



**Hewlett Packard  
Enterprise**

# **HPE 3PAR OS 3.3.1 GA, EGA, MU1, EMU1, MU2 Release Notes**

## **Abstract**

This document describes the features and issues included in HPE 3PAR OS 3.3.1 GA, EGA, MU1, EMU1, and MU2, and is intended for use by Hewlett Packard Enterprise customers, partners and field representatives.

Part Number: QL226-99868a  
Published: May 2018  
Edition: 2

## Notices

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

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

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

## Acknowledgments

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

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

Adobe<sup>®</sup> and Acrobat<sup>®</sup> are trademarks of Adobe Systems Incorporated.

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

UNIX<sup>®</sup> is a registered trademark of The Open Group.

## 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.

# Contents

<b>HPE 3PAR OS 3.3.1 GA Release Notes.....</b>	<b>5</b>
Upgrade Considerations.....	5
Supported Platforms.....	5
Notes.....	5
Components .....	5
HPE 3PAR OS 3.3.1 GA Release Notes.....	8
What's New in the OS.....	8
Modifications to the HPE 3PAR OS.....	12
Known Issues with the OS.....	27
HPE 3PAR 3.3.1 File Persona GA Release Notes.....	38
Modifications to File Persona.....	38
Known Issues with File Persona.....	41
HPE 3PAR 3.3.1 CLI GA Release Notes.....	46
Installation Notes for the CLI.....	46
Supported Operating Systems.....	47
What's New in the CLI.....	47
Modifications to the CLI.....	50
HPE 3PAR 3.3.1 CIM API GA Release Notes.....	56
What's New with the CIM API and SNMP Software .....	56
Modifications to the 3PAR CIM API.....	57
HPE 3PAR 3.3.1 WSAPI GA Release Notes.....	58
What's New with the Web Services API Software .....	58
Modifications to the 3PAR Web Services API.....	60
 <b>HPE 3PAR OS 3.3.1 EGA Release Notes.....</b>	 <b>62</b>
Online Upgrade Considerations.....	62
Affected components.....	62
Modifications .....	62
Verification.....	68
 <b>HPE 3PAR OS 3.3.1 MU1 Release Notes.....</b>	 <b>71</b>
Upgrade Considerations.....	71
Supported Platforms.....	71
Notes.....	71
HPE 3PAR OS 3.3.1 MU1 Release Notes.....	72
What's New in the OS.....	72
Modifications to the HPE 3PAR OS.....	72
Known Issues with the OS.....	83
Modifications to File Persona.....	88
HPE 3PAR OS 3.3.1 CLI Release Notes.....	88
What's New in the CLI.....	88
Modifications to the CLI.....	89
HPE 3PAR OS 3.3.1 MU1 CIM API Release Notes.....	91
Modifications to the 3PAR CIM API.....	91
HPE 3PAR WSAPI 3.3.1 MU1 Release Notes.....	92
What's New with the Web Services API Software .....	92
Modifications to the 3PAR Web Services API.....	93

<b>HPE 3PAR OS 3.3.1 EMU1 Release Notes.....</b>	<b>94</b>
Upgrade Considerations.....	94
Supported Platforms.....	94
Components.....	95
Modifications to the OS.....	96
 <b>HPE 3PAR OS 3.3.1 MU2 Release Notes.....</b>	 <b>99</b>
Update Considerations.....	99
Supported Platforms.....	99
HPE 3PAR 3.3.1 MU2 Release Notes.....	99
Patches Included in This Release.....	99
Modifications to the HPE 3PAR OS.....	100
Known Issues with the OS.....	112
HPE 3PAR 3.3.1 File Persona MU2 Release Notes.....	116
What's New in File Persona.....	116
HPE 3PAR 3.3.1 MU2 CLI Release Notes.....	117
Changed Commands.....	117
HPE 3PAR 3.3.1 MU2 CIM API Release Notes.....	117
What's New in the CIM API.....	117
HPE 3PAR 3.3.1 MU2 Web Services API Release Notes.....	118
What's New with the Web Services API Software .....	118
HPE 3PAR 3.3.1 VASA/VVol MU2 Release Notes.....	118
What's New in the VASA/VVol.....	118
 <b>Component Versions .....</b>	 <b>120</b>
 <b>Websites.....</b>	 <b>123</b>
 <b>Support and other resources.....</b>	 <b>124</b>
Accessing Hewlett Packard Enterprise Support.....	124
Accessing updates.....	124
Customer self repair.....	125
Remote support.....	125
Warranty information.....	125
Regulatory information.....	126
Documentation feedback.....	126

# 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

*Table Continued*

Component	Version
Console Menu	3.3.1.215
Event Manager	3.3.1.215
Internal Test Tools	3.3.1.215
LD Check Tools	3.3.1.215
Network Controller	3.3.1.215
Node Disk Scrubber	3.3.1.215
PD Scrubber	3.3.1.215
Per-Node Server	3.3.1.215
Persistent Repository	3.3.1.215
Powerfail Tools	3.3.1.215
Preserved Data Tools	3.3.1.215
Process Monitor	3.3.1.215
Software Updater	3.3.1.215
TOC Server	3.3.1.215
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

*Table Continued*

Component	Version
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

*Table Continued*

Component	Version
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 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**

DAC Description	DAC Part #	3PAR StoreServ Platforms and Required DAC OS Support					Speed/ Protocols Supported
		7000	10000	8000	9000 and 20000		
<b>HPE 3COM (H3C)</b>							
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

*Table Continued*

		3PAR StoreServ Platforms and Required DAC OS Support				
DAC Description	DAC Part #	7000	10000	8000	9000 and 20000	Speed/Protocols Supported
HPE x240 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
<b>HPE Procurve</b>						
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
<b>HPE StoreFabric</b>						
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

*Table Continued*

3PAR StoreServ Platforms and Required DAC OS Support						
DAC Description	DAC Part #	7000	10000	8000	9000 and 20000	Speed/ Protocols Supported
<b>HPE Blade System</b>						
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.

*Table Continued*

**Conditions of occurrence:** When using SPOCC to complete 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

*Table Continued*

**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 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 `showcpvg` 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 `showcpvg` output may indicate an alert is set after the underlying condition has been resolved.

