



AssuredSAN Event Descriptions Reference Guide

Abstract

This guide is for reference by storage administrators to help troubleshoot storage-system issues. It describes event messages that may be reported during system operation and specifies any actions recommended in response to an event.

Copyright © 2014 Dot Hill Systems Corp. All rights reserved. Dot Hill Systems Corp., Dot Hill, the Dot Hill logo, AssuredSAN, AssuredSnap, AssuredCopy, AssuredRemote, EcoStor, SimulCache, R/Evolution, and the R/Evolution logo are trademarks of Dot Hill Systems Corp. All other trademarks and registered trademarks are proprietary to their respective owners.

The material in this document is for information only and is subject to change without notice. While reasonable efforts have been made in the preparation of this document to assure its accuracy, changes in the product design can be made without reservation and without notification to its users.



Adobe PostScript

Contents

About this guide	4
Intended audience	4
Prerequisites	4
Document conventions and symbols	4
Event descriptions	6
Introduction	6
Events and event messages	6
Event format in this guide	6
Resources for diagnosing and resolving problems	6
Event descriptions	7
Troubleshooting steps for leftover disk drives	71
Using the trust command	71
Power supply faults and recommended actions	72
Events sent as indications to SMI-S clients	73
Glossary	74

About this guide

This guide describes events that the AssuredSAN™ storage systems may report and recommended actions to take in response to those events. It also gives more details for troubleshooting leftover disks and warnings for usage of the `trust` command.

Intended audience

This guide is intended for storage system administrators and service personnel.

Prerequisites

Prerequisites for using this product include knowledge of:

- Network administration
- Storage system configuration
- Storage area network (SAN) management and direct attach storage (DAS)
- Fibre Channel, Serial Attached SCSI (SAS), Internet SCSI (iSCSI), and Ethernet protocols
- RAID technology

Before you begin to follow procedures in this guide, you must have already installed enclosures and learned of any late-breaking information related to system operation, as described in the Setup Guide and in Release Notes.

Document conventions and symbols

Table 1 Document conventions

Convention	Element
Blue text	Cross-reference links
Blue, underlined text	Email addresses
Blue, underlined text	Website addresses
Bold text	<ul style="list-style-type: none">• Keys that are pressed• Text typed into a GUI element, such as a box• GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes
<i>Italic</i> text	Text emphasis
Monospace text	<ul style="list-style-type: none">• File and directory names• System output• Code• Commands, their arguments, and argument values
<i>Monospace, italic</i> text	<ul style="list-style-type: none">• Code variables• Command variables
Monospace, bold text	Emphasized of file and directory names, system output, code, and text typed at the command line

△ **CAUTION:** Indicates that failure to follow directions could result in damage to equipment or data.

📄 **IMPORTANT:** Provides clarifying information or specific instructions.

NOTE: Provides additional information.

 **TIP:** Provides helpful hints and shortcuts.

1 Event descriptions

Introduction

This guide is for reference by storage administrators and technical support personnel to help troubleshoot storage-system issues. It describes event messages that may be reported during system operation and specifies any actions recommended in response to an event.

This guide applies to AssuredSAN 3004, 4004, and Ultra48 Series storage systems, and to AssuredSAN 2333, 2002 Series, and 3000 Series storage systems that have been updated to the most recent firmware available. It describes all event codes that exist as of publication. Depending on your system model and firmware version, some events described in this guide may not apply to your system. The event descriptions should be considered as explanations of events that you do see. They should not be considered as descriptions of events that you should have seen but did not. In such cases those events probably do not apply to your system.

In this guide:

- The term *disk group* refers to either a vdisk for linear storage or a virtual disk group for virtual storage.
- The term *pool* refers to either a single vdisk for linear storage or a virtual pool for virtual storage.

Events and event messages

When an event occurs in a storage system, an event message is recorded in the system's event log and, depending on the system's event notification settings, may also be sent to users (using email) and host-based applications (via SNMP or SMI-S).

Each event has a numeric code that identifies the type of event that occurred, and has one of the following severities:

- **Critical:** A failure occurred that may cause a controller to shut down. Correct the problem *immediately*.
- **Error:** A failure occurred that may affect data integrity or system stability. Correct the problem as soon as possible.
- **Warning:** A problem occurred that may affect system stability but not data integrity. Evaluate the problem and correct it if necessary.
- **Informational:** A configuration or state change occurred, or a problem occurred that the system corrected. No immediate action is required. In this guide, this severity is abbreviated as "Info."

An event message may specify an associated error code or reason code, which provides additional detail for technical support. Error codes and reason codes are outside the scope of this guide.

Event format in this guide

This guide lists events by event code and severity, where the most severe form of an event is described first. Events are listed in the following format.

Event code

Severity Event description.

Recommended actions

- If the event indicates a problem, actions to take to resolve the problem.

Resources for diagnosing and resolving problems

For further information about diagnosing and resolving problems, see:

- The troubleshooting chapter and the LED descriptions appendix in your product's Setup Guide
- The topics about verifying component failure in your product's FRU Installation and Replacement Guide

For a summary of storage events and corresponding SMI-S indications, see [Events sent as indications to SMI-S clients](#) on page 73.

Event descriptions

1

Warning If the indicated disk group is RAID 6, it is operating with degraded health due to the failure of two disks.

If the indicated disk group is not RAID 6, it is operating with degraded health due to the failure of one disk.

The disk group is online but cannot tolerate another disk failure.

If a dedicated spare (linear only) or global spare of the proper type and size is present, that spare is used to automatically reconstruct the disk group. Events 9 and 37 are logged to indicate this. For linear disk groups, if no usable spare disk is present, but an available disk of the proper type and size is present and the dynamic spares feature is enabled, that disk is used to automatically reconstruct the disk group and event 37 is logged.

Recommended actions

- If no spare was present and the dynamic spares feature (linear only) is disabled (event 37 was NOT logged), replace the failed disk with one of the same type and the same or greater capacity. The new disk will be used to automatically reconstruct the disk group. Confirm this by checking that events 9 and 37 are logged.
- Otherwise, reconstruction automatically started and event 37 was logged. Replace the failed disk and configure the replacement as a dedicated (linear only) or global spare for future use.
- For continued optimum I/O performance, the replacement disk should have the same or better performance.
- Confirm that all failed disks have been replaced and that there are sufficient spare disks configured for future use.

3

Error The indicated disk group went offline.

One disk failed for RAID 0 or NRAID, three disks failed for RAID 6, or two disks failed for other RAID levels. The disk group cannot be reconstructed. This is not a normal status for a disk group unless you have done a manual dequarantine.

For virtual disk groups, when a disk failure occurs the data in the disk group that uses that disk will be automatically migrated to another available disk group if space is available, so no user data is lost. Data will be lost only if multiple disk failures occur in rapid succession so there is not enough time to migrate the data, or if there is insufficient space to fit the data in another tier, or if failed disks are not replaced promptly by the user.

Recommended actions

- The CLI `trust` command may be able to recover some of the data in the disk group. See the CLI help for the `trust` command. It is recommended that you contact technical support for assistance in determining if the `trust` operation is applicable to your situation and for assistance in performing it.
- If you choose to not use the `trust` command, perform these steps:
 - Replace the failed disk or disks. (Look for event 8 in the event log to determine which disks failed and for advice on replacing them.)
 - Delete the disk group (`remove disk-groups` CLI command).
 - Re-create the disk group (`add disk-group` CLI command).
- To prevent this problem in the future, use a fault-tolerant RAID level, configure one or more disks as spare disks, and replace failed disks promptly.

4

Info. The indicated disk had a bad block which was corrected.

Recommended actions

- Monitor the error trend and whether the number of errors approaches the total number of bad-block replacements available.

6

Warning A failure occurred during initialization of the indicated disk group. This was probably caused by the failure of a disk drive. The initialization may have completed but the disk group probably has a status of FTDN (fault tolerant with a down disk), CRIT (critical), or OFFL (offline), depending on the RAID level and the number of disks that failed.

Recommended actions

- Look for another event logged at approximately the same time that indicates a disk failure, such as event 55, 58, or 412. Follow the recommended actions for that event.

Info. Disk group creation failed immediately. The user was given immediate feedback that it failed at the time they attempted to add the disk group.

Recommended actions

- No action is required.

7

Error In a testing environment, a controller diagnostic failed and reports a product-specific diagnostic code.

Recommended actions

- Perform failure analysis.

8

Warning One of the following conditions has occurred:

- A disk that was part of a disk group is down. The indicated disk in the indicated disk group failed and the disk group probably has a status of FTDN (fault tolerant with a down disk), CRIT (critical), or OFFL (offline), depending on the RAID level and the number of disks that failed. If a spare is present and the disk group is not offline, the controller automatically uses the spare to reconstruct the disk group. Subsequent events indicate the changes that happen to the disk group. When the problem is resolved, event 9 is logged.
- Reconstruction of a disk group failed. The indicated disk was being used as the target disk for reconstructing the indicated disk group. While the disk group was reconstructing, another disk in the disk group failed and the status of the disk group went to OFFL (offline). The indicated disk has a status of LEFTOVR (leftover).

Recommended actions

- If a disk that was part of a disk group is down:
 - If the indicated disk failed because of excessive media errors, imminent disk failure, possible hardware failure, or a disk that is not supported, replace the disk.
 - If the indicated disk failed because a user forced the disk out of the disk group, or for an unknown reason, and the associated disk group is offline or quarantined, contact technical support; otherwise, clear the disk's metadata to reuse the disk.
 - If the indicated disk failed because a previously detected disk is no longer present, reinsert the disk or insert a replacement disk. If the disk then has a status of leftover (LEFTOVR), clear the metadata to reuse the disk. If the associated disk group is critical, event 1 will also be logged; see the recommended actions for that event. If the associated disk group is offline or quarantined, contact technical support.

- If reconstruction of a disk group failed:
 - If the associated disk group is online, clear the indicated disk's metadata so that the disk can be re-used.
 - If the associated disk group is offline, the CLI `trust` command may be able to recover some or all of the data in the disk group. However, trusting a partially reconstructed disk may lead to data corruption. See the CLI help for the `trust` command. It is recommended that you contact technical support for assistance in determining if the `trust` operation is applicable to your situation and for assistance in performing it.
 - If the associated disk group is offline and you do not want to use the `trust` command, perform these steps:
 - Delete the disk group (`remove disk-groups` CLI command).
 - Clear the indicated disk's metadata so the disk can be re-used (`clear disk-metadata` CLI command).
 - Replace the failed disk or disks. (Look for other instances of event 8 in the event log to determine which disks failed.)
 - Re-create the disk group (`add disk-group` CLI command).
- If you replace a disk, the replacement disk must be the same type (SAS SSD, enterprise SAS, or midline SAS) and the same or greater capacity. For continued optimum I/O performance, the replacement disk should have performance that is the same as or better than the one it is replacing.

9

Info. The indicated spare disk has been used in the indicated disk group to bring it back to a fault-tolerant status.

Disk group reconstruction starts automatically. This event indicates that a problem reported by event 8 is resolved.

Recommended actions

- No action is required.

16

Info. The indicated disk has been designated a global spare.

Recommended actions

- No action is required.

18

Info. Disk group reconstruction completed.

Recommended actions

- No action is required.

19

Info. A rescan has completed.

Recommended actions

- No action is required.

20

Info. Storage Controller firmware update has completed.

Recommended actions

- No action is required.

21

Error Disk group verification completed. Errors were found but not corrected.

Recommended actions

- Perform a disk group scrub to find and correct the errors.

Warning Disk group verification did not complete because of an internally detected condition such as a failed disk.

If a disk fails, data may be at risk.

Recommended actions

- Resolve any non-disk hardware problems, such as a cooling problem or a faulty controller module, expansion module, or power supply.
- Check whether any disks in the disk group have logged SMART events or unrecoverable read errors.
 - If so, and the disk group is a non-fault-tolerant RAID level (RAID 0 or non-RAID), copy the data to a different disk group and replace the faulty disks.
 - If so, and the disk group is a fault-tolerant RAID level, replace the faulty disks. Before replacing a disk, confirm that a reconstruction is not currently running on the disk group. It is also recommended to make a full backup of all the data in the disk group before replacing disks. If more than one disk in the disk group has errors, replace the disks one at a time and allow reconstruction to complete after each disk is replaced.

Info. Disk group verification failed immediately, was aborted by a user, or succeeded.

Recommended actions

- No action is required.

23

Info. Disk group creation has started.

Recommended actions

- No action is required.

25

Info. Disk group statistics were reset.

Recommended actions

- No action is required.

28

Info. Controller parameters have been changed.

This event is logged when general configuration changes are made. For example, utility priority, remote notification settings, user interface passwords, and network port IP values. This event is not logged when changes are made to disk group or volume configuration.

Recommended actions

- No action is required.

31

Info. The indicated disk is no longer a global or dedicated spare.

Recommended actions

- No action is required.

32

Info. Disk group verification has started.

Recommended actions

- No action is required.

33

Info. Controller time/date has been changed.

This event is logged before the change happens, so the timestamp of the event shows the old time. This event may occur often if NTP is enabled

Recommended actions

- No action is required.

34

Info. The controller configuration has been restored to factory defaults.

Recommended actions

- For an FC controller, restart it to make the default loop ID take effect.

37

Info. Disk group reconstruction has started. When complete, event 18 is logged.

Recommended actions

- No action is required.

39

Warning The sensors monitored a temperature or voltage in the warning range. When the problem is resolved, event 47 is logged for the component that logged event 39.

If the event refers to a disk sensor, disk behavior may be unpredictable in this temperature range.

Check the event log to determine if more than one disk has reported this event.

- If multiple disks report this condition there could be a problem in the environment.
- If one disk reports this condition, there could be a problem in the environment or the disk has failed.

Recommended actions

For a 2U12 or 2U24 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5–40° C (41° F–104° F).
- Check for any obstructions to the airflow.
- Check that there is a module or blank plate in every module slot in the enclosure.
- If none of the above explanations apply, replace the disk or controller module that logged the error.

For a 2U48 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5–35° C (41° F–95° F).
- Check for any obstructions to the airflow.
- Check that the drawers are closed and there is a module or blank plate in every module slot in the enclosure.
- If none of the above explanations apply, replace the disk or controller module that logged the error.

40

Error The sensors monitored a temperature or voltage in the failure range. When the problem is resolved, event 47 is logged for the component that logged event 40.

Recommended actions

For a 2U12 or 2U24 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5–40° C (41° F–104° F).
- Check for any obstructions to the airflow.
- Check that there is a module or blank plate in every module slot in the enclosure.
- If none of the above explanations apply, replace the disk or controller module that logged the error.

For a 2U48 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5–35° C (41° F–95° F).
- Check for any obstructions to the airflow.
- Check that the drawers are closed and there is a module or blank plate in every module slot in the enclosure.
- If none of the above explanations apply, replace the disk or controller module that logged the error.

41

Info. The indicated disk has been designated a spare for the indicated disk group.

Recommended actions

- No action is required.

43

Info. The indicated disk group has been deleted.

Recommended actions

- No action is required.

44

Warning The controller contains cache data for the indicated volume but the corresponding disk group is not online.

Recommended actions

- Determine the reason that the disks comprising the disk group are not online.
- If an enclosure is down, determine corrective action.
- If the disk group is no longer needed, you can clear the orphan data. This will result in lost data.
- If the disk group is missing and was not intentionally removed, see [Resources for diagnosing and resolving problems](#) on page 6.

47

Info. An error detected by the sensors has been cleared. This event indicates that a problem reported by event 39 or 40 is resolved.

Recommended actions

- No action is required.

48

Info. The indicated disk group has been renamed.

Recommended actions

- No action is required.

49

Info. A lengthy SCSI maintenance command has completed. (This typically occurs during disk firmware update.)

Recommended actions

- No action is required.

50

Error A correctable ECC error occurred in cache memory more than 10 times during a 24-hour period, indicating a probable hardware fault.

Recommended actions

- Replace the controller module that logged this event.

Warning A correctable ECC error occurred in cache memory.

This event is logged with Warning severity to provide information that may be useful to technical support, but no action is required now. It will be logged with Error severity if it is necessary to replace the controller module.

Recommended actions

- No action is required.

51

Error An uncorrectable ECC error occurred in cache memory more than once during a 48-hour period, indicating a probable hardware fault.

Recommended actions

- Replace the controller module that logged this event.

Warning An uncorrectable ECC error occurred in cache memory.

This event is logged with Warning severity to provide information that may be useful to technical support, but no action is required now. It will be logged with Error severity if it is necessary to replace the controller module.

Recommended actions

- No action is required.

52

Info. Disk group expansion has started.

This operation can take days, or weeks in some cases, to complete. Allow adequate time for the expansion to complete.

When complete, event 53 is logged.

Recommended actions

- No action is required.

53

Warning Too many errors occurred during disk group expansion to allow the expansion to continue.

Recommended actions

- If the expansion failed because of a disk problem, replace the disk with one of the same type (SAS SSD, enterprise SAS, or midline SAS) and the same or greater capacity. For continued optimum I/O performance, the replacement disk should have performance that is the same as or better than the one it is replacing. If disk group reconstruction starts, wait for it to complete and then retry the expansion.

