Hewlett Packard Enterprise

HPE 3PAR OS 3.3.1 MU2 Patch 39 Release Notes

Abstract

This document describes the HPE 3PAR OS 3.3.1 MU2 patch release details. Hewlett Packard Enterprise recommends this patch for all systems running HPE 3PAR OS 3.3.1 MU2 with File Persona.

Part Number: QL226- 99998b
Published: October 2018
Edition: 1

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.

Acknowledgements

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

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

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

UNIX® is a registered trademark of The Open Group.

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

HPE 3PAR OS 3.3.1 MU2 P39 Release Notes

Hewlett Packard Enterprise recommends installing patches in the same sequence as they are released, unless instructed otherwise.

Install this patch on arrays using File Persona with HPE 3PAR OS 3.3.1 MU2 P39. For instructions on updating your specific software see, **HPE 3PAR OS and Service Processor Software Update Guide** (**https://support.hpe.com/hpsc/doc/public/display?docId=a00040316en_us**).

Installation recommendation

Highly recommended

Prerequisites

- SP prerequisite: 5.0.3.0 + latest patch
- OS prerequisites: OS-3.3.1 MU2

Patch details

Patch ID:	P39
Synopsis:	Required patch to support File Persona version 1.5 with 3.3.1 MU2.
Date:	September 18, 2018
Description:	See the Release Notes text.
Affected Packages:	tpd-cli, tpd-libauth, tpd-libvmsvc, tpd-vmmgr, tpd-fs, tpd-prerevert
Obsoletes:	None
Requires:	OS-3.3.1.410-MU2
Build Version:	3.3.1.468
Supports Revert:	No
Patches Partially Superseded:	OS-3.3.1.410-P36
Patches Obsolete by Combination:	None
Notes:	None

What's new?

Remote Copy auto failover for FPGs

When you add FPGs to a Remote Copy Group that uses the AutoFailover policy, the FPGs are automatically activated on the secondary system if the primary system fails.

Remote Copy manual failover/failback for FPGs

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

File lock compliance mode

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

File access auditing enhancements

Authentication improvements

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

  Specify multiple LDAP servers to ensure resilience in the case of a single server failure.
- **Local user mapping**

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

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

Major version on-disk upgrade

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

SMB v1 protocol control

As per security best practices, SMB v1 defaults to disabled for newly configured File Persona instances. For existing File Persona instances, the administrator can disable SMB v1 after upgrade, and after confirming that no clients require this type of access. Administrators can also re-enable this support on newly configured instances if needed for legacy support.

Network diagnostics

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

General performance improvements

What's New in the CLI?

Command	Description
`checkhealth`	Includes NFS health information for each node when available.
`controlport`	Introduces the following new subcommands: **`fs ping`** and **`fs traceroute`** Enhances network diagnostics. Initiate from nodes that are enabled with file services.

Table Continued

Command	Description
createfpg	Introduces the following new subcommands: **-rcopygroup** Allows easier association of an FPG to a Remote Copy Group. **-recover** Provides load balancing of recovered FPGs similar to newly created FPGs.
createfsgroup	Changes gid grouping details: • 0-99 – Globally allocated by the Debian project and used for system accounts. The root account, for instance, is GID 0. • 100-1000000000 – Normal groups. • 65534 – The nobody group, with no rights or permissions.
createfshare	Adds an -audit option to the nfs subcommand. Separates File Access Auditing options from the rest of the share options.
createfsuser	Changes uid grouping details: • 0-99 – Globally allocated by the Debian project and used for system accounts. The root account, for instance, is UID 0. • 100-1000000000 – Normal user accounts. • 65534 – The nobody account, with no rights or permissions.
removefpg	Updates help and error messages to clarify that you cannot remove an FPG if it contains a compliance-enabled VFS or File Store.
removefsarchive	Updates help and error messages to clarify that compliance-enabled stores cannot remove retention and files using CLI.
setfpg	Introduces the following new subcommands: **-rcopygroup** Allows easier association of an FPG to a Remote Copy Group. **-version** Allows upgrading of the on-disk version to a specific version.
setfs	Introduces the following new subcommands: **ldap** Supports a new -cloneservers option, to allow configuration of additional LDAP servers for redundancy, in case one server is temporarily unavailable. **smb** Supports two new options, -enablesmb1 and -enablesmb1ad, to control whether SMB v1 protocol is allowed. Defaults to allowed after upgrade from previous versions. Defaults to disallowed when configuring file services for the first time on this version.

Table Continued

Command	Description
`setfsaudit`	Modifies the following subcommand: **`logpol`** Allows you to specify both the `-size` and `-time` options, and to initiate log rotation when detecting the first criteria. Updates help and warning messages to clarify upcoming removal of `replica`. Adds the following subcommand: **`tz`** Configures local time zone information to include with audit logs. Also includes a `-clear` option to clear this configuration.
`setfshare`	Adds an `-audit` option to the `nfs` subcommand. Separates File Access Auditing options from the rest of the share options.
`setsys`	Updates help and error messages to clarify that you cannot disable the `ComplianceOfficerApproval` if compliance-enabled FPGs exist.
`setuser`	Updates help to clarify that users with the role of Compliance Officer (CO) should not have a browse role in a domain.
`showfsquota`	Changes the capacity unit for quotas from MB to MiB.
`showfs`	Introduces the following subcommands: **`-nfs`** Displays NFS health and global settings information per node. **`-supportedversion`** Displays all upgradable on-disk versions. Adds a new field in the `-ldap` display to show cloned server configurations. Adds new values to the `-smb` display to show for SMB v1 protocol support.
`showfsaudit`	Introduces the following new subcommand: **`tz`** Displays the time zone for audit log files.

Modifications

Issue ID: 89753

Issue summary: No alert was raised when the FPG approached or exceeded the supported limits.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: A file system with 250 million or more files can start to become slower and less responsive. No alert is raised to warn the user that the file count could be the issue.

Symptoms: Gradual decreased performance of FPG I/O without any indication or alert as to why.

Table Continued

Conditions of occurrence: When an FPG has more than or is approaching 250 million files.

Impact: Low

Customer circumvention: Monitor the number of files in each FPG periodically using the `showfpg -d` command. View the details of the FPG in SSMC. If the file count is approaching 250 million files, create a new FPG to receive new writes.

Customer recovery steps: Migrate files to a new FPG and then remove files from the existing FPG until the file count is below the 250 million file limit.

Issue ID: 96608

Issue summary: Several changes are included to address temporary interruptions when accessing SMB shares.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: Various conditions are causing the SMB stack to restart (self-heal), and is resulting in `fcollect` support collections.

Symptoms: SMB share clients may notice momentary unresponsiveness in read or write access to files on the file share.

Conditions of occurrence: Accessing SMB shares.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue ID: 97625

Issue summary: Management of file services was temporarily unavailable. The `showfs` command reported a "Starting" state.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: Under conditions where authentication providers such as LDAP and Active Directory are not responding in a timely manner, management of file services could become temporarily unavailable.

Symptoms: The `showfs` command could display a "Starting" state, even though the data services are not actually restarting.

Table Continued

Conditions of occurrence: This could happen when a customer is bound to an LDAP for authentication. If the LDAP access is slow due to lot traffic or heavy network, then it negatively impacts the name resolution. Consequently, it impacts the file services manageability.

Impact: Medium

Customer circumvention: Ensure that any configured authentication provider is healthy and responsive.

Customer recovery steps: None

Issue IDs: 97644

Issue summary: The system time used by the file services and the Active Directory was not synchronized, which resulted in unsuccessful file share creation and access.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: The system time for file services is not in the acceptable time threshold range (anything more or less than 5 minutes), in relation to Active Directory. Due to the time drift, the file services are no longer joined to the Active Directory domain. As a result, authentication and name resolution is unsuccessful for the file services. The problem persists until the system time for file services and Active Directory are synchronized.

Symptoms: All management operations or access to file shares that are configured to use the Active Directory authentication are unsuccessful. Creating a file share or accessing an existing one using MMC is unsuccessful too.

Conditions of occurrence:

The file services time drifts ahead from the Active Directory domain. As a result, file services are no longer joined to the Active Directory. This is because the NTP time is not set properly before the file services get started.