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

*Table Continued*

**Conditions of occurrence:**

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

**Impact:** Low

**Customer circumvention:** Issue is resolved in 3.3.1

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

**Issue IDs:** 153893

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

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

**Affected software versions:** 3.2.1, 3.2.2

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

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

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

**Impact:** High

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

**Customer recovery steps:** None

**Issue IDs:** 156155

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

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

**Affected software versions:** 3.2.2

*Table Continued*

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

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

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

**Impact:** Low

**Customer circumvention:** None

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

**Issue IDs:** 158195

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

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

**Affected software versions:** 3.2.2

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

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

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

**Impact:** Low

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

**Customer recovery steps:** None

**Issue IDs:** 159520

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

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

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

*Table Continued*

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

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

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

**Impact:** Low

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

**Customer recovery steps:** None

**Issue IDs:** 160406

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

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

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

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

**Symptoms:** The array becomes unresponsive and requires restart.

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

**Impact:** High

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

**Customer recovery steps:** None

**Issue 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 plugi 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.

*Table Continued*

**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:** `srdataac` 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 `srdataac` 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 `srdataac` log file.

**Conditions of occurrence:** Excessive writing to `srdataac` 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

*Table Continued*

**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.

*Table Continued*

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

**Affected software versions:** 3.2.2

**Issue description:** An "infinite" loop in `srdac` 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 nodes.

**Symptoms:** The array will restart three times.

**Conditions of occurrence:**

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

**Impact:** High

*Table Continued*

**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 node down recovery.

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

**Affected software versions:** 3.3.1.215

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

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

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

**Impact:** Medium

**Customer circumvention:** Install P01.

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

## Known Issues with the OS

**Issue IDs:** 94331

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

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

**Affected software versions:**

**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

*Table Continued*

**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

*Table Continued*

**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

<p><b>Issue IDs:</b> 192368</p> <p><b>Issue summary:</b> <code>cachesvr</code> process memory consumption may cause other processes to stop.</p> <p><b>Affected platforms:</b> StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000</p> <p><b>Affected software versions:</b> 3.2.1.GA - 3.2.1.MU3, 3.2.2.GA - 3.2.2.MU3</p> <p><b>Issue description:</b> Over time the <code>cachesvr</code> process on the cluster master node may exhaust free memory, causing other user processes to halt. When this occurs, the affected process will restart and may continue to halt until the <code>cachesvr</code> process is restarted. Once the <code>cachesvr</code> 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.</p> <p><b>Symptoms:</b> <code>cachesvr</code> process memory size grows over time and causes other process to halt with the message "Unable to allocate xxxxxxxx bytes."</p> <p><b>Conditions of occurrence:</b> The issue is most likely to occur on systems which have large configurations and which execute frequent array management interactions.</p> <p><b>Impact:</b> Medium</p> <p><b>Customer circumvention:</b> None</p> <p><b>Customer recovery steps:</b> None</p>
<p><b>Issue IDs:</b> 193758</p> <p><b>Issue summary:</b> Large number of <code>updatevv</code> operations could lead to rare and unexpected IO stalls.</p> <p><b>Affected platforms:</b> StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000</p> <p><b>Affected software versions:</b> 322GA-322MU4, 3.3.1</p> <p><b>Issue description:</b> A large number of <code>updatevv</code> operations could lead to rare and unexpected IO stalls.</p> <p><b>Symptoms:</b> IO stalls could be encountered on StoreServ which goes through frequent and large number of <code>updatevv</code> operations.</p> <p><b>Conditions of occurrence:</b> Frequent and intense <code>updatevv</code> operations on snapshot volumes.</p> <p><b>Impact:</b> Medium</p> <p><b>Customer circumvention:</b> Reduce the frequency of events leading to intense <code>updatevv</code> operations.</p> <p><b>Customer recovery steps:</b> None</p>
<p><b>Issue IDs:</b> 193846</p> <p><b>Issue summary:</b> <code>tunesys</code> does not apply the <code>-fulldiskpct</code> or <code>-chunkpct</code> options to the intra-node phase when active-active PDs are present.</p>

*Table Continued*

**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

*Table Continued*

**Affected software versions:** 3.3.1.GA

**Issue description:** An issue has been reported with the CLI `setcpgr` 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 `setcpgr` 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 `setcpgr` is used from the CLI. (This issue does not affect changing the CPG settings via the SSMC.)

**Customer recovery steps:** Use `setcpgr` 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

*Table Continued*

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

**Customer recovery steps:** Use `setcpgr` 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 nodes."

*Table Continued*

**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.
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.	<p>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:</p> <ol style="list-style-type: none"> <li>1. Go through explorer.</li> <li>2. Create the directory.</li> <li>3. Share the directory once it is created.</li> </ol>

*Table Continued*

Issue ID	Summary	Description
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 <code>showfpg -d</code> command.	The "Files" value in the <code>showfpg -d</code> 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.
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	<p>"Unknown error 528" error message on NFSv3 during <code>setfacl</code>.</p> <p>Unknown error 528 may be encountered when using Network File System (NFS) version 3 (NFSv3) to set file permissions using the <code>setfacl</code> utility or from access contention handling when accessing the file Access Control List (ACL).</p>	<p>Unknown error 528 may be encountered when using Network File System (NFS) version 3 (NFSv3) to set file permissions using the <code>setfacl</code> utility or from access contention handling when accessing the file Access Control List (ACL).</p> <p>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:</p> <p><a href="https://www.ietf.org/rfc/rfc1813.txt">https://www.ietf.org/rfc/rfc1813.txt</a></p>	<p>To prevent this issue, user must either utilize a client that complies with the NFSv3 specification for retries, or do not use <code>setfacl</code> via script or utility that would allow multiple operations to occur in a short period of time.</p> <p>To recover from this issue, retry the failed operation. Several retries may be needed during periods of heavy <code>setfacl</code> call load.</p>
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.
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 <code>noac</code> option during NFS mount would help address these situations of incorrect file attribute cache handling. But using the <code>noac</code> option will have a significant performance impact, and it is recommend to use it only for those applications which exhibit these issues.