Info. Disk group expansion either completed, failed immediately, or was aborted by a user.

Recommended actions

- If the expansion failed because of a disk problem, replace the disk with one of the same type (SAS SSD, enterprise SAS, or midline SAS) and the same or greater capacity. For continued optimum I/O performance, the replacement disk should have performance that is the same as or better than the one it is replacing. If disk group reconstruction starts, wait for it to complete and then retry the expansion.

55

Warning The indicated disk reported a SMART event.

A SMART event indicates impending disk failure.

Recommended actions

- Resolve any non-disk hardware problems, especially a cooling problem or a faulty power supply.
- If the disk is in a disk group that uses a non-fault-tolerant RAID level (RAID 0 or non-RAID), copy the data to a different disk group and replace the faulty disk.
- If the disk is in a disk group that uses a fault-tolerant RAID level, replace the faulty disk. Before replacing the disk, confirm that a reconstruction is not currently running on the disk group. It is also recommended to make a full backup of all the data in the disk group before replacing disks. If more than one disk in the disk group has reported SMART events, replace the disks one at a time and allow reconstruction to complete after each disk is replaced.

56

Info. A controller has powered up or restarted.

Recommended actions

- No action is required.

58

Error A disk drive detected a serious error, such as a parity error or disk hardware failure.

Recommended actions

- Replace the failed disk with one of the same type (SAS SSD, enterprise SAS, or midline SAS) and the same or greater capacity. For continued optimum I/O performance, the replacement disk should have performance that is the same as or better than the one it is replacing.

Warning A disk drive reset itself due to an internal logic error.

Recommended actions

- The first time this event is logged with Warning severity, if the indicated disk is not running the latest firmware, update the disk firmware.
- If this event is logged with Warning severity for the same disk more than five times in one week, and the indicated disk is running the latest firmware, replace the disk with one of the same type (SAS SSD, enterprise SAS, or midline SAS) and the same or greater capacity. For continued optimum I/O performance, the replacement disk should have performance that is the same as or better than the one it is replacing.

Info. A disk drive reported an event.

Recommended actions

- No action is required.

59

Warning The controller detected a parity event while communicating with the indicated SCSI device. The event was detected by the controller, not the disk.

Recommended actions

- If the event indicates that a disk or an expansion module is bad, replace the indicated device.

Info. The controller detected a non-parity error while communicating with the indicated SCSI device. The error was detected by the controller, not the disk.

Recommended actions

- No action is required.

61

Error The controller reset a disk channel to recover from a communication error. This event is logged to identify an error trend over time.

Recommended actions

- If the controller recovers, no action is required.
- View other logged events to determine other action to take.

62

Warning The indicated global or dedicated spare disk has failed.

Recommended actions

- Replace the disk with one of the same type (SAS SSD, enterprise SAS, or midline SAS) and the same or greater capacity. For continued optimum I/O performance, the replacement disk should have performance that is the same as or better than the one it is replacing.
- If the failed disk was a global spare, configure the new disk as a global spare.
- If the failed disk was a dedicated spare, configure the new disk as a dedicated spare for the same disk group.

65

Error An uncorrectable ECC error occurred in cache memory on startup.

The controller is automatically restarted and its cache data are restored from the partner controller's cache.

Recommended actions

- Replace the controller module that logged this event.

68

Info. The controller that logged this event is shut down, or both controllers are shut down.

Recommended actions

- No action is required.

71

Info. The controller has started or completed failing over.

Recommended actions

- No action is required.

72

Info. After failover, recovery has either started or completed.

Recommended actions

- No action is required.

73

Info. The two controllers are communicating with each other and cache redundancy is enabled.

Recommended actions

- No action is required.

74

Info. The FC loop ID for the indicated disk group was changed to be consistent with the IDs of other disk groups. This can occur when disks that constitute a disk group are inserted from an enclosure having a different FC loop ID.

This event is also logged by the new owning controller after disk group ownership is changed.

Recommended actions

- No action is required.

75

Info. The indicated volume's LUN (logical unit number) has been unassigned because it conflicts with LUNs assigned to other volumes. This can happen when disks containing data for a mapped volume have been moved from one storage system to another.

Recommended actions

- If you want hosts to access the volume data in the inserted disks, map the volume with a different LUN.

76

Info. The controller is using default configuration settings. This event occurs on the first power up, and might occur after a firmware update.

Recommended actions

- If you have just performed a firmware update and your system requires special configuration settings, you must make those configuration changes before your system will operate as before.

77

Info. The cache was initialized as a result of power up or failover.

Recommended actions

- No action is required.

78

Warning The controller could not use an assigned spare for a disk group because the spare's capacity is too small.

This occurs when a disk in the disk group fails, there is no dedicated spare available and all global spares are too small or, if the dynamic spares feature is enabled, all global spares and available disks are too small, or if there is no spare of the correct type. There may be more than one failed disk in the system.

Recommended actions

- Replace each failed disk with one of the same type (SAS SSD, enterprise SAS, or midline SAS) and the same or greater capacity. For continued optimum I/O performance, the replacement disk should have performance that is the same as or better than the one it is replacing.
- Configure disks as dedicated spares or global spares.

- For a dedicated spare, the disk must be of the same type as the other disks in the disk group and at least as large as the smallest-capacity disk in the disk group, and it should have the same or better performance.
- For a global spare, it is best to choose a disk that is as big as or bigger than the largest disk of its type in the system and of equal or greater performance. If the system contains a mix of disk types (SAS SSD, enterprise SAS, or midline SAS), there should be at least one global spare of each type (unless dedicated spares are used to protect every disk group of a given type).

79

Info. A trust operation has completed for the indicated disk group.

Recommended actions

- Be sure to complete the trust procedure as documented in the CLI help for the `trust` command.

80

Info. The controller enabled or disabled the indicated parameters for one or more disks.

Recommended actions

- No action is required.

81

Info. The current controller has unkilld the partner controller. The other controller will restart.

Recommended actions

- No action is required.

83

Info. The partner controller is changing state (shutting down or restarting).

Recommended actions

- No action is required.

84

Warning The current controller that logged this event forced the partner controller to fail over.

Recommended actions

- Download the debug logs from your storage system and contact technical support. A service technician can use the debug logs to determine the problem.

86

Info. Host-port or disk-channel parameters have been changed.

Recommended actions

- No action is required.

87

Warning The mirrored configuration retrieved by this controller from the partner controller has a bad cyclic redundancy check (CRC). The local flash configuration will be used instead.

Recommended actions

- Restore the default configuration by using the `restore defaults` command, as described in the CLI Reference Guide.

88

Warning The mirrored configuration retrieved by this controller from the partner controller is corrupt. The local flash configuration will be used instead.

Recommended actions

- Restore the default configuration by using the `restore defaults` command, as described in the CLI Reference Guide.

89

Warning The mirrored configuration retrieved by this controller from the partner controller has a configuration level that is too high for the firmware in this controller to process. The local flash configuration will be used instead.

Recommended actions

- The current controller that logged this event probably has down-level firmware. Update the firmware in the down-level controller. Both controllers should have the same firmware versions.

When the problem is resolved, event 20 is logged.

90

Info. The partner controller does not have a mirrored configuration image for the current controller, so the current controller's local flash configuration is being used.

This event is expected if the other controller is new or its configuration has been changed.

Recommended actions

- No action is required.

91

Error In a testing environment, the diagnostic that checks hardware reset signals between controllers in Active-Active mode failed.

Recommended actions

- Perform failure analysis.

95

Error Both controllers in an Active-Active configuration have the same serial number. Non-unique serial numbers can cause system problems. For example, WWNs are determined by serial number.

Recommended actions

- Remove one of the controller modules and insert a replacement, then return the removed module to be reprogrammed.

96

Info. Pending configuration changes that take effect at startup were ignored because customer data might be present in cache.

Recommended actions

- If the requested configuration changes did not occur, make the changes again and then use a user-interface command to shut down the Storage Controller and then restart it.

103

Info. The name has been changed for the indicated volume.

Recommended actions

- No action is required.

104

Info. The size has been changed for the indicated volume.

Recommended actions

- No action is required.

105

Info. The default LUN (logical unit number) has been changed for the indicated volume.

Recommended actions

- No action is required.

106

Info. The indicated volume has been added to the indicated pool.

Recommended actions

- No action is required.

107

Error A serious error has been detected by the controller. In a single-controller configuration, the controller will restart automatically. In an Active-Active configuration, the partner controller will kill the controller that experienced the error.

Recommended actions

- Download the debug logs from your storage system and contact technical support. A service technician can use the debug logs to determine the problem.

108

Info. The indicated volume has been deleted from the indicated pool.

Recommended actions

- No action is required.

109

Info. The statistics for the indicated volume have been reset.

Recommended actions

- No action is required.

110

Info. Ownership of the indicated disk group has been given to the other controller.

Recommended actions

- No action is required.

111

Info. The link for the indicated host port is up.

This event indicates that a problem reported by event 112 is resolved. For a system with FC ports, this event also appears after loop initialization.

Recommended actions

- No action is required.

112

Warning The link for the indicated host port has unexpectedly gone down.

Recommended actions

- Look for corresponding event 111 and monitor excessive transitions indicating a host-connectivity or switch problem. If this event occurs more than 8 times per hour, it should be investigated.
- This event is probably caused by equipment outside of the storage system, such as faulty cabling or a faulty switch.
- If the problem is not outside of the storage system, replace the controller module that logged this event.

Info. The link for the indicated host port has gone down because the controller is starting up.

Recommended actions

- No action is required.

114

Info. The link for the indicated disk-channel port is down. Note that events 114 and 211 are logged whenever a user-requested rescan occurs and do not indicate an error.

Recommended actions

- Look for corresponding event 211 and monitor excessive transitions indicating disk problems. If more than 8 transitions occur per hour, see [Resources for diagnosing and resolving problems](#) on page 6.

116

Error After a recovery, the partner controller was killed while mirroring write-back cache data to the controller that logged this event. The controller that logged this event restarted to avoid losing the data in the partner controller's cache, but if the other controller does not restart successfully, the data will be lost.

Recommended actions

- To determine if data might have been lost, check whether this event was immediately followed by event 56, closely followed by event 71. The failover indicates that the restart did not succeed.

117

Warning This controller module detected or generated an error on the indicated host channel.

Recommended actions

- Restart the Storage Controller that logged this event.
- If more errors are detected, check the connectivity between the controller and the attached host.
- If more errors are generated, shut down the Storage Controller and replace the controller module.

118

Info. Cache parameters have been changed for the indicated volume.

Recommended actions

- No action is required.

127

Warning The controller has detected an invalid disk dual-port connection. This event indicates that a controller host port is connected to an expansion port instead of to a port on a host or a switch.

Recommended actions

- Disconnect the host port and expansion port from each other and connect them to the proper devices.

136

Warning Errors detected on the indicated disk channel have caused the controller to mark the channel as degraded.

Recommended actions

- Determine the source of the errors on the indicated disk channel and replace the faulty hardware.

When the problem is resolved, event 189 is logged.

139

Info. The Management Controller (MC) has powered up or restarted.

Recommended actions

- No action is required.

140

Info. The Management Controller is about to restart.

Recommended actions

- No action is required.

141

Info. This event is logged when the IP address used for management of the system has been changed by a user or by a DHCP server (if DHCP is enabled). This event is also logged during power up or failover recovery, even when the address has not changed.

Recommended actions

- No action is required.

152

Warning The Management Controller (MC) has not communicated with the Storage Controller (SC) for 15 minutes and may have failed.

This event is initially logged as Informational severity. If the problem persists, this event is logged a second time as Warning severity and the MC is automatically restarted in an attempt to recover from the problem. Event 156 is then logged.

Recommended actions

- If this event is logged only one time as Warning severity, no action is required.
- If this event is logged more than one time as Warning severity, do the following:
 - If you are now able to access the management interfaces of the controller that logged this event, do the following:
 - Check the version of the controller firmware and update to the latest firmware if needed.
 - If the latest firmware is already installed, the controller module that logged this event probably has a hardware fault. Replace the module.
 - If you are **not** able to access the management interfaces of the controller that logged this event, do the following:
 - Shut down that controller and reseal the module.
 - If you are then able to access the management interfaces, check the version of the controller firmware and update to the latest firmware if needed.
 - If the problem recurs, replace the module.

Info. The Management Controller (MC) has not communicated with the Storage Controller (SC) for 160 seconds. If communication is restored in less than 15 minutes, event 153 is logged. If the problem persists, this event is logged a second time as Warning severity.

NOTE: It is normal for this event to be logged as Informational severity during firmware update.

Recommended actions

- Check the version of the controller firmware and update to the latest firmware if needed.
- If the latest firmware is already installed, no action is required.

153

Info. The Management Controller (MC) has re-established communication with the Storage Controller (SC).

Recommended actions

- No action is required.

154

Info. New firmware has been loaded in the Management Controller (MC).

Recommended actions

- No action is required.

155

Info. New loader firmware has been loaded in the Management Controller (MC).

Recommended actions

- No action is required.

156

Warning The Management Controller (MC) has been restarted from the Storage Controller (SC) for the purpose of error recovery.

Recommended actions

- See the recommended actions for event 152, which is logged at approximately the same time.

Info. The Management Controller (MC) has been restarted from the Storage Controller (SC) in a normal case, such as when initiated by a user.

Recommended actions

- No action is required.

157

Error A failure occurred when trying to write to the Storage Controller (SC) flash chip.

Recommended actions

- Replace the controller module that logged this event.

158

Error A correctable ECC error occurred in Storage Controller CPU memory more than once during a 12-hour period, indicating a probable hardware fault.

Recommended actions

- Replace the controller module that logged this event.

Warning A correctable ECC error occurred in Storage Controller CPU memory.

This event is logged with Warning severity to provide information that may be useful to technical support, but no action is required now. It will be logged with Error severity if it is necessary to replace the controller module.

Recommended actions

- No action is required.

161

Info. One or more enclosures do not have a valid path to an enclosure management processor (EMP).

All enclosure EMPs are disabled.

Recommended actions

- Download the debug logs from your storage system and contact technical support. A service technician can use the debug logs to determine the problem.

162

Warning The host WWNs (node and port) previously presented by this controller module are unknown. In a dual-controller system this event has two possible causes:

- One or both controller modules have been replaced or moved while the system was powered off.
- One or both controller modules have had their flash configuration cleared (this is where the previously used WWNs are stored).

The controller module recovers from this situation by generating a WWN based on its own serial number.

Recommended actions

- If the controller module was replaced or someone reprogrammed its FRU ID data, verify the WWN information for this controller module on all hosts that access it.

163

Warning The host WWNs (node and port) previously presented by the partner controller module, which is currently offline, are unknown.

This event has two possible causes:

- The online controller module reporting the event was replaced or moved while the system was powered off.
- The online controller module had its flash configuration (where previously used WWNs are stored) cleared.

The online controller module recovers from this situation by generating a WWN based on its own serial number for the other controller module.

Recommended actions

- If the controller module was replaced or someone reprogrammed its FRU ID data, verify the WWN information for the other controller module on all hosts that access it.

166

Warning The RAID metadata level of the two controllers does not match, which indicates that the controllers have different firmware levels.

Usually, the controller at the higher firmware level can read metadata written by a controller at a lower firmware level. The reverse is typically not true. Therefore, if the controller at the higher firmware level failed, the surviving controller at the lower firmware level cannot read the metadata in disks that have failed over.

Recommended actions

- If this occurs after a firmware update, it indicates that the metadata format changed, which is rare. Update the controller with the lower firmware level to match the firmware level in the other controller.

167

Warning A diagnostic test at controller bootup detected an abnormal operation, which might require a power cycle to correct.

Recommended actions

- Download the debug logs from your storage system and contact technical support. A service technician can use the debug logs to determine the problem.

168

Error The indicated SES alert condition was detected in the indicated enclosure. This event is logged as Error severity when one of the power supplies in an enclosure has no power supplied to it or when a hardware failure is detected.

Recommended actions

For a 2U12 or 2U24 enclosure:

- Check that all modules in the enclosure are fully seated in their slots and that their latches are locked.
- If the reported problem is with a power supply, perform these checks:
 - Check that each power supply module has its switch turned on (if equipped with a switch).
 - Check that each power cable is firmly plugged into both the power supply and a functional electrical outlet.
- If the reported problem is with a temperature sensor or fan or power supply, perform these checks:
 - Check that all of the enclosure's fans are running.
 - Check that the ambient temperature is not too warm. The enclosure operating range is 5°–40°C (41°–104°F).
 - Check for any obstructions to the airflow.
 - Check that there is a module or blank plate in every module slot in the enclosure.
- If none of the above resolve the issue, the indicated FRU has probably failed and should be replaced. The failed FRU will probably have an amber LED lit.

When the problem is resolved, event 169 is logged.

For a 2U48 enclosure:

- Check that all modules in the enclosure are fully seated in their slots and that their latches are locked.
- If the reported problem is with a power supply, perform these checks:
 - Check that each power supply module has its switch turned on (if equipped with a switch).
 - Check that each power cable is firmly plugged into both the power supply and a functional electrical outlet.
- If the reported problem is with a temperature sensor or fan or power supply, perform these checks:
 - Check that all of the enclosure's fans are running.
 - Check that the ambient temperature is not too warm. The enclosure operating range is 5°–35°C (41°–95°F).

