

# Quantum<sup>®</sup>

## QXS 12G Events Description Guide

# 12, 24, and 84-Drive Systems



QXS 12G Events Description Guide, 6-68653-01 Rev B, September 2019, Product of USA.

#### Copyright Statement

© 2019 Quantum Corporation. All rights reserved.

Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law.

#### Trademark Statement:

Artico, Be Certain (and the Q brackets design), DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, DXiV-Series, FlexTier, Lattus, the Q logo, the Q Quantum logo, Q-Cloud, Quantum (and the Q brackets design), the Quantum logo, Quantum Be Certain (and the Q brackets design), Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo (and design), vmPRO, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners.

Products mentioned herein are for identification purposes only and may be registered trademarks or trademarks of their respective companies. All other brand names or trademarks are the property of their respective owners.

Quantum specifications are subject to change.



# Contents

<b>About This Guide</b> .....	<b>xi</b>
Intended audience .....	xi
Prerequisites .....	xi
Related Documentation .....	xii
Document conventions and symbols .....	xiii
.....	<b>1</b>
<b>Event Descriptions 1-99</b> .....	<b>1</b>
Introduction .....	1
Events and Event Messages .....	2
Event Format in this Guide .....	2
Event code - Severity .....	2
Resources for Diagnosing and Resolving Problems .....	2
Event Descriptions 1-19 .....	3
1 .....	3
3 .....	3
4 .....	4
6 .....	4
7 .....	4
8 .....	5
9 .....	6
16 .....	6
18 .....	6
19 .....	7
Event Descriptions 20-39 .....	7
20 .....	7
21 .....	7
23 .....	8
25 .....	8
28 .....	8
31 .....	9
32 .....	9
33 .....	9
34 .....	9
37 .....	9
39 .....	10
Event Descriptions 40-59 .....	10
40 .....	10
41 .....	11

43	11
44	11
47	11
48	11
49	12
50	12
51	12
52	13
53	13
55	13
56	14
58	14
59	15
Event Descriptions 60-79	15
61	15
62	15
65	16
68	16
71	16
72	16
73	16
74	17
75	17
76	17
77	17
78	18
79	18
Event Descriptions 80-99	18
80	18
81	19
83	19
84	19
86	19
87	19
88	20
89	20
90	20
91	20
95	21
96	21
<b>Event Descriptions 100-199</b>	<b>23</b>
Event Descriptions 100-119	23
103	23
104	23
105	24
106	24
107	24
108	24
109	24
110	25
111	25
112	25
114	25
116	26
117	26
118	26
Event Descriptions 120-139	26

127	26
136	27
139	27
Event Descriptions 140-159	27
140	27
141	27
152	28
153	28
154	29
155	29
156	29
157	29
158	30
Event Descriptions 160-179	30
161	30
162	30
163	31
166	31
167	31
170	32
171	32
172	32
173	33
174	33
175	34
176	34
177	34
Event Descriptions 180-199	34
181	34
182	35
183	35
185	35
186	35
187	36
188	36
189	36
190	36
191	36
192	37
193	37
194	37
195	37
198	38
199	38
<b>Event Descriptions 200-299</b>	<b>39</b>
Event Descriptions 200-219	39
200	39
201	40
202	40
203	40
204	40
205	41
206	41
207	41
208	43
209	43
210	43

211	44
216	44
217	44
218	44
219	45
Event Descriptions 220-239	45
232	45
233	45
235	45
236	46
237	46
238	46
239	47
Event Descriptions 240-259	47
240	47
241	47
242	47
243	48
245	48
246	48
247	48
248	49
249	49
250	49
251	49
253	49
255	50
259	50
Event Descriptions 260-279	50
260	50
261	50
262	50
263	51
266	51
267	51
268	51
269	52
270	52
271	52
273	52
274	53
275	54
Event Descriptions 280-299	54
298	54
299	54
<b>Event Descriptions 300-399</b>	<b>55</b>
Event Descriptions 300-319	55
300	55
301	55
302	56
303	56
304	56
305	56
306	56
307	57
309	57
310	57