*Table Continued*

Issue ID	Summary	Description	Corrective Action
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)."	<p>To prevent possible name collisions and confusion with names stored in ACLs, the following is recommended:</p> <ol style="list-style-type: none"> <li>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.</li> <li>2. Set both of the two name fields to the same name when creating a user.</li> </ol>
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.
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.

*Table Continued*

Issue ID	Summary	Description	Corrective Action
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.	<p>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.</p> <p>The error following error message is displayed:</p> <p>"Windows cannot copy file &lt;Local Windows path&gt; to location &lt;Share path&gt;. This error may be caused by network problems or insufficient security rights. DETAIL - Access is denied."</p>	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 <code>setfs</code> command. To avoid this issue, avoid stopping any node within 10 minutes of configuring additional nodes.
93279	Spurious <code>monitor.startprocess.ok</code> event reported.	Occasionally, an event with the identifier <code>monitor.startprocess.ok</code> 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 <code>stopfs</code> command followed by the <code>startfs -enable</code> 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 <code>stopfs</code> command following by the <code>startfs -enable</code> 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 <code>showfsnap</code> 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 <code>showfsnap</code> 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 <code>showfs</code> 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 <code>showfsquota</code> commands can cause system to respond slowly or cause subsequent commands to fail.	When LDAP server is unavailable (LDAP is configured), executing the <code>showfsquota</code> 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.	<p>Perform the following steps to rediscover attached VTLs:</p> <ol style="list-style-type: none"> <li>1. Execute following command on the HPE 3PAR CLI: <pre>showfsndmp -vtl vtldevices</pre> <p>It will list VTL device IPs similar to the following:</p> <pre>VtlDeviceIp 1.1.1.1 1.1.1.2</pre> </li> <li>2. Execute following command by providing all above IPs separated by commas: <pre>setfsndmp vtl +1.1.1.1,1.1.1.2</pre> </li> </ol> <p>All VTLs will be rediscovered.</p>

## 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
<code>checkhealth</code>	New <code>-d</code> option
<code>checkvv</code>	New <code>-compr_dryrun</code> option
<code>controlsr</code>	New subcommands <code>setperiod</code> and <code>setretention</code>
<code>createfpg</code>	Max size 64 TiB
<code>createfshare</code>	New subcommand <code>ftp</code>
<code>createfststore</code>	New mandatory <code>-secmode</code> option
<code>creategroupsv</code>	New <code>-addto</code> and <code>-match</code> option
<code>creategroupsvvcopy</code>	New <code>-compr</code> and <code>-deup</code> compression options
<code>createsched</code>	<code>importvv</code> now allowed, command limit 1023 bytes
<code>createsralertcrit</code>	Additional space categories, New <code>%_average</code> condition comparisons; Added <code>SYSSPACE</code> type
<code>createsv</code>	New <code>-addto</code> option
<code>createvv</code>	Added three new policies for host DIF support; extended <code>-f</code> option to skip DIF policy change warning message; Compression changes
<code>growfpg</code>	Max size 64 TiB
<code>histpd</code>	New <code>-devsvtime</code> option

*Table Continued*

Command	Description
importvv	New <code>-compr</code> and <code>-dedup</code> compression options
locatecage	Support locate commands on HPE 3PAR StoreServ 8000 Storage system
removedomain	Added <code>-pat</code> option
removedomainset	Added <code>-pat</code> option
removefshare	New subcommand <code>ftp</code>
removehost	Added <code>-pat</code> option
removehostset	Added <code>-pat</code> option
removevvset	Added <code>-pat</code> option
setfpg	New <code>-upgrade</code> option
setfs	New subcommand <code>usermap</code>
setfsav	New <code>-quar_file</code> ; SOPHOS added to <code>-vendor</code>
setfshare	New subcommand <code>ftp</code>
setfstore	New <code>-secop_errsuppress</code> and <code>-secmode</code> options
setrcopygroup	New policy <code>mt_pp</code>
setrcopytarget	New subcommand <code>autotunelinks</code>
setsralertcrit	Allows more changes, Merges SSD100 and SSD150 metrics
setsys	Added <code>OverprovRatioLimit</code> , <code>OverprovRatioWarning</code> , <code>allowR5OnFCDrives</code> , <code>DisableCompr</code> , <code>AllowWrtbackUpgrade</code> , and <code>AllowWrtbackSingleNode</code>
setvv	New policies: <code>3par_host_dif</code> , <code>std_host_dif</code> , <code>no_host_dif</code>
showcpg	New <code>-listcols</code> and <code>-showcols</code> , output format changes
showfs	New <code>-usermap</code> option
showfsarchive	New <code>-importfile</code> , <code>-export</code> options and subcommand <code>export</code>
showfshare	New subcommand <code>ftp</code>

*Table Continued*

Command	Description
showfstore	Output changes
showhost	Output changes for <code>-agent</code>
showiscsisession	New <code>-d</code> option
showld	New <code>-ck</code> option
shownode	New <code>-pci type "combo"</code>
showportdev	New <code>-d</code> option for subcommand <code>tzone</code> , new subcommand <code>uns</code>
showsys	New <code>-vvspace</code> option
showtask	Limit increased to 2000
showuserconn	Output for <code>-d</code> lists memory
showvlun	New <code>-pathsum</code> columns
showvv	New <code>showvv -pol</code> output for host DIF settings; New compression output changes, changes to output of <code>showvv -s</code> and <code>showvv -d</code>
sr*	New <code>-compareby</code> option
srcpgspace	Compression output changes
srhistvlun	VVol filtering
srrgiodensity	Added <code>-totpct</code> option
srstatvlun	New <code>-vlun</code> , VVol filtering
srvvspace	VVol filtering. Compression output changes.
statpd	Added <code>-devsvtime</code> option
tunesys	New <code>-force</code> , <code>-slsz</code> , <code>-slth</code> , <code>-compactmb</code> , <code>-cleanwait</code> , <code>-maxnodetasks</code> and <code>-ss</code>

## Modifications to the CLI

<b>Issue IDs:</b> 79971
<b>Issue summary:</b> <code>checkhealth</code> doesn't detect degraded SFPs in converged network adapters (CNAs).