Impact: Medium

Customer circumvention: Use the same external time server (NTP) configuration as the Active Directory.

Customer recovery steps: To address the issue, upgrade to the latest HPE 3PAR OS 3.3.1 MU2. With the upgrade installed, a 3PAR alert is generated when the time between File Persona and the Active Directory drifts.

Issue ID: 98756

Issue summary: Share access was lost on one of the array controller nodes. The following message was displayed `Access Denied`.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Table Continued

Issue description: Various conditions can cause the SMB stack to restart (self-heal), and this can result in `fcollect` support collections.

Symptoms: SMB clients cannot access shares on a controller node even though the other controller nodes have it.

Conditions of occurrence: Clients accessing SMB shares on a controller node.

Impact: Medium

Customer circumvention: None

Customer recovery steps: Perform a failover FPG to another controller node and then run the `stopfs <node>` followed by the `startfs -enable <node>` commands.

Issue ID: 98997

Issue summary: The `httpd` monitoring service is running on both controller nodes (4-nodes) after performing an FPG failover. This is happening even though the Object Access API file shares are not available on both controller nodes of the cluster

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: The `httpd` monitoring service is running on both controller nodes (4-nodes) after performing an FPG failover. Port 80 or Port 443 open on unexpected interfaces.

Symptoms: Port 80 or Port 443 open on unexpected interfaces.

Conditions of occurrence: An FPG failover to another controller node containing Object Access API file shares was unsuccessful.

Impact: Low

Customer circumvention: None

Customer recovery steps: None

Issue ID: 99290

Issue summary: A new Active Directory request was unsuccessful when a secure channel connection to the Active Directory was reset.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: When a secure channel connection to an Active Directory is reset, there are chances that the authentication requests could be unsuccessful.

Table Continued

Symptoms: Customer intermittently loses a secure channel connection to Active Directory. All new Active Directory requests do not complete successfully.

Conditions of occurrence: The Active Directory domain services get restarted while file services are in use.

Impact: High

Customer circumvention: Do not disrupt the network connection between the SMB client and the Active Directory server when an Active Directory request is executed.

Customer recovery steps: Reduce the frequency of open, close, and delete operations.

Issue ID: 99851

Issue summary: FPG became unavailable just before the NFS file share got full.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: FPG could become unavailable just before an NFS file share gets full. Even a retry could give the same result. Consequently, activating the FPG may no longer be possible without further intervention.

Symptoms: NFS share becomes unavailable during writing when the FPG is nearly full and snapshots are concurrently in use.

Conditions of occurrence: A snapshot was taken in an `fsfull` condition

Impact: High

Customer circumvention: The issue can be avoided if the files are created and accessed with the same combination of cases of letters in the name

Customer recovery steps: None

Issue ID: 100011

Issue summary: Under high I/O load, the following message is displayed `The thread pool's task queue is full, limit: 75`.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: Sometimes management commands are unresponsive due to a high I/O load.

Symptoms: File Persona management commands become unavailable till the internal File Persona management service detects the issue and recovers from the situation automatically.

Table Continued

Conditions of occurrence: High I/O load.

Impact: Medium

Customer circumvention: Run fewer parallel instances of the `robocopy` command.

Customer recovery steps: None

Issue ID: 101659

Issue summary: Higher than expected CPU load on active file services although no SMB file share was configured.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: Active Directory was not used in the setup, rather LDAP was the directory server. Hence, the File Persona authentication order does not have an entry for the Active Directory.

Symptoms: Higher than expected CPU load, as seen in the `statfs -cpu` command output.

Conditions of occurrence: The mechanism that searches for groups and users when using an LDAP provider is not efficient for large number of groups and users. Too many resources are consumed.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue ID: 101885

Issue summary: The offline FSCK utility took longer than expected to complete.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: An offline FSCK executed by support could take an unusual amount of time to complete when an FPG has a large number of files.

Symptoms: The offline FSCK utility was taking longer than expected to complete.

Conditions of occurrence: Offline FSCK was taking longer than expected to complete.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue ID: 102366

Issue summary: The user was unable to traverse the root of shares by default.

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

Affected software versions: All versions 3.3.1 GA or later and earlier than 3.3.1 MU2 P39

Issue description: In the NTFS security mode, users were unable to traverse the root of the share by default. The default ACE for the root of the share was modified to allow for traversal by the `Everyone` user to address this issue.

Conditions of occurrence: When creating shares using default permissions.

Impact: Low

Customer circumvention: Create a share at the File Store root and modify the Share Folder ACL if this level of default access is not desired.

Customer recovery steps: None

Issue ID: 102927

Issue summary: The `mkdir` operation at the client's side was unsuccessful. The following message was displayed `retry operation` when simultaneous cache clear management commands were issued.

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

Affected software versions: All versions 3.3.1 GA or later and earlier than 3.3.1 MU2 P39

Issue description: The `mkdir` operation at the client's side could be unsuccessful. A message such as `retry operation` will be displayed when simultaneous cache clear management commands are issued. The `setfs auth -clearcache` command is issued at the same time as a `mkdir` request.

Symptoms: Directory creation is unsuccessful and is returning a retry message.

Conditions of occurrence: Executing the `setfs auth -clearcache` command simultaneously with the `mkdir` operation causes the `setfs auth -clearcache` command to be unsuccessful.

Impact: Medium

Customer circumvention: Do not issue the `setfs auth -clearcache` command during I/O operations.

Customer recovery steps: Redo the directory creation.

Issue ID: 103728

Issue summary: Unreachable file services system.

Table Continued

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

Affected software versions: All versions 3.3.1 GA or later and earlier than 3.3.1 MU2 P39

Issue description: File services became temporarily unavailable on multiple controller nodes during an SMB share load test.

Symptoms: An unreachable system which cannot be pinged.

Conditions of occurrence: I/O load during an SMB share load test.

Impact: High

Customer circumvention: None

Customer recovery steps: None

Issue ID: 103877

Issue summary: False NFS health check failure alerts were generated.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: False NFS health check alerts were generated. There was no impact on file access. There was no pattern to the occurrence of the alerts. Sometimes they occurred within a period of hours and sometimes over multiple days.

Symptoms: Alerts similar to the following example were generated:

```
Time     : 2017-08-17 17:03:53.40 JST
Node     : 0
Seq      : 66125
Class    : Alert
Severity : Major
Type     : NFS Share
Component: sw_fs_fstore_share_nfs:3702654806126534045:nfs-resource-health-check
Message  : File Services NFS Share:3702654806126534045:nfs-resource-health-check Failed  (FAILED)
```

Conditions of occurrence: Normal operation

Impact: Low

Customer circumvention: None

Customer recovery steps: None

Issue ID: 104334

Issue summary: The `HPEidmapd` daemon got restarted unexpectedly when the `Authenticated^Users@NT^AUTHORITY` object was not found.

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

Table Continued

Affected software versions: All versions 3.3.1 MU1 P07 or later and earlier than 3.3.1 MU2 P39

Issue description: The `Authenticated^Users@NT^AUTHORITY` object could not be found when the Name Service Switch interface was used. The `HPEidmapd` daemon could get restarted.

Symptoms: Alert indicating that the `HPEidmapd` service has restarted.

Conditions of occurrence:

- Name Service Switch configuration
- Query of the `Authenticated^Users@NT^AUTHORITY` object was not found.

Impact: High

Customer circumvention: None

Customer recovery steps: None

Issue ID: 104469

Issue summary: An SMB directory folder was not created successfully.

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

Affected software versions: All versions 3.3.1 GA or later and earlier than 3.3.1 MU2 P39

Issue description: An SMB directory folder was not created successfully when a share was mapped as a user from another domain.

Symptoms: An SMB directory folder was not created successfully.

Conditions of occurrence: This issue occurs when a client, who is a member of a domain, logs in as a non-admin user and tries to map a drive as a user from another domain with a primary group "Domain Users" and attempts to create folders and files.

Impact: Low

Customer circumvention: None

Customer recovery steps: None

Issue ID: 105799

Issue summary: Online FSCK restarted the file services on the controller node where the operation was active.

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

Affected software versions: All versions 3.3.1 GA or later and earlier than 3.3.1 MU2 P39

Issue description: Online FSCK mapped the same part of the file in the memory multiples times, consequently straining the memory.