- Check for any obstructions to the airflow.
- Check that the drawers are closed and there is a module or blank plate in every module slot in the enclosure.
- If the reported problem is a processor fault, schedule preventive maintenance. Contact customer support for more details.
- If none of the above resolve the issue, the indicated FRU has probably failed and should be replaced. The failed FRU will probably have an amber LED lit.

When the problem is resolved, event 169 is logged.

Warning The indicated SES alert condition was detected in the indicated enclosure.

Recommended actions

For a 2U12 or 2U24 enclosure:

- Check that all modules in the enclosure are fully seated in their slots and that their latches are locked.
- If the reported problem is with a power supply, perform these checks:
 - Check that each power supply module has its switch turned on (if equipped with a switch).
 - Check that each power cable is firmly plugged into both the power supply and a functional electrical outlet.
- If the reported problem is with a temperature sensor or fan or power supply, perform these checks:
 - Check that all of the enclosure's fans are running.
 - Check that the ambient temperature is not too warm. The enclosure operating range is 5°–40°C (41°–104°F).
 - Check for any obstructions to the airflow.
 - Check that there is a module or blank plate in every module slot in the enclosure.
- If none of the above resolve the issue, the indicated FRU has probably failed and should be replaced. The failed FRU will probably have an amber LED lit.

When the problem is resolved, event 169 is logged.

For a 2U48 enclosure:

- Check that all modules in the enclosure are fully seated in their slots and that their latches are locked.
- If the reported problem is with a power supply, perform these checks:
 - Check that each power supply module has its switch turned on (if equipped with a switch).
 - Check that each power cable is firmly plugged into both the power supply and a functional electrical outlet.
- If the reported problem is with a temperature sensor or fan or power supply, perform these checks:
 - Check that all of the enclosure's fans are running.
 - Check that the ambient temperature is not too warm. The enclosure operating range is 5°–35°C (41°–95°F).
 - Check for any obstructions to the airflow.
 - Check that the drawers are closed and there is a module or blank plate in every module slot in the enclosure.
 - If the reported problem is a processor fault, schedule preventive maintenance. Contact customer support for more details.
- If none of the above resolve the issue, the indicated FRU has probably failed and should be replaced. The failed FRU will probably have an amber LED lit.

When the problem is resolved, event 169 is logged.

Info. The indicated SES alert condition was detected in the indicated enclosure.

Recommended actions

- No action is required.

169

Info. The indicated SES alert condition has been cleared in the indicated enclosure. This event indicates that a problem reported by event 168 is resolved.

Recommended actions

- No action is required.

170

Info. The last rescan detected that the indicated enclosure was added to the system.

Recommended actions

- No action is required.

171

Info. The last rescan detected that the indicated enclosure was removed from the system.

Recommended actions

- No action is required.

172

Warning The indicated disk group has been quarantined because not all of its disks are accessible. While the disk group is quarantined, any attempt to access the volumes in the disk group from a host will fail. If all of the disks become accessible, the disk group will be dequarantined automatically with a resulting status of FTOL (fault tolerant and online). If not all of the disks become accessible but enough become accessible to allow reading from and writing to the disk group, the disk group will be dequarantined automatically with a resulting status of FTDN (fault tolerant with a down disk) or CRIT (critical). If a spare disk is available, reconstruction will begin automatically. When the disk group has been removed from quarantine, event 173 is logged. For a more detailed discussion of quarantine, see the WBI help for the Tools > Dequarantine Vdisk panel (linear only) or the CLI help for the `dequarantine` command.

△ CAUTION:

- Avoid using the manual dequarantine operation as a recovery method when event 172 is logged because this causes data recovery to be more difficult or impossible.
 - If you clear unwritten cache data while a disk group is quarantined or offline, that data will be permanently lost.
-

Recommended actions

- If event 173 has subsequently been logged for the indicated disk group, no action is required. The disk group has already been removed from quarantine.
- Otherwise, perform the following actions:
 - Check that all enclosures are powered on.
 - Check that all disks and I/O modules in every enclosure are fully seated in their slots and that their latches are locked.
 - Reseat any disks in the quarantined disk group that are reported as missing or failed in the user interface. (Do NOT remove and reinsert disks that are not members of the disk group that is quarantined.)
 - Check that the SAS expansion cables are connected between each enclosure in the storage system and that they are fully seated. (Do NOT remove and reinsert the cables because this can cause problems with additional disk groups.)
 - Check that no disks have been removed from the system unintentionally.
 - Check for other events that indicate faults in the system and follow the recommended actions for those events. But, if the event indicates a failed disk and the recommended action is to replace the disk, do NOT replace the disk at this time because it may be needed later for data recovery.

- If the disk group is still quarantined after performing the above steps, shut down both controllers and then power down the entire storage system. Power it back up, beginning with any disk enclosures (expansion enclosures), then the controller enclosure.
- If the disk group is still quarantined after performing the above steps, contact technical support.

173

Info. The indicated disk group has been removed from quarantine.

Recommended actions

- No action is required.

174

Info. Enclosure or disk firmware update has succeeded, been aborted by a user, or failed.

If the firmware update fails, the user will be notified about the problem immediately and should take care of the problem at that time, so even when there is a failure, this event is logged as Informational severity.

Recommended actions

- No action is required.

175

Info. The network-port Ethernet link has changed status (up or down) for the indicated controller.

Recommended actions

- If this event is logged indicating the network port is up shortly after the Management Controller (MC) has booted up (event 139), no action is required.
- Otherwise, monitor occurrences of this event for an error trend. If this event occurs more than 8 times per hour, it should be investigated.
 - This event is probably caused by equipment outside of the storage system, such as faulty cabling or a faulty Ethernet switch.
 - If this event is being logged by only one controller in a dual-controller system, swap the Ethernet cables between the two controllers. This will show whether the problem is outside or inside the storage system.
 - If the problem is not outside of the storage system, replace the controller module that logged this event.

176

Info. The error statistics for the indicated disk have been reset.

Recommended actions

- No action is required.

177

Info. Cache data was purged for the indicated missing volume.

Recommended actions

- No action is required.

181

Info. One or more configuration parameters associated with the Management Controller (MC) have been changed, such as configuration for SNMP, SMI-S, email notification, and system strings (system name, system location, etc.).

Recommended actions

- No action is required.

182

Info. All disk channels have been paused. I/O will not be performed on the disks until all channels are unpaused.

Recommended actions

- If this event occurs in relation to disk firmware update, no action is required. When the condition is cleared, event 183 is logged.
- If this event occurs and you are not performing disk firmware update, see [Resources for diagnosing and resolving problems](#) on page 6.

183

Info. All disk channels have been unpaused, meaning that I/O can resume. An unpaused initiates a rescan, which when complete is logged as event 19.

This event indicates that the pause reported by event 182 has ended.

Recommended actions

- No action is required.

185

Info. An enclosure management processor (EMP) write command has completed.

Recommended actions

- No action is required.

186

Info. Enclosure parameters have been changed by a user.

Recommended actions

- No action is required.

187

Info. The write-back cache has been enabled.

Event 188 is the corresponding event that is logged when write-back cache is disabled.

Recommended actions

- No action is required.

188

Info. Write-back cache has been disabled.

Event 187 is the corresponding event that is logged when write-back cache is disabled.

Recommended actions

- No action is required.

189

Info. A disk channel that was previously degraded or failed is now healthy.

Recommended actions

- No action is required.

190

Info. The controller module's supercapacitor pack has started charging.

This change met a condition to trigger the auto-write-through feature, which has disabled write-back cache and put the system in write-through mode. When the fault is resolved, event 191 is logged to indicate that write-back mode has been restored.

Recommended actions

- If event 191 is not logged within 5 minutes after this event, the supercapacitor has probably failed and the controller module should be replaced.

191

Info. The auto-write-through trigger event that caused event 190 to be logged has been resolved.

Recommended actions

- No action is required.

192

Info. The controller module's temperature has exceeded the normal operating range.

This change met a condition to trigger the auto-write-through feature, which has disabled write-back cache and put the system in write-through mode. When the fault is resolved, event 193 is logged to indicate that write-back mode has been restored.

Recommended actions

- If event 193 has not been logged since this event was logged, the over-temperature condition probably still exists and should be investigated. Another over-temperature event was probably logged at approximately the same time as this event (such as event 39, 40, 168, 307, 469, 476, or 477). See the recommended actions for that event.

193

Info. The auto-write-through trigger event that caused event 192 to be logged has been resolved.

Recommended actions

- No action is required.

194

Info. The Storage Controller in the partner controller module is not up.

This indicates that a trigger condition has occurred that has caused the auto-write-through feature to disable write-back cache and put the system in write-through mode. When the fault is resolved, event 195 is logged to indicate that write-back mode has been restored.

Recommended actions

- If event 195 has not been logged since this event was logged, the other Storage Controller is probably still down and the cause should be investigated. Other events were probably logged at approximately the same time as this event. See the recommended actions for those events.

195

Info. The auto-write-through trigger event that caused event 194 to be logged has been resolved.

Recommended actions

- No action is required.

198

Info. A power supply has failed.

This indicates that a trigger condition has occurred that has caused the auto-write-through feature to disable write-back cache and put the system in write-through mode. When the fault is resolved, event 199 is logged to indicate that write-back mode has been restored.

Recommended actions

- If event 199 has not been logged since this event was logged, the power supply probably does not have a health of OK and the cause should be investigated. Another power-supply event was probably logged at approximately the same time as this event (such as event 168). See the recommended actions for that event.

199

Info. The auto-write-through trigger event that caused event 198 to be logged has been resolved.

Recommended actions

- No action is required.

200

Info. A fan has failed.

This indicates that a trigger condition has occurred that has caused the auto-write-through feature to disable write-back cache and put the system in write-through mode. When the fault is resolved, event 201 is logged to indicate that write-back mode has been restored.

Recommended actions

- If event 201 has not been logged since this event was logged, the fan probably does not have a health of OK and the cause should be investigated. Another fan event was probably logged at approximately the same time as this event (such as event 168). See the recommended actions for that event.

201

Info. The auto-write-through trigger event that caused event 200 to be logged has been resolved.

Recommended actions

- No action is required.

202

Info. An auto-write-through trigger condition has been cleared, causing write-back cache to be re-enabled. The environmental change is also logged at approximately the same time as this event (event 191, 193, 195, 199, 201, and 241.)

Recommended actions

- No action is required.

203

Warning An environmental change occurred that allows write-back cache to be enabled, but the auto-write-back preference is not set. The environmental change is also logged at approximately the same time as this event (event 191, 193, 195, 199, 201, or 241).

Recommended actions

- Manually enable write-back cache.

204

Error An error occurred with either the CompactFlash card (NV device) or the transport mechanism. The system may attempt to recover itself.

The CompactFlash card is used for backing up unwritten cache data when a controller goes down unexpectedly, such as when a power failure occurs. This event is generated when the Storage Controller (SC) detects a problem with the CompactFlash as it is booting up.

Recommended actions

- Restart the Storage Controller that logged this event.
- If this event is logged again, shut down the Storage Controller and replace the controller module.

Warning The system has started and found an issue with the CompactFlash card (NV device). The system will attempt to recover itself.

The CompactFlash card is used for backing up unwritten cache data when a controller goes down unexpectedly, such as when a power failure occurs. This event is generated when the Storage Controller (SC) detects a problem with the CompactFlash as it is booting up.

Recommended actions

- Restart the Storage Controller that logged this event.
- If this event is logged again, shut down the Storage Controller and replace the controller module.

Info. The system has come up normally and the CompactFlash card (NV device) is in a normal expected state. This event will be logged as an Error or Warning event if any user action is required.

Recommended actions

- No action is required.

205

Info. The indicated volume has been mapped or unmapped.

Recommended actions

- No action is required.

206

Info. Disk group scrub has started.

The scrub checks disks in the disk group for the following types of errors:

- Data parity errors for a RAID 3, 5, 6, or 50 disk group.
- Mirror verify errors for a RAID 1 or RAID 10 disk group.
- Media errors for all RAID levels including RAID 0 and non-RAID disk groups.

When errors are detected, they are automatically corrected.

When the scrub is complete, event 207 is logged.

Recommended actions

- No action is required.

207

Error Disk group scrub completed and found an excessive number of errors in the indicated disk group.

This event is logged as Error severity when more than 100 parity or mirror mismatches are found and corrected during a scrub or when 1 to 99 parity or mirror mismatches are found and corrected during each of 10 separate scrubs of the same disk group.

For non-fault-tolerant RAID levels (RAID 0 and non-RAID), media errors may indicate loss of data.

Recommended actions

- Resolve any non-disk hardware problems, such as a cooling problem or a faulty controller module, expansion module, or power supply.
- Check whether any disks in the disk group have logged SMART events or unrecoverable read errors.
 - If so, and the disk group is a non-fault-tolerant RAID level (RAID 0 or non-RAID), copy the data to a different disk group and replace the faulty disks.
 - If so, and the disk group is a fault-tolerant RAID level, replace the faulty disks. Before replacing a disk, confirm that a reconstruction is not currently running on the disk group. It is also recommended to make a full backup of all the data on the disk group before replacing disks. If more than one disk in the disk group has errors, replace the disks one at a time and allow reconstruction to complete after each disk is replaced.

Warning Disk group scrub did not complete because of an internally detected condition such as a failed disk.

If a disk fails, data may be at risk.

Recommended actions

- Resolve any non-disk hardware problems, such as a cooling problem or a faulty controller module, expansion module, or power supply.
- Check whether any disks in the disk group have logged SMART events or unrecoverable read errors.
 - If so, and the disk group is a non-fault-tolerant RAID level (RAID 0 or non-RAID), copy the data to a different disk group and replace the faulty disks.
 - If so, and the disk group is a fault-tolerant RAID level, replace the faulty disks. Before replacing a disk, confirm that a reconstruction is not currently running on the disk group. It is also recommended to make a full backup of all the data in the disk group before replacing disks. If more than one disk in the disk group has errors, replace the disks one at a time and allow reconstruction to complete after each disk is replaced.

Info. Disk group scrub completed or was aborted by a user.

This event is logged as Informational severity when fewer than 100 parity or mirror mismatches are found and corrected during a scrub.

For non-fault-tolerant RAID levels (RAID 0 and non-RAID), media errors may indicate loss of data.

Recommended actions

- No action is required.

208

Info. A scrub-disk job has started for the indicated disk. The result will be logged with event 209.

Recommended actions

- No action is required.

209

Error A scrub-disk job logged with event 208 has completed and found one or more media errors, SMART events, or hard (non-media) errors. If this disk is used in a non-fault-tolerant disk group, data may have been lost.

Recommended actions

- Replace the disk with one of the same type (SAS SSD, enterprise SAS, or midline SAS) and the same or greater capacity. For continued optimum I/O performance, the replacement disk should have performance that is the same as or better than the one it is replacing.

Warning A scrub-disk job logged with event 208 has been aborted by a user, or has reassigned a disk block. These bad-block replacements are reported as "other errors". If this disk is used in a non-fault-tolerant disk group, data may have been lost.

Recommended actions

- Monitor the error trend and whether the number of errors approaches the total number of bad-block replacements available.

Info. A scrub-disk job logged with event 208 has completed and found no errors, or a disk being scrubbed (with no errors found) has been added to a disk group, or a user has aborted the job.

Recommended actions

- No action is required.

210

Info. All snapshots have been deleted for the indicated parent volume (or snap pool for linear volumes only).

Recommended actions

- No action is required.

211

Warning SAS topology has changed. No elements are detected in the SAS map. The message specifies the number of elements in the SAS map, the number of expanders detected, the number of expansion levels on the native (local controller) side and on the partner (partner controller) side, and the number of device PHYs.

Recommended actions

- Perform a rescan to repopulate the SAS map.
- If a rescan does not resolve the problem, then shut down and restart both Storage Controllers.
- If the problem persists, see [Resources for diagnosing and resolving problems](#) on page 6.

Info. SAS topology has changed. The number of SAS expanders has increased or decreased. The message specifies the number of elements in the SAS map, the number of expanders detected, the number of expansion levels on the native (local controller) side and on the partner (partner controller) side, and the number of device PHYs.

Recommended actions

- No action is required.

212

Info. All master volumes associated with the indicated snap pool have been deleted.

Recommended actions

- No action is required.

213

Info. The indicated standard volume has been converted to a master volume, or the indicated master volume has been converted to a standard volume.

Recommended actions

- No action is required.

214

Info. The creation of snapshots is complete. The number of snapshots is indicated.

Additional events give more information for each snapshot.

Recommended actions

- No action is required.

215

Info. A previously created batch of snapshots is now committed and ready for use. The number of snapshots is indicated.

Additional events give more information for each snapshot.

Recommended actions

- No action is required.

216

Info. An uncommitted snapshot has been deleted. Removal of the indicated snapshot completed successfully.

Recommended actions

- No action is required.

217

Error A supercapacitor failure occurred in the controller.

Recommended actions

- Replace the controller module that logged this event.

218

Warning The supercapacitor pack is near end of life.

Recommended actions

- Replace the controller module reporting this event.

