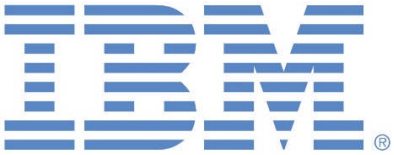

IBM FlashSystem 840

Firmware Version 1.1.2.7

Release Notes

October 8, 2014

Contents

1.0 Applicable systems

This release is intended for the following systems:

- IBM® FlashSystem™ 840, machine type 9840, model AE1
- IBM FlashSystem 840, machine type 9843, model AE1

2.0 Product resources

IBM FlashSystem 840 product information resources guide users through the various features and components of the storage system, including usage and troubleshooting guides. To read about this storage system and learn how to use or troubleshoot, see the IBM Knowledge Center for IBM FlashSystem 840 at www.ibm.com/support/knowledgecenter/ or visit the IBM Redbooks® website at www.redbooks.ibm.com for the *IBM FlashSystem 840 Product Guide*.

3.0 Bug severity legend

The following explains the bug severity ranking used in Section 4.1 for key fixes, as well as Section 7.0 for the release history:

S1: Highest	Recommend upgrade for all users as soon as possible.
S2: Medium	Recommend upgrade for all users at the next scheduled maintenance window.
S3: Average	Recommend upgrade at the next scheduled maintenance window for users experiencing these issues. All other users may wish to upgrade at the next scheduled maintenance window.
S4: Low	Upgrade at the next scheduled maintenance window. May be performed at the discretion of the user if the issue is having a negative impact.
S5: Lowest	Upgrade is not necessary. This would include a mostly cosmetic or minor annoyance fix.

4.0 Latest changes

After initial configuration of the hardware is complete, IBM strongly recommends that you make sure that your IBM FlashSystem firmware is up-to-date. Visit IBM Fix Central using the link below to see if any updates are available for your system.

4.1 Remediated security vulnerabilities

The following security vulnerabilities have been remediated:

- The "Shellshock" (also known as "bash bug" and "bashdoor") vulnerability, as detailed in the security bulletin entitled “Vulnerabilities in Bash affect IBM FlashSystem 840 and V840 (CVE-2014-6271, CVE-2014-7169, CVE-2014-7186, CVE-2014-7187, CVE-2014-6277, CVE-2014-6278)” at the following URL: http://www.ibm.com/support/docview.wss?uid=ssg1S1004932
- A NSS & NSPR vulnerability, as detailed in the security bulletin entitled “Six (6) Vulnerabilities in Network Security Services (NSS) & Netscape Portable Runtime (NSPR) affect IBM FlashSystem 840 and V840 (CVE-2013-1740, CVE-2014-1490, CVE-2014-1492, CVE-2014-1544, CVE-2014-1545)” at the following URL: http://www.ibm.com/support/docview.wss?uid=ssg1S1004930
- An OpenSSL vulnerability, as detailed in the security bulletin entitled “Seven (7) Vulnerabilities in OpenSSL affect IBM FlashSystem 840 and V840s (CVEs)” at the following URL: http://www.ibm.com/support/docview.wss?uid=ssg1S1004931
- A vulnerability in IBM Java™ SDK which affects IBM FlashSystem 840 and V840, as detailed in the security bulletin entitled “Vulnerability in IBM Java SDK affects IBM FlashSystem 840 and V840 (CVE-2014-4263)” at the following URL: http://www.ibm.com/support/docview.wss?uid=ssg1S1004929

4.2 Issues fixed

The following are fixes that are included in this release and the severity level for each fix. Use these severity levels as described in Section 3.0 to aid in your decision to upgrade.

S1

31892 - Remediate PSIRT Advisory 2106: MCP affected by Open Source - 7 issues for OpenSSL.
31893 - Remediate PSIRT Advisory 2093: MCP affected by Open Source - 2 issue(s) for glibc.
31975 - Remediate PSIRT Advisories 2209 and 2211 ("bash bug" vulnerability).

Note: If your system is on a firmware release prior to 1.1.2.2, check Section 6, the Release history. There may be critical fixes that need to be applied to your system.