Table Continued

Symptoms: File services restarts on the controller node during online FSCK operation.

Conditions of occurrence: If a snapshot is taken on a file store and the files get modified at irregular intervals, then the snapshots are sparse.

Impact: High

Customer circumvention: None

Customer recovery steps: None

Issue ID: 105838

Issue summary: When Microsoft Management Console (MMC) is used to display open files, MMC is displaying files for all FPGs owned by the File Persona node instead of files for a specific VFS

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: MMC is opening files for an entire FPG instead of the file share that is managed.

Symptoms: Inaccurate per file share open file count in MMC

Conditions of occurrence: Microsoft Management Console (MMC) is used to display open files.

Impact: Low

Customer circumvention: None

Customer recovery steps: None

Issue ID: 106099

Issue summary: Microsoft Project could not open files. The following state for the `SMB Create` attribute was displayed `STATUS_SHARING_VIOLATION`.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: Microsoft Project sometimes does not open files properly. The following state for the `SMB Create` attribute is displayed `STATUS_SHARING_VIOLATION`

Symptoms: Microsoft Project is not opening files properly.

Table Continued

Conditions of occurrence:

- An SMB threads creates a file and enforces a particular shared lock mode
- An SMB thread tries to enforce compatible shared mode locks

Impact: Medium

Customer circumvention: None

Customer recovery steps: If the file is accessed by the SMB protocol, then turn-off the share mode enforcement

Issue ID: 106597

Issue summary: Failover was unsuccessful with the following message `Error handling umount notification before for host`. The Virtual File Server IP address was still deactivated.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: Rolling back and `umount` was unsuccessful.

Symptoms: During protocol deactivation, an unsuccessful `umount` did not initiate a roll back process and brought the protocol service down. This caused extended data unavailability for the impacted Virtual File Server.

Conditions of occurrence: Deactivation of an FPG under I/O load.

Impact: High

Customer circumvention: None

Customer recovery steps: Issue a new request to deactivate the FPG until the request is successful, then reactivate the FPG.

Issue ID: 107555

Issue summary: Online FSCK was unsuccessful with the following message `Failed to add missing entries for tag`.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: Online FSCK could be unsuccessful when namespace connectivity is checked for link mismatch. Consequently, there could be filesystem activation issues.

Symptoms: FPG is not activated successfully after running Online FSCK

Table Continued

Conditions of occurrence: Running Online FSCK after deleting directories in between multiple snapshots.

Impact: High

Customer circumvention: None

Customer recovery steps: None

Issue ID: 109148

Issue summary: The client IP list of an existing NFS file share could not be modified after File Persona upgrade.

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

Affected software versions: 3.3.1 MU1/EMU1 with P07, P08 , or P19

Issue description: After File Persona is upgraded, the command to modify the client IP list of pre-existing NFS file shares was unsuccessful with the message `NFS Modify Export failed with error: Modifying export path is not allowed`.

Symptoms: After File Persona is upgraded, the command to modify the client IP list of an existing NFS file shares is unsuccessful as shown below:

```
cli% setfshare nfs -clientip +<client IP> -fstore <file store> <vfs> <share name>
NFS Modify Export failed with error: Modifying export path is not allowed.
```

Conditions of occurrence: Upgrade with NFS file shares configured.

Impact: Low

Customer circumvention: None

Customer recovery steps: Delete and recreate the share.

Issue ID: 109538

Issue summary: Excessive SMB status calls were received although no SMB file shares were exported from the system.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: When excessive SMB status calls can be received even though no SMB file shares are exported from the system.

Symptoms: Performance of the controller node degrades due to excessive SMB status calls.

Table Continued

Conditions of occurrence: File services enabled.

Impact: Low

Customer circumvention: None

Customer recovery steps: None

Issue ID: 109769

Issue summary: Degradation in performance when copying large files (such as `.iso`) to an SMB file share.

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

Affected software versions: All versions 3.3.1 GA or later and earlier than 3.3.1 MU2 P39

Issue description: Degradation in performance when copying large files (such as `.iso`) to an SMB file share.

Symptoms: Slow I/O, disk fragmentation

Conditions of occurrence: Writing large files to an SMB share.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Issue ID: 110892

Issue summary: After upgrading from 1.4.2, the `srvsvc` container indicates a failed state and the system generates a log file.

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

Affected software versions: All versions

Issue description: At startup, if FP was joined to a Windows domain, LSASS has to re-join and/or enumerate domains. LWIO calls NetLogon, which uses the customer DNS to resolve the Hostname to an IP. If the DNS has no entry for hostname, or has an incorrect reverse lookup pointer, the DNS response is unpredictable. In this case LWIO releases the allocated memory, resulting in double free memory.

Symptoms: HP-SMB cannot join the domain and the HP-SMB stack restarts.

Table Continued

Conditions of occurrence:

- Customer DNS does not have a reverse lookup zone.
- Hostname does not have a record pointer in the reverse lookup zone.
- Hostname has a record pointer in the reverse lookup zone but it is different from the Hostname (canonical name).

Impact: High

Customer circumvention:

Verify DNS entries for correct configuration.

Customer recovery steps: None

Issue ID: 112990

Issue summary: The Server Message Block (SMB) service health monitor restarts the SMB service because of extremely slow system response. This restart interrupts existing SMB client sessions.

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

Affected software versions: All versions between 3.3.1 MU1 P19 and 3.3.1 MU2 P39

Issue description: On a system with a heavy load, system calls that take more than 20 seconds can cause the SMB stack to restart, which interrupts SMB client sessions.

Symptoms: During times of heavy use on the cluster, SMB clients experience session interruptions. These interruptions recover automatically, and can appear as a momentary pause in connectivity.

Conditions of occurrence: Heavy I/O loads or very slow authentication responses can cause the SMB service health monitor to restart the SMB service.

Impact: Medium

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 114353

Issue summary: After generating `lwsm` log files, File Services stop running.

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

Affected software versions: HPE 3PAR 3.3.1 MU1 and later.

Issue description: When attempting to join a domain, the name resolution using Active Directory return unexpected results and causes abnormal termination of HP-SMB stack. File Services (in particular SMB) stop running.

Table Continued

Symptoms: Customer encounters an abnormal termination from HP-SMB stack.

Conditions of occurrence:

1. Insufficient or incorrect Active Directory DNS configuration.
2. Joining domain with HP-SMB server.

Impact: Medium

Customer circumvention: Properly configure the Active Directory DNS.

Customer recovery steps: Upgrade the system to the latest patch.

Issue ID: 116735

Issue summary: An attempt to access a file or folder is unsuccessful.

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

Affected software versions: HPE 3PAR OS 3.3.1 MU1 plus File Persona limited availability versions; HPE 3PAR 3.3.1 MU2 plus File Persona limited availability versions.

Issue description: As part of authenticating a user and authorizing access to files and folders, the domain must resolve the user ID. An offline domain within the AD forest, coupled with a disabled rfc2307 mode, can prevent access to a file or folder.

Symptoms: An attempt to access a file or folder is unsuccessful.

Conditions of occurrence: 1) The system is joined to an Active Directory (AD) domain. 2) The AD domain exists in a forest with one or more one-way trusts. 3) At least one domain in the forest is offline. 4) rfc2307 mode is disabled.

Impact: Medium

Customer circumvention: Ensure that all domains in the forest are healthy. Use bi-directional trusts.

Customer recovery steps: Bring the unhealthy domain back online.

Issue ID: 116805

Issue summary: Access denied is returned when attempting to access files and subdirectories with the local Guest account.

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

Affected software versions: HPE 3PAR OS 3.3.1 MU1 plus File Persona limited availability versions; HPE 3PAR 3.3.1 MU2 plus File Persona limited availability versions.

Table Continued

Issue description: When the file store contains files or directories with ownership and/or access permissions for the local user Guest (Guest@LOCAL_CLUSTER), the issue is that when the Local user Guest account is enabled and Guest attempts to access those files or directories, “access denied” failures occur.

Symptoms: An attempt to access a file or folder is unsuccessful.

Symptom 1:

When the local Guest account is enabled:

When Local user Guest creates a directory or file from an SMB client, when the parent directory has create permissions for Guest, and a CREATOR_OWNER ACE exists, the directory/file will be created, but Guest will get “access denied” when trying to access the newly created file or directory that should be owned by Guest.