311	57
312	58
313	58
314	58
315	58
316	58
317	59
319	59
Event Descriptions 320-339	59
322	59
Event Descriptions 340-359	60
352	60
353	60
354	60
356	61
357	61
358	61
359	62
Event Descriptions 360-379	62
360	62
361	62
362	62
363	63
364	63
365	63
Event Descriptions 380-399	64
<b>Event Descriptions 400-499</b>	<b>65</b>
Event Descriptions 400-419	65
400	65
401	65
402	66
412	66
Event Descriptions 420-439	66
Event Descriptions 440-459	67
442	67
443	67
454	67
455	67
456	68
457	68
458	68
459	68
Event Descriptions 460-479	68
460	68
461	69
462	69
463	70
464	70
465	70
466	70
467	70
468	71
469	71
470	71
473	72
474	72
476	72

477	72
478	73
479	73
Event Descriptions 480-499	73
480	73
481	74
482	74
483	74
484	74
485	75
486	76
487	76
488	76
489	76
490	77
491	77
492	77
493	77
495	77
496	78
<b>Event Descriptions 500-599</b>	<b>79</b>
Event Descriptions 500-519	79
501	79
502	79
503	80
504	80
505	80
506	80
507	81
508	81
509	81
510	81
511	82
512	82
513	82
514	82
515	83
516	83
517	83
518	83
Event Descriptions 520-539	84
521	84
522	84
523	84
524	84
525	85
526	85
527	85
528	86
529	86
530	86
531	86
533	87
Event Descriptions 540-559	87
542	87
543	87
545	88

546	88
548	88
549	88
550	89
551	89
552	90
553	91
554	92
555	93
556	94
557	95
Event Descriptions 560-579	96
562	96
563	96
565	96
566	97
568	97
571	97
572	98
573	98
574	98
575	98
576	99
577	99
578	99
579	100
Event Descriptions 580-599	100
580	100
581	100
582	101
584	101
585	101
586	101
587	102
590	102
593	102
594	102
595	103
598	104
599	104
<b>Event Descriptions 600-699</b>	<b>105</b>
Event Descriptions 600-619	105
602	105
603	105
604	106
605	106
606	107
607	107
608	107
609	107
610	108
611	108
612	109
613	109
615	110
616	110
617	110

618	111
619	111
626	111
627	111
628	112
Event Descriptions 620-699	112
<b>Additional Information</b>	<b>113</b>
Deprecated Events	113
Removed Events	113
Using the trust Command	113
Single Disk/Drive Failed	114
Resynchronize Leftover Disks/Drives	114
Events Sent as Indications to SMI-S Clients	114



# About This Guide

---

This guide provides information for the following 12-G QXS systems:

- QXS-312 12-G: 12-Drive (2-Port: FC or iSCSI)
- QXS-324 12-G: 24-Drive (2-Port: FC or iSCSI)
- QXS-412 12-G: 12-Drive (4-Port: FC or iSCSI)
- QXS-424 12-G: 24-Drive (4-Port: FC or iSCSI)
- QXS-484 12-G: 84-Drive (4-Port: FC or iSCSI)

---

## Intended audience

This guide is intended for storage customers and technicians.

This guide describes events that QXS 12-G 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.

---

**NOTE:** For information about product features and terms, see your product's *QXS 12G Disk Management Utility User Guide* or *QXS 12G CLI Reference Guide*.

---

---

## Prerequisites

Prerequisites for planning, installing, and using this product include knowledge of:

- Servers and computer networks
- Network administration
- Storage system installation and configuration
- Storage area network (SAN) management and direct attach storage (DAS)
- Converged Network Controllers (CNCs)
- Fibre Channel (FC) protocols
- Serial Attached SCSI (SAS) protocol

- Internet SCSI (iSCSI) protocol
- Ethernet protocol

---

## Related Documentation

Refer to [Table 1](#) for related 12G QXS documentation.

**Table 1** Related Documentation

<b>For Information About</b>	<b>See</b>
Enhancements, known issues, and late-breaking information not included in product documentation	<i>QXS 12G Release Notes</i>
Overview of hardware installation	<i>QXS 12G Quick Start Guide</i>
Product hardware installation and maintenance	<i>QXS 12G Hardware Installation and Maintenance Guide</i>
Obtaining and installing a license to use licensed features	<i>QXS 12G Licensing Guide</i>
Using the web interface to configure and manage the product	<i>QXS 12G Disk Management Utility User Guide</i>
Event codes and recommended actions	<i>QXS 12G Event Descriptions Reference Guide</i>
Using the command-line interface (CLI) to configure and manage the product	<i>QXS 12G CLI Reference Guide</i>
Cabinet information, QXS 12-G specifications, and environment and requirements	<i>QXS 12G Site Planning Guide</i>
Regulatory compliance and safety and disposal information	<i>* Product Regulatory Compliance and Safety</i>

\*Printed document included with product.