219

Info. Utility priority has been changed by a user.

Recommended actions

- No action is required.

220

Info. Roll back of data in the indicated master volume to data in the indicated snapshot has been started by a user.

Recommended actions

- No action is required.

221

Info. Snapshot reset has completed.

Recommended actions

- No action is required.

222

Info. The policy for the snap pool has been changed by a user. A policy specifies the action for the system to automatically take when the snap pool reaches the associated threshold level.

Recommended actions

- No action is required.

223

Info. The threshold level for the snap pool has been changed by a user. Each snap pool has three threshold levels that notify you when the snap pool is reaching decreasing capacity. Each threshold level has an associated policy that specifies system behavior when the threshold is reached.

Recommended actions

- No action is required.

224

Info. Roll back of data in the indicated master volume to data in the indicated snapshot has completed.

Recommended actions

- No action is required.

225

Error A copy-on-write failure occurred when copying data from the indicated master volume to a snapshot.

Due to a problem accessing the snap pool, the write operation could not be completed to the disk. Data are left in cache.

Recommended actions

- Delete all snapshots for the master volume and then convert the master volume to a standard volume.

226

Error Roll back for the indicated master volume failed to start due to inability to initialize the snap pool.

The roll back is in a suspended state.

Recommended actions

- Make sure the snap pool and the pool on which this volume exists are online. Restart the roll-back operation.

227

Error Failed to execute roll back for a particular LBA (logical block address) range of the indicated parent volume.

Recommended actions

- Restart the roll-back operation.

228

Error Roll back for the indicated master volume failed to end due to inability to initialize the snap pool.

The roll back is in a suspended state.

Recommended actions

- Make sure the snap pool and the pool on which this volume exists are online. Restart the roll-back operation.

229

Warning The indicated snap pool has reached its warning threshold.

Recommended actions

- You can expand the snap pool or delete snapshots.

230

Warning The indicated snap pool has reached its error threshold.

When the error threshold is reached, the system automatically takes the action set in the policy for this threshold level. The default policy for the error threshold is to auto-expand the snap pool.

Recommended actions

- You can expand the snap pool or delete snapshots.

231

Warning The indicated snap pool has reached its critical threshold.

When the critical threshold is reached, the system automatically takes the action set in the policy for this threshold level. The default policy for the critical threshold is to delete all snapshots in the snap pool.

Recommended actions

- If the policy is to halt writes, then you must free up space in the snap pool by deleting snapshots.
- For other policies, no action is required.

232

Warning The maximum number of enclosures allowed for the current configuration has been exceeded.

The platform does not support the number of enclosures that are configured. The enclosure indicated by this event has been removed from the configuration.

Recommended actions

- Reconfigure the system.

233

Warning The indicated disk type is invalid and is not allowed in the current configuration.

All disks of the disallowed type have been removed from the configuration.

Recommended actions

- Replace the disallowed disks with ones that are supported.

234

Error The indicated snap pool is unrecoverable and can therefore no longer be used.

Recommended actions

- All the snapshots associated with this snap pool are invalid and you may want to delete them. However, the data in the master volume can be recovered by converting it to a standard volume.

235

Error An enclosure management processor (EMP) detected a serious error.

Recommended actions

- Replace the indicated controller module or expansion module.

Info. An enclosure management processor (EMP) reported an event.

Recommended actions

- No action is required.

236

Error A special shutdown operation has started. These special shutdown types indicate an incompatible feature.

Recommended actions

- Replace the indicated controller module with one that supports the indicated feature.

Info. A special shutdown operation has started. These special shutdown types are used as part of the firmware-update process.

Recommended actions

- No action is required.

237

Error A firmware update attempt was aborted because of either general system health issue(s), or unwritable cache data that would be lost during a firmware update.

Recommended actions

- Resolve before retrying a firmware update. For health issues, issue the `show system` CLI command to determine the specific health issue(s). For unwritten cache data, use the `show unwritable-cache` CLI command.

Info. A firmware update has started and is in progress. This event provides details of the steps in a firmware-update operation that may be of interest if you have problems updating firmware.

Recommended actions

- No action is required.

238

Warning An attempt to install a licensed feature failed due to an invalid license.

Recommended actions

- Check the license for what is allowed for the platform, make corrections as appropriate, and reinstall.

239

Warning A timeout occurred while flushing the CompactFlash.

Recommended actions

- Restart the Storage Controller that logged this event.
- If this event is logged again, shut down the Storage Controller and replace the controller module.

240

Warning A failure occurred while flushing the CompactFlash.

Recommended actions

- Restart the Storage Controller that logged this event.
- If this event is logged again, shut down the Storage Controller and replace the controller module.

241

Info. The auto-write-through trigger event that caused event 242 to be logged has been resolved.

Recommended actions

- No action is required.

242

Error The controller module's CompactFlash card has failed.

This change met a condition to trigger the auto-write-through feature, which has disabled write-back cache and put the system in write-through mode. When the fault is resolved, event 241 is logged to indicate that write-back mode has been restored.

Recommended actions

- If event 241 has not been logged since this event was logged, the CompactFlash probably does not have health of OK and the cause should be investigated. Another CompactFlash event was probably logged at approximately the same time as this event (such as event 239, 240, or 481). See the recommended actions for that event.

243

Info. A new controller enclosure has been detected. This happens when a controller module is moved from one enclosure to another and the controller detects that the midplane WWN is different from the WWN it has in its local flash.

Recommended actions

- No action is required.

245

Info. An existing disk channel target device is not responding to SCSI discovery commands.

Recommended actions

- Check the indicated target device for bad hardware or bad cable, then initiate a rescan.

246

Warning The coin battery is not present, is not properly seated, or has reached end-of-life.

The battery provides backup power for the real-time (date/time) clock. In the event of a power failure, the date and time will revert to 1980-01-01 00:00:00.

Recommended actions

- Replace the controller module that logged this event.

247

Warning The FRU ID EEPROM for the indicated field replaceable unit (FRU) cannot be read. FRU ID data might not be programmed.

FRU ID data includes the worldwide name, serial numbers, firmware and hardware versions, branding information, etc. This event is logged once each time a Storage Controller (SC) is started for each FRU that is not programmed.

Recommended actions

- Return the FRU to have its FRU ID data reprogrammed.

248

Info. A valid feature license was successfully installed. See event 249 for details about each licensed feature.

Recommended actions

- No action is required.

249

Info. After a valid license is installed, this event is logged for each licensed feature to show the new license value for that feature. The event specifies whether the feature is licensed, whether the license is temporary, and whether the temporary license is expired.

Recommended actions

- No action is required.

250

Warning A license could not be installed.

The license is invalid or specifies a feature that is not supported on your product.

Recommended actions

- Review the readme file that came with the license. Verify that you are trying to install the license in the system that the license was generated for.

251

Info. A volume-copy operation has started for the indicated source volume.

If the source volume is a master volume, you can remount it.

If the source volume is a snapshot, do not remount it until the copy is complete (as indicated by event 268).

Recommended actions

- No action is required.

252

Info. Data written to the indicated snapshot after it was created has been deleted. The snapshot now represents the state of the parent volume when the snapshot was created.

Recommended actions

- No action is required.

253

Info. A license was uninstalled.

Recommended actions

- No action is required.

255

Info. The PBCs across controllers do not match as PBC from controller A and PBC from controller B are from different vendors. This may limit the available configurations.

Recommended actions

- No action is required.

256

Info. The indicated snapshot has been prepared but is not yet committed.

This can occur when a snapshot is taken by an application, such as the VSS hardware provider, that is timing-sensitive and needs to take a snapshot in two stages.

After the snapshot is committed and event 258 is logged, the snapshot can be used.

Recommended actions

- No action is required.

- 257**
Info. The indicated snapshot has been prepared and committed and is ready for use.
Recommended actions
- No action is required.
- 258**
Info. The indicated snapshot has been committed and is ready for use.
Recommended actions
- No action is required.
- 259**
Info. In-band CAPI commands have been disabled.
Recommended actions
- No action is required.
- 260**
Info. In-band CAPI commands have been enabled.
Recommended actions
- No action is required.
- 261**
Info. In-band SES commands have been disabled.
Recommended actions
- No action is required.
- 262**
Info. In-band SES commands have been enabled.
Recommended actions
- No action is required.
- 263**
Warning The indicated spare disk is missing. Either it was removed or it is not responding.
Recommended actions
- Replace the disk with one of the same type (SAS SSD, enterprise SAS, or midline SAS) and the same or greater capacity.
 - Configure the disk as a spare.
- 266**
Info. A volume-copy operation for the indicated master volume has been aborted by a user.
Recommended actions
- No action is required.

267

Error While cleaning up resources in metadata at the end of a volume-copy operation, the firmware found at least one error for the indicated volume.

Recommended actions

- Make sure that the disk groups and disks associated with the volume copy do not have problems (health OK, status FTOL or UP) and then retry the volume copy.

268

Info. A volume-copy operation for the indicated volume has completed.

Recommended actions

- No action is required.

269

Error A partner firmware upgrade attempt aborted because of either general system health issue(s) or unwritable cache data that would be lost during a firmware update.

Recommended actions

- Resolve before retrying a firmware update. For health issues, issue the `show system` CLI command to determine the specific health issue(s). For unwritten cache data, use the `show unwritable-cache` CLI command.

Info. A partner firmware update operation has started. This operation is used to copy firmware from one controller to the other to bring both controllers up to the same version of firmware.

Recommended actions

- No action is required.

270

Warning Either there was a problem reading or writing the persistent IP data from the FRU ID SEEPROM, or invalid data were read from the FRU ID SEEPROM.

Recommended actions

- Check the IP settings (including iSCSI host-port IP settings for an iSCSI system), and update them if they are incorrect.

271

Info. The storage system could not get a valid serial number from the controller's FRU ID SEEPROM, either because it couldn't read the FRU ID data, or because the data in it are not valid or have not been programmed. Therefore, the MAC address is derived by using the controller's serial number from flash. This event is only logged one time during bootup.

Recommended actions

- No action is required.

272

Info. Expansion of the indicated snap pool has started.

Recommended actions

- No action is required.

273

Info. PHY fault isolation has been enabled or disabled by a user for the indicated enclosure and controller module.

Recommended actions

- No action is required.

274

Warning The indicated PHY has been disabled, either automatically or by a user. Drive PHYs are automatically disabled for empty disk slots or if a problem is detected. The following reasons indicate a likely hardware fault:

- Disabled because of error count interrupts
- Disabled because of excessive PHY change counts
- PHY is ready but did not pass COMINIT

Recommended actions

- If none of the reasons listed above, no action is required.
- If any of the reasons listed above is indicated and the event occurs shortly after the storage system is powered up, do the following:
 - Shut down the Storage Controllers. Then turn off the power for the indicated enclosure wait a few seconds, and turn it back on.
 - If the problem recurs and the event message identifies a disk slot, replace the disk in that slot.
 - If the problem recurs and the event message identifies a module, do the following:
 - If the indicated PHY type is Egress, replace the cable in the module's egress port.
 - If the indicated PHY type is Ingress, replace the cable in the module's ingress port.
 - For other indicated PHY types or if replacing the cable does not fix the problem, replace the indicated module.
 - If the problem persists, check for other events that may indicate faulty hardware, such as an event indicating an over-temperature condition or power supply fault, and follow the recommended actions for those events.
 - If the problem still persists, the fault may be in the enclosure midplane. Replace the chassis-and-midplane FRU.
- If any of the reasons listed above is indicated and this event is logged shortly after a failover, user-initiated rescan, or restart, do the following:
 - If the event message identifies a disk slot, reseal the disk in that slot.
 - If the problem persists after reseating the disk, replace the disk.
 - If the event message identifies a module, do the following:
 - If the indicated PHY type is Egress, replace the cable in the module's egress port.
 - If the indicated PHY type is Ingress, replace the cable in the module's ingress port.
 - For other indicated PHY types or if replacing the cable does not fix the problem, replace the indicated module.
 - If the problem persists, check for other events that may indicate faulty hardware, such as an event indicating an over-temperature condition or power supply fault, and follow the recommended actions for those events.
 - If the problem still persists, the fault may be in the enclosure midplane. Replace the chassis-and-midplane FRU.

275

Info. The indicated PHY has been enabled.

Recommended actions

- No action is required.

298

Warning The controller's real-time clock (RTC) setting is invalid.

This event will most commonly occur after a power loss if the real-time clock battery has failed. The time may have been set to a time that is up to 5 minutes before the power loss occurred, or it may have been reset to 1980-01-01 00:00:00.

Recommended actions

- Check the system date and time. If either is incorrect, set them to the correct date and time.
- Also look for event 246 and follow the recommended action for that event.

When the problem is resolved, event 299 is logged.

299

Info. The controller's RTC setting was successfully recovered.

This event will most commonly occur after an unexpected power loss.

Recommended actions

- No action is required, but if event 246 is also logged, follow the recommended action for that event.

300

Info. CPU frequency has changed to high.

Recommended actions

- No action is required.

301

Info. CPU frequency has changed to low.

Recommended actions

- No action is required.

302

Info. DDR memory clock frequency has changed to high.

Recommended actions

- No action is required.

303

Info. DDR memory clock frequency has changed to low.

Recommended actions

- No action is required.

304

Info. The controller has detected I²C errors that may have been fully recovered.

Recommended actions

- No action is required.

305

Info. A serial number in Storage Controller (SC) flash memory was found to be invalid when compared to the serial number in the controller-module or midplane FRU ID EEPROM. The valid serial number has been recovered automatically.

Recommended actions

- No action is required.

306

Info. The controller-module serial number in Storage Controller (SC) flash memory was found to be invalid when compared to the serial number in the controller-module FRU ID EEPROM. The valid serial number has been recovered automatically.

Recommended actions

- No action is required.

307

Critical A temperature sensor on a controller FRU detected an over-temperature condition that caused the controller to shut down.

Recommended actions

For a 2U12 or 2U24 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5° C–40° C (41° F–104° F).
- Check for any obstructions to the airflow.
- Check that there is a module or blank plate in every module slot in the enclosure.

If none of the above explanations apply, replace the controller module that logged the error.

For a 2U48 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5° C–35° C (41° F–95° F).
- Check for any obstructions to the airflow.
- Check that the drawers are closed and there is a module or blank plate in every module slot in the enclosure.

If none of the above explanations apply, replace the controller module that logged the error.

309

Info. Normally when the Management Controller (MC) is started, the IP data is obtained from the midplane FRU ID EEPROM where it is persisted. If the system is unable to write it to the EEPROM the last time it changed, a flag is set in flash memory. This flag is checked during startup, and if set, this event is logged and the IP data that is in flash memory is used. The only time that this would not be the correct IP data would be if the controller module was swapped and then whatever data are in the controller's flash memory is used.

Recommended actions

- No action is required.

310

Info. After a rescan, back-end discovery and initialization of data for at least one EMP (Enclosure Management Processor) has completed. This event is not logged again when processing completes for other EMPs in the system.

Recommended actions

- No action is required.

311

Info. This event is logged when a user initiates a ping of a host via the iSCSI interface.

Recommended actions

- If the ping operation failed, check connectivity between the storage system and the remote host.

312

Info. This event is used by email messages and SNMP traps when testing notification settings. This event is not recorded in the event log.

Recommended actions

- No action is required.

313

Error The indicated controller module has failed. This event can be ignored for a single-controller configuration.

Recommended actions

- If this is a dual-controller system, replace the failed controller module. The module's Fault/Service Required LED will be illuminated (not blinking).

314

Error The indicated FRU has failed or is not operating correctly. This event follows some other FRU-specific event indicating a problem.

Recommended actions

- To determine whether the FRU needs to be replaced, see the topic about verifying component failure in your product's FRU Installation and Replacement Guide.

315

Critical The controller module is incompatible with the enclosure.

The controller will automatically shut down. If two incompatible controllers are inserted at the same time or booted at the same time, one controller will crash and the other will hang. This behavior is expected and prevents data loss.

Recommended actions

- Move the controller module to a compatible enclosure.

316

Warning The temporary license for a feature has expired.

Any components created with the feature remain accessible but new components cannot be created.

Recommended actions

- To continue using the feature, purchase a permanent license.

Info. The temporary license for a feature will expire in 10 days. Any components created with the feature will remain accessible but new components cannot be created.

Recommended actions

- To continue using the feature after the trial period, purchase a permanent license.

317

Error A serious error has been detected on the Storage Controller's disk interface. The controller will be killed by its partner.

Recommended actions

- Visually trace the cabling between the controller modules and expansion modules.
- If the cabling is OK, replace the controller module that logged this event.
- If the problem recurs, replace the expansion module that is connected to the controller module.

319

Warning The indicated available disk has failed.

Recommended actions

- Replace the disk with one of the same type (SAS SSD, enterprise SAS, or midline SAS) and the same or greater capacity. For continued optimum I/O performance, the replacement disk should have performance that is the same as or better than the one it is replacing.

322

Warning The controller has an older Storage Controller (SC) version than the version used to create the CHAP authentication database in the controller's flash memory.

The CHAP database cannot be read or updated. However, new records can be added, which will replace the existing database with a new database using the latest known version number.