S2

There are no severity level 2 issues found in this release.

S3

31666 - T3 recovery attempts to restore the iSCSI host twice and fails.

S4

There are no severity level 4 issues found in this release.

S5

There are no severity level 5 issues found in this release.

4.3 Known issues

The following sections discuss known issues that exist for this release and the workarounds available for each issue. Contact Support if you encounter any issues.

4.3.1 FCoE interface

31515, “FCoE interface links sometimes fail to come online on system power on or after mkarray command” is a current known issue. This issue can be corrected by resetting the interface using the following command:

```
svctask maintenance -action reset -canister canister_id -adapter adapter_id
```

To determine the IDs of the appropriate canister and adapter, issue the `lsportfc` command to locate the offline ports.

4.3.2 Upgrade directory

In firmware versions earlier than 1.2.1.7, a known issue exists where the upgrade directory in which upgrade files are stored after being uploaded does not automatically delete old files. In versions 1.2.x.x, update files are automatically removed from the system after 30 minutes. Each firmware update package is around 350 MB in size. The directory limit is 1.6GB, which means that the size limit is 4 packages. If the limit is reached, the update will fail with the following error message:

“The upload of the package IBM9840_INSTALL_X.X.X.X-xx.xx.tgz.gpg failed.”

A simple workaround exists for this issue and can be executed through the use of the CLI. Previously uploaded packages can also be seen in the GUI through the “Update package” drop down. However, in order to clear the older files and resolve the issue, you must use the CLI. If this issue is the reason for the update failure, you should see four existing files listed in the output. In order to enable the update, you must clear the files from the directory. To do this, issue the following command:

```
cleardumps –prefix /home/admin/upgrade
```


To ensure that the upgrade directory is now cleared, issue the following command:

```
lsdumps –prefix /home/admin/upgrade
```

4.4 Current supported specifications

SCSI-SAM-3	SCSI Architecture Model – v3
SCSI-SPC-3	SCSI Primary Commands - v3
SCSI-SBC-2	SCSI Block Commands – v2
SCSI-FCP-3	Fibre Channel Protocol for SCSI - v3
SCSI-SRP	SCSI RDMA Protocol - v1
FC-PH-3	Fibre Channel Physical and Signaling Interface - v3
FC-AL-2	Fibre Channel Arbitrated Loop – v2
IBTA-1.2	InfiniBand Trade Association Architecture Specification - v1.2

Note: In order to perform concurrent maintenance on canisters and batteries, the following procedures should be used.

4.5 Canister concurrent maintenance demonstration

In order to understand how to perform concurrent maintenance (CM), you should understand the difference between a canister and a node. A canister is one of two physical components at the rear of the enclosure, containing interface cards, fan modules, USB connectors, and Ethernet management ports. Facing the enclosure, the canister on the left is Canister 1, and the one on the right is Canister 2. The node is a logical processing unit in the enclosure. At any one time a single node in the system manages configuration activity. This configuration node manages a cache of the configuration information that describes the system configuration and provides a focal point for configuration commands. If the configuration node fails, the other node in the system takes over its responsibilities. The configuration node may be physically located in either Canister 1 or Canister 2. It is, however, important to note that node ID is not tied to a Canister ID.

The CM demonstration follows these steps:

1. To determine which canister to replace, look at the back of the FlashSystem 840. Canister 1 is on the left, as seen from the rear. Canister 2 is on the right. See Figure 1 below.