Symptom 2:

When the local Guest account is enabled:

If there are existing files owned by local user Guest in the file store on a previous release (e.g. 3.2.2-MU4), after upgrade to a later release with this problem (e.g. 3.3.1-MU1): Guest is be able to log on but is not able to access the directories/files previously owned and accessible by Guest, getting “access denied” errors.

Conditions of occurrence:

1) Local Guest account is enabled.

2) Files and/or directories exist that are owned by local user Guest or have ACLs that should make them accessible by Guest.

Impact: Medium

Table Continued

Customer circumvention:

Do not upgrade the on-disk version of the file system to 12.2 or greater if there are files with ownership or permissions by the Guest local user, as the permissions metadata may become corrupted. Wait until after upgrading to the latest release before upgrading the on-disk version.

Do not enable the local Guest account before upgrading to the latest release.

Customer recovery steps: Bring the unhealthy domain back online.

1. If the customer has not enabled the Guest account and/or has not created files owned by Guest, this is not an issue. No recovery is necessary.

2. If the customer does enable the Guest account but the customer has not upgraded the on-disk version prior to upgrading to the latest release, no recovery is necessary.

3. If the customer had files or directories owned or with permissions by Guest@LOCAL_CLUSTER, and the on-disk version was upgraded to 12.2 or greater before upgrading to the latest release, the metadata may have been corrupted, and access by the Guest user may continue to fail even after the release upgrade.

This can be fixed after upgrading to the latest release by performing the following steps:

a) Map a share to the file store from a Windows client as the domain Administrator.

b) Reset the ownership and/or ACLs of relevant files to Guest. For example:

c) Recursively change the ownership at the parent directory owned by Guest to be owned by the domain Administrator.

For example, in Windows 10, Windows Explorer Security Tab=>Advanced Security Settings=>Owner Change (Check “Replace all child object permission entries with inheritable permission entries from this object”)

d) Recursively change the ownership at the same parent directory back to Guest

Issue ID: 117510

Issue summary: With File Persona P07 and later, file system does not mount after fail over. System returns the message: `ftx_capsule_play_redo_on_page(): bs_access() of bfset 2 file-tag 261786531 failed(-1032) dirTag : 2`

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

Affected software versions: HPE 3PAR OS 3.3.1 MU1 P7 and P8

Issue description:

Reloading the file extents without checking whether the metadata of file changed during the recovery (part of mount process).

Symptoms: With File Persona P07 and later, file system fails to mount after fail over and returns the message: `ftx_capsule_play_redo_on_page(): bs_access() of bfset 2 file-tag 261786531 failed(-1032) dirTag : 2`

System recovery of file system during becomes excessive.

Table Continued

Conditions of occurrence: Create a fragmented filesystem. Create large number of files so that bmt will have more extents. If the system is crashed and is getting mounted it has to first recovered for metadata consistency which will take long time.

Impact: Medium

Customer circumvention: None.

Customer recovery steps:

Run domain activate on the system and wait for completion of the recovery process. Remount the system.

Issue ID: 118971

Issue summary: SMB protocol access and user authentication services are unavailable on startup.

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

Affected software versions: HPE 3PAR OS 3.3.1 MU1 plus File Persona limited availability versions; HPE 3PAR 3.3.1 MU2 plus File Persona limited availability versions.

Issue description: Each time the protocol and authentication services start, they enumerate the domain controllers in the forest and add them to a list, even if the domain controllers are already listed. The duplicate entries to not immediately cause an issue, but if the list becomes sufficiently long, the services no longer start successfully. This issue is more likely to occur in systems connected to large Active Directory forests.

Symptoms: High CPU utilization. Loss of access to SMB protocol and authentication services.

Conditions of occurrence: Restart of protocol services several times, accelerated in proportion to the number of domain controllers in the forest.

Impact: High

Customer circumvention: None.

Customer recovery steps: None.

Issue ID: 200909

Issue summary: File Services did not automatically start on the controller node even though the controller node started (during upgrade or otherwise).

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: When a controller node starts up, enabling file services on the controller node requires that all volumes in use by file services are in a ready state. If it takes longer than expected for these volumes to reach a ready state, file services may be left in a `shutoff` state.

Table Continued

Symptoms: Even though the controller node was started, the `showfs` command reported that file services on the controller node are in a `shutoff` state. The following was displayed:

```
cli% showfs
Node FSNode State     Active InCluster -Version- ---N:S:P--- BondMode MTU
   0 Yes    Upgrading No     No        -         0:4:2,0:4:1        -   -
   1 Yes    Shutoff   No     No        -         1:4:2,1:4:1        -   -
   2 No     Unknown   No     No        -         -                  -   -
   3 No     Unknown   No     No        -         -                  -   -
-------------------------------------------------------------------------------
   4 total
If this occurs during an upgrade, an error like the following may be reported:
Rebooting node 1.
Waiting for node to go down....
Waiting for node 1 to rejoin the cluster.....................................
Checking if the system is healthy enough to proceed...
Waiting for File Services update to complete................................
File Services update did not complete after 600 secs. (Unable to connect with server.
```

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

Impact: Medium

Customer circumvention: None

Customer recovery steps: Use the `startfs -enable <node>` command for the node when file services are in a `shutoff` state. If the upgrade is in progress, and the controller node is still reporting `Upgrading`, resume the upgrade process.

Issue ID: 211445

Issue summary: File Persona upgrade was unsuccessful because the File Persona node did not start after a StoreServ node upgrade.

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

Affected software versions: All versions earlier than 3.3.1 MU2 P39

Issue description: The networking configuration for File Persona was in an unexpected state. This caused the upgrade to stall.

Symptoms: The `showfs` command shows a node in a `shutoff` state after an upgrade.

Conditions of occurrence: File services are enabled.

Impact: Medium

Customer circumvention: None

Customer recovery steps: None

Known issues

Issue ID: 109129

Issue summary: The system produces the following error when `domainSID` is not configured on the LDAP server: `Failed to obtain Domain SID from Ldap Server`

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

Affected software versions: All versions

Issue description: Unsuccessful bind when using POSIX to bind either a newly configured LDAP server or an LDAP server without `domainSID` configured.

Symptoms: `Setfs -ldap` command is not successful and returns the message: `Failed to obtain Domain SID from Ldap Server`.

Conditions of occurrence: Using POSIX to bind either a newly configured LDAP server or an LDAP server without `domainSID` configured.

Impact: Medium

Customer circumvention: Customers binding to an LDAP provider using POSIX schema must configure `domainSID` for the LDAP server (see, *HPE 3PAR File Persona User Guide*).

Customer recovery steps: None.

Issue ID: 110635

Issue summary: Online upgrade was unsuccessful with the following message `Failed to failover the filesystem(s).`

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

Affected software versions: All versions

Issue description: When the system is under an excessive load with many delayed ACKs, an online upgrade could be unsuccessful.

Symptoms: Online upgrade was unsuccessful.

Conditions of occurrence: Upgrade is performed when the system is under excessive load with lots of delayed ACKs.

Impact: Medium

Customer circumvention: Before performing an upgrade, use the `statcmp` or `srstatcmp` commands to ensure that there are no delayed ACKs reported. If there are delayed ACKs reported, reduce the load on the system before proceeding with an upgrade.

Customer recovery steps: None

Issue ID: 110992

Issue summary: User mapping to a local provider cannot be configured for Active Directory users or groups with a UID or GID value greater than 1,000,000,000.

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

Affected software versions: 3.3.1 MU2 P39 and later

Issue description: If there are pre-existing files with owners or groups with a UID or GID value greater than 1,000,000,000, the owners and groups of those files will not be mapped to a local user or group, Adding local users or groups with a UID or GID value greater than 1,000,000,000 is not supported.

Symptoms: If the `AD<=>Local` user mapping is configured for users/groups, files or directories with owners/groups with a UID/GID greater than 1,000,000,000 are no longer accessible.

Conditions of occurrence:

- There are existing files or directories with users or groups with a UID or GID greater than 1,000,000,000
- User mapping is enabled and configured for the `AD<=>Local` user mapping for the same users or groups

Impact: Medium

Customer circumvention: If the user or group UID or GID to be migrated is above the supported maximum value of 1,000,000,000 for the local provider, use LDAP mapping instead of local mapping.