---

# Document conventions and symbols

Table 2 Document conventions

Convention	Element
Blue text	Cross-reference links and e-mail addresses
<a href="#">Blue, underlined text</a>	Web site addresses
Bold text	<ul style="list-style-type: none"><li>• Key names</li><li>• Text typed into a GUI element, such as into a box</li><li>• GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes</li></ul>
<i>Italic text</i>	Text emphasis
Monospace text	<ul style="list-style-type: none"><li>• File and directory names</li><li>• System output</li><li>• Code</li><li>• Text typed at the command-line</li></ul>
<i>Monospace, italic text</i>	<ul style="list-style-type: none"><li>• Code variables</li><li>• Command-line variables</li></ul>
<b>Monospace, bold text</b>	Emphasis of file and directory names, system output, code, and text typed at the command-line


---

**WARNING!** Indicates that failure to follow directions could result in bodily injury.

---

**△ 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.

---





# Event Descriptions 1-99

This chapter provides the following information:

- [Introduction](#)
- [Events and Event Messages](#)
- [Event Format in this Guide](#)
- [Resources for Diagnosing and Resolving Problems](#)
- [Event Descriptions 1-19](#)
- [Event Descriptions 20-39](#)
- [Event Descriptions 40-59](#)
- [Event Descriptions 60-79](#)
- [Event Descriptions 80-99](#)

---

## 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 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).

---

**TIP:** A best practice is to enable notifications to be sent for events having severity Warning and above.

---

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."
- **Resolved:** A condition that caused an event to be logged has been resolved.

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 chapter in the *QXS 12G Hardware Installation and Maintenance Guide*.

For a summary of storage events and corresponding SMI-S indications, see [Events Sent as Indications to SMI-S Clients](#) on page 114.

---

# Event Descriptions 1-19

## 1

### Warning

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

- 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.

If a 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 event 37 was not logged, a spare of the proper type and size was not available for reconstruction. Replace the failed disk with one of the same type and the same or greater capacity and, if necessary, designate it as a spare. 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 in the Performance tier, 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.
  - Contact technical support for help to determine if the trust operation applies to your situation and for help to perform 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 (CLI remove disk-groups command).
- Re-create the disk group (CLI add disk-group 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

Either:

- Disk group creation completed successfully.
- 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.

## 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).
- An SSD that was part of a disk group has reported that it has no life remaining. 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.

### Recommended Actions

- If a disk that was part of a disk group is down:
  - If the indicated disk failed for one of these reasons—excessive media errors, imminent disk failure, possible hardware failure, disk is not supported, too many controller-recoverable errors, illegal request, due to being degraded, or due to being too slow—replace the disk with one of the same type (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 indicated disk failed because a user forced the disk out of the disk group, RAID-6 initialization failed, or for an unknown reason:
    - If 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 of the same type (SSD, enterprise SAS, or midline SAS) and the same or greater capacity as the one that was in the slot. 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 disk then has a status of leftover (LEFTOVR), clear the metadata to reuse the disk.
    - 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. Contact technical support for help to determine if the trust operation applies to your situation and for help to perform 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 (CLI remove disk-groups command).
  - Clear the indicated disk's metadata so the disk can be re-used (CLI clear disk-metadata 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 (CLI add disk-group command).
- If an SSD that was part of a disk group has reported that it has no life remaining, replace the disk with one of the same type 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

### Error

Disk group reconstruction completed with errors.

When a disk fails, reconstruction is performed using a spare disk. However, this operation failed. 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.

#### Recommended actions

- If you do not have a backup copy of the data, take a backup.
- 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.

## Info

Disk group reconstruction completed.

For the ADAPT disk group that completed partially, either there is no available spare space, or the spare space cannot be used because of ADAPT fault-tolerant requirements.

### Recommended Actions

- No action is required.

## 19

## Info

A rescan has completed.

### Recommended Actions

- No action is required.

---

# Event Descriptions 20-39

## 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, check the current state of the disk group. If it is not FTOL then back up the data as data may be at risk. If it is FTOL then replace the indicated disk. If more than one disk in the same disk group has logged a SMART event, back up the data and replace each disk one at a time. In virtual storage it may be possible to remove the affected disk group, which will drain its data to another disk group, and then re-add the disk group.