Recommended actions

- Upgrade the controller firmware to a version whose SC is compatible with the indicated database version.
 - If no records were added, the database becomes accessible and remains intact.
 - If records were added, the database becomes accessible but contains only the new records.

352

Info. Expander Controller (EC) assert data or stack-dump data are available.

Recommended actions

- No action is required.

353

Info. Expander Controller (EC) assert data and stack-dump data have been cleared.

Recommended actions

- No action is required.

354

Warning SAS topology has changed on a host port. At least one PHY has gone down. For example, the SAS cable connecting a controller host port to a host has been disconnected.

Recommended actions

- Check the cable connection between the indicated port and the host.
- Monitor the log to see if the problem persists.

Info. SAS topology has changed on a host port. At least one PHY has become active. For example, the SAS cable connecting a controller host port to a host has been connected.

Recommended actions

- No action is required.

355

Warning The controller module's debug button was found to be stuck in the On position during boot up.

Recommended actions

- If the button remains stuck, replace the controller module.

356

Warning This event can only result from tests that are run in the manufacturing environment.

Recommended actions

- Follow the manufacturing process.

357

Warning This event can only result from tests that are run in the manufacturing environment.

Recommended actions

- Follow the manufacturing process.

358

Critical All PHYs are down for the indicated disk channel. The system is degraded and is not fault tolerant because all disks are in a single-ported state.

Recommended actions

- Turn off the power for the controller enclosure, wait a few seconds, and turn it back on.
- If the condition doesn't persist (event 359 has been logged for the indicated channel), no further action is required.
- If the condition persists, this indicates a hardware problem in one of the controller modules or in the controller enclosure midplane. For help identifying which FRU to replace, see [Resources for diagnosing and resolving problems](#) on page 6.

Warning Some, but not all, PHYs are down for the indicated disk channel.

Recommended actions

- Monitor the log to see whether the condition persists.
- If the condition doesn't persist (event 359 has been logged for the indicated channel), no further action is required.
- If the condition persists, this indicates a hardware problem in one of the controller modules or in the controller enclosure midplane. For help identifying which FRU to replace, see [Resources for diagnosing and resolving problems](#) on page 6, or the CLI help for the `show events` command.

359

Info. All PHYs that were down for the indicated disk channel have recovered and are now up.

Recommended actions

- No action is required.

360

Info. The speed of the indicated disk PHY was renegotiated.

Recommended actions

- No action is required.

361

Critical, Error, or Warning The scheduler experienced a problem with the indicated schedule.

Recommended actions

- Take appropriate action based on the indicated problem.

Info. A scheduled task was initiated.

Recommended actions

- No action is required.

362

Critical, Error, or Warning The scheduler experienced a problem with the indicated task.

Recommended actions

- Take appropriate action based on the indicated problem.

Info. The scheduler experienced a problem with the indicated task.

Recommended actions

- No action is required.

363

Error When the Management Controller (MC) is restarted, firmware versions that are currently installed are compared against those in the bundle that was most recently installed. When firmware is updated, it is important that all components are successfully updated or the system may not work correctly. Components checked include the CPLD, Expander Controller (EC), Storage Controller (SC), and MC.

Recommended actions

- Reinstall the firmware bundle.

Info. When the Management Controller (MC) is restarted, firmware versions that are currently installed are compared against those in the bundle that was most recently installed. If the versions match, this event is logged as Informational severity. Components checked include the CPLD, Expander Controller (EC), Storage Controller (SC), and MC.

Recommended actions

- No action is required.

364

Info. The broadcast bus is running as generation 1.

Recommended actions

- No action is required.

365

Error An uncorrectable ECC error occurred in Storage Controller CPU memory more than once, indicating a probable hardware fault.

Recommended actions

- Replace the controller module that logged this event.

Warning An uncorrectable ECC error occurred in Storage Controller CPU memory.

This event is logged with Warning severity to provide information that may be useful to technical support, but no action is required now. It will be logged with Error severity if it is necessary to replace the controller module.

Recommended actions

- No action is required.

400

Info. The indicated log has filled to a level at which it needs to be transferred to a log-collection system.

Recommended actions

- No action is required.

401

Warning The indicated log has filled to a level at which diagnostic data will be lost if not transferred to a log-collection system.

Recommended actions

- Transfer the log file to the log-collection system.

402

Error The indicated log has wrapped and has started to overwrite its oldest diagnostic data.

Recommended actions

- Investigate why the log-collection system is not transferring the logs before they are overwritten. For example, you might have enabled managed logs without configuring a destination to send logs to.

412

Warning One disk in the indicated RAID-6 disk group failed. The disk group is online but has a status of FTDN (fault tolerant with a down disk).

If a dedicated spare (linear only) or global spare of the proper type and size is present, that spare is used to automatically reconstruct the disk group. Events 9 and 37 are logged to indicate this. If no usable spare disk is present, but an available disk of the proper type and size is present and the dynamic spares feature is enabled, that disk is used to automatically reconstruct the disk group and event 37 is logged.

Recommended actions

- If no spare was present and the dynamic spares feature (linear only) is disabled (event 37 was NOT logged), replace the failed disk with one of the same type and the same or greater capacity. The new disk will be used to automatically reconstruct the disk group. Confirm this by checking that events 9 and 37 are logged.
- Otherwise, reconstruction automatically started and event 37 was logged. Replace the failed disk and configure the replacement as a dedicated (linear only) or global spare for future use.
- For continued optimum I/O performance, the replacement disk should have the same or better performance.
- Confirm that all failed disks have been replaced and that there are sufficient spare disks configured for future use.

413

Info. Created a replication set with the indicated primary volume.

Recommended actions

- No action is required.

414

Error Failed to create the indicated replication set for the indicated volume.

This operation is not permitted if the specified volume is already in a replication set or is not a master volume.

Recommended actions

- If the volume is a master volume and is not in a replication set, retry the operation.

415

Info. Deleted the indicated replication set.

Recommended actions

- No action is required.

416

Error Failed to delete the indicated replication set.

This can occur if an invalid identifier was specified for the replication set, or if the specified primary volume is not in the local system.

Recommended actions

- Repeat the deletion using a valid replication-set identifier, or on the local system for the primary volume.

417

Info. The indicated snapshot was automatically deleted to make space for a new snapshot or for a remote snapshot proxy volume, or while changing the secondary volume to be the primary volume.

The indicated snapshot was automatically deleted, as determined by the snap pool's policy.

Recommended actions

- No action is required.

418

Warning A replication operation cannot complete because it needs to create a proxy volume and a replication snapshot in the secondary pool, but the maximum number of volumes exists for that pool or its owning controller and the pool contains no suitable snapshot to automatically delete.

This event is logged in the secondary volume's system only.

Recommended actions

- To enable the replication operation to continue, delete at least one unneeded volume from the destination pool or from another pool owned by the same controller.
- After performing the above action, if the replication fails for the same reason and becomes suspended, events 431 and 418 will be logged. Repeat the above action and resume the replication.
- To allow additional volumes to be created in the future (standard volumes, replication volumes, or snapshots), delete any unneeded volumes.

419

Info. Started to add the indicated secondary volume to the indicated replication set.

Recommended actions

- No action is required.

420

Error Failed to add the indicated secondary volume to the indicated replication set.

This can occur for several reasons, such as:

- The volume is already a replication volume.
- The volume is not local to the system.
- The communication link is busy or experienced an error.
- The volume is not the same size as the existing volume or is no longer in the set.
- The volume record is not up to date.
- Replication is not licensed or the license limit would be exceeded.

Recommended actions

- If any of the above problems exist, resolve them. Then repeat the add operation with a valid volume.

421

Info. Completed adding the indicated secondary volume to the indicated replication set.

Recommended actions

- No action is required.

422

Info. Completed removing the indicated secondary volume from the indicated replication set.

Recommended actions

- No action is required.

423

Error Failed to remove the indicated volume from the indicated replication set.

This can occur for several reasons, such as:

- The volume record is not found.
- The volume record is not yet available.
- A primary volume conflict exists.
- You cannot delete the volume from a remote system.
- You cannot remove the volume because it is the primary volume.

Recommended actions

- If any of the above problems exist, resolve them. Then repeat the remove operation with a valid volume.

424

Info. Completed modifying parameters for the indicated secondary volume in the indicated replication set.

Recommended actions

- No action is required.

425

Info. Started a replication to the indicated secondary volume.

Recommended actions

- No action is required.

426

Info. Completed a replication to the indicated secondary volume.

Recommended actions

- No action is required.

427

Warning A communication error occurred when sending information between storage systems.

Recommended actions

- Check your network or fabric for abnormally high congestion or connectivity issues.

428

Info. A user suspended a replication to the indicated secondary volume.

Recommended actions

- No action is required.

429

Info. A user resumed a replication to the indicated secondary volume.

Recommended actions

- No action is required.

430

Info. A user aborted a replication to the indicated secondary volume.

Recommended actions

- No action is required.

431

Error Replication to the indicated volume has been suspended due to an error detected during the replication process. User intervention is required to resume replication. This can occur for several reasons, such as:

- The cache request was aborted.
- The cache detected that the source or target volume is offline.
- The cache detected a media error.
- The snap pool is full.
- The communication link is busy or experienced an error.
- The snapshot being used for the replication is invalid.
- There was a problem establishing proxy communication.

Recommended actions

- If the reported problem is with a primary volume, back up as much of the volume as possible.
- Resolve the error and then resume the replication.

432

Error Aborted a replication due to an error in the indicated secondary volume.

Recommended actions

- Verify that the secondary volume is valid and that the system where the volume resides is accessible.

433

Info. Skipped a replication to the indicated secondary volume.

Recommended actions

- No action is required.

434

Warning A replication collided with an ongoing replication to the indicated secondary volume.

A replication was in progress for the replication set when a new replication was requested. The new replication has been queued.

Recommended actions

- This can be a normal operation, but in some cases this can indicate a problem. Ensure that: there are no network issues; there is sufficient bandwidth between the primary and secondary systems; the interval between replications is set to a sufficient amount of time to allow replications to complete. Having too many replications queued can result in some replications not completing.

435

Warning Failed to initialize the indicated replication set.

This can occur because:

- Firmware in the remote system is incompatible with firmware in the local system.
- The primary volume and secondary volume involved in replication cannot communicate.

Recommended actions

- Update the firmware on one or both systems so they are running the same version.
- Check your network or fabric for abnormally high congestion or connectivity issues.

436

Warning Firmware in the remote system is incompatible with firmware in the local system so they cannot communicate with each other to perform replication operations.

Recommended actions

- Update the firmware on one or both systems so they are running the same version.

437

Info. Started to change the primary volume for the indicated replication set to the indicated volume.

Recommended actions

- No action is required.

438

Info. Completed changing the primary volume for the indicated replication set to the indicated volume.

The primary volume of the indicated replication set has been changed by a user to a different volume in the replication set.

Recommended actions

- No action is required.

439

Error Failed to change the primary volume to the indicated volume for the indicated replication set. This can occur for several reasons, such as:

- The volume is not in the replication set.
- Configuration tag or configuration data not found.
- The retry limit has been reached.

Recommended actions

- Verify that the specified volume is part of the replication set.
- Verify that there are no network issues preventing communication between the local and remote storage systems.

440

Warning Retrying a replication due to an error in the indicated secondary volume.

This can occur for several reasons, such as:

- The cache request was aborted.
- The cache detected that the source or target volume is offline.
- The cache detected a media error.
- The snap pool is full.
- The communication link is busy or experienced an error.
- The snapshot being used for the replication is invalid.
- There was a problem establishing proxy communication.

The replication is being automatically retried according to policies in place. If the issue is resolved before retries are exhausted, the replication will continue on its own. Otherwise, it will go into a suspended state unless the policy is set up to retry forever.

Recommended actions

- If any of the above problems exist, resolve them.

441

Error Failed to forward an add-volume request for a volume in a replication set. The secondary volume cannot be fully added to the replication set.

Recommended actions

- Remove the indicated secondary volume from the replication set.

442

Warning Power-On Self Test (POST) diagnostics detected a hardware error in a UART chip.

Recommended actions

- Replace the controller module that logged this event.

444

Warning A snap pool reached a capacity threshold and the associated Auto Expand policy failed because there is not enough available space in the disk group.

Recommended actions

- Increase the available space in the disk group either by expanding the disk group or by removing any unneeded volumes.

Info. A snap pool reached a capacity threshold and the associated policy completed successfully. For example, the snap pool was expanded successfully, or the oldest snapshot was deleted, or all snapshots were deleted. If the policy is Delete Oldest Snapshot, the serial number of the deleted snapshot is reported.

Recommended actions

- No action is required.

449

Info. Roll back was aborted for the indicated master volume.

This can occur if a roll back is in progress and a user selects to roll back a different volume, which will abort the first roll back and start a new roll back. A user can't explicitly abort a roll back because that would corrupt the parent volume.

Recommended actions

- No action is required.

450

Warning A remote volume's status changed from online to offline.

This can occur for several reasons, such as:

- The communication link is busy or experienced an error.
- The local initiator experienced an error.

Recommended actions

- Verify that there are no network issues preventing communication between the local and remote storage systems.

451

Info. A remote volume's status changed from offline to online.

Recommended actions

- No action is required.

452

Info. The indicated volume has been detached from the indicated replication set.

The volume can now be physically moved to another storage system.

Recommended actions

- No action is required.

453

Info. The indicated volume has been reattached to the indicated replication set.

Recommended actions

- No action is required.

454

Info. A user changed the drive-spin-down delay for the indicated disk group to the indicated value.

Recommended actions

- No action is required.

455

Warning The controller detected that the configured host-port link speed exceeded the capability of an FC SFP. The speed has been automatically reduced to the maximum value supported by all hardware components in the data path.

Recommended actions

- Replace the SFP in the indicated port with an SFP that supports a higher speed.

456

Warning The system's IQN was generated from the default OUI because the controllers could not read the OUI from the midplane FRU ID data during startup. If the IQN is wrong for the system's branding, iSCSI hosts might be unable to access the system.

Recommended actions

- If event 270 with status code 0 is logged at approximately the same time, restart the Storage Controllers.

457

Info. The indicated virtual pool was created.

Recommended actions

- No action is required.

458

Info. Disk groups were added to the indicated virtual pool.

Recommended actions

- No action is required.

459

Info. Removal of the indicated disk group(s) was started.

When this operation is complete, event 470 is logged.

Recommended actions

- No action is required.

460

Error The indicated disk group is missing from the indicated virtual pool.

This may be caused by missing disk drives or unconnected or powered-off enclosures.

Recommended actions

- Ensure that all disks are installed and all enclosures are connected and powered on. When the problem is resolved, event 461 is logged.

461

Info. The indicated disk group that was missing from the indicated virtual pool was recovered.

This event indicates that a problem reported by event 460 is resolved.

Recommended actions

- No action is required.

462

Warning The indicated virtual pool exceeded its high threshold for allocated pages and the virtual pool is overcommitted.

There are three thresholds, two of which are user-settable. The third and highest setting is set automatically by the controller and cannot be changed. This event is logged with Warning severity if the high threshold is exceeded and the virtual pool is overcommitted. Overcommitted means that the total committed size of all virtual volumes exceeds the physical space in the virtual pool. If the storage usage drops below a threshold, event 463 is logged.

Recommended actions

- You should immediately take steps to reduce storage usage or add capacity.

463

Info. The indicated virtual pool has dropped below one of its thresholds for allocated pages.

This event indicates that a condition reported by event 462 is no longer applicable.

Recommended actions

- No action is required.

464

Warning A user inserted an unsupported cable or SFP into the indicated controller host port.

Recommended actions

- Replace the cable or SFP with a supported type.

465

Info. A user removed an unsupported cable or SFP from the indicated controller host port.

Recommended actions

- No action is required.

466

Info. The indicated virtual pool was deleted.

Recommended actions

- No action is required.

467

Info. Addition of the indicated disk group completed successfully.

Recommended actions

- No action is required.

468

Info. FPGA temperature has returned to the normal operating range and the speed of buses connecting the FPGA to downstream adapters has been restored. The speed was reduced to compensate for an FPGA over-temperature condition.

This event indicates that a problem reported by event 469 is resolved.

Recommended actions

- No action is required.

469

Warning The speed of buses connecting the FPGA to downstream adapters has been reduced to compensate for an FPGA over-temperature condition.

The storage system is operational but I/O performance is reduced.

Recommended actions

For a 2U12 or 2U24 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5°C–40°C (41°F–104°F).
- Check for any obstructions to the airflow.
- Check that there is a module or blank plate in every module slot in the enclosure.
- If none of the above explanations apply, replace the controller module that logged the error.

For a 2U48 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5°C–35°C (41°F–95°F).
- Check for any obstructions to the airflow.
- Check that the drawers are closed and there is a module or blank plate in every module slot in the enclosure.
- If none of the above explanations apply, replace the controller module that logged the error.

When the problem is resolved, event 468 is logged.

470

Info. Removal of the indicated disk group(s) completed successfully.

Recommended actions

- No action is required.

471

Info. A replication was queued because the indicated secondary volume is detached from the replication set.

Recommended actions

- No action is required.
- To enable the replication to proceed, reattach the secondary volume and then resume the replication.

472

Error A replication failed to start on the indicated secondary volume.