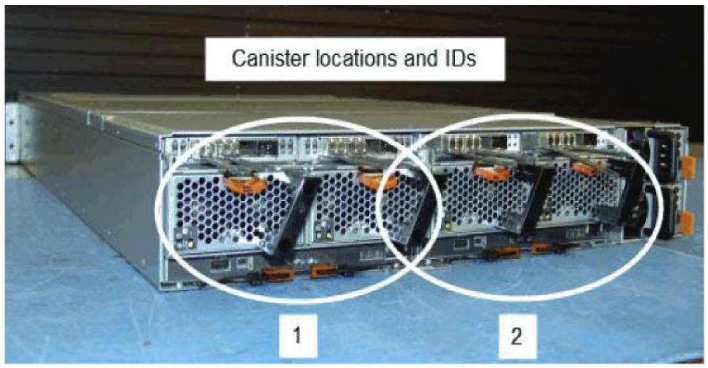

Figure 1. FlashSystem 840 rear view of canisters with canister IDs identified

2. From the Command-line Interface (CLI), run the `lsnodecanister` command to determine which node is the “config_node” as well as to view the “node_name” of the canister to be removed. In the following output example, node 1 is the configuration node and the node that we wish to remove.

```
id name  UPS_serial_number WWNN status ... config_node ...
1  node1 XXXXXXXXXXXXXXXX       online ...  yes...
2  node2 XXXXXXXXXXXXXXXX       online ...  no      ...
```

3. If the node you wish to remove is the configuration node, you must first force a failover to occur. If the node you wish to remove is not the configuration node, skip this step.
    a. Make sure that the non-configuration node has a node status of "active." Issue `sainfo lsservicenodes` to determine the node status.
    b. Place the node into a service state using the “panel_name” of the current configuration node listed in the output from the command listed in step a. Issue the following command to place the node into service state: `satask startservice –force` *`panel_name`*
    c. If you are logged in to the cluster IP, your SSH session will be terminated. Log back in to the cluster IP. Verify that the failover changed the configuration node and that the node you wish to remove is now in service state, as shown in the following example.

```
id name     UPS_serial_number WWNN   status ... config_node ...
1  node1    XXXXXXXXXXXXXXXX        online ... no          ...
2  node2    XXXXXXXXXXXXXXXX        online ... yes         ...
```

    d. Ensure that the non-configuration node is listed as "active” under the “node_status” field after issuing `sainfo lsservicenodes`. The following is the expected output:

```
panel_name . . . cluster_name node_id node_name relation node_status . . .
01-2       . . . Cluster_56    1       node1      local    Service 690 . . .
01-1       . . . Cluster_56    2       node2      partner  Active . . .
```

    Note: In the example above, the node with the node name “node1” will be removed. The cluster has failed over to the node named “node2,” which is now the configuration node, and node1 is now in service state.

4. Use the value under the “node_name” field to specify that you will shut down the node in service state in the following command: `stopsystem –node` *`node_id`*
5. Wait for the power status LED on the back of the canister to flash slowly. This LED activity means that the node has powered down.
6. View the latest events in the management GUI under Status Alerts or Monitoring → Events and click “Run Fix” to walk through the Directed Maintenance Procedure (DMP) for the canister replacement.

4.6 Battery concurrent maintenance demonstration

The following steps show a demonstration of how CM is performed on a system battery:

1. Remove a single battery.
2. View the latest events in the management GUI under Status Alerts or Monitoring → Events and click “Run Fix” to walk through the DMP for battery replacement.

5.0 Upgrading firmware

Use the following sections to perform code upgrades for your systems to the current release.

Warning: Please read all of the instructions below before upgrading.

5.1 Release overview

If you are upgrading to this release and your system is healthy, you can perform a Concurrent Code Upgrade (CCU). A CCU is a non-disruptive upgrade and is the preferred upgrade method. For general instructions on performing upgrades, refer to the FlashSystem Knowledge Center using the following URL: http://www.ibm.com/support/knowledgecenter/

From this link, you can access the appropriate FlashSystem page through the following path:

System Storage → Flash Storage → Flash high availability systems →IBM FlashSystem 840 → Upgrading

5.2 Supported upgrade paths

The following upgrade paths are supported for this release:

1.1.0.3 → 1.1.2.7
1.1.0.7 → 1.1.2.7
1.1.1.1 → 1.1.2.7
1.1.1.2 → 1.1.2.7
1.1.1.3 → 1.1.2.7
1.1.1.4 → 1.1.2.7
1.1.2.2 → 1.1.2.7
1.1.2.5 → 1.1.2.7
1.1.2.6 → 1.1.2.7

5.3 Preparing to upgrade

CCU is a non-disruptive upgrade, which means that the system remains online throughout the process and that you can continue to access data normally. As a precaution, it is recommended that the upgrade occur during a time of reduced traffic. During the upgrade, the interface adapters in each canister are taken offline temporarily in order to be upgraded. This might impact performance or throughput. The impact is more noticeable under heavy load conditions. With a properly configured multi-path configuration, access to your data is maintained at all times.