*Table Continued*

**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 node slots.

**Impact:** Medium

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

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

**Issue IDs:** 136799

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

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

**Affected software versions:** 3.2.1 (MU3)

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

*Table Continued*

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

**Conditions of occurrence:** CLI connection stall.

**Impact:** Medium

**Customer circumvention:** None

**Customer recovery steps:** Quit unresponsive CLI connection process.

**Issue IDs:** 138748

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

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

**Affected software versions:** 3.2.1 (MU2)

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

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

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

**Impact:** Low

**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

*Table Continued*

**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.

*Table Continued*

**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

*Table Continued*

<p><b>Issue description:</b> <code>admithw</code> reconfigures all Fibre Channel ports, including RC ports, that are in a "free" state to host connection mode.</p> <p><b>Symptoms:</b> A possible loss of RC ports used during HPE 3PAR OS or hardware upgrade when <code>admithw</code> is executed.</p> <p><b>Conditions of occurrence:</b> Having RC in use, but temporary free or disconnected, during <code>admithw</code> execution.</p> <p><b>Impact:</b> High</p> <p><b>Customer circumvention:</b> Guarantee that before executing <code>admithw</code>, all FC ports, including RC ports, are properly connected and not showing as <code>free</code> in <code>showport</code>.</p> <p><b>Customer recovery steps:</b> Reconfigure any incorrectly configured RC port back to Remote Copy mode.</p>
--

<p><b>Issue IDs:</b> 179378</p>
<p><b>Issue summary:</b> Users with edit or higher permissions are able to use <code>updatevv</code> on virtual volumes in their domains.</p> <p><b>Affected platforms:</b> StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000</p> <p><b>Affected software versions:</b> All versions before 3.3.1</p> <p><b>Issue description:</b> Previously, a super-user would have to issue the command <code>setuseracl &lt;username&gt; updatevv &lt;virtual volume name&gt;</code> to allow a non-super user to utilize the <code>updatevv</code> 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 <code>updatevv</code> without requiring a super-user issue the <code>setuseracl</code> command.</p> <p><b>Symptoms:</b> When a non-super user, issues the command <code>updatevv &lt;virtual volume name&gt;</code> the user will get a "permission denied" message, given the command <code>setuseracl</code> was not issued for them.</p> <p><b>Conditions of occurrence:</b> The user does not have edit or higher permissions for the domain to which the virtual volume belongs.</p> <p><b>Impact:</b> Low</p> <p><b>Customer circumvention:</b> None</p> <p><b>Customer recovery steps:</b> None</p>

<p><b>Issue IDs:</b> 184028</p>
<p><b>Issue summary:</b> WSAPI audit trail support: <code>tpdtcl</code> needs to put original request IP/port info in the <code>eventlog</code> and <code>showuserconn</code>.</p> <p><b>Affected platforms:</b> StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000</p> <p><b>Affected software versions:</b> 3.2.1, 3.2.2</p>

*Table Continued*

**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 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

<p><b>Issue IDs:</b> 145085</p> <p><b>Issue summary:</b> A cimserver IndicationSubscription cannot be deleted.</p> <p><b>Affected platforms:</b> StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000</p> <p><b>Affected software versions:</b> 3.2.1, 3.2.2</p> <p><b>Issue description:</b> CIM_IndicationFilter instances that exist only in the root/tpd but not interop namespace cannot be enumerated and deleted.</p> <p><b>Symptoms:</b> The cimserver API will return a NOT FOUND error when attempting to delete a CIM_IndicationSubscription.</p> <p><b>Conditions of occurrence:</b> CIM_IndicationFilter is created in root/tpd namespace only.</p> <p><b>Impact:</b> Low</p> <p><b>Customer circumvention:</b> None</p> <p><b>Customer recovery steps:</b> Create the exact same CIM_IndicationFilter in interop namespace also.</p>
<p><b>Issue IDs:</b> 161149</p> <p><b>Issue summary:</b> Volumes created with CreateStorageVolumeFromStoragePoolWithTemplate do not use the snapshot CPG specified by the storage setting.</p> <p><b>Affected platforms:</b> StoreServ 7000, StoreServ 8000, StoreServ 10000, StoreServ 20000</p> <p><b>Affected software versions:</b> 3.2.2</p> <p><b>Issue description:</b> The snapshot CPG specified by the TPD_StorageSetting template is not configured for volumes created with the CIM API call CreateStorageVolumeFromStoragePoolWithTemplate.</p>

*Table Continued*

**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 node will not boot after clean shutdown when 2nd node has a bad voltage regulator.

**Affected platforms:** StoreServ 7000

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

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

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

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

**Impact:** High

**Customer circumvention:** None

**Customer recovery steps:** None

**Issue IDs:**199218

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

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

**Affected software versions:**3.3.1 GA

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

**Symptoms:** Longer than normal host service times on VLUNS.

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

**Impact:** High

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

**Customer recovery steps:** Use standard recovery for host timeouts.

**Issue IDs:**200023

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

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

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

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

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

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

**Impact:** Low

**Customer circumvention:** None

**Customer recovery steps:** None

**Issue IDs:**200464

**Issue summary:** The command `updatevv -removeandcreate` 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 -removeandcreate`, may skip A VV while adding it in Virtual Volume Set (VVSet).