This can occur when the secondary system is not licensed for replication (for example, a temporary license expired).

Recommended actions

- To perform replication, ensure that the secondary system has a valid replication license.

473

Info. The indicated volume is using more than its threshold percentage of its virtual pool.

This is an indication that the storage usage crossed the user-specified threshold for this volume. If the storage usage drops below the threshold, event 474 is logged.

Recommended actions

- No action is required. How this information is utilized is left to the discretion of the user.

474

Info. The indicated volume is no longer using more than its threshold percentage of its virtual pool. This event indicates that the condition reported by event 473 is no longer applicable.

Recommended actions

- No action is required.

475

Info. A replication was queued because the indicated secondary volume is offline.

Recommended actions

- No action is required.
- To enable the replication to proceed, reattach the secondary volume and then resume the replication.

476

Warning The CPU temperature exceeded the safe range so the CPU entered its self-protection state. IOPS were reduced.

The storage system is operational but I/O performance is reduced.

Recommended actions

For a 2U12 or 2U24 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5°C–40°C (41°F–104°F).
- Check for any obstructions to the airflow.
- Check that there is a module or blank plate in every module slot in the enclosure.
- If none of the above explanations apply, replace the controller module that logged the error.

For a 2U48 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5°C–35°C (41°F–95°F).
- Check for any obstructions to the airflow.
- Check that the drawers are closed and there is a module or blank plate in every module slot in the enclosure.
- If none of the above explanations apply, replace the controller module that logged the error.

When the problem is resolved, event 478 is logged.

477

Info. The CPU temperature exceeded the normal range so the CPU speed was reduced. IOPS were reduced. The storage system is operational but I/O performance is reduced.

Recommended actions

For a 2U12 or 2U24 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5°C–40°C (41°F–104°F).
- Check for any obstructions to the airflow.
- Check that there is a module or blank plate in every module slot in the enclosure.
- If none of the above explanations apply, replace the controller module that logged the error.

For a 2U48 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5°C–35°C (41°F–95°F).
- Check for any obstructions to the airflow.
- Check that the drawers are closed and there is a module or blank plate in every module slot in the enclosure.
- If none of the above explanations apply, replace the controller module that logged the error.

When the problem is resolved, event 478 is logged.

478

Info. A problem reported by event 476 or 477 is resolved.

Recommended actions

- No action is required.

479

Error The controller reporting this event was unable to flush data to or restore data from non-volatile memory.

This mostly likely indicates a CompactFlash failure, but it could be caused by some other problem with the controller module. The Storage Controller that logged this event will be killed by its partner controller, which will use its own copy of the data to perform the flush or restore operation.

Recommended actions

- If this is the first time this event has been logged, restart the killed Storage Controller.
- If this event is then logged again, replace the controller module.

480

Error An IP address conflict was detected for the indicated iSCSI port of the storage system. The indicated IP address is already in use.

Recommended actions

- Contact your data-network administrator to help resolve the IP address conflict.

481

Error The periodic monitor of CompactFlash hardware detected an error. The controller was put in write-through mode, which reduces I/O performance.

Recommended actions

- Restart the Storage Controller that logged this event.
- If this event is logged again, shut down the Storage Controller and replace the controller module.

482

Warning One of the PCIe buses is running with fewer lanes than it should.

This event is the result of a hardware problem that has caused the controller to use fewer lanes. The system works with fewer lanes, but I/O performance is degraded.

Recommended actions

- Replace the controller module that logged this event.

483

Error An invalid expansion-module connection was detected for the indicated disk channel. An egress port is connected to an egress port, or an ingress port is connected to an incorrect egress port.

Recommended actions

- Visually trace the cabling between enclosures and correct the cabling.

484

Warning No compatible spares are available to reconstruct this disk group if it experiences a disk failure. Only disk groups that have dedicated or suitable global spares will start reconstruction automatically.

This situation puts data at increased risk because it will require user action to configure a disk as a dedicated or global spare before reconstruction can begin on the indicated disk group if a disk in that disk group fails in the future.

If the last global spare has been deleted or used for reconstruction, ALL disk groups that do not have at least one dedicated spare are at increased risk. Note that even though there may be global spares still available, they cannot be used for reconstruction of a disk group if that disk group uses larger-capacity disks or a different type of disk. Therefore, this event may be logged even when there are unused global spares. If the dynamic spares feature is enabled (linear only), this event will be logged even if there is an available disk that may be used for reconstruction.

Recommended actions

- Configure disks as dedicated spares or global spares.
 - For a dedicated spare, the disk must be of the same type as the other disks in the linear disk group and at least as large as the smallest-capacity disk in the linear disk group, and it should have the same or better performance.
 - For a global spare, it is best to choose a disk that is as big as or bigger than the largest disk of its type in the system and of equal or greater performance. If the system contains a mix of disk types (SAS SSD, enterprise SAS, or midline SAS), there should be at least one global spare of each type (unless dedicated spares are used to protect every disk group of a given type, which will only apply to a linear storage configuration).

485

Warning The indicated disk group was quarantined to prevent writing invalid data that may exist in the controller that logged this event.

This event is logged to report that the indicated disk group has been put in the quarantined offline state (status of QTOF) to prevent loss of data. The controller that logged this event has detected (via information saved in the disk group metadata) that it may contain outdated data that should not be written to the disk group. Data may be lost if you do not follow the recommended actions carefully. This situation is typically caused by removal of a controller module without shutting it down first, then inserting a different controller module in its place. To avoid having this problem occur in the future, always shut down the Storage Controller in a controller module before removing it. This situation may also be caused by failure of the CompactFlash card, as indicated by event 204.

Recommended actions

- If event 204 is logged, follow the recommended actions for event 204.
- If event 204 is NOT logged, perform the following recommended actions:
 - If event 486 is not logged at approximately the same time as event 485, reinsert the removed controller module, shut it down, then remove it again.
 - If events 485 and 486 are both logged at approximately the same time, wait at least 5 minutes for the automatic recovery process to complete. Then sign in and confirm that both controller modules are operational. (You can determine if the controllers are operational with the `show controllers` CLI command or with the WBI.) In most cases, the system will come back up and no further action is required. If both controller modules do not become operational in 5 minutes, data may have been lost. If both controllers are not operational, follow this recovery process:
 - Remove the controller module that first logged event 486.

- Turn off the power for the controller enclosure, wait a few seconds, then turn it back on.
- Wait for the controller module to restart, then sign in again.
- Check the status of the disk groups. If any of the disk groups have a status of quarantined offline (QTOF), dequarantine those disk groups.
- Reinsert the previously removed controller module. It should now restart successfully.

486

Warning A recovery process was initiated to prevent writing invalid data that may exist in the controller that logged this event.

The controller that logged this event has detected (via information saved in the disk group metadata) that it may contain outdated data that should not be written to the disk groups. The controller will log this event, restart the partner controller, wait 10 seconds, then kill itself. The partner controller will then unkill this controller and mirror the correct cache data to it. This procedure will, in most cases, allow all data to be correctly written without any loss of data and without writing any outdated data.

Recommended actions

- Wait at least 5 minutes for the automatic recovery process to complete. Then sign in and confirm that both controller modules are operational. (You can determine if the controllers are operational with the `show redundancy-mode` CLI command.) In most cases, the system will come back up and no action is required.
- If both controller modules do not become operational in 5 minutes, see the recommended actions for event 485, which will be logged at approximately the same time.

487

Info. Historical performance statistics were reset.

Recommended actions

- No action is required.

488

Info. Creation of a group of volumes started.

Recommended actions

- No action is required.

489

Info. Creation of a group of volumes completed.

Recommended actions

- No action is required.

490

Info. Creation of a group of volumes failed.

Recommended actions

- No action is required.

491

Info. Creation of a local group of volumes started.

Recommended actions

- No action is required.

492

Info. A local group of volumes was dissolved.

Recommended actions

- No action is required.

493

Info. A local group of volumes was modified.

Recommended actions

- No action is required.

494

Info. Reinitialization of a snap pool has completed.

Recommended actions

- No action is required.

495

Warning The algorithm for best-path routing selected the alternate path to the indicated disk because the I/O error count on the primary path reached its threshold.

The controller that logs this event indicates which channel (path) has the problem. For example, if the B controller logs the problem, the problem is in the chain of cables and expansion modules connected to the B controller module.

Recommended actions

- If this event is consistently logged for only one disk in an enclosure, perform the following actions:
 - Replace the disk.
 - If that does not resolve the problem, the fault is probably in the enclosure midplane. Replace the chassis-and-midplane FRU for the indicated enclosure.
- If this event is logged for more than one disk in an enclosure or disks in multiple enclosures, perform the following actions:
 - Check for disconnected SAS cables in the bad path. If no cables are disconnected, replace the cable connecting to the ingress port in the most-upstream enclosure with reported failures. If that does not resolve the problem, replace other cables in the bad path, one at a time until the problem is resolved.
 - If that does not resolve the problem, replace the expansion modules that are in the bad path. Begin with the most-upstream module that is in an enclosure with reported failures. If that does not resolve the problem, replace other expansion modules (and the controller module) upstream of the affected enclosure(s), one at a time until the problem is resolved.
 - If that does not resolve the problem, the fault is probably in the enclosure midplane. Replace the chassis-and-midplane FRU of the most-upstream enclosure with reported failures. If that does not resolve the problem and there is more than one enclosure with reported failures, replace the chassis-and-midplane FRU of the other enclosures with reported failures until the problem is resolved.

496

Warning An unsupported disk type was found.

Recommended actions

- Replace the disk with a supported type.

497

Info. A disk copyback operation started. The indicated disk is the source disk.

When a disk fails, reconstruction is performed using a spare disk. When the failed disk is replaced, the data that was reconstructed in the spare disk (and any new data that was written to it) is copied to the disk in the slot where the data was originally located. This is known as slot affinity. For the copyback operation, the reconstructed disk is called the source disk, and the newly replaced disk is called the destination disk. All of the data is copied from the source disk to the destination disk and the source disk then becomes a spare disk again.

Recommended actions

- No action is required.

498

Warning A disk copyback operation failed.

When a disk fails, reconstruction is performed using a spare disk. When the failed disk is replaced, the data that was reconstructed in the spare disk (and any new data that was written to it) is copied to the disk in the slot where the data was originally located. However, this copyback operation failed.

Recommended actions

- If the event's Additional Information says that some of the data in the source disk is unreadable (uncorrectable media error) and that data cannot be reconstructed from the other disks in the disk group, do the following:
 - If you do not have a backup copy of the data in the disk group, make a backup.
 - Note the configuration of the disk group, such as its size and host mappings.
 - Look for another event logged at approximately the same time that indicates a disk failure, such as event 8, 55, 58, or 412. Follow the recommended actions for that event.
 - Delete the disk group.
 - Re-create the disk group.
 - Restore the data from the backup.
- If the event's Additional Information says the failure may be because the disk that was inserted as a replacement for the failed disk is also faulty or because the source disk for the copyback is faulty, or may be caused by a fault in the midplane of the enclosure that the disks are inserted into, do the following:
 - Look for another event logged at approximately the same time that indicates a disk failure, such as event 8, 55, 58, or 412. Follow the recommended actions for that event.
 - If the problem then recurs for the same slot, replace the chassis-and-midplane FRU.

Info. A disk copyback operation completed.

Recommended actions

- No action is required.

499

Info. A disk copyback operation started. The indicated disk is the destination disk.

When a disk fails, reconstruction is performed using a spare disk. When the failed disk is replaced, the data that was reconstructed in the spare disk (and any new data that was written to it) is copied to the disk in the slot where the data was originally located. This is known as slot affinity. For the copyback operation, the reconstructed disk is called the source disk, and the newly replaced disk is called the destination disk. All of the data is copied from the source disk to the destination disk and the source disk then becomes a spare disk again.

Recommended actions

- No action is required.

500

Info. A disk copyback operation completed. The indicated disk was restored to being a spare.

When a disk fails, reconstruction is performed using a spare disk. When the failed disk is replaced, the data that was reconstructed in the spare disk (and any new data that was written to it) is copied to the disk in the slot where the data was originally located. This is known as slot affinity. For the copyback operation, the reconstructed disk is called the source disk, and the newly replaced disk is called the destination disk. All of the data is copied from the source disk to the destination disk and the source disk then becomes a spare disk again.

Recommended actions

- No action is required.

501

Error The enclosure hardware is not compatible with the I/O module firmware.

The Expander Controller firmware detected an incompatibility with the midplane type. As a preventive measure, disk access was disabled in the enclosure.

Recommended actions

- Update the storage system to the latest firmware.

502

Warning The indicated SSD has 5% or less of its life remaining.

This event will be logged again as the device approaches and reaches its end of life.

Recommended actions

- Be sure you have a spare SSD of the same type and capacity available.
- If a spare is available, it is recommended to replace the SSD now.

Info. The indicated SSD has 20% or less of its life remaining.

This event will be logged again with a severity of Warning as the SSD further approaches its end of life.

Recommended actions

- Obtain a replacement SSD of the same type and capacity if you do not already have one available.

504

Info. Service debug access to the system has been enabled or disabled by a user. Allowing service debug access may have security implications. After the diagnosis is complete you may want to disallow such access.

Recommended actions

- No action is required.

505

Warning The indicated virtual pool was created with a size smaller than 500 GB, which can lead to unpredictable behavior.

The storage system may not perform correctly.

Recommended actions

- Add disk groups to the virtual pool to increase the size of the pool.

506

Info. Addition of the indicated disk group was started.
When this operation is complete, event 467 is logged.

Recommended actions

- No action is required.

507

Info. The link speed of the indicated disk does not match the link speed that the enclosure is capable of.
This event is logged when the auto-negotiated link speed is less than the maximum speed that the enclosure supports. The disk is functional, but I/O performance is reduced. This event may be logged for one disk channel or for both disk channels.

Recommended actions

- If the disk is a member of a non-fault-tolerant disk group (RAID 0 or non-RAID), move the data to a different disk group.
- Replace the disk with one of the same type (SAS SSD, enterprise SAS, or midline SAS) and the same or greater capacity. For continued optimum I/O performance, the replacement disk should have performance that is the same as or better than the one it is replacing.

508

Error The indicated virtual pool went offline. All of its volumes also went offline.
All data in the virtual pool has been lost. This condition can be caused by corrupt or inaccessible virtual pool metadata.

Recommended actions

- Check for other events that indicate faults in the system and follow the recommended actions for those events.
- Re-create the virtual pool.
- Restore the data from a backup, if available.

509

Error The metadata volume for the indicated virtual pool went offline. Volume mappings and persistent reservations are inaccessible or lost.

Recommended actions

- Check for other events that indicate faults in the system and follow the recommended actions for those events.
- Create new mappings for the volumes. Persistent reservations will be restored by host systems automatically.

510

Info The FDE lock key has been set or changed by a user.

Recommended actions

- Be sure to record the lock key passphrase and the new lock ID.

511

Info The FDE import lock key has been set by a user.
This is normally used to import into the system an FDE disk that was locked by another system.

Recommended actions

- Ensure that the imported disk(s) are integrated into the system.

512

Info

The system was set to the FDE secured state by a user.

Full Disk Encryption is now enabled. Disks removed from this system will not be readable unless they are imported into another system.

Recommended actions

- No action is required.

513

Info

The system was set to the FDE repurposed state by a user.

All disks have been repurposed and set to their initial factory states. FDE is no longer enabled on the system.

Recommended actions

- No action is required.

514

Info

The FDE lock key and import key were cleared by a user.

I/O operations may continue as long as the system is not restarted.

Recommended actions

- If the system is restarted and access to data is intended, the lock key must be reinstated.

515

Info

An FDE disk was repurposed by a user.

The disk was reset to its original factory state.

Recommended actions

- No action is required.

516

Error

An FDE disk has been placed in the unavailable state.

The related event message 518, which indicates that a disk operation failed, may provide additional information.

Recommended actions

- See the recommended action specified in the event message.

517

Info

A disk that was formerly in the FDE unavailable state is no longer unavailable.

The disk has returned to normal operations.

Recommended actions

- No action is required.

518

Error

An FDE disk operation has failed.

This event provides detail about the operation that failed.

Recommended actions

- The disk may need to be removed, imported, repurposed, or replaced.

521

Error An error occurred while accessing the midplane SEEPROM to store or fetch Full Disk Encryption keys. The midplane's memory is used to store the FDE lock key.

Recommended actions

- The midplane may need to be replaced if the error persists.

524

Error A temperature or voltage sensor reached a critical threshold.

A sensor monitored a temperature or voltage in the critical range. When the problem is resolved, event 47 is logged for the component that logged event 524.

If the event refers to a disk sensor, disk behavior may be unpredictable in this temperature range.

Check the event log to determine if more than one disk has reported this event.

- If multiple disks report this condition there could be a problem in the environment.
- If one disk reports this condition, there could be a problem in the environment or the disk has failed.

Recommended actions

For a 2U12 or 2U24 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5–40° C (41° F–104° F).
- Check for any obstructions to the airflow.
- Check that there is a module or blank plate in every module slot in the enclosure.
- If none of the above explanations apply, replace the disk or controller module that logged the error.