Customer recovery steps: Remove the `AD<=>Local` user mapping rules for all affected users or groups.

Issue ID: 112689

Issue summary: Unauthorized NFS clients got access to an NFS share.

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

Affected software versions: All versions

Issue description: NFS clients could encounter problems when mounting a nested directory from an export. The issue occurs when the export has a different client-access list when compared to the parent directory.

Symptoms: NFSv3 mount request succeeds for an export operation for a client that is restricted from accessing the export. In other words, the client is not part of the client-access list for the export.

Table Continued

Conditions of occurrence:

1. An export is configured on the parent directory with a client-access list as "*" and the nested export that gets created after having a restricted access to "10.x.x.22"
2. The nested export gets successfully mounted from "10.x.x.22" (restricted) till the NFSv3/NFSv4 mount of the parent export is attempted and data is accessed from the same client.
3. The issue occurs when the client-access list is specified in the subnetwork as well.

Impact: Medium

Customer circumvention: It is recommended to set a same client-access list for both the parent and child exports (nested) to avoid undesired NFSv3 access.

Customer recovery steps: Modify the parent and nested exports to have the same set of client-access lists.

Issue ID: 115337

Issue summary: The `setfs ad` command completes successfully, but the `showfs -ad` command indicates that the File Services nodes are not joined to Active Directory.

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

Affected software versions: 3.3.1 MU2 P39 and later

Issue description: An attempt to join a system using a previously working Active Directory provider appears to complete successfully. However, because the Active Directory provider no longer appears in the stacking order, the connection is unsuccessful. The system does not provide notification.

Symptoms: The `setfs ad` command completes successfully but the `showfs -ad` command indicates that the File Services nodes are not joined to Active Directory.

Conditions of occurrence: The system has been joined to Active Directory previously, but the Active Directory provider is not currently in the stacking order.

Impact: Medium

Customer circumvention: Use `showfs -auth` to verify that ActiveDirectory is listed in the stacking order before attempting to join the system to Active Directory using the `setfs ad` command.

Customer recovery steps:

Add the Active Directory provider back to the stacking order with the `setfs auth` command. Because Active Directory was previously joined, the command produces the following output:

```
 cli% showfs -ad
Domain Name  : <domainFQDN>
NetBIOS Name : <domainNetBIOS>
Forest       : <forestFQDN>
Status       : Online
```

Issue ID: 115360

Issue summary: The system returns a misleading message that `Another AD task is already running` after a successful request to join Active Directory

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

Affected software versions: 3.3.1 MU2 P39 and later

Issue description: Because the active File Persona node is not the lowest numbered node, an attempt to join a system to Active Directory returns an unexpected message, even though the request completed successfully.

Symptoms: The system displays an unexpected message after an attempt to join the system to Active Directory.

Conditions of occurrence: The Active node for File Persona is not the lowest numbered node when attempting to join the system to Active Directory.

Impact: Low

Customer circumvention: None.

Customer recovery steps:

Ignore the message and use the `showfs -ad` command to confirm a successful join. The command returns a response similar to the following:

```
cli% showfs -ad
Domain Name  : <domainFQDN>
NetBIOS Name : <domainNetBIOS>
Forest       : <forestFQDN>
Status       : Online
```

Issue ID: 117638

Issue summary: The file services management service for a node indicates the node state as `starting` and does not proceed to a `running` state. This is a rare occurrence.

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

Affected software versions: All

Issue description: In a rare scenario, the file services management service for a node continues to report the node state as `starting` and does not change to a `running` state. Active data services at the time of transition continue to run, but the system does not allow configuration changes to this node.

Symptoms: The `showfs` command shows a Starting state for one but not all of the nodes in the system

Conditions of occurrence: The node has been up and running file services for more than a month without interruption.

Table Continued

Impact: Medium

Customer circumvention: None.

Customer recovery steps:

If a single node in the system reports a starting state continuously and for an extended period of time (greater than 30 minutes), stop and restart the node using the `stopfs <node>` command, followed by a `startfs –enable <node>` command.

Issue ID: 120330

Issue summary: Lost AD join after restart, including during upgrade or while switching domains.

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

Affected software versions: 3.3.1 MU2 P39 and later

Issue description: Upgrading or switching domains can leave the HP-SMB in an unknown state. An attempt to join the nodes returns an incorrect message saying the nodes are already joined.

If HP-SMB can't connect to the customer Domain Controller at start-up (network disruption or unavailable DNS/DC), then domain/trust enumeration or refreshing a Kerberos Ticket with DC does not succeed. In this case, even if HP-SMB is joined to the domain, it cannot provide authentication/authorization services.

Symptoms: After upgrading, the system loses share access, and attempts to join the domain result in a message that nodes are already joined. A domain query of join status repots that the nodes are not joined to the domain.

Conditions of occurrence: Either DNS or the DC is not available at start-up and HP-SMB is already joined to a Domain.

Impact: Low

Customer circumvention: Make sure that the network infrastructure is healthy.

Customer recovery steps:

None.

Issue ID: 120389

Issue summary: Cannot join new domain after leaving existing domain.

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

Affected software versions: 3.3.1 MU2 P39 and later

Issue description: Joining Windows Domain (for new installation) or switching from Windows Domain A to Windows Domain B (for existing installation) or upgrade (File Persona was already joined), does not succeed. Either the new Windows DC FQDN is not registered in DNS or DNS does not have a reverse lookup pointer to the DC FQDN.

Table Continued

Symptoms: Leaving a domain and attempting to join a different domain is unsuccessful and returns a `Status 31` error. Repeated attempts to rejoin are also unsuccessful. When attempting to join a new Windows domain, or re-establish a Kerberos ticket time for a previously established join, File Persona sometimes cannot resolve the DC FQDN. Even if File Persona can resolve the request, HP-SMB sometimes cannot provide authentication.

Conditions of occurrence:

- The Active node for File Persona is not the lowest numbered node when attempting to join the system to Active Directory.
- DNS does not have a reverse pointer to DC FQDN.
- Previously, nodes were successfully joined to a domain.
- Leaving the joined domain and joining an entirely different domain after updating the DNS configuration.

Impact: Low

Customer circumvention: Make sure that the DNS contains correct forward and reverse entries to DC.

Customer recovery steps:

Fix the DNS then retry Join Windows Domain operation.

Issue ID: 121097

Issue summary: After upgrading to 3PAR OS 3.3.1 MU2 P7, customers that use NFS with a local authentication provider observe access issues if there are inconsistent group memberships between clients and the group membership configured by the File Persona administrator.

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

Affected software versions: HPE 3PAR OS 3.3.1 MU2 P7 and later

Issue description: After upgrading to 3PAR OS 3.3.1 MU2 P7, customers that use NFS with a local authentication provider observe access issues if there are inconsistent group memberships between clients and the group membership configured by the File Persona administrator.

This occurs because of security improvements and the removal of the NFS 16 group limitation introduced in 3.3.1 MU2 Patch 7. These changes implemented server side group enumeration for NFS with AUTH_SYS. Because the group membership enumeration occurs on the server, the administrator can control the group membership for each local user and ensure proper access enforcement.

Symptoms: Unexpected user access issues after upgrade to 3.3.1 MU2 P7.

Conditions of occurrence: After upgrade to 3.3.1 MU2 P7, users report access issues.

Impact: Low

Table Continued

Customer circumvention: To avoid the user access issue, Hewlett Packard Enterprise recommends that the administrator check the output of `showfsuser -d` to verify that the secondary group membership for each user matches the configured group membership. If there is a disparity, the administrator can use `setfsuser <username> -grplist <grplist>` to address it.

Customer recovery steps: N/A

Issue ID: 143701

Issue summary: All thinly provisioned FPGs deactivated when the backing CPG was almost out of space.

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

Affected software versions: All versions

Issue description: When a CPG backing an FPG is almost out of space (as could be seen in the `free` column of the `showspace -cpg <cpgName>` command output), the system tries to automatically deactivate all thinly provisioned FPGs. The FPGs use a CPG to ensure that the FPGs do not experience a write error due to lack of space for additional allocations when backing the VVs. However, the alert that indicates that the FPG has been deactivated does not clearly show lack of space in the CPG.