Note: With this release, a RAID certify will run upon completion of CCU to verify RAID stripes. This process takes a few hours depending on workload and can impact performance during this time.

In order to ensure a successful, non-disruptive upgrade, you should verify that your interface ports are all online and all the system hardware is functioning normally. Ideally, you should have the following:

- All host interfaces should be online. An active multi-path configuration is required to ensure no loss of access during the upgrade.
- Both batteries should be online and charged at least 85%. Use the CLI command `lsenclosurebattery` or the management GUI under Monitoring → Systems to verify battery charge.
- If using encryption, ensure both USB keys are inserted during CCU.
- All hardware should be online and functioning normally. There should be no unfixed alerts in the event log (see the exceptions below).

Important: Before you begin the upgrade, we recommend that you perform a backup of your data and a backup of the FlashSystem configuration. To back up the configuration, log into the cluster management IP address and issue the following command using admin-level authority:

```
svcconfig backup
```

Optionally, you can copy the configuration backup file from the FlashSystem to your workstation using secure copy (scp) on Linux or PuTTY secure copy (pscp.exe) on Windows as in the following examples:

(Using Linux)

```
scp superuser@cluster_ip:/dumps/svc.config.backup.* .
```

(Using Windows)

```
pscp -unsafe superuser@cluster_ip:/dumps/svc.config.backup.* .
```

Note: Do not ignore the periods shown above at the end of each command. In addition, replacement of italicized descriptions within angle brackets with appropriate information is required.

5.4 Performing the upgrade

It is highly recommended that the upgrade be performed using the web-based cluster management interface known as the management GUI. Access this option under General tab in the GUI. Using the management GUI, you will be prompted to install and run the latest version of the software upgrade test utility, which is designed to detect and warn of various conditions that prevent a successful upgrade. For more detailed instructions and information on using the upgrade utility from the CLI or the GUI, see the upgrade utility release notes available on Fix Central.

The upgrade can also be performed using the `applysoftware` command using the CLI. This requires that you manually upload the latest release to the `/upgrade` directory on the cluster management node.

5.5 Troubleshooting

Use the following sections to troubleshoot problems that may occur during the upgrade process.

5.5.1 Stalled upgrade

If the upgrade takes more than two hours to complete, it may have stalled. Upgrade status is viewed by issuing `lssoftwareupgradestatus` CLI command. In most cases, this can be resolved by aborting the upgrade and reattempting the upgrade after the system downgrades to its original level. To abort the upgrade, issue the `applysoftware -abort` CLI command or click the “Stop Upgrade” button in the GUI, as seen in Figure 1 below.

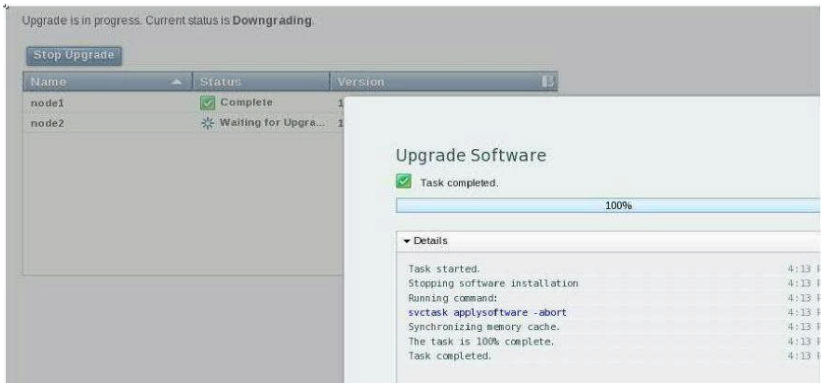

Figure 2. Aborting the upgrade after it has stalled

After the system is downgraded, you can reattempt your upgrade from the GUI or CLI. If the upgrade stalls repeatedly or if you have alerts which cannot be cleared, contact IBM Support. Support contact information is provided in Section 6.0 of this document.