For a 2U48 enclosure:

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The enclosure operating range is 5–35° C (41° F–95° F).
- Check for any obstructions to the airflow.
- Check that the drawers are closed and there is a module or blank plate in every module slot in the enclosure.
- If none of the above explanations apply, replace the disk or controller module that logged the error.

527

Error Expander Controller (EC) firmware is incompatible with the enclosure.

As a preventative measure, the Expander Controller (EC) disabled all PHYs and reported the short enclosure status page in the supported diagnostic list.

Recommended actions

- Upgrade the controller module to the latest supported bundle version.

528

Error Expander Controller firmware detected that the partner Expander Controller (EC) firmware is incompatible with the enclosure.

As a preventative measure, the Expander Controller (EC) disabled all PHYs and reported the short enclosure status page in the supported diagnostic list.

Recommended actions

- Upgrade the partner controller module to the latest supported bundle version.

529

Error The local Expander Controller (EC) is incompatible with the enclosure.

As a preventative measure, the Expander Controller (EC) disabled all PHYs and reported the short enclosure status page in the supported diagnostic list.

Recommended actions

- Replace the controller module with one that is compatible with the enclosure.

530

Error The local Expander Controller (EC) firmware detected a level of incompatibility with the partner Expander Controller (EC) . This incompatibility could be due to unsupported hardware or firmware.

As a preventative measure, the local Expander Controller (EC) is holding the partner Expander Controller (EC) in a reset loop.

Recommended actions

- Remove the partner controller module from the enclosure. Boot the partner controller module in single-controller mode in a separate enclosure (without the controller module that reported this event). Load the latest compatible bundle version. If the version fails to load, replace the partner controller module.

531

Error The indicated controller module was unable to recover from a stall. The system will need to be manually recovered.

Recommended actions

- Download the debug logs from your storage system and contact technical support. A service technician can use the debug logs to determine the problem.

Warning The indicated controller module detected a stall. The system will perform corrective actions.

Recommended actions

- No action is required.

533

Error This event provides details about the result of the MC test of the indicated component.

If the test succeeded, the message says the component is present and operational. If the test failed, the message says the component is unavailable.

Recommended actions

- If the event indicates the test failed, replace the controller module that logged this event.

Info. This event provides details about the result of the MC test of the indicated component.

Recommended actions

- No action is required.

545

Warning A controller module is connected to a legacy enclosure midplane, resulting in degraded performance.

Recommended actions

- To achieve better performance, replace the enclosure's legacy chassis-and-midplane FRU with the latest version of the FRU.

546

Error The controller that logged this event killed the partner controller which has an incompatible host port configuration.

Recommended actions

- Replace the killed controller module with a controller module that has the same host port configuration as the surviving controller module.

548

Warning Disk group reconstruction failed. When a disk fails, reconstruction is performed using a spare disk. However, this operation failed. This event can be logged for two reasons:

- Some of the data in the other disk(s) in the disk group is unreadable (uncorrectable media error), so part of the data cannot be reconstructed.
- The disk that was used as a replacement for the failed disk is also faulty. This failure may also be caused by a fault in the midplane of the enclosure that the disks are inserted into.

Recommended actions

- If the event was logged for the first reason above, do the following:
 - If you do not have a backup copy of the data in the disk group, make a backup.
 - Note the configuration of the disk group, such as its size and host mappings.
 - Look for another event logged at approximately the same time that indicates a disk failure, such as event 8, 55, 58, or 412. Follow the recommended actions for that event.
 - Remove the disk group.
 - Re-add the disk group.
 - Restore the data from the backup.
- If the event was logged for the second reason above, do the following:
 - Look for another event logged at approximately the same time that indicates a disk failure, such as event 8, 55, 58, or 412. Follow the recommended actions for that event.
 - If the problem then recurs for the same slot, replace the chassis-and-midplane FRU.

549

Critical The indicated controller module detected that it recovered from an internal processor fault.

Recommended actions

- Replace the controller module.

550

Critical The read data path between the Storage Controller and the disk drives was detected to be unreliable. The Storage Controller took action to correct this.

Recommended actions

- Replace the controller.

562

Info. Virtual pool statistics have been reset.

Recommended actions

- No action is required.

Troubleshooting steps for leftover disk drives

Storage systems use metadata on disk drives to identify disk group members and identify other disk members of the disk group.

Disks enter a Leftover state for several reasons:

- Disk spin up was not completed before a controller polled the disk. When the controller queries the disk and finds the disk is not in a ready state, the controller may place the disk into a Leftover state.
- Improper power-on sequences.
- Firmware upgrade (due to a timing issue).
- Failover taking longer than expected.
- The disk is swapped from another system, or removed and reinserted in the enclosure.

Metadata on a disk identifies the disk as being a member of a disk group. Improperly clearing the metadata from a disk may cause permanent data loss.

△ **CAUTION:** Clearing metadata from a leftover disk should be done with extreme care. Only clear metadata if you are certain the disk has never been associated with a disk group in this system or contains no data. This situation most often occurs when inserting a previously used disk into a live system or moving a disk between two systems.

Never clear metadata from a disk if any disk group in the system is in an Offline, Quarantined, or inaccessible state. Do not clear metadata from a disk if you are unsure this is the correct step to take. Clearing metadata from a disk permanently clears all data from the disk. In these types of situations, a backup of data should be done if possible.

Using the trust command

The CLI `trust` command should only be used as a last step in a disaster recovery situation. This command has the potential to cause permanent data loss and unstable operation of the disk group. If a disk group with a single disk is in a leftover or failed condition, the `trust` command should never be used. The `trust` command should only be used if the disk group is in an Offline state.

If a single disk in a disk group has failed or been placed into a Leftover state due to errors, reintegrating the disk into the same or a different disk group has the potential to cause data loss. A disk that has failed or been placed into a Leftover state due to multiple errors should be replaced with a new disk. Assign the new disk back to the disk group as a spare and allow reconstruction to complete in order to return the disk group to a fault tolerant state

The `trust` command attempts to resynchronize leftover disks in order to make any leftover disk an active member of the disk group again. The user might need to take this step when a disk group is offline because there is no data backup, or as a last attempt to try to recover the data on a disk group. In this case, `trust` may work, but only as long as the leftover disk continues to operate. When the "trusted" disk group is back online, backup all data on the disk group and verify all data to ensure it is valid. The user then needs to delete the trusted disk group, add a new disk group, and restore data from the backup to the new disk group.

📌 **IMPORTANT:** Using `trust` on a disk group is only a disaster-recovery measure. The disk group has no tolerance for additional failures and should never be put back into a production environment.

△ **CAUTION:** Before trusting a disk group, carefully read the cautions and procedures for using the `trust` command in the CLI Reference Guide and online help.

Once the `trust` command has been issued on a disk group, further troubleshooting steps may be limited towards disaster recovery. If you are unsure of the correct action to take, contact technical support for further assistance.

Power supply faults and recommended actions

Table 1 Power supply faults and recommended actions

Fault	Recommended action
Power supply fan warning or failure, or power supply warning or failure. Event code 168.	<ul style="list-style-type: none"> • If the front bezel of your enclosure contains an air filter, remove the bezel and determine if the air filter needs cleaning or replacement. • Check that all of the fans are working using the WBI. • Make sure that no slots are left open for more than two minutes. If you need to replace a module, leave the old module in place until you have the replacement, or use a blank cover to close the slot. Leaving a slot open negatively affects the airflow and might cause the unit to overheat. • Make sure that the power supply modules are properly seated in their slots and that their latches are locked.
Power supply module status is listed as failed or you receive a voltage event notification. Event code 168.	<ul style="list-style-type: none"> • For each power supply module with a switch, check that the switch is turned on. • Check that the power cables are firmly plugged into both power supply and into an appropriate and functional electrical outlet. • Replace the power supply module.
Power LED is off.	Same as above.
Voltage/Fan Fault/Service Required LED is on.	Replace the power supply module.

Events sent as indications to SMI-S clients

If the storage system's SMI-S interface is enabled, the system will send events as indications to SMI-S clients so that SMI-S clients can monitor system performance. For information about enabling the SMI-S interface, see the chapter about configuring the system in your product's RAIDar User Guide or Storage Management Guide.

The event categories below pertain to FRU assemblies and certain FRU components.

Table 2 Events and corresponding SMI-S indications

FRU/Event category	Corresponding SMI-S class	Operation status values that would trigger alert conditions
Controller	DHS_Controller	Down, Not Installed, OK
Hard Disk Drive	DHS_DiskDrive	Unknown, Missing, Error, Degraded, OK
Fan	DHS_PSUFan	Error, Stopped, OK
Power Supply	DHS_PSU	Unknown, Error, Other, Stressed, Degraded, OK
Temperature Sensor	DHS_OverallTempSensor	Unknown, Error, Other, Non-Recoverable Error, Degraded, OK
Battery/SuperCap	DHS_SuperCap	Unknown, Error, OK
FC Port	DHS_FCPort	Stopped, OK
SAS Port	DHS_SASTargetPort	Stopped, OK
iSCSI Port	DHS_ISCSIEthernetPort	Stopped, OK

Glossary

2U12	An enclosure that is two rack units in height and can contain 12 disks.
2U24	An enclosure that is two rack units in height and can contain 24 disks.
2U48	An enclosure that is two rack units in height and can contain 48 disks.
Additional Sense Code/Additional Sense Code Qualifier	See ASC/ASCQ.
Advanced Encryption Standard	See AES.
AES	Advanced Encryption Standard. A specification for the encryption of data using a symmetric-key algorithm.
Air Management Sled	See AMS.
Air Management Solution	See AMS.
allocated page	A page of storage-pool space that has been allocated to a volume to store data.
allocation rate	The rate, in pages per minute, at which a pool is allocating pages to its volumes because they need more space to store data.
ALUA	Asymmetric Logical Unit Access.
AMS	For a 2U12 or 2U24 enclosure, Air Management Sled. A drive blank designed to fill an empty disk slot in an enclosure to maintain optimum airflow through the chassis. For a 2U48 enclosure, Air Management Solution. A plastic insert designed to fill an empty disk bay (four disk slots) within a drawer to maintain optimum airflow through the chassis.
array	See storage system.
ASC/ASCQ	Additional Sense Code/Additional Sense Code Qualifier. Information on sense data returned by a SCSI device.
ATS	Automated tiered storage. A paged-storage feature that automatically uses the appropriate tier of disks to store data based on how frequently the data is accessed. This enables higher-cost, higher-speed disks to be used only for frequently needed data, while infrequently needed data can reside in lower-cost, lower-speed disks.
automated tiered storage	See ATS.
auto-write-through	See AWT.
available disk	A disk that is not a member of a disk group, is not configured as a spare, and is not in the leftover state. It is available to be configured as a part of a disk group or as a spare. See <i>also</i> compatible disk, dedicated spare, dynamic spare, and global spare.
AWT	Auto-write-through. A setting that specifies when the RAID controller cache mode automatically changes from write-back to write-through.
base volume	A virtual volume that is not a snapshot of any other volume, and is the root of a snapshot tree.
CAPI	Configuration Application Programming Interface. A proprietary protocol used for communication between the Storage Controller and the Management Controller in a controller module. CAPI is always enabled.
CHAP	Challenge-Handshake Authentication Protocol.
chassis	The sheetmetal housing of an enclosure.
child volume	The snapshot of a parent volume in a snapshot tree.
chunk size	The amount of contiguous data that is written to a disk group member before moving to the next member of the disk group.

CIM	Common Information Model. The data model for WBEM. It provides a common definition of management information for systems, networks, applications and services, and allows for vendor extensions.
CIM Query Language	See CQL.
CIMOM	Common Information Model Object Manager. A component in CIM that handles the interactions between management applications and providers.
CNC	Converged Network Controller. A controller module whose host ports can be set to operate in FC or iSCSI mode, using qualified SFP and cable options. Changing the host-port mode is also known as changing the ports' personality.
comma separated values	See CSV.
Common Information Model	See CIM.
Common Information Model Object Manager	See CIMOM.
compatible disk	A disk that can be used to replace a failed member disk of a disk group because it both has enough capacity and is of the same type (enterprise SAS, for example) as the disk that failed. See also available disk, dedicated spare, dynamic spare, and global spare.
complex programmable logic device	See CPLD.
Configuration Application Programming Interface	See CAPI.
controller A (or B)	A short way of referring to controller module A (or B).
controller enclosure	An enclosure that contains one or two controller modules.
controller module	A FRU that contains the following subsystems and devices: a Storage Controller processor; a Management Controller processor; a SAS expander and Expander Controller processor; management interfaces; cache protected by a supercapacitor pack and nonvolatile memory (CompactFlash); host, expansion, network, and service ports; and midplane connectivity. In a controller enclosure, the upper controller module is designated A and the lower one is designated B.
converged network controller	See CNC.
Coordinated Universal Time	See UTC.
CPLD	Complex programmable logic device. An electronic component used to build reconfigurable digital circuits. It can replace large numbers of logic gates.
CQL	CIM Query Language.
CRC	Cyclic Redundancy Check. A mathematical algorithm that, when implemented in software or hardware, can be used to detect errors in data.
CSV	Comma separated values. A format to store tabular data in plain-text form.
Cyclic Redundancy Check	See CRC.
DAS	Direct Attached Storage. A dedicated storage device that connects directly to a host without the use of a switch.
Data Encryption Standard	See DES.
deallocation rate	The rate, in pages per minute, at which a pool is deallocating pages from its volumes because they no longer need the space to store data.

dedicated spare	A disk that is reserved for use by a specific linear disk group to replace a failed disk. See <i>also</i> available disk, compatible disk, dynamic spare, and global spare.
default mapping	Host-access settings that apply to all initiators that are not explicitly mapped to that volume using different settings. See <i>also</i> explicit mapping and masking.
DES	Data Encryption Standard. An algorithm for the encryption of electronic data.
DHCP	Dynamic Host Configuration Protocol. A network configuration protocol for hosts on IP networks.
Direct Attached Storage	See DAS.
disk group	A set of disk drives that is configured to use a specific RAID type and provides storage capacity for a pool. See <i>also</i> linear disk group and virtual disk group.
Distributed Management Task Force	See DMTF.
DMTF	Distributed Management Task Force. An industry organization that develops and maintains standards for system management.
drain	Moving active volume data from a virtual disk group to other disk-group members within the same pool.
drawer	In a 2U48 enclosure, a FRU that holds up to 16 disks installed in groups of four. Each enclosure has three drawers.
drive enclosure	See expansion enclosure. See <i>also</i> JBOD.
drive spin down	See DSD.
DRM	Disaster recovery management. Storage-system firmware features that, when the Site Replication Adapter (SRA) feature is enabled, support the use of VMware's Site Recovery Manager to automate disaster-recovery failover and failback tasks. See <i>also</i> SRA.
DSD	Drive spin down. A power-saving feature that monitors disk activity in the storage system and spins down inactive disks based on user-selectable policies. Drive spin down is not applicable to disks in virtual pools.
dual-port disk	A disk that is connected to both controllers so it has two data paths, achieving fault tolerance.
Dynamic Host Configuration Protocol	See DHCP.
dynamic spare	An available compatible disk that is automatically assigned, if the dynamic spares option is enabled, to replace a failed disk in a disk group with a fault-tolerant RAID level. See <i>also</i> available disk, compatible disk, dedicated spare, and global spare.
EC	Expander Controller. A processor (located in the SAS expander in each controller module and expansion module) that controls the SAS expander and provides SES functionality. See <i>also</i> EMP, MC, and SC.
EMP	Enclosure management processor. An EC subsystem that provides SES data such as temperature, power supply and fan status, and the presence or absence of disks.
enclosure	A physical storage device that contains I/O modules, disk drives, and other FRUs.
enclosure management processor	See EMP.
Expander Controller	See EC.
expansion enclosure	An enclosure that contains one or two expansion modules. Expansion enclosures can be connected to a controller enclosure to provide additional storage capacity. See <i>also</i> JBOD.
expansion module	A FRU that contains the following subsystems and devices: a SAS expander and EC processor; host, expansion, and service ports; and midplane connectivity. In a drive enclosure, the upper expansion module is designated <i>A</i> and the lower one is designated <i>B</i> .

