Affected platforms: All StoreServ

Affected software versions: 3.3.1 MU2

Issue description: The license information key for `thinProvisioning` is incorrectly spelled as `thinProvisioing`.

Symptoms: An incorrectly spelled field.

Conditions of occurrence: Always occurs.

Impact: Low

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 220990

Issue summary: WSAPI does not use updated user privilege if user privilege has been changed.

Affected platforms: All StoreServ

Affected software versions: 3.3.1

Issue description: WSAPI does not use updated user privilege if the user is removed or recreated with different user privileges, or if the user privilege is modified.

Symptoms: WSAPI uses old user privilege is a user privilege has been modified.

Conditions of occurrence:

1. User creates WSAPI credential with old privilege; and sends WSAPI request with that credential.
2. User privilege is modified.
3. WSAPI will serve upcoming requests from that user with the user's old privilege.

Impact: Medium

Customer circumvention: Restart WSAPI server if a user privilege has been modified.

Customer recovery steps: Restart WSAPI server.

Issue ID: 227775

Issue summary: WSAPI does not accept some special characters in password

Affected platforms: All StoreServ

Affected software versions: 3.3.1

Issue description: WSAPI returns `internal server error` for any request with a session key which was created with a user with special characters in the password. The special characters could be **`$`** or **`[`**.

Symptoms: WSAPI returns `internal server error`.

Conditions of occurrence: User password contains characters **`$`** or **`[`**.

Impact: Low

Customer circumvention: Change the password to avoid the characters **`$`** or **`[`**.

Customer recovery steps: None.

Issue IDs: 243195

Issue summary: WSAPI does not allow to modify **`disableCompr`** system parameter.

Affected platforms: All StoreServ

Affected software versions: 3.3.1

Issue description: WSAPI does not allow to set **`disableCompr`** system parameter to **`true`**. Instead use the CLI command, **`setsys DisableCompr`** to modify this parameter.

Symptoms: WSAPI will generate error when setting **`disableCompr`** to true.

Conditions of occurrence: Use HTTP PUT request to modify the **`disableCompr`** system parameter.

Impact: Medium

Customer circumvention: Do not use WSAPI to set **`disableCompr`** system parameter.

Customer recovery steps: Use the CLI command **`setsys DisableCompr`**, to set the parameter.

Component Versions

Table 3: Components and Versions

Component	Version
CLI Server	3.3.1.460 (MU3)
CLI Client	3.3.1.460
System Manager	3.3.1.460 (MU3)
Kernel	3.3.1.460 (MU3)
TPD Kernel Code	3.3.1.460 (MU3)
CIM Server	3.3.1.460 (MU3)

Table Continued

Component	Version
WSAPI Server	3.3.1.460 (MU3)
Console Menu	3.3.1.460 (MU3)
Event Manager	3.3.1.460 (MU3)
Internal Test Tools	3.3.1.460 (MU3)
LD Check Tools	3.3.1.460 (MU3)
Network Controller	3.3.1.460 (MU3)
Node Disk Scrubber	3.3.1.460 (MU3)
PD Scrubber	3.3.1.460 (MU3)
Per-Node Server	3.3.1.460 (MU3)
Persistent Repository	3.3.1.460 (MU3)
Powerfail Tools	3.3.1.460 (MU3)
Preserved Data Tools	3.3.1.460 (MU3)
Process Monitor	3.3.1.460 (MU3)
Software Updater	3.3.1.460 (MU3)
TOC Server	3.3.1.460 (MU3)
VV Check Tools	3.3.1.460 (MU3)
Upgrade Check Scripts	180829.U015
File Persona	1.5.2.8-20180817 (MU3)
SNMP Agent	1.13.0
SSH	7.5p1-5
VASA Provider	3.0.18 (MU3)
Firmware Database	3.3.1.460 (MU3)
Drive Firmware	3.3.1.460 (MU3)
UEFI BIOS	05.04.04 (MU3)
MCU Firmware (OKI)	4.9.01 (MU3)

Table Continued

Component	Version
MCU Firmware (STM)	5.4.00 (MU3)
Cage Firmware (DC1)	4.44 (MU3)
Cage Firmware (DC2)	2.64 (MU3)
Cage Firmware (DC3)	08 (MU3)
Cage Firmware (DC4)	2.64 (MU3)
Cage Firmware (DCN1)	4082 (MU3)
Cage Firmware (DCN2)	4082 (MU3)
Cage Firmware (DCS1)	4082 (MU3)
Cage Firmware (DCS2)	4082 (MU3)
Cage Firmware (DCS5)	2.88 (MU3)
Cage Firmware (DCS6)	2.88 (MU3)
Cage Firmware (DCS7)	4082 (MU3)
Cage Firmware (DCS8)	4082 (MU3)
QLogic QLA4052C HBA Firmware	03.00.01.77 (MU3)
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	05.07.36
Emulex LP11002 HBA Firmware	02.82.x10
Emulex LPe12002 HBA Firmware	02.10.x08
Emulex LPe12004 HBA Firmware	02.10.x08
Emulex LPe16002 HBA Firmware	11.1.220.10
Emulex LPe16004 HBA Firmware	11.1.220.10
3PAR FC044X HBA Firmware	200A8