Symptoms: An alert like the following would indicate that the CPG is nearly out of space:

```
Id          : 417
State       : New
Message Code: 0x027001d
Time        : 2018-05-18 14:42:59 MDT
Severity    : Critical
Type        : CPG growth failure non-admin
Message     : CPG FC_r6 SD and/or user space could not grow due to
unavailability of free space. CPG grow attempted using degraded availability
parameters (-ha mag) also failed.
              One or more alerts like the following would indicate that the
FPGs have been deactivated:
Id          : 418
State       : New
Message Code: 0x0720001
Catalog-Key : filesystem-event:filesystem.cmd.unmount
Time        : 2018-05-18 14:43:17 MDT
Severity    : Informational
Type        : File Provisioning Group
Message     : File Provisioning Group:<id>:<fpgName> Normal  (UNMOUNTED)
Details     : FPG Event: Unmounted FPG <fpgName> on host <nodeName>.
```

Conditions of occurrence: CPG space is exhausted but thinly provisioned FPGs are still in use.

Impact: Medium

Table Continued

Customer circumvention: Whenever thin provisioning is used, pay close attention to the remaining capacity in the CPGs. Respond to all alerts by adding capacity to the CPG or otherwise free up space in the CPG. An early indicator that CPGs dependent on a given storage class will soon run out of capacity could be seen with the following type of alert:

```
Id          : 452
State       : New
Message Code: 0x0270010
Time        : 2018-05-18 14:26:49 MDT
Severity    : Major
Type        : FC raw space allocation 85% alert
Message     : Total FC raw space usage at 10453G (above 91% of total 11442G).
```

This alert could be resolved by adding space of the given storage class.

Customer recovery steps: Resolve the CPG space issue and then manually activate the FPG again using the `setfpg` command.

Issue ID: 199108

Issue summary: When a patch was unsuccessfully applied to the `tpd-fs` package, the correct status did not get reflected in the overall upgrade status.

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

Affected software versions: All patches including the `tpd-fs` package

Issue description: When an update of the `tpd-fs` package is still in progress, but the update was unsuccessful, the overall status update could be reported as successful. The version of the file services on the controller nodes does not all match the version as reported by the `showfs` command.

Symptoms: Unsuccessful upgrade is reported in the verbose progress, but the final status of the upgrade is reported as successful. This could look something like the following:

```
[2017-01-23T19:57:59 HKT] Failed to upgrade node1fs.
[2017-01-23T19:57:59 HKT] current_version=1.2.2.6-20161017, expected_version=1.2.3.2-20161117
[2017-01-23T19:58:00 HKT] Finalizing File Services. command=curl command to finalize
[2017-01-23T19:58:00 HKT] exit_status=0
[2017-01-23T19:58:02 HKT] File Services is healthy. (count=1)
[2017-01-23T19:58:02 HKT] File Services upgrade completed successfully.
```

After the upgrade, the output of the `showfs` command displays different versions across the set of nodes.

Conditions of occurrence: Application of a patch that contains an update to the `tpd-fs` package.

Impact: Medium

Customer circumvention: None

Customer recovery steps: To correct the inconsistency in versions across the controller nodes, contact HPE Support.

Issue ID: 211324

Issue summary: A fully provisioned FPG could unexpectedly deactivate when the backing CPG is almost out of space.

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

Affected software versions: All versions

Issue description: When a CPG backing an FPG is nearly out of space (as can be seen in the `free` column of the `showspace -cpg <cpgName>` command output), the system tries to automatically deactivate all the thinly provisioned FPGs. These FPGs use the CPG to ensure that the FPG does not experience a write error due to lack of space for additional allocations for backing VVs. In this case, fully provisioned FPGs could get deactivated.

Symptoms: An alert like the following would indicate that the CPG is almost out of space:

```
Id           : 417
State        : New
Message Code: 0x027001d
Time         : 2018-05-18 14:42:59 MDT
Severity     : Critical
Type         : CPG growth failure non-admin
Message      : CPG FC_r6 SD and/or user space could not grow due to
unavailability of free space. CPG grow attempted using degraded availability
parameters (-ha mag) also failed.
An alert like the following would indicate that the FPG has been deactivated:
Id           : 418
State        : New
Message Code: 0x0720001
Catalog-Key : filesystem-event:filesystem.cmd.unmount
Time         : 2018-05-18 14:43:17 MDT
Severity     : Informational
Type         : File Provisioning Group
Message      : File Provisioning Group:<id>:<fpgName> Normal  (UNMOUNTED)
Details      : FPG Event: Unmounted FPG <fpgName> on host <nodeName>.
```

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

Impact: Medium

Table Continued

Customer circumvention: Whenever thin provisioning is used, it is critical that the administrator pays close attention to the remaining capacity of the CPGs. Respond to alerts by adding capacity to the CPG or otherwise free up space in the CPG. An early indicator that of a CPG that is dependent on a given storage class and is soon running out of capacity could be seen with the following type of alert:

```
Id          : 452
State       : New
Message Code: 0x0270010
Time        : 2018-05-18 14:26:49 MDT
Severity    : Major
Type        : FC raw space allocation 85% alert
Message     : Total FC raw space usage at 10453G (above 91% of total 11442G).
```

If this alert is resolved by adding space of the given storage class, this issue could be avoided

Customer recovery steps: Resolve the CPG space issue and then manually activate the FPG again using the `setfpg` command.

Issue ID: 213022

Issue summary: Online or offline upgrade from 3.2.1 MU3 to 3.3.1 MU2 (or later) was unsuccessful.

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

Affected software versions: 3.3.1 MU2 or later

Issue description: Online or offline upgrade from 3.2.1 MU3 to 3.3.1 MU2 (or later) is unsuccessful if 3.2.1 MU3 P20 was not installed.

Symptoms: Online or offline upgrade from 3.2.1 MU3 to 3.3.1 MU2 (or later) is unsuccessful.

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

Impact: High

Customer circumvention: Install 3.2.1 MU3 P20 before attempting the upgrade.

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

Issue ID: 228325

Issue summary: FPG growth in a Remote Copy Group in a failed over state was successful but the FPG was no longer present in the Remote Copy Group.

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

Affected software versions: All versions

Issue description: When a Remote Copy Group is in a failed over state, FPG commands such as `growfpg` should not be issued. The FPG will successfully grow but the FPG might be taken out from the Remote Copy Group.

Table Continued

Symptoms: FPG growth was successful but the FPG was taken out from the Remote Copy Group. The following message is displayed:

```
Failed to admit new volumes to Remote Copy Group: Unable to restart RC Group
<rcGroupName> on target <rcTarget>: Error: Could not be started on target
<rcTarget>. Attempt to start group <rcGroupName> failed: Secondary group for
group <rcGroupName> on target system <rcTarget> doesn't match primary group.
```

Conditions of occurrence: Attempt to grow an FPG while an associated Remote Copy Group is in a failed over state.

Impact: Medium

Customer circumvention: Perform a recover operation on the Remote Copy Group before attempting to grow the FPG.

Customer recovery steps: Manually grow a VV in the secondary group to match the primary group and then restart the Remote Copy Group.

Issue ID: 228736

Issue summary: File services could not be started on a controller node.

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

Affected software versions: All versions

Issue description: An internal service gets in a state where file services are not allowed on a controller node any longer.

Symptoms: When the `startfs –enable <node>` command is issued, output such as the following is seen in the task. File services remain unavailable.

```
2018-01-09 14:40:07 GMT Updated     Executing "startfsen_task -enable <node>" as 3:10968
…
2018-01-09 14:40:09 GMT Error       error: Failed to attach interface
2018-01-09 14:40:09 GMT Error       error: cannot fork child process: Cannot allocate memory
```

Conditions of occurrence: Restarting file services after the controller node has been up and running for a long time.

Impact: Medium

Customer circumvention: None

Customer recovery steps: Run the `shutdownnode check <node>` command to ensure the node can be rebooted safely. Restart the impacted controller node with the `shutdownnode reboot <node>` command. Alternatively, contact HPE Support for assistance in getting the service back in a correct state to avoid a controller node reboot.

Issue ID: 229706

Issue summary: The `growfpg` command was unsuccessful with the following message: `Error: command 'blockresize' requires <path>`.