## 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 designated as a 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 required.

## Info

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

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The controller enclosure operating range is 5°C to 35°C (41°F to 95°F). The expansion enclosure operating range is 5°C to 40°C (41°F to 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.

---

## Event Descriptions 40-59

## 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

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The controller enclosure operating range is 5°C to 35°C (41°F to 95°F). The expansion enclosure operating range is 5°C to 40°C (41°F to 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.

## 41

### Info

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

#### Recommended Actions

- No action required.

## 43

### Info

The indicated disk group has been deleted.

#### Recommended Actions

- No action 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](#), page 2.

## 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 required.

## 48

### Info

The indicated disk group has been renamed.

#### Recommended Actions

- No action required.

## 49

### Info

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

#### Recommended Actions

- No action 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 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 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 (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 (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, check the current state of the disk group. If it is not FTOL then back up the data as data may be at risk. If it is FTOL then replace the indicated disk. If more than one disk in the same disk group has logged a SMART event, back up the data and replace each disk one at a time. In virtual storage it may be possible to remove the affected disk group, which will drain its data to another disk group, and then re-add the disk group.

## 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 (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 (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.

---

## Event Descriptions 60-79

## 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 dedicated spare disk or global spare disk has failed.

#### Recommended actions

- Replace the disk with one of the same type (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 (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 (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.

---

## Event Descriptions 80-99

## 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.





## Event Descriptions 100-199

---

This chapter provides the following information:

- [Event Descriptions 100-119](#)
- [Event Descriptions 120-139](#)
- [Event Descriptions 140-159](#)
- [Event Descriptions 160-179](#)
- [Event Descriptions 180-199](#)

---

### Event Descriptions 100-119

#### 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. This can affect host mappings.

#### 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](#).

## 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 (Storage Controller booted up), closely followed by event 71 (failover started). 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.

---

## Event Descriptions 120-139

## 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.

---

## Event Descriptions 140-159

## 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.

---

## Event Descriptions 160-179

## 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.

## 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 was quarantined for one of the following reasons:

- Not all of its disks are accessible. While the disk group is quarantined, in linear storage any attempt to access its volumes in the disk group from a host will fail. In virtual storage, all volumes in the pool will be forced read-only. If all of the disks become accessible, the disk group will be dequarantined automatically with a resulting status of FTOL. If not all of the disks become accessible but enough become accessible to allow reading from and writing to the disk group, it will be dequarantined automatically with a resulting status of FTDN or CRIT. 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 dequarantine, see the disk management utility or CLI documentation.

---

**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.

---

- It contains data in a format that is not supported by this system. The controller does not support linear disk groups.

#### Recommended Actions

- If the disk group was quarantined because not all of its disks are accessible:
  - 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.
- If the disk group was quarantined because it contains data in a format that is not supported by this system:
  - Recover full support and manageability of the quarantined disk groups and volumes by replacing your controllers with controllers that support this type of disk group.
  - If you are sure that the data on this disk group is not needed, simply remove the disk

## 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.

---

## Event Descriptions 180-199

## 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](#).

## 183

**Info**

All disk channels have been unpaused, meaning that I/O can resume. An unpauses 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 cash is disabled.

#### Recommended Actions

- No action is required.

## 188

### Info

Write-back cache has been disabled.

Event 187 is the corresponding even 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.



## Chapter 3

# Event Descriptions 200-299

---

This chapter provides the following information:

- [Event Descriptions 200-219](#)
- [Event Descriptions 220-239](#)
- [Event Descriptions 240-259](#)
- [Event Descriptions 260-279](#)
- [Event Descriptions 280-299](#)

---

## Event Descriptions 200-219

### 200

#### Info

A fan 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 NV device itself 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 CompactFlash.
- 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 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 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, check the current state of the disk group. If it is not FTOL then back up the data as data may be at risk. If it is FTOL then replace the indicated disk. If more than one disk in the same disk group has logged a SMART event, back up the data and replace each disk one at a time. In virtual storage it may be possible to remove the affected disk group, which will drain its data to another disk group, and then re-add the disk group.

---

**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, check the current state of the disk group. If it is not FTOL then back up the data as data may be at risk. If it is FTOL then replace the indicated disk. If more than one disk in the same disk group has logged a SMART event, back up the data and replace each disk one at a time. In virtual storage it may be possible to remove the affected disk group, which will drain its data to another disk group, and then re-add the disk group.

## 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 (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 when using virtual storage.