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

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

**Impact:** High

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

**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	<a href="#">HPE 3PAR OS 3.3.1 Patch 01 Release Notes</a>
Patch 02	P02 provides several quality improvements.	None	<a href="#">HPE 3PAR OS 3.3.1 Patch 02 Release Notes</a>

### Known Issues with the OS

**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 nodes."

*Table Continued*

**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.

## 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

<p><b>Issue ID:</b> 179894</p> <p><b>Issue summary:</b> 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.</p> <p><b>Affected platforms:</b> All StoreServ</p> <p><b>Affected software versions:</b> 3.3.1 GA, 3.3.1 EGA, and all previous versions</p> <p><b>Issue description:</b> Disks exhibiting certain types of correctable errors will not be identified early for replacement.</p> <p><b>Symptoms:</b> HVIPC disk models report unusually high numbers of correctable errors, leading to eventual disk replacement.</p> <p><b>Conditions of occurrence:</b> On StoreServ 10000 with HVIPC drives installed, higher than normal correctable errors may be observed, leading to eventual disk replacement.</p> <p><b>Impact:</b> Medium</p> <p><b>Customer circumvention:</b> Perform maintenance when disks require replacement.</p> <p><b>Customer recovery steps:</b> None</p>
<p><b>Issue ID:</b> 184101</p> <p><b>Issue summary:</b> Occasionally peer motion volume migration from 3par array to 3par array does not complete.</p> <p><b>Affected platforms:</b> All StoreServ</p> <p><b>Affected software versions:</b> 3.2.2 GA – 3.2.2 MU4</p> <p><b>Issue description:</b> Peer motion import would return error string <code>Name -srctpg is too long, should be less than 5 characters Error: bad rv argument</code>. This was due to a misinterpretation of a unusual mode page.</p> <p><b>Symptoms:</b> Peer motion migrations would fail.</p> <p><b>Conditions of occurrence:</b> Edge case in data handling, when certain internal fields were set by source array describing the volume to be migrated.</p> <p><b>Impact:</b> Low</p> <p><b>Customer circumvention:</b> Convert the volumes to fully provisioned before migration.</p> <p><b>Customer recovery steps:</b> Retry Migration.</p>
<p><b>Issue ID:</b> 193352</p> <p><b>Issue summary:</b> High volume of fixed events in the event log.</p>

*Table Continued*

**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 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.

*Table Continued*

**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:a11_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

*Table Continued*

**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 -removeandcreate` command on the vvset, a Virtual Volume(s) is missing from the output. Additionally, when the `updatevv -removeandcreate` 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 -removeandcreate` on vvset.

**Impact:** High

**Customer circumvention:** Do not create a VLUN set from vvset. Rather create an individual VLUN for 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

*Table Continued*

**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

*Table Continued*

**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.

*Table Continued*

**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

*Table Continued*

**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.

<p><b>Issue ID:</b> 205406</p> <p><b>Issue summary:</b> Remote Copy disaster recovery operation did not complete, leaving the Remote Copy groups in an unexpected (inconsistent) state.</p> <p><b>Affected platforms:</b> All StoreServ</p> <p><b>Affected software versions:</b> 3.1.1 GA - 3.3.1 GA, 3.3.1 EGA</p> <p><b>Issue description:</b> 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.</p> <p><b>Symptoms:</b> The CLI command <code>showrcopy</code> 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.</p> <p><b>Conditions of occurrence:</b> Performing a Remote Copy disaster recovery operation while a region move is in progress.</p> <p><b>Impact:</b> Medium</p> <p><b>Customer circumvention:</b> Wait until all region moves are complete before performing Remote Copy disaster recovery.</p> <p><b>Customer recovery steps:</b> Use <code>setcopygroup</code> command with appropriate options to restore the Remote Copy groups to a normal state.</p>
<p><b>Issue ID:</b> 206188</p> <p><b>Issue summary:</b> FC Multi-Queue feature was not enabled on 16GB FC HBA after an array update.</p> <p><b>Affected platforms:</b> StoreServ 8000, StoreServ 20000, StoreServ 20000 R2</p> <p><b>Affected software versions:</b> 3.3.1 GA, 3.3.1 EGA</p> <p><b>Issue description:</b> 3PAR 3.3.1 OS upgrade from any version of 3.2.2 or 3.2.1 required additional node reboot after completion of OS upgrade before the Multi-Queue feature is enabled on the LPe16002 or LPe16004 16G FC ports.</p> <p><b>Symptoms:</b> 3PAR array performance may be less than expected.</p> <p><b>Conditions of occurrence:</b> Upgrading the HPE 3PAR OS from 3.2.2 or 3.2.1 to 3.3.1 GA or 3.3.1 EGA.</p> <p><b>Impact:</b> Medium</p> <p><b>Customer circumvention:</b> None</p> <p><b>Customer recovery steps:</b> Reboot each node once after the 3PAR 3.3.1GA OS upgrade is complete.</p>
<p><b>Issue ID:</b> 207547</p> <p><b>Issue summary:</b> Remote Copy read failure results in unexpected controller node restart.</p>

*Table Continued*

**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	<a href="#">HPE 3PAR OS 3.2.1 MU5 Patch 59 Release Notes</a>
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	<a href="#">HPE 3PAR OS 3.2.1 MU5 Patch 71 Release Notes</a>
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	<a href="#">HPE 3PAR OS 3.2.2 MU4 Patch 74 Release Notes</a>
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	<a href="#">HPE 3PAR OS 3.2.2 MU3 Patch 70 Release Notes</a>
HPE 3PAR OS 3.2.2 MU4 Patch 80	Patch 80 provides several quality improvements.	OS-3.2.2.612-P76	<a href="#">HPE 3PAR OS 3.2.2 MU4 Patch 80 Release Notes</a>
HPE 3PAR OS 3.2.2 MU4 Patch 84	Patch 84 provides several quality improvements.	OS-3.2.2.612-P76	<a href="#">HPE 3PAR OS 3.2.2 MU4 Patch 84 Release Notes</a>
HPE 3PAR OS 3.3.1 GA/EGA Patch 04	Patch 04 provides improvements for slow disks and virtual volume management.	None	<a href="#">HPE 3PAR OS 3.3.1 Patch 04 Release Notes</a>

## Known Issues with the OS

<b>Issue ID:</b> 181445
<b>Issue summary:</b> After an unexpected array restart, the normal consistency checks performed on Virtual Volumes may report as <b>not_started, needs_check</b> .
<b>Affected platforms:</b> All StoreServ
<b>Affected software versions:</b> 3.2.1 GA - 3.3.1 MU1
<b>Issue description:</b> Automatic <b>checkvv</b> at restart time corrects any metadata issues found, but does not start the VV. Manual intervention of running <b>checkvv</b> is required to have the volume start.
<b>Symptoms:</b> Virtual Volumes reporting status as <b>Not_started,needs_check</b> .