explicit mapping	Access settings for an initiator to a volume that override the volume's default mapping. See <i>also</i> default mapping and masking.
extrinsic methods	Methods which are particular to a specific class in SMI-S. See <i>also</i> intrinsic methods.
failback	See recovery.
failover	In an active-active configuration, failover is the act of temporarily transferring ownership of controller resources from an offline controller to its partner controller, which remains operational. The resources include pools, volumes, cache data, host ID information, and LUNs and WWNs. See recovery.
FC	Fibre Channel interface protocol.
FC-AL	Fibre Channel Arbitrated Loop. The FC topology in which devices are connected in a one-way loop.
FDE	Full disk encryption. A method by which you can secure the data residing on a system. See <i>also</i> lock key, passphrase, repurpose, and SED.
Fibre Channel Arbitrated Loop	See FC-AL.
field-programmable gate array	See FPGA.
field-replaceable unit	See FRU.
FPGA	Field-programmable gate array. An integrated circuit designed to be configured after manufacturing.
FRU	Field-replaceable unit. A part that can be removed and replaced by the user or support technician without having to send the product to a repair facility.
full disk encryption	See FDE.
global spare	A compatible disk that is reserved for use by any disk group with a fault-tolerant RAID level to replace a failed disk. See <i>also</i> available disk, compatible disk, dedicated spare, and dynamic spare.
HBA	Host bus adapter. A device that facilitates I/O processing and physical connectivity between a host and the storage system.
host	(v3) A user-defined group of initiators that represents a server or switch. (v2) An external port that the storage system is attached to. The external port may be a port in an I/O adapter in a server, or a port in a network switch. Product interfaces use the terms host and initiator interchangeably.
host bus adapter	See HBA.
host group	A user-defined group of hosts for ease of management, such as for mapping operations.
host port	A port on a controller module that interfaces to a host computer, either directly or through a network switch.
image ID	A globally unique serial number that identifies the point-in-time image source for a volume. All volumes that have identical image IDs have identical data content, whether they be snapshots or stand-alone volumes.
initiator	(v3) An external port to which the storage system is connected. The external port may be a port in an I/O adapter in a server, or a port in a network switch. (v2) See host.
I/O Manager	A MIB-specific term for a controller module.
I/O module	See IOM.
intrinsic methods	Methods inherited from CIM and present in all classes such as <code>getClass</code> , <code>createInstance</code> , <code>enumerateInstances</code> , and <code>associatorNames</code> in SMI-S.
IOM	Input/output module. An IOM can be either a controller module or an expansion module.
IQN	iSCSI Qualified Name.
iSCSI	Internet SCSI interface protocol.

iSNS	Internet Storage Name Service.
JBOD	“Just a bunch of disks.” See drive enclosure.
jumbo frame	In an iSCSI network, a frame that can contain 9000 bytes for large data transfers. A normal frame can contain 1500 bytes.
large form factor	See LFF.
LBA	Logical Block Address. The address used for specifying the location of a block of data.
leftover	The state of a disk that the system has excluded from a disk group because the timestamp in the disk’s metadata is older than the timestamp of other disks in the disk group, or because the disk was not detected during a rescan. A leftover disk cannot be used in another disk group until the disk’s metadata is cleared; for information and cautions about doing so, see documentation topics about clearing disk metadata.
LFF	Large form factor.
linear	The storage-class designation for logical components such as volumes that do not use paged-storage technology to virtualize data storage. The linear method stores user data in sequential, fully allocated physical blocks, using a fixed (static) mapping between the logical data presented to hosts and the physical storage where it is stored.
linear disk group	A set of disk drives that is configured to use a specific RAID type. The number of disks that a linear disk group can contain is determined by its RAID level. Any supported RAID level can be used. When a linear disk group is created, a linear pool with the same name is also created to represent the volume-containment properties of the disk group. See <i>also</i> linear pool.
linear pool	A container for volumes that is composed of one linear disk group.
LIP	Loop Initialization Primitive. An FC primitive used to determine the loop ID for a controller.
lock key	A system-generated value that manages the encryption and decryption of data on FDE-capable disks. See <i>also</i> FDE and passphrase.
logical block address	See LBA.
Logical Unit Number	See LUN.
loop	See FC-AL.
Loop Initialization Primitive	See LIP.
LUN	Logical Unit Number. A number that identifies a mapped volume to a host system.
MAC address	Media Access Control Address. A unique identifier assigned to network interfaces for communication on a network.
Management Controller	See MC.
Management Information Base	See MIB.
map/mapping	Settings that specify whether a volume is presented as a storage device to a host system, and how the host system can access the volume. Mapping settings include an access type (read-write, read-only, or no access), controller host ports through which initiators may access the volume, and a LUN that identifies the volume to the host system. See <i>also</i> default mapping, explicit mapping, and masking.
masking	A volume-mapping setting that specifies no access to that volume by hosts. See <i>also</i> default mapping and explicit mapping.
master volume	A volume that is enabled for snapshots and has an associated snap pool.
MC	Management Controller. A processor (located in a controller module) that is responsible for human-computer interfaces, such as the WBI, and computer-computer interfaces, such as SNMP, and interacts with the Storage Controller. See <i>also</i> EC and SC.
Media Access Control Address	See MAC address.

metadata	Data in the first sectors of a disk drive that stores all disk-, disk-group-, and volume-specific information including disk group membership or spare identification, disk group ownership, volumes and snapshots in the disk group, host mapping of volumes, and results of the last media scrub.
MIB	Management Information Base. A database used for managing the entities in SNMP.
mount	To enable access to a volume from a host OS. <i>See also</i> host, map/mapping, and volume.
network port	The Ethernet port on a controller module through which its Management Controller is connected to the network.
network time protocol	<i>See</i> NTP.
NTP	Network time protocol.
object identifier	<i>See</i> OID.
OID	Object Identifier. In SNMP, an identifier for an object in a MIB.
orphan data	<i>See</i> unwritable cache data.
overcommit	A setting that controls whether a virtual pool is allowed to have volumes whose total size exceeds the physical capacity of the pool.
overcommitted	The amount of storage capacity that is allocated to volumes exceeds the physical capacity of the storage system.
page	A range of contiguous LBAs in a virtual disk group.
paged storage	A method of mapping logical host requests to physical storage that maps the requests to virtualized "pages" of storage that are in turn mapped to physical storage. This provides more flexibility for expanding capacity and automatically moving data than the traditional, linear method in which requests are directly mapped to storage devices. Paged storage is also called virtual storage.
parent volume	A volume that has snapshots (can be either a base volume or a base snapshot volume). The parent of a snapshot is its immediate ancestor in the snapshot tree.
partner firmware update	<i>See</i> PFU.
passphrase	A user-created password that allows users to manage lock keys in an FDE-capable system. <i>See also</i> FDE and lock key.
PCBA	Printed circuit board assembly.
PFU	Partner firmware update. The automatic update of the partner controller when the user updates firmware on one controller.
PGR	Persistent group reservations.
PHY	One of two hardware components that form a physical connection between devices in a SAS network that enables transmission of data.
physical layer	<i>See</i> PHY.
point-to-point	Fibre Channel Point-to-Point topology in which two ports are directly connected.
pool	<i>See</i> linear pool and virtual pool.
POST	Power-on self test. Tests that run immediately after a device is powered on.
Power-On Self Test	<i>See</i> POST.
power supply unit	<i>See</i> PSU.
primary volume	The volume that is the source of data in a replication set and that can be mapped to hosts. For disaster recovery purposes, if the primary volume goes offline, a secondary volume can be designated as the primary volume. The primary volume exists in a primary disk group in the primary (or local) storage system.
proxy volume	A virtual volume in the local system that represents a volume in a remote system. Proxy volumes are used internally by the controllers to perform actions such as transferring replication data.
PSU	Power supply unit. The power supply FRU.

quick rebuild	A feature for virtual storage that reduces the time that user data is less than fully fault-tolerant after a disk failure in a disk group. The quick-rebuild process rebuilds only data stripes that contain user data. Data stripes that have not been allocated to user data are rebuilt in the background.
RAID head	See controller enclosure.
read cache	A special disk group using SSDs that can be added to a virtual pool for the purpose of speeding up read access to data stored on spinning disks elsewhere in the pool. Read cache is also referred to as read flash cache.
read flash cache	See read cache.
recovery	In an active-active configuration, recovery is the act of returning ownership of controller resources to a controller (which was offline) from its partner controller. The resources include volumes, cache data, host ID information, and LUNs and WWNs. <i>See also</i> failover.
remote replication	Asynchronous (batch) replication of block-level data from a volume in a primary system to a volume in one or more secondary systems by creating a replication snapshot of the primary volume and copying the snapshot data to the secondary systems via Fibre Channel or iSCSI links. The capability to perform remote replication is a licensed feature (AssuredRemote).
remote syslog support	See syslog.
replication image	A conceptual term for replication snapshots that have the same image ID in primary and secondary systems. These synchronized snapshots contain identical data and can be used for disaster recovery.
replication-prepared volume	A volume created for the purpose of being the secondary volume in a replication set. Replication-prepared volumes are automatically created by the WBI Replication Setup Wizard, or they can be created manually in the CLI or the WBI.
replication set	Associated primary and secondary volumes that are enabled for replication and that typically reside in two physically or geographically separate storage systems. <i>See</i> primary volume and secondary volume.
replication snapshot	A special type of snapshot, created by the remote replication feature, that preserves the state of data of a replication set's primary volume as it existed when the snapshot was created. For a primary volume, the replication process creates a replication snapshot on both the primary system and, when the replication of primary-volume data to the secondary volume is complete, on the secondary system. Replication snapshots are unmappable and are not counted toward a license limit, although they are counted toward the system's maximum number of volumes. A replication snapshot can be exported to a regular, licensed snapshot. <i>See also</i> replication sync point.
replication sync point	The state of a replication snapshot whose corresponding primary or secondary snapshot exists and contains identical data. For a replication set, four types of sync point are identified: the only replication snapshot that is copy-complete on any secondary system is the "only sync point"; the latest replication snapshot that is copy-complete on any secondary system is the "current sync point"; the latest replication snapshot that is copy-complete on all secondary systems is the "common sync point"; a common sync point that has been superseded by a new common sync point is an "old common sync point."
repurpose	A method by which all data on a system or disk is erased in an FDE-capable system. Repurposing unsecures the system and disks without needing the correct passphrase. <i>See also</i> FDE and passphrase.
RFC	Read flash cache. <i>See</i> read cache.
SAS	Serial Attached SCSI interface protocol or disk-drive architecture.
SC	Storage Controller. A processor (located in a controller module) that is responsible for RAID controller functions. The SC is also referred to as the RAID controller. <i>See also</i> EC and MC.
SCSI Enclosure Services	See SES.

secondary volume	The volume that is the destination for data in a replication set and that is not accessible to hosts. For disaster recovery purposes, if the primary volume goes offline, a secondary volume can be designated as the primary volume. The secondary volume exists in a secondary disk group in a secondary (or remote) storage system. The contents of a secondary volume are in a constant state of flux and are not in a consistent state while a replication is in process. Only snapshots that are associated with a secondary volume are data consistent.
secret	For use with CHAP, a password that is shared between an initiator and a target to enable authentication.
secure hash algorithm	See SHA.
secure shell	See SSH.
Secure Sockets Layer	See SSL.
SED	Self-encrypting drive. A disk drive that provides hardware-based data encryption and supports use of the storage system's Full Disk Encryption feature. <i>See also</i> FDE.
SEEPROM	Serial electrically erasable programmable ROM. A type of nonvolatile (persistent if power removed) computer memory used as FRU ID devices.
Self-Monitoring Analysis and Reporting Technology	See SMART.
serial electrically erasable programmable ROM	See SEEPROM.
Service Location Protocol	See SLP.
SES	SCSI Enclosure Services. The protocol that allows the initiator to communicate with the enclosure using SCSI commands.
SFCB	Small Footprint CIM Broker.
SFF	Small form factor. A type of disk drive.
SHA	Secure Hash Algorithm. A cryptographic hash function.
SLP	Service Location Protocol. Enables computers and other devices to find services in a local area network without prior configuration.
Small Footprint CIM Broker	See SFCB.
small form factor	See SFF.
SMART	Self-Monitoring Analysis and Reporting Technology. A monitoring system for disk drives that monitors reliability indicators for the purpose of anticipating disk failures and reporting those potential failures.
SMI-S	Storage Management Initiative - Specification. The SNIA standard that enables interoperable management of storage networks and storage devices. The interpretation of CIM for storage. It provides a consistent definition and structure of data, using object-oriented techniques.
snap pool	A volume that stores data that is specific to snapshots of an associated master volume, including copy-on-write data and data written explicitly to the snapshots. A snap pool cannot be mapped.
snapshot	A "virtual" volume that preserves the state of a master volume's data as it existed when the snapshot was created. Data associated with a snapshot is recorded in both the master volume and in its associated snap pool. A snapshot can be mapped and written to. The capability to create snapshots is a licensed feature (AssuredSnap). Snapshots that can be mapped to hosts are counted against the snapshot-license limit, whereas transient and unmappable snapshots are not.

snapshot tree	A group of virtual volumes that are interrelated due to creation of snapshots. Since snapshots can be taken of existing snapshots, volume inter-relationships can be thought of as a "tree" of volumes. A tree can be 254 levels deep. See <i>also</i> base volume, child volume, parent volume, and source volume.
SNIA	Storage Networking Industry Association. An association regarding storage networking technology and applications.
source volume	A volume that has snapshots. Used as a synonym for parent volume.
sparse snapshot	A type of point-in-time copy that preserves the state of data at an instant in time by storing only those blocks that are different from an already existing full copy of the data.
SRA	Storage Replication Adapter. A host-based software component that allows VMware's Site Recovery Manager to manage the storage-system firmware's disaster recovery management (DRM) features, automating disaster-recovery failover and failback tasks. The SRA uses the CLI XML API to control the storage system. See <i>also</i> DRM.
SSD	Solid-state drive.
SSH	Secure Shell. A network protocol for secure data communication.
SSL	Secure Sockets Layer. A cryptographic protocol that provides security over the internet.
standard volume	A volume that can be mapped to initiators and presented as a storage device to a host system, but is not enabled for snapshots.
Storage Controller	See SC.
Storage Management Initiative - Specification	See SMI-S.
Storage Networking Industry Association	See SNIA.
storage system	A controller enclosure with at least one connected drive enclosure. Product documentation and interfaces use the terms storage system and system interchangeably.
syslog	A protocol for sending event messages across an IP network to a logging server.
thin provisioning	A feature that allows actual storage for a virtual volume to be assigned as data is written, rather than storage being assigned immediately for the eventual size of the volume. This allows the storage administrator to overcommit physical storage, which in turn allows the connected host system to operate as though it has more physical storage available than is actually allocated to it. When physical resources fill up, the storage administrator can add storage capacity on demand.
tier	A homogeneous set of disk drives, typically of the same capacity and performance level, that comprise one or more disk groups in the same pool. Tiers differ in their performance, capacity, and cost characteristics, which forms the basis for the choices that are made with respect to which data is placed in which tier. The predefined tiers are: <ul style="list-style-type: none"> • Performance, which uses SAS SSDs (high speed, low capacity) • Standard, which uses enterprise-class spinning SAS disks (lower speed, higher capacity) • Archive, which uses midline spinning SAS disks (low speed, high capacity).
tier migration	The automatic movement of blocks of data, associated with a single volume, between tiers based on the access patterns that are detected for the data on that volume.
tray	See enclosure.
UCS Transformation Format - 8-bit	See UTF-8.
ULP	Unified LUN Presentation. A RAID controller feature that enables a host system to access mapped volumes through any controller host port. ULP incorporates Asymmetric Logical Unit Access (ALUA) extensions.
undercommitted	The amount of storage capacity that is allocated to volumes is less than the physical capacity of the storage system.
Unified LUN Presentation	See ULP.

unmount	To remove access to a volume from a host OS.
unwritable cache data	Cache data that has not been written to disk and is associated with a volume that no longer exists or whose disks are not online. If the data is needed, the volume's disks must be brought online. If the data is not needed it can be cleared, in which case it will be lost and data will differ between the host system and disk. Unwritable cache data is also called orphan data.
UTC	Coordinated Universal Time. The primary time standard by which the world regulates clocks and time. It replaces Greenwich Mean Time.
UTF-8	UCS transformation format - 8-bit. A variable-width encoding that can represent every character in the Unicode character set used for the CLI and WBI interfaces.
v2	The legacy interface for managing linear storage. This is the default for a system that has been upgraded from a previous release.
v3	The new interface for managing virtual and linear storage. This is the default for a new installation.
vdisk	A virtual disk comprising the capacity of one or more disks. The number of disks that a vdisk can contain is determined by its RAID level. See linear disk group.
vdisk spare	See dedicated spare.
virtual	The storage-class designation for logical components such as volumes that use paged-storage technology to virtualize data storage. See paged storage.
virtual disk	See vdisk.
virtual disk group	A set of disk drives that is configured to use a specific RAID type. The number of disks that a linear disk group can contain is determined by its RAID level. A virtual disk group can use RAID 1, 5, 6, or 10. A virtual disk group can be added to a new or existing virtual pool. See also virtual pool.
virtual pool	A container for volumes that is composed of one or more virtual disk groups.
volume	A logical representation of a fixed-size, contiguous span of storage that is presented to host systems for the purpose of storing data.
volume copy	An independent copy of the data in a linear volume. The capability to create volume copies is a licensed feature (AssuredCopy) that makes use of snapshot functionality.
volume group	A user-defined group of volumes for ease of management, such as for mapping operations.
WBEM	Web-Based Enterprise Management. A set of management and internet standard technologies developed to unify the management of enterprise computing environments.
web-based interface/web-browser interface	See WBI.
WBI	Web-browser interface. The primary interface for managing the system. A user can enable the use of HTTP, HTTPS for increased security, or both.
Web-Based Enterprise Management	See WBEM.
World Wide Name	See WWN.
World Wide Node Name	See WWNN.
World Wide Port Name	See WWPN.
WWN	World Wide Name. A globally unique 64-bit number that identifies a device used in storage technology.
WWNN	World Wide Node Name. A globally unique 64-bit number that identifies a device.
WWPN	World Wide Port Name. A globally unique 64-bit number that identifies a port.