#### 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” (page 7).

### 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.

## 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.

---

## Event Descriptions 220-239

## 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.

## 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 CLI show system command to determine the specific health issue(s). For unwritten cache data, use the CLI show unwritable-cache 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 CompactFlash.
- If this event is logged again, shut down the Storage Controller and replace the controller module.

---

## Event Descriptions 240-259

## 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 CompactFlash.
- 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 SEEPROM 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.

Do not mount either volume until the copy is complete (as indicated by event 268).

#### 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.

## 259

### Info

In-band CAPI commands have been disabled.

#### Recommended Actions

- No action is required.

---

## Event Descriptions 260-279

## 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 (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

A volume-copy operation completed with a failure.

This event has two variants:

- 1 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).
- 2 Possible causes are the pool running out of available space and crossing the high threshold, volumes being unavailable, or general I/O errors.

#### Recommended Actions

- For variant 1: No action is required.
- For variant 2: Look for other events logged at approximately the same time that indicate a pool space or volume failure. Follow the recommended actions for those events.

## 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 CLI *show system* command to determine the specific health issue(s). For unwritten cache data, use the CLI *show unwritable-cache* 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.

## 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.

## 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 persists, the fault may be in the enclosure midplane. Replace the chassis 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 persists, the fault may be in the enclosure midplane. Replace the chassis FRU.

## 275

### Info

The indicated PHY has been enabled.

#### Recommended Actions

- No action is required.

---

## Event Descriptions 280-299

## 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.



## Chapter 4

# Event Descriptions 300-399

---

This chapter provides the following information:

- [Event Descriptions 300-319](#)
- [Event Descriptions 320-339](#)
- [Event Descriptions 340-359](#)
- [Event Descriptions 360-379](#)
- [Event Descriptions 380-399](#)

---

## Event Descriptions 300-319

### 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 SEEPROM. 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 SEEPROM. 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

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The controller enclosure operating range is 5°C to 35°C (41°F to 95°F). The expansion enclosure operating range is 5°C to 40°C (41°F to 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.

## 309

### Info

Normally when the Management Controller (MC) is started, the IP data is obtained from the midplane FRU ID SEEPROM where it is persisted. If the system is unable to write it to the SEEPROM 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 Hardware Installation and Maintenance Guide.

## 315

### Critical

This IOM is incompatible with the enclosure in which it is inserted.

#### Recommended Actions

- Replace this IOM with an IOM that is compatible with this 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 (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.

---

## Event Descriptions 320-339

## 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.

---

## Event Descriptions 340-359

### 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.

## 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 event 359 has been logged for the indicated channel, indicating the condition no longer exists, 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](#).

### 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 event 359 has been logged for the indicated channel, indicating the condition no longer exists, 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](#).

## 359

### Info

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

#### Recommended Actions

- No action is required.

---

## Event Descriptions 360-379

## 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.

---

# Event Descriptions 380-399

Currently there are no events between 380 and 399.



## Chapter 5

# Event Descriptions 400-499

---

This chapter provides the following information:

- [Event Descriptions 400-419](#)
- [Event Descriptions 420-439](#)
- [Event Descriptions 440-459](#)
- [Event Descriptions 460-479](#)
- [Event Descriptions 480-499](#)

---

## Event Descriptions 400-419

### 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 event 37 was not logged, a spare of the proper type and size was not available for reconstruction. Replace the failed disk with one of the same type and the same or greater capacity and, if necessary, designate it as a spare. 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 spare (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.

---

## Event Descriptions 420-439

Currently there are no events between 420 and 439.

---

## Event Descriptions 440-459

### 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.

### 443

#### Error

The firmware for the indicated enclosure is not supported in this configuration.

The firmware for the indicated enclosure does not support this enclosure for use as an expansion chassis. Its firmware supports this enclosure only as a direct attached JBOD.

#### Recommended Actions

- Replace the indicated enclosure. It is not supported.

### 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.

---

## Event Descriptions 460-479

## 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

### Error

The indicated virtual pool reached its storage limit.

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.

### 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.

### Info

The indicated virtual pool exceeded one of its thresholds for allocated pages.

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

- No action is required for the low and mid thresholds. However, you may want to determine if your storage usage is growing at a rate that will result in the high threshold being crossed in the near future. If this will occur, either take steps to reduce storage usage or purchase additional capacity.