5.5.2 Failures during upgrade

You may get a battery or quorum alert during upgrade due to required reconfiguration. These alerts need to be cleared once the CCU process is complete. Follow the instructions provided by the fix procedures in the Events view of the GUI. In many cases, these issues can be resolved by reseating the power supply or battery as instructed by the GUI DMP. After events are fixed, they may be visible from the Events view of the management GUI if the filter is set to "Show All," but they should no longer appear in the Recommended Actions, Unfixed Messages, or Alerts views. If you see unfixed battery or quorum alerts after an upgrade is complete, contact IBM Support.

If the upgrade has failed or stopped due to a hardware failure, you will see the “Hardware Failed” status as presented in Figure 2 below.

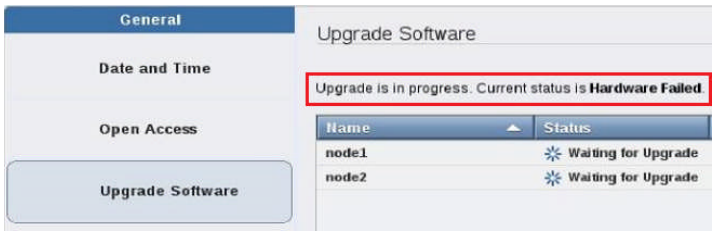

Figure 3. Viewing the upgrade status

If you suspect a hardware failure, you can issue the `lssoftwareupgradestatus` command to confirm the state of your system. This command shows that the system is in a "hardware_failed" state, and that the event log contains a "System upgrade suspended" event. You may resume the upgrade by issuing the `applysoftware -resume –force` command for the following conditions:

- PSU Unsupported events
- Battery fault type 1 faults that are fixed and online according to CLI command `lsenclosurebattery`
- Fan events

If the upgrade cannot be resumed or you have other alerts which cannot be cleared, contact IBM Support. Upon completion of upgrade to 1.1.2.7, a RAID certify will run to verify RAID stripes. This process takes a few hours depending on workload and can impact performance during this time.

6.0 Contact information

Call IBM at 1-800-IBM-SERV (1-800-426-7378). To find contact information for a specific region, visit the IBM directory of worldwide contacts at http://www.ibm.com/planetwide/.

7.0 Release history

The following sections include a list of all fixes and improvements for previous releases.

Release 1.1.2.6

29687, 29694, 30973, 30992 - Improved Robustness of Hot Canister Pulls. (S1)
30524 - T3 failed with Error code 78 (S2)
30999, 31602 - T3 recovery fails due to quorum data offset by 32 bytes after system power cycle (S2)
31308 - Direct attach FC P2P results in incorrect credit and port unfairness, also 31346 (S2)
31267 - FC-4 target reset not working correctly (S2)
31106 - SCSI III Persistent Reservations (without APTPL bit set) do not persist across node failover (S2)
29021 - System is now resilient to canister failures during a rebuild. (S2)
29814 - Intermittent 114 Enclosure Battery fault type 1 (S2)
29881 - MGMT to XBAR link integrity improvement (S2)
28767 - Battery inlet temperature sensor no longer reports impossible values. (S2)
28767 - Battery inlet temperature sensor no longer reports impossible values. (S2)
29124 - PMOM: Add mask bits for PSU error reporting to software (S2)
30806 - "chsystem -alias" command does not work, causing incorrect UUIDs after T3 (S2)
30937 - T3 does not recover vdisk mapping SCSI ids (LUN) (S2)
30143 - System report with canister powered off can crash remaining canister (S2)
30740 - Changing flashcard positions while the system is powered off leads to data loss (S2)
29508 - Interface's PCI links have issues while data running, should fail interface not both xbars (S2)
29917 - Fixed issue with FRU replacement upgrades causing nodes to come up with 841 node error (S2)
26338 - Fixed DMA stall due to wear leveling moves (S2)
30754 - Fixed erroneous interface timeout that can occur approximately 497 days after boot. (S2)
30920 - Interface with PCI errors can incorrectly take down access to entire system (S2)
29515 - Fixed issue that caused a RAID controller failure on canister when hot removing the other canister while running and replacing it within 15 seconds. (S2)
31251 - Non config SYS panic due to battery information misread (S3)
31149 - Remove "Battery not charging when it should be" log message (S3)
27191 - Correctly read Power Supply’s VPD which fixes an un-clearable error in the event log. (S3)
29028 - Added hysteresis to temperature monitoring to avoid strong reactions to temporary temperature spikes due to sensor inexactness. (S3)
29443 - Only an active management node can power off flashcards; previously the redundant node could cause a power off if it detected bad batteries, unfortunately the user may not be properly notified of the power off since it was being sourced by the redundant node. (S3)
28698 - Concurrent code updates are more resilient to interface programming errors. (S3)
28924 - Interface connectivity alerts greatly improved and clarified. (S3)
27922 - Power supply temperature notices no longer flood system logs. (S3)
28622 - Flash data retention logic now correctly notes system power off time to better correct for stored data on flash cards powered off for extended periods of time. (S3)
29539 - ftdc - system report - seeing samnet issues in certain cases (S3)
29970 - lsibportcandidate does not show any IB ports (S3)
30448 - svctask chnodehw asserts when used to fix node error 841 (S3)
30872 - lsiogrphost command caused a node restart on non-config node (S3)
29819 - Non-config Intel complex dies during satask snap (S3)
29946 - Concurrent replacement of Canister 2 caused EventID=085071 error 1039 (S3)
29822 - Fan failure not cleared on node failover (S3)