*Table Continued*

**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.

*Table Continued*

**Impact:** Low

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

**Issue ID:** 204959

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

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

**Affected software versions:** All versions

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

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

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

**Impact:** Low

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

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

**Issue ID:** 211785

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.

*Table Continued*

**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 `admit hw` 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 `admit hw` 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 `admit hw` 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.

*Table Continued*

**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 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

*Table Continued*

**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 `sdmatack`. 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 `sdmatack` kicks off these omitted snapshots will be missed.

**Symptoms:** Should `sdmatack` 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 `sdmatack` 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 `sdmatack` 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
<code>addsnmpmgr</code>	New <code>-notify</code> option
<code>createcert</code>	Add 4 syslog Services
<code>createfshare</code>	New <code>-audit</code> option

*Table Continued*

Command	Description
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
showsapisession	New type and -filter
SR commands	Add percentile, per_group, per_time, only_compareby to summary option

## Modifications to the CLI

<b>Issue ID:</b> 163864
<b>Issue summary:</b> Enables additional commands in the audit user environment.
<b>Affected platforms:</b> All StoreServ
<b>Affected software versions:</b> 3.1.3 GA to 3.3.1 GA/EGA
<b>Issue description:</b> This enhancement enables <code>itables -I</code> and <code>netstat -avntp</code> in the audit user environment.

*Table Continued*

**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**.

*Table Continued*

**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

*Table Continued*

**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 requests.

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: `cpgspace`data, `volumespace`data, `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 WSPA 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	<a href="#">HPE 3PAR OS 3.3.1 MU1 Patch 09 Release Notes</a>
Patch 11	Patch 11 improves SSMC connectivity when LDAP is used.	None	<a href="#">HPE 3PAR OS 3.3.1 MU1 Patch 11 Release Notes</a>
Patch 18	Patch 18 adds quality improvements including OS upgrade and node down recovery.	Obsoletes P14 and P17	<a href="#">HPE 3PAR OS 3.3.1 MU1 Patch 18 Release Notes</a>

3PAR OS 3.3.1 EMU1 also includes the following modifications:

<p><b>Issue ID:</b> 214315</p> <p><b>Issue summary:</b> In some environments, the gFC driver might deliver a false positive detection of IO resource shortage. This is resolved.</p> <p><b>Affected platforms:</b> All StoreServ</p> <p><b>Affected software versions:</b> 3.3.1.GA, 3.3.1.EGA, 3.3.1.MU1</p> <p><b>Issue description:</b> 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.</p> <p><b>Symptoms:</b> The target port types show as <code>free</code> from the CLI command <code>showport</code>.</p> <p><b>Conditions of occurrence:</b> Occurs with 16 Gb FC adapters.</p> <p><b>Impact:</b> High</p> <p><b>Customer circumvention:</b> None.</p> <p><b>Customer recovery steps:</b> None.</p>
<p><b>Issue ID:</b> 215674</p> <p><b>Issue summary:</b> 3PAR 16Gb array ports may auto-negotiate to switches at 8Gb instead of 16Gb.</p> <p><b>Affected platforms:</b> All StoreServ</p> <p><b>Affected software versions:</b> 3.3.1.GA, 3.3.1.MU1</p> <p><b>Issue description:</b> 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.</p> <p><b>Symptoms:</b> 16 GB HBA connecting at 8GB to the Fibre Channel Switch.</p> <p><b>Conditions of occurrence:</b> TTS is disabled on the switch port.</p> <p><b>Impact:</b> Medium</p>