- If the high threshold is crossed, you should promptly 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

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The controller enclosure operating range is 5°C to 35°C (41°F to 95°F). The expansion enclosure operating range is 5°C to 40°C (41°F to 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.

When the problem is resolved, event 468 is logged.

## 470

### Warning

Removal of the indicated disk group(s) completed with failure.

Removal of disk groups can fail for several reasons, and the specific reason for this failure is included with the event. Removal most often fails because there is no longer room in the remaining pool space to move data pages off of the disks in the disk group.

#### Recommended Actions

- Resolve the issue specified by the error messaging included with this event and re-issue the request to remove the disk group.

### Info

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

#### Recommended Actions

- No action is required.

## 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.

## 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

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The controller enclosure operating range is 5°C to 35°C (41°F to 95°F). The expansion enclosure operating range is 5°C to 40°C (41°F to 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.

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

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The controller enclosure operating range is 5°C to 35°C (41°F to 95°F). The expansion enclosure operating range is 5°C to 40°C (41°F to 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.

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

---

## Event Descriptions 480-499

## 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 CompactFlash.
- 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 or global 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, 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 (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 CLI `show controllers` command or with the disk management utility.) 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 CLI `show redundancy-mode` 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 volume group started.

#### Recommended Actions

- No action is required.

## 489

### Info

Creation of a volume group completed.

#### Recommended Actions

- No action is required.

## 490

### Info

Creation of a volume group failed.

#### Recommended Actions

- No action is required.

## 491

### Info

Creation of a volume group started.

#### Recommended Actions

- No action is required.

## 492

### Info

The volumes in a volume group were ungrouped.

#### Recommended Actions

- No action is required.

## 493

### Info

A volume group was modified.

#### 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 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 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 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.

### Warning

An unsupported disk vendor was found.

#### Recommended Actions

- Replace the disk with a disk that is supported by your system vendor.



# Event Descriptions 500-599

---

This chapter provides the following information:

- [Event Descriptions 500-519](#)
- [Event Descriptions 520-539](#)
- [Event Descriptions 540-559](#)
- [Event Descriptions 560-579](#)
- [Event Descriptions 580-599](#)

---

## Event Descriptions 500-519

### 501

#### Error

The chassis 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/drive access was disabled in the chassis.

#### Recommended Actions

- Update the storage system to the latest firmware.

### 502

#### Warning

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

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

#### Recommended Actions

- Update the storage system to the latest firmware.

## 503

### Info

The Intelligent BackEnd Error Monitor (IBEEM) has discovered that continuous errors are being reported for the indicated PHY.

IBEEM logged this event after monitoring the PHY for 30 minutes.

#### Recommended Actions

- No action is required.

## 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 (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.

---

# Event Descriptions 520-539

## 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.

## 522

### Warning

A scrub-disk-group job encountered an error at the indicated logical block address.

The event message always includes the disk group name and the logical block address of the error within that disk group. If the block with an error falls within the LBA range used by a volume, the event message also includes the volume name and the LBA within that volume.

#### Recommended Actions

- Examine event 207 that was logged previously to this event. Follow the recommended actions for that event.

## 523

### Info

This event provides additional details associated with a scrub-disk-group job, expanding on the information in event 206, 207, or 522.

#### Recommended Actions

- Follow the recommended actions for the associated event.

## 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**

- Check that the storage system's fans are running.
- Check that the ambient temperature is not too warm. The controller enclosure operating range is 5°C to 35°C (41°F to 95°F). The expansion enclosure operating range is 5°C to 40°C (41°F to 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.

## 525

**Info**

A drawer has been stopped by a user.

The drawer has powered down and may be safely removed. A rescan must complete before the updated drawer information will be available.

**Recommended Actions**

- Restart the drawer using the start drawer command, or remove the drawer for replacement.

## 526

**Info**

A drawer has been started by a user.

The drawer has powered up. Disks in the drawer may take a few minutes to spin up. A rescan must complete before the updated drawer information will be available.

**Recommended Actions**

- No action is required.

## 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.

---

# Event Descriptions 540-559

## 542

## Critical

A data block was fenced by the controller due to lost data.

Event 543 will also be logged to describe volume information for the fenced data block.

### Recommended Actions

- Perform recovery procedures, which may include restoring from backups.

## 543

## Critical

A data block in a volume was fenced by the controller due to lost data.

This event describes volume information for a fenced data block. It is logged in conjunction with event 542, which describes disk group and disk information for the data block.