30323 - Fixed canister power off issues during battery concurrent upgrades (S3)
29705, 30140, 30261 - Fixed drive encryption issues after a concurrent update (S3)
29534 - Improved automatic reset recovery of interfaces (S3)
29818 - Improved error handling of interfaces during automatic internal hardware upgrades (S3)
29848 - Improved error handling of RAID Controllers during automatic internal hardware upgrades (S3)
29957 - Improved software upgrade test utility to check Drives and PSU states (S3)
30051 - Fixed quorum communication timeouts during concurrent updates (S3)
30119 - Improved checks for interface failures during RAID Controller concurrent updates (S3)
30135 - Handle batteries with bad manufacturing set up during upgrades (S3)
30240 - Add better handling of single canister upgrades (S3)
30284 - Improve handling of PSoC upgrade failures (S3)
28487 - ecmon kills volume manager SOMETIMES on boot, killing node during CCL (S3)
25822 - Write CPLD mask register on PMOM reconfigure to prevent canister power off during CCL (S3)
29505, 29979 -- Fix I2C issues on battery module (S3)
29833 - Increase fault tolerance on reads to the battery module (S3)
30350 - Fix issue with not high charging current when both batteries are offline. (S3)
29472 - Added support for UNTAGGED task attribute. This had been removed in a previous release as obsolete, but is needed for interoperability support. **(**S3)
29628 - Fixed interface cards failure when both batteries are pulled out and replaced while data is running. (S3)
30259 - A resolve transport ID conflicts with certain versions of SVC (S3)
29712 - Fixed issue with sending FLOGI requests with non-zero S_ID values when moving Fibre Channel cables between switch ports. (S3)
29864 - XBAR link failure to GBE or Interface following CCL of XBAR can cause Interface heartbeat failure (S3)
24802 - The PTPL_A bit in the PERSISTENT RESERVE IN command is now reported correctly with the REPORT CAPABILITIES service action. (S3)
29651 - quick (< 1 second) battery reseat caused quorum error (S4)
30416 - 2 instances of rsyslogd running on Texan (S4)
29545 - Fix packet to orca fpgas times out during canister insertion (S4)
30444 - Flashcard log messages not being serviced from log buffer (S4)
30704 - Reseating canisters one after another cause samnet timeouts (S4)
29427 - Made upgrade more resilient to node failovers (S4)
29729, 29787 - Improved battery state checking during concurrent updates (S4)
30295 - Flashcard "CCE" -- Enable the SEM core to check and, with FW assist, correct configuration memory (S4)
30219 - Fix issue with system not getting out of service 657 state when both batteries are removed from system (S4)
29894 - Management FPGA is not always patched on booth (S4)
28949 - Average latency performance shown in GUI no longer averages unused links (which report “0” and cause better latency numbers when averaged). (S4)
28948 - Repeated SCSI RESERVE6 commands no longer cause failures. (S4)
28818 - When system is powered off, but plugged in, power supply fans will only ramp up when the batteries are charging. (S4)
29186 - Corrected a variety of displayed URLs to point to valid web sites; in particular, the code upgrade check tool can now be downloaded using the embedded URL link. (S4)
29036 - Can now recover an offline array into a degraded state; for example if you powered off a system and removed the spare and a data flashcard, you can now add a single card and recover the array. (S4)
27761 - FC link breaks will now wait 10 seconds before cancelling all open exchanges and logging out all connected initiators for that port. (S4)
29549 - Power controller monitoring no longer result in false battery failures. (S4)