*Table Continued*

**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
```

**Known Issues with the OS****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>

## HPE 3PAR 3.3.1 MU2 Release Notes

### 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	<a href="https://support.hpe.com/hpsc/doc/public/display?docId=a00027067en_us">https://support.hpe.com/hpsc/doc/public/display?docId=a00027067en_us</a>
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	<a href="https://support.hpe.com/hpsc/doc/public/display?docId=a00026783en_us">https://support.hpe.com/hpsc/doc/public/display?docId=a00026783en_us</a>
HPE 3PAR OS 3.3.1 Patch 21	Quality improvements to SD metadata, compression, deduplication and others.	OS-3.3.1.269-P18	<a href="https://support.hpe.com/hpsc/doc/public/display?docId=a00040475en_us">https://support.hpe.com/hpsc/doc/public/display?docId=a00040475en_us</a>
HPE 3PAR OS 3.3.1 MU1 Patch 24	Corrects an issue with alert processing via the SP.	None.	<a href="https://support.hpe.com/hpsc/doc/public/display?docId=a00040750en_us">https://support.hpe.com/hpsc/doc/public/display?docId=a00040750en_us</a>
HPE 3PAR OS 3.3.1 Patch 25	Provides several critical quality improvements.	OS-3.3.1.269-P24	<a href="https://support.hpe.com/hpsc/doc/public/display?docId=a00043628en_us">https://support.hpe.com/hpsc/doc/public/display?docId=a00043628en_us</a>

## Modifications to the HPE 3PAR OS

The following issues have been addressed in this release.

<b>Issue ID:</b> 182665
<b>Issue summary:</b> Physical Disks (PD) may lose both paths nearly simultaneously.
<b>Affected platforms:</b> All StoreServ
<b>Affected software versions:</b> 3.2.2 MU4, 3.2.2 MU6, 3.3.1 GA-MU1
<b>Issue description:</b> PD may lose both paths with a check condition of 06/29, which can lead to host I/O timeouts.
<b>Symptoms:</b> Both paths are lost on PDs.
<b>Conditions of occurrence:</b> Normal operations.
<b>Impact:</b> Medium
<b>Customer circumvention:</b> None.
<b>Customer recovery steps:</b> None.

<p><b>Issue ID:</b> 187217</p> <p><b>Issue summary:</b> In rare situations, a single controller node may unexpectedly restart during internal region moves during controller node reintegration or <code>tunevv</code>.</p> <p><b>Affected platforms:</b> All StoreServ</p> <p><b>Affected software versions:</b> 3.1.2 MU3, 3.2.1, 3.2.2, 3.3.1 GA, 3.3.1 MU1</p> <p><b>Issue description:</b> During controller node integration into the cluster or when <code>tunevv</code> is running, a single controller node may unexpectedly restart if the if an inconsistency in the metadata for the tuned volume is encountered.</p> <p><b>Symptoms:</b> A controller node unexpectedly restarts while running <code>tunevv</code>.</p> <p><b>Conditions of occurrence:</b> The CLI command <code>tunevv</code> is running or a controller node is attempting to join the cluster.</p> <p><b>Impact:</b> Medium</p> <p><b>Customer circumvention:</b> Avoid running <code>tunevv</code>.</p> <p><b>Customer recovery steps:</b> None. The array recovers itself.</p>
<p><b>Issue ID:</b> 190961</p> <p><b>Issue summary:</b> The array unexpectedly restarts when internal operations are performed on virtual volumes.</p> <p><b>Affected platforms:</b> All StoreServ</p> <p><b>Affected software versions:</b> 3.3.1 GA - MU1</p> <p><b>Issue description:</b> The array unexpectedly restarts when VV close operations, which occur, for example, when <code>checkvv -offline</code> is running, are being performed on VVs while the System Manager transfers mastership to another controller node.</p> <p><b>Symptoms:</b> The array unexpectedly restarts when a VV is transitioning to an offline state.</p> <p><b>Conditions of occurrence:</b> The controller node running the System Manager is either intentionally rebooted or unexpectedly restarts while processing VV close operations.</p> <p><b>Impact:</b> High</p> <p><b>Customer circumvention:</b> None.</p> <p><b>Customer recovery steps:</b> None.</p>
<p><b>Issue ID:</b> 193779</p> <p><b>Issue summary:</b> Unexpected controller node restarts occur on arrays with four or more controller nodes and SSD drives.</p>

*Table Continued*

**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.

*Table Continued*

**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

*Table Continued*

**Issue description:** Applying the incorrect license changes the existing W19/WWNBASE ID. This causes the WWPNs of the HBAs to change and a 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 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

*Table Continued*

**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

*Table Continued*

**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

*Table Continued*

**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 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.

*Table Continued*

**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.

*Table Continued*

**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 -no1d** option for tunes.

**Customer recovery steps:** None.

<p><b>Issue ID:</b> 229075</p> <p><b>Issue summary:</b> A compressed virtual volume experiences SD metadata inconsistencies.</p> <p><b>Affected platforms:</b> StoreServ 8000, StoreServ 9000, StoreServ 20000, StoreServ 20000R2</p> <p><b>Affected software versions:</b> 3.3.1 GA-MU1</p> <p><b>Issue description:</b> A compressed virtual volume may experience SD metadata inconsistencies while I/O is occurring to the volume and compression garbage collection is also running.</p> <p><b>Symptoms:</b> The CLI command <code>showvv</code> on a compressed VV shows <code>sd_metadata_inconsistent</code>.</p> <p><b>Conditions of occurrence:</b> Use of compressed volumes.</p> <p><b>Impact:</b> High</p> <p><b>Customer circumvention:</b> None.</p> <p><b>Customer recovery steps:</b> The <code>sd_meta_inconsistent</code> state must be cleared by running <code>checkvv -fixsd</code>.</p>
<p><b>Issue ID:</b> 230334</p> <p><b>Issue summary:</b> Controller nodes become unresponsive while removing or updating volumes.</p> <p><b>Affected platforms:</b> All StoreServ</p> <p><b>Affected software versions:</b> 3.2.2, 3.3.1</p> <p><b>Issue description:</b> Controller nodes become unresponsive while removing or updating volumes which leads to the cluster manager removing the controller node from the cluster.</p> <p><b>Symptoms:</b> Controller nodes or the entire array unexpectedly restart.</p> <p><b>Conditions of occurrence:</b> 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.</p> <p><b>Impact:</b> High</p> <p><b>Customer circumvention:</b> Reduce the I/O load before performing removals or updates of VVs within the VV family.</p> <p><b>Customer recovery steps:</b> None.</p>
<p><b>Issue ID:</b> 232878</p> <p><b>Issue summary:</b> <code>setvvols -remove</code> does not complete successfully when the Remote Copy licenses are not installed.</p> <p><b>Affected platforms:</b> All StoreServ</p>

*Table Continued*

**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`.

## 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 unstated.

**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 `-removeandcreate` option.

**Impact:** High

**Customer circumvention:** Use the `-removeandcreate` 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

*Table Continued*

**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 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 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.

*Table Continued*