### Recommended Actions

- Perform recovery procedures, which may include restoring from backups.

## 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 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. In this case, the reconstruction operation failed because unreadable data (uncorrectable media error) exists in at least one other disk in the disk group. Because of this, a portion of the data cannot be reconstructed.

#### Recommended Actions

- 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 to a new disk group.

## 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.

## 551

### Error

An EMP reported one of the following for a power supply unit (PSU):

- The PSU is unable to communicate with the EMP.
- The PSU in an enclosure does not have power supplied to it or has a hardware failure.
- The PSU is running with corrupted firmware.

#### Recommended Actions

- If the EMP is unable to communicate with the indicated PSU:
  - Wait for at least 10 minutes and check if the error resolves.
  - If the error persists, check that all modules in the enclosure are fully seated in their slots and that their latches, if any, are locked.
  - If this does not resolve the issue, note down the PSU. Ensure the partner PSU is not degraded. If the partner PSU is degraded, contact technical support.
  - If the partner PSU is not degraded, remove and reinsert the indicated PSU.
  - If this does not resolve the issue, the indicated FRU has probably failed and should be replaced.
- If one of the PSUs in an enclosure does not have power supplied to it or has a hardware failure:
  - Check that the indicated PSU is fully seated in its slot and that the PSU's latches, if any, are locked.
  - Check that each PSU has its switch turned on (if equipped with a switch).
  - Check that each power cable is firmly plugged into both the PSU and a functional electrical outlet.
  - If none of the above resolves the issue, the indicated PSU has probably failed and should be replaced.
- If a PSU is running with corrupted firmware:
  - The indicated PSU has failed and should be replaced.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

### Warning

An EMP reported that a power supply unit (PSU) has been uninstalled.

#### Recommended Actions

- Check that the indicated PSU is in the indicated enclosure.
- If the PSU is not in the enclosure, install a PSU immediately.

- If the PSU is in the enclosure, ensure that the power supply is fully seated in its slot and that its latch is locked.
- If none of the above resolves the issue, the indicated FRU has failed and should be replaced.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Resolved

A SES alert for a power supply in the indicated enclosure has been resolved.

### Recommended Actions

- No action is required.

## 552

## Error

An EMP reported an alert condition.

- A hardware failure has been detected and all fans in the indicated FRU have failed.
- The fan is unable to communicate with the EMP.

### Recommended Actions

- If a hardware failure has been detected and all fans in the indicated FRU have failed.
  - Inspect the system health information to determine which FRU contains the affected fans. Event 551 or 558 should give further information on the containing FRUs.
  - Replace the containing FRUs.
- If the fan is unable to communicate with the EMP.
  - Wait for at least 10 minutes and check if the error resolves.
  - If the error persists, check that all modules in the enclosure are fully seated in their slots and that their latches, if any, are locked.
  - If this does not resolve the issue, note down the FRU. Ensure the partner FRU is not degraded. If the partner FRU is degraded, contact technical support.
  - If the partner FRU is not degraded, remove and reinsert the indicated FRU.
  - If this does not resolve the issue, the indicated FRU has probably failed and should be replaced.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Warning

An EMP reported one of the following:

- A fan in the indicated FRU has been uninstalled.
- A fan in the indicated FRU has failed and fan redundancy for the FRU has been lost.

### Recommended Actions

If a fan in the indicated FRU has been uninstalled.

- Check that the indicated FRU is in the indicated enclosure.
- If the FRU is not in the enclosure, install the appropriate FRU immediately.
- If the FRU is in the enclosure, ensure that the FRU is fully seated in its slot and that its latch is locked.

- If none of the above resolves the issue, the indicated FRU has failed and should be replaced.

If a fan in the indicated FRU has failed and fan redundancy for the FRU has been lost:

- The indicated FRU has failed and should be replaced.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Resolved

A SES alert for a fan in the indicated enclosure has been resolved.

### Recommended Actions

- No action is required.

## 553

## Error

A temperature sensor reported an alert condition.

- A temperature sensor is outside critical temperature threshold in the indicated FRU.
- The temperature sensor is not able to communicate with the EMP.