Table Continued

Component	Version
LSI 9201-16e HBA Firmware	17.11.03
LSI 9205-8e HBA Firmware	17.11.03
LSI 9300-8e HBA Firmware	10.10.03

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	3P03
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	3P03

Table Continued

HCBF1800S5xeN010	10K	1.8TB	Y	Y	Y	Y	N	3P05
STHB1800S5xeN010	10K	1.8TB	Y	Y	Y	Y	N	3P03
AREA0400S5xnNTRI	SSD	400GB	N	Y	Y	Y	Y	3P01
AREX0400S5xnNTRI	SSD	400GB	Y	Y	Y	Y	Y	3P03
DDYE0400S5xnNMRI	SSD	400GB	N	Y	Y	Y	Y	3P03
AREA0480S5xnNTRI	SSD	480GB	N	Y	N	N	N	3P01
DOPE0480S5xnNMRI	SSD	480GB	N	Y	Y	Y	N	3P0A
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	3P07
DDYE1920S5xnNMRI	SSD	1.92TB	N	Y	Y	Y	Y	3P03
DOPE1920S5xnNMRI	SSD	1.92TB	N	Y	Y	Y	N	3P0A
AREA3840S5xnNTRI	SSD	3.84TB	N	Y	Y	Y	Y	3P00

Table Continued

AREX3840S5xnFTRI	SSD	3.84TB	N	Y	Y	Y	Y	3P03
AREX3840S5xnNTRI	SSD	3.84TB	N	Y	Y	Y	Y	3P07
DDYE3840S5xnNMRI	SSD	3.84TB	N	Y	Y	Y	Y	3P03
DOPM3840S5xnNMRI	SSD	3.84TB	N	Y	Y	Y	N	3P07
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	3P03
AREA15T4S5xnFTRI	SSD	15.3TB	N	N	Y	Y	Y	3P01
AREA15T4S5xnNTRI	SSD	15.3TB	N	N	Y	Y	Y	3P04
AREX0480S5xnNTRI	SSD	480GB	N	Y	Y	Y	N	3P07
AREX0920S5xnNTRI	SSD	920GB	N	Y	N	N	N	3P07
AREX0920S5xnFTRI	SSD	920GB	N	Y	Y	Y	N	3P03
AREX1920S5xnFTRI	SSD	1.92TB	N	Y	Y	Y	N	3P03

Table Continued

HCFP0600S5xeN010	10K	600GB	Y	Y	Y	Y	N	3P00
HCFP1200S5xeN010	10K	1.2TB	Y	Y	Y	Y	N	3P00
HCFP1800S5xeN010	10K	1.8TB	Y	Y	Y	Y	N	3P00
HCFP1200S5xeF010	10K	1.2TB	Y	Y	Y	Y	N	3P00
SSKB0600S5xeN010	10K	600GB	Y	Y	Y	Y	N	3P01
SSKB1200S5xeN010	10K	1.2TB	Y	Y	Y	Y	N	3P01
SSKB1800S5xeN010	10K	1.8TB	Y	Y	Y	Y	N	3P01
SSKB1200S5xeF010	10K	1.2TB	Y	Y	Y	Y	N	3P01

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.

Websites

Website	Link
Hewlett Packard Enterprise Information Library	**www.hpe.com/info/enterprise/docs**
Hewlett Packard Enterprise Support Center	**www.hpe.com/support/hpesc**
Contact Hewlett Packard Enterprise Worldwide	**www.hpe.com/assistance**
Subscription Service/Support Alerts	**www.hpe.com/support/e-updates**
Software Depot	**www.hpe.com/support/softwaredepot**
Customer Self Repair	**www.hpe.com/support/selfrepair**
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**

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:

www.hpe.com/support/selfrepair

Remote support

Remote support is available with supported devices as part of your warranty, Care Pack Service, 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.

For more information and device support details, go to the following website:

www.hpe.com/info/insightremotesupport/docs

Documentation feedback

To help improve Aruba documentation, please send documentation suggestions, comments, and notice of any errors 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.

Warranty and regulatory information

For important safety, environmental, and regulatory information, see *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at **www.hpe.com/support/Safety-Compliance-EnterpriseProducts**.

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