**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 File Persona MU2 Release Notes

### What's New in File Persona

- Increased scalability.
- 256 TiB per node pair up to 1 PiB total file capacity per system.
- Improved support for Data Migration.
- Stores native IDs instead of names on disk when using the latest On-Disk Version (12.2).
- File Access Auditing framework.
- Integrates with external ISV applications for data governance regarding file access history.
- File Lock Compliance Mode.
- Increased control over management of data retained using the File Lock feature, by means of a secondary approval of management operations by a user in a Compliance Officer role.
- Support for Kaspersky as an Antivirus ISV.
- User driven snapshot recovery for protocols other than SMB.

- Improved caching for LDAP authentication.
- General performance improvements.

## 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.

### Changed Commands

Command	Description
<code>setcim</code>	New <code>-pol</code> options <code>tls_strict</code> and <code>no_tls_strict</code>
<code>setwsapi</code>	New <code>-pol</code> options <code>tls_strict</code> and <code>no_tls_strict</code>
<code>showcim</code>	New <code>tls_strict</code> and <code>no_tls_strict</code> policies in <code>showcim -pol</code>
<code>showwsapi</code>	New Policy field in <code>showwsapi -d</code>

## 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 3.3.1 VASA/VVol MU2 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 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  
AES128-GCM-SHA256

ECDH-RSA-AES128-SHA256  
AES128-SHA256

# Component Versions

**Table 3: Components and Versions**

Component	Version
Maintenance Update	3.3.1.410 (MU2)
CLI Server	3.3.1.410 (MU2)
CLI Client	3.3.1.410
System Manager	3.3.1.410 (MU2)
Kernel	3.3.1.410 (MU2)
TPD Kernel Code	3.3.1.410 (MU2)
CIM Server	3.3.1.410 (MU2)
WSAPI Server	3.3.1.410 (MU2)
Console Menu	3.3.1.410 (MU2)
Event Manager	3.3.1.410 (MU2)
Internal Test Tools	3.3.1.410 (MU2)
LD Check Tools	3.3.1.410 (MU2)
Network Controller	3.3.1.410 (MU2)
Node Disk Scrubber	3.3.1.410 (MU2)
PD Scrubber	3.3.1.410 (MU2)
Per-Node Server	3.3.1.410 (MU2)
Persistent Repository	3.3.1.410 (MU2)
Powerfail Tools	3.3.1.410 (MU2)
Preserved Data Tools	3.3.1.410 (MU2)
Process Monitor	3.3.1.410 (MU2)
Software Updater	3.3.1.410 (MU2)
TOC Server	3.3.1.410 (MU2)

*Table Continued*

Component	Version
VV Check Tools	3.3.1.410 (MU2)
Upgrade Check Scripts	180507.U013
File Persona	1.4.2.40-20171006 (MU2)
SNMP Agent	1.10.0
SSH	6.0p1-4+deb7u5
VASA Provider	3.0.17 (MU2)
Firmware Database	3.3.1.410 (MU2)
Drive Firmware	3.3.1.410 (MU2)
UEFI BIOS	05.02.54 (MU2)
MCU Firmware (OKI)	4.8.60 (MU2)
MCU Firmware (STM)	5.3.17 (MU2)
Cage Firmware (DC1)	4.44 (MU2)
Cage Firmware (DC2)	2.64 (MU2)
Cage Firmware (DC3)	08 (MU2)
Cage Firmware (DC4)	2.64 (MU2)
Cage Firmware (DCN1)	4082 (MU2)
Cage Firmware (DCN2)	4082 (MU2)
Cage Firmware (DCS1)	4082 (MU2)
Cage Firmware (DCS2)	4082 (MU2)
Cage Firmware (DCS5)	2.79 (MU2)
Cage Firmware (DCS6)	2.79 (MU2)
Cage Firmware (DCS7)	4082 (MU2)
Cage Firmware (DCS8)	4082 (MU2)
QLogic QLA4052C HBA Firmware	03.00.01.77 (MU2)
QLogic QLE8242 CNA Firmware	04.15.27

*Table Continued*

Component	Version
QLogic 260x HBA FC Firmware	174.03.70
QLogic 27xx/268x HBA FC Firmware	174.03.70
QLogic 83xx HBA FCoE Firmware	08.01.05
QLogic 8300 HBA iSCSI Firmware	05.07.35
Emulex LP11002 HBA Firmware	02.82.x10
Emulex LPe12002 HBA Firmware	02.10.x03
Emulex LPe12004 HBA Firmware	02.10.x03
Emulex LPe16002 HBA Firmware	11.1.220.10
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

# Websites

## **General websites**

**Hewlett Packard Enterprise Information Library**

**[www.hpe.com/info/EIL](http://www.hpe.com/info/EIL)**

**Single Point of Connectivity Knowledge (SPOCK) Storage compatibility matrix**

**[www.hpe.com/storage/spock](http://www.hpe.com/storage/spock)**

**Storage white papers and analyst reports**

**[www.hpe.com/storage/whitepapers](http://www.hpe.com/storage/whitepapers)**

For additional websites, see **[Support and other resources](#)**.

# 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](http://www.hpe.com/support/hpesc)

### Hewlett Packard Enterprise Support Center: Software downloads

[www.hpe.com/support/downloads](http://www.hpe.com/support/downloads)

### Software Depot

[www.hpe.com/support/softwaredepot](http://www.hpe.com/support/softwaredepot)

- To subscribe to eNewsletters and alerts:  
[www.hpe.com/support/e-updates](http://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](http://www.hpe.com/support/AccessToSupportMaterials)



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

---

## Customer self repair

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

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

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

## Remote support

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

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

### Remote support and Proactive Care information

#### HPE Get Connected

[www.hpe.com/services/getconnected](http://www.hpe.com/services/getconnected)

#### HPE Proactive Care services

[www.hpe.com/services/proactivecare](http://www.hpe.com/services/proactivecare)

#### HPE Proactive Care service: Supported products list

[www.hpe.com/services/proactivecaresupportedproducts](http://www.hpe.com/services/proactivecaresupportedproducts)

#### HPE Proactive Care advanced service: Supported products list

[www.hpe.com/services/proactivecareadvancedsupportedproducts](http://www.hpe.com/services/proactivecareadvancedsupportedproducts)

### Proactive Care customer information

#### Proactive Care central

[www.hpe.com/services/proactivecarecentral](http://www.hpe.com/services/proactivecarecentral)

#### Proactive Care service activation

[www.hpe.com/services/proactivecarecentralgetstarted](http://www.hpe.com/services/proactivecarecentralgetstarted)

## Warranty information

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

[www.hpe.com/support/Safety-Compliance-EnterpriseProducts](http://www.hpe.com/support/Safety-Compliance-EnterpriseProducts)

### Additional warranty information

#### HPE ProLiant and x86 Servers and Options

[www.hpe.com/support/ProLiantServers-Warranties](http://www.hpe.com/support/ProLiantServers-Warranties)

## HPE Enterprise Servers

[www.hpe.com/support/EnterpriseServers-Warranties](http://www.hpe.com/support/EnterpriseServers-Warranties)

## HPE Storage Products

[www.hpe.com/support/Storage-Warranties](http://www.hpe.com/support/Storage-Warranties)

## HPE Networking Products

[www.hpe.com/support/Networking-Warranties](http://www.hpe.com/support/Networking-Warranties)

# Regulatory information

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

[www.hpe.com/support/Safety-Compliance-EnterpriseProducts](http://www.hpe.com/support/Safety-Compliance-EnterpriseProducts)

## Additional regulatory information

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

[www.hpe.com/info/reach](http://www.hpe.com/info/reach)

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

[www.hpe.com/info/ecodata](http://www.hpe.com/info/ecodata)

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

[www.hpe.com/info/environment](http://www.hpe.com/info/environment)

# Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback ([docsfeedback@hpe.com](mailto: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.