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

Affected software versions: All versions

Issue description: The `growfpg` command could be unsuccessful when an FPG grows.

Symptoms: When this issue occurs, the `growfpg` task does not complete successfully and returns a message like:

```
Unable to grow FPG <fpgName>. Error: Failed to refresh grown VV...error: command 'blockresize' requires <path> option
```

Conditions of occurrence: Occurs when growing an FPG.

Impact: Medium

Customer circumvention: None

Customer recovery steps: Issue the `growfpg -recover_storage` command and then retry the original `growfpg` command.

Issue ID: 230966

Issue summary: During a Remote Copy failover, if the FPG name was longer than 18 characters, the FPG could not be activated.

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

Affected software versions: All versions

Issue description: The maximum length of VV name that could be used by File Persona is 23 characters. The VV name format generated when using the `creatfpg` or `growfpg` commands is `<fpgName>.N`. When creating an FPG, the name is limited to 21 characters. When an FPG is configured for Remote Copy, a suffix is typically added to the secondary volumes like "`.r`" or "`_sec`". This additional suffix could cause the VV name to exceed 23 characters, causing an unsuccessful activation of the FPG on the secondary system.

Symptoms: When an FPG was activated by using the `createfpg -recover` command on the secondary system, the following message was displayed:

```
Error: Unable to get minor number of device '/dev/tpddev/vvb/<vvid>': No such file or directory
```

Conditions of occurrence: Attempting to recover an FPG, where the VV name is 24 characters or longer.

Impact: Medium

Table Continued

Customer circumvention: After setting up Remote Copy for an FPG, confirm that the VV names on the secondary system are 23 characters or shorter. If the names are longer than 23 characters, rename them to names that are shorter than 23 characters.

Customer recovery steps: Rename the affected volumes on the secondary system to 23 characters or shorter. Reissue the `createfpg -recover` command with the volumes for the impacted FPG.

Issue ID: 232375

Issue summary: FPG did not deactivate automatically if the backing CPG was almost out of space.

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

Affected software versions: All versions

Issue description: When a CPG backing an FPG is almost out of space (as can be seen in the `free` column of the `showspace -cpg <cpgName>` command output), the system automatically deactivates all the thinly provisioned FPGs. The FPGs use the CPG to ensure that the FPGs do not experience a write error due to lack of space for additional allocations for backing VVs. If the FPG has active I/O, the deactivation request could be unsuccessful. The FPG could remain active even though there is little space left in the CPG. Continuing to write to an FPG in this state results in an alert such as the following:

```
FPG Event: FPG <fpgName> domain <domainID> has become unavailable. Reason:
ADE filesystem requested failover of FPG <fpgName>, segment 1, reason: disk
I/O error.
```

If the above alert gets displayed, contact HPE Support to run a filesystem check (FSCK) before the FPG can be reactivated.

Table Continued

Symptoms: An alert like the following could indicate that the CPG is nearly out of space:

```
Id          : 417
State       : New
Message Code: 0x027001d
Time        : 2018-05-18 14:42:59 MDT
Severity    : Critical
Type        : CPG growth failure non-admin
Message     : CPG FC_r6 SD and/or user space could not grow due to
unavailability of free space. CPG grow attempted using degraded availability
parameters (-ha mag) also failed.
If the deactivation of an FPG has also failed, an alert like the following
would be present:
Id          : 454
State       : New
Message Code: 0x00e000a
Time        : 2018-05-18 14:43:24 MDT
Severity    : Minor
Type        : Task failed
Message     : Task 26383 (type "background_command", name "setfpg_task") has
failed (Task Failed). Please see task status for details.
The failed task will have details like the following:
   Id Type               Name        Status Phase Step -------
StartTime------- ------FinishTime------- -Priority- -User--
26383 background_command setfpg_task failed   ---  --- 2018-05-18 14:43:03
MDT 2018-05-18 14:43:24 MDT n/a        3parsvc

Detailed status:
2018-05-18 14:43:03 MDT Created     task.
2018-05-18 14:43:03 MDT Updated     Executing "setfpg_task" as 0:20955
2018-05-18 14:43:04 MDT Updated     Deactivating FPG <fpgName>
2018-05-18 14:43:24 MDT Error       Failed to deactivate <fpgName>: RPC error
umount error : filesystem mountpoint  directory is busy. Filesystem <fpgName>
host <nodeName>.
2018-05-18 14:43:24 MDT Error       FPG <fpgName> was not activated on host
<backupNodeName>.
2018-05-18 14:43:24 MDT Error       Task exited with status 1
2018-05-18 14:43:24 MDT Failed      Could not complete task.
```

Conditions of occurrence: The CPG space is exhausted but an FPG has active I/O, which prevents successful deactivation.

Impact: High

Table Continued

Customer circumvention: Pay close attention to the remaining capacity in the CPGs. Whenever thin provisioning is used, it is critical that the administrator pays close attention Respond to alerts by adding capacity to the CPG or free up space in the CPG. An alert of the following type will be displayed:

```
Id           : 452
State        : New
Message Code: 0x0270010
Time         : 2018-05-18 14:26:49 MDT
Severity     : Major
Type         : FC raw space allocation 85% alert
Message      : Total FC raw space usage at 10453G (above 91% of total 11442G).
```

This alert could be resolved by adding space of the given storage class.

Customer recovery steps: Halt all I/O to and from the FPG and manually deactivate the FPG with the `setfpg` command. Once the CPG space issue has been taken care of, manually activate the FPGs by using the `setfpg` command. If the `FPG Event: FPG <fpgName> domain <domainID> has become unavailable` alert has already been raised, contact HPE Support for assistance on how to run a filesystem check to get the FPG online back again.

Issue ID: 233101

Issue summary: When a Remote Copy Group containing FPGs fails over, more than a permitted number of FPGs are listed on the target.

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

Affected software versions: 3.3.1 MU2 P39 and later

Issue description: File Persona supports 16 FPGs per node pair. It is important to take this limit into account when configuring Remote Copy. If there are more FPGs present between the two sites and the FPGs failover, the FPGs cannot be activated. One additional FPG beyond the limit is recovered, but is left in a deactivated state.

Symptoms: After a Remote Copy failover, not all FPGs are listed. Only one FPG is listed in a deactivated state. In the following example, a 2-node system that supports 16 FPGs is shown:

```
cli% showfpg
                          ------(GiB)-------
FPG           --Mountpath--- --Size-- Available ActiveStates -DefaultCpg- --------VVs-------- State    Version
src_node1fs_0 /src_node1fs_0     0.00      0.00 DEACTIVATED  fs_cpg       src_node1fs_0.1_sec degraded 12.3
tgt_node0fs_0 /tgt_node0fs_0  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_0.1     normal   12.3
tgt_node0fs_1 /tgt_node0fs_1  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_1.1     normal   12.3
tgt_node0fs_2 /tgt_node0fs_2  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_2.1     normal   12.3
tgt_node0fs_3 /tgt_node0fs_3  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_3.1     normal   12.3
tgt_node0fs_4 /tgt_node0fs_4  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_4.1     normal   12.3
tgt_node0fs_5 /tgt_node0fs_5  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_5.1     normal   12.3
tgt_node0fs_6 /tgt_node0fs_6  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_6.1     normal   12.3
tgt_node0fs_7 /tgt_node0fs_7  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node0fs_7.1     normal   12.3
tgt_node1fs_0 /tgt_node1fs_0  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_0.1     normal   12.3
tgt_node1fs_1 /tgt_node1fs_1  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_1.1     normal   12.3
tgt_node1fs_2 /tgt_node1fs_2  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_2.1     normal   12.3
tgt_node1fs_3 /tgt_node1fs_3  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_3.1     normal   12.3
tgt_node1fs_4 /tgt_node1fs_4  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_4.1     normal   12.3
tgt_node1fs_5 /tgt_node1fs_5  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_5.1     normal   12.3
tgt_node1fs_6 /tgt_node1fs_6  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_6.1     normal   12.3
tgt_node1fs_7 /tgt_node1fs_7  1024.00   1023.32 ACTIVATED    fs_cpg       tgt_node1fs_7.1     normal   12.3
-----------------------------------------------------------------------------------------------------------------
           17 total          16384.00  16373.12
```

Table Continued

Conditions of occurrence: Remote Copy configured with FPGs and more than the supported number of FPGs for the secondary site configured between the two sites.