29451 - Average latency performance shown in GUI no longer averages unused links (which report “0” and cause better latency numbers when averaged). (S4)
29020 - Temperature statistics now being collected. (S4)
29392 - Battery goes into fault type 1 temporarily due to issues with battery gas gauge communication. (S4)
28978 - Now correctly detect failed RAID controller instead of incorrectly failing flashcards. (S4)
29012 - Battery charge voltage doesn't get set properly due to issues with battery gas gauge communication. (S4)
28823 - Canister concurrent maintenance now supported. (S4)
28948 - Repeated SCSI RESERVE6 commands no longer cause failures. (S4)
28818 - When system is powered off, but plugged in, power supply fans will only ramp up when the batteries are charging. (S4)
28760 - Non-concurrent upgrades done via CLI now are correctly represented by the GUI. (S4)
28902 - Interface “Links Degraded” messages can now be correctly cleared. (S4)
29186 - Corrected a variety of displayed URLs to point to valid web sites; in particular, the code upgrade check tool can now be downloaded using the embedded URL link. (S4)
28960 - Upon power-on, flashcard now correctly report power status, thus preventing initial flashcard arrays being incorrectly created without all members. (S4)
28903 - Interface “Links Degraded” message downgraded from “alert” to “warning”. (S4)
31099 - Switching rebuild from one spare to another fails all interfaces. (S4)
31104 - Run array certify on CCL completion to check RAID stripe integrity (S4)
31157 - Flashcard failure during certify fails interfaces (S4)
31262 - Increase drive certify time out (S4)
29893 - Displayed Battery charge percentage calculation not correct (S5)
30575 - Panel name renames as serial number when clustered (S5)
29645 - Allow CLI access from the serial port user (S5)
29838 - svcinfo lsnode shows partner node status is 'offline' instead of 'service' preventing obtaining partner snap (S5)
29639 - sainfo lsservicenodes error_data column showing incorrect string (S5)
29274 - Status of Ethernet port does not get updated in lsportip (S5)
29829 - HWERRLOG fills up and stops rather than rotates (S5)
30159 - System will generate Call Home PMR but will NOT heartbeat (S5)
29891 - Added new events for automatic internal hardware upgrades during boot (S5)
30023 - Added a manual way to bypass the 30 min multipath wait during upgrades (S5)
29386 - Fixed issue with concurrent upgrade lssoftwareupgradestatus command percent_complete field decreasing (S5)
29120 - Improved support maintenance command (S5)
29342 - Improved PSoC error handling during upgrades (S5)
24943 - Persist flash program erase count across flashcard initialization (S5)
26915 - Correctly calculate the amount of time required for a battery to charge when using "sainfo lsservicestatus" (S5)
29525 - Add system power statistics to call home data (S5)
29538 - GUI stat "Latency" renamed "Internal Latency" to correctly represent the statistic. (S5)
30004 - svc_snap -dumpall returns CMMVC5741E The filter value [] is not valid (S5)

8.0 Copyright notice

9.0 Revision history

The following table outlines the revision history of this document.

Document version	Date	Revision details
2.0	October 26, 2015	Entire document – formatting changed, wording improved. Section 4.3 – Known issue added.