Impact: Low

Customer circumvention: Do not configure more than the supported number of FPGs at the secondary site between the two sites when enabling Remote Copy.

Customer recovery steps: To revert to the supported number of FPGs, use the `removefpg –forget` command on the deactivated FPG. Failback the Remote Copy Group to the primary site. Configure fewer FPGs for Remote Copy or add controller nodes to support the aggregate set of FPGs between the sites.

Issue ID: 233575

Issue summary: An FPG AutoFailover was unsuccessful when the Compliance feature was enabled on the primary array but not enabled on the target or secondary array.

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

Affected software versions: 3.3.1 MU2 P39 and later

Issue description: An FPG AutoFailover could be unsuccessful when the Compliance feature is enabled on the primary array but is not enabled on the target or secondary array. The issue occurs when a primary array that is compliance-enabled has a Compliance policy applied to one or more FPGs. During failover, the FPG that is compliance-enabled will not succeed as expected. As the FPGs are recovered serially, all the FPGs that come after the compliance-enabled FPG will not get enabled on the target array.

Symptoms: A `createfpg -recover` operation did not complete successfully and returned the following status in the task:

```
Could not activate FPG <fpgName> as it has compliance enabled object(s). To
activate this FPG, enable ComplianceOfficerApproval using "setsys" and try
again.
```

Conditions of occurrence: The Compliance feature is enabled on the primary array but not on the secondary array. Some but not all FPGs are compliance-enabled.

Impact: Medium

Customer circumvention: Ensure that the Compliance feature is enabled on both the source and target arrays.

Customer recovery steps: Enable FPGs that are not compliance-enabled manually using the `createfpg -recover set:<FPG-name>` command on the secondary array.

Issue ID: 233780

Issue summary: When failback after a failover was triggered using a pushbutton failover using SSMC, not all FPGs were activated successfully.

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

Table Continued

Affected software versions: 3.3.1 MU2 P39 and later

Issue description: When multiple Remote Copy Groups are failed over or failed back in parallel, some of the FPGs might not be recoverable and activated successfully. They will be left in a deactivated state.

Symptoms: One or more FPGs could be in a deactivated state on the arrays to which the FPGs are intended to be activated. The error message when trying to manually activate such an FPG could look like the following:

```
Failed to activate <fpgName>: node <nodeName> has no access to the volumes
[[<vvId>]] used by filesystems [<fpgName>] (suggested action - check volume
access on the specified nodes)
```

Conditions of occurrence: Failover or failback of multiple Remote Copy Groups containing FPGs.

Impact: Medium

Customer circumvention: Failover or failback one Remote Copy Group at a time.

Customer recovery steps: Recover the missing FPGs manually using `removefpg -forget <FPG-name>` command. Follow the task to completion using `waittask -v <taskid>`. Recover the FPG by calling the `createfpg -recover set:<FPG-name>` command and following the task to completion.

Issue ID: 235843

Issue summary: The `checkvv -dedup_dryrun <vvName>` command was not permitted when the VV belonged to an activated FPG.

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

Affected software versions: All versions 3.2.2 MU2 or later

Issue description: If data is being written on an FPG, do not run `checkvv` on the activated FPG. Even though the `-dedup_dryrun` command does not modify the volume, it is not allowed on an activated FPG.

Symptoms: The following message is displayed when attempting to run the `checkvv -dedup_dryrun <vvName>` command:

```
<fpgName>.<N> has an associated activated FPG. Cannot run checkvv on VV of an
activated FPG. Please deactivate FPG and re-run checkvv
```

Conditions of occurrence: Attempting to run a `dedup dryrun check` on a VV that is part of an activated FPG.

Impact: Medium

Customer circumvention: Select only one of the following options:

- Deactivate the FPG. Run the check and reactivate the FPG. This will make the FPG unavailable to clients.
- Create a Virtual Copy of the FPG volumes. Run the `checkvv` command on the volumes. Delete the Virtual Copy.

Customer recovery steps: None

Issue ID: 243445

Issue summary: An FPG recovery operation returns a message that the operation was not successful.

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

Affected software versions: All versions 3.2.2 MU2 or later.

Issue description: An FPG recovery operation targeted to a volume that is already associated with an activated FPG returns a message that the operation was not successful. There is no impact to the activated FPG when File Services are in a `running` state on all nodes in the system. When File Services is in a `starting` state on one of the nodes, the FPG can become unavailable.

Symptoms: After attempting a recover operation on a volume, the system returns a message that the operation was not successful, and that the FPG using that volume is unavailable.

Conditions of occurrence: This issue occurs when an FPG is in an activated state and the request to recover an FPG uses one of the volumes of the activated FPG.

Impact: High

Customer circumvention: Use the `showfpg` command to show whether the volume belongs to an FPG.

If the volume does belong to a listed FPG, do not attempt a recover operation. Use the recover operation only when volumes belonging to FPGs were forgotten using the `removefpg -forget` option.

Customer recovery steps: Make sure that File Services are in a `running` state on all of the nodes in the system.

Deactivate the FPG, if possible, and then reactivate it.

Affected components

Applying an HPE 3PAR OS patch can cause a restart of the affected OS components. Restarting components generates events and alerts. This behavior is expected. The system continues to serve data, but a restart can interrupt existing CLI or SSMC sessions.

Component	Version
File Persona	1.5.1.900-20180914 (P39)
CLI Server	3.3.1.452 (P36)
CLI Client	3.3.1.452

Verification

To verify the installation of the latest patch from an interactive CLI session, enter the CLI command **`showversion -a -b`**.

NOTE: Because the CLI Client version in the SP code is a fixed number, the `showversion` command response sometimes shows a different version when running the command from the SP, versus the output from any other system.

```
showversion -a -b output:

Release version 3.3.1.410 (MU2)
Patches:  P30,P32,P36,P38,P39

Component Name                    Version
CLI Server                        3.3.1.452 (P36)
CLI Client                        3.3.1.452
System Manager                    3.3.1.454 (P38)
Kernel                            3.3.1.410 (MU2)
TPD Kernel Code                   3.3.1.410 (MU2)
TPD Kernel Patch                  3.3.1.454 (P38)
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)
VV Check Tools                    3.3.1.410 (MU2)
Upgrade Check Scripts             180906.U641
File Persona                      1.5.1.900-20180914 (P39)
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
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

Single Point of Connectivity Knowledge (SPOCK) Storage compatibility matrix

www.hpe.com/storage/spock

Storage white papers and analyst reports

www.hpe.com/storage/whitepapers

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

Accessing Hewlett Packard Enterprise Support

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

  http://www.hpe.com/assistance

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

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

Information to collect

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

Accessing updates

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

  Hewlett Packard Enterprise Support Center

  www.hpe.com/support/hpesc

  Hewlett Packard Enterprise Support Center: Software downloads

  www.hpe.com/support/downloads

  Software Depot

  www.hpe.com/support/softwaredepot

- To subscribe to eNewsletters and alerts:

  www.hpe.com/support/e-updates

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

  www.hpe.com/support/AccessToSupportMaterials

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

Customer self repair

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

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

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

Remote support

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

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

Remote support and Proactive Care information

HPE Get Connected

www.hpe.com/services/getconnected

HPE Proactive Care services

www.hpe.com/services/proactivecare

HPE Proactive Care service: Supported products list

www.hpe.com/services/proactivecaresupportedproducts

HPE Proactive Care advanced service: Supported products list

www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central

www.hpe.com/services/proactivecarecentral

Proactive Care service activation

www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

To view the warranty information for your product, see the links provided below:

HPE ProLiant and IA-32 Servers and Options

www.hpe.com/support/ProLiantServers-Warranties

HPE Enterprise and Cloudline Servers

www.hpe.com/support/EnterpriseServers-Warranties

HPE Storage Products

www.hpe.com/support/Storage-Warranties

HPE Networking Products

www.hpe.com/support/Networking-Warranties

Regulatory information

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

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

Additional regulatory information

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

www.hpe.com/info/reach

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

www.hpe.com/info/ecodata

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

www.hpe.com/info/environment

Documentation feedback

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