### Recommended Actions

- If temperature sensor is outside critical temperature threshold in the indicated FRU.
  - Check that the ambient temperature is not too warm. For the normal operating range, see your product's Hardware Installation and Maintenance Guide.
  - Check for any obstructions to the airflow.
  - Check that all modules in the enclosure are fully seated in their slots and that their latches, if any, are locked.
  - Check that all fans in the enclosure are running.
  - 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 temperature sensor is not able to communicate with the EMP.
  - Wait for at least 10 minutes and check if the error resolves.
  - If the error persists, check that all modules in the enclosure are fully seated in their slots and that their latches, if any, are locked.
  - If this does not resolve the issue, note down the FRU. Ensure the partner FRU is not degraded. If the partner FRU is degraded, contact technical support.
  - For all FRU types except the enclosure, if the partner FRU is not degraded, remove and reinsert the indicated FRU.
  - If the indicated FRU is the enclosure, set up a preventive maintenance window and power cycle the enclosure at that time.
  - If this does not resolve the issue, the indicated FRU has probably failed and should be replaced.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Warning

A temperature sensor is not within normal operating temperature thresholds but is within safe operating limits; or, a temperature sensor has been uninstalled.

### Recommended Actions

- If a temperature sensor has exceeded the normal operating range but is within safe operating limits.
  - Check that the ambient temperature is not too warm. For the normal operating range, see your product's Hardware Installation and Maintenance Guide.
  - Check for any obstructions to the airflow.
- If a temperature sensor has been uninstalled:
  - Check that the indicated FRU is in the indicated enclosure.
- If the FRU is not in the enclosure, install the FRU immediately.
- If the FRU is in the enclosure, ensure that the FRU is fully seated in its slot and that its latches, if any, are locked.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Resolved

A SES alert for a temperature sensor in the indicated enclosure has been resolved.

### Recommended Actions

- No action is required.

## 554

## Error

A voltage sensor reported an alert condition.

- A voltage sensor is outside a critical voltage threshold in the indicated FRU.
- A voltage sensor is not able to communicate with the EMP.

### Recommended Actions

- If a voltage sensor is outside a critical voltage threshold in the indicated FRU.
  - Check that all modules in the enclosure are fully seated in their slots and that their latches, if any, are locked.
  - If this does not resolve the issue, the indicated FRU has probably failed and should be replaced.
- If the voltage sensor is not able to communicate with the EMP.

- Wait for at least 10 minutes and check if the error resolves.
- If the error persists, check that all modules in the enclosure are fully seated in their slots and that their latches, if any, are locked.
- If this does not resolve the issue, note down the FRU. Ensure the partner FRU is not degraded. If the partner FRU is degraded, contact technical support.
- For all FRU types except the enclosure, if the partner FRU is not degraded, remove and reinsert the indicated FRU.
- If the indicated FRU is the enclosure, set up a preventive maintenance window and power cycle the enclosure at that time.
- If this does not resolve the issue, the indicated FRU has probably failed and should be replaced.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Warning

A voltage sensor is not within the normal operating range but is within safe operating limits; or, a voltage sensor has been removed.

### Recommended Actions

If a voltage sensor has exceeded the normal operating range but is within safe operating limits:

- Check that all modules in the enclosure are fully seated in their slots and that their latches are locked.
- If this does not resolve the issue, the indicated FRU has probably failed and should be replaced.

If a voltage sensor has been removed:

- Check that the indicated FRU is in the indicated enclosure.
- If the FRU is not in the enclosure, install the FRU immediately.
- If the FRU is in the enclosure, ensure that the FRU is fully seated in its slot and that its latches are locked.
- If this does not resolve the issue, the indicated FRU has probably failed and should be replaced.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Resolved

A SES alert for a voltage sensor in the indicated enclosure has been resolved.

### Recommended Actions

- No action is required.

## 555

## Error

The local Expander Controller firmware has detected a level of incompatibility with the partner Expander Controller firmware or hardware. As a preventive measure, the local Expander Controller may disable all the PHYs.

### Recommended Actions

- Check that both the Expander Controllers have the correct firmware revision.
- If both Expander Controllers have different firmware versions, upgrade the partner controller module to the appropriate firmware that is compatible with the enclosure.

- If this does not resolve the problem, replace the partner controller module.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Warning

An expander in a controller module, expansion module, or drawer is mated but is not responding; or, an expander in an expansion module has been removed.

### Recommended Actions

- Check that the indicated FRU is in the indicated enclosure.
- If the FRU is not in the enclosure, install the appropriate FRU immediately.
- If the FRU is in the enclosure, ensure that the FRU is fully seated in its slot and that its latches, if any, are locked.
- If none of the above resolves the issue, the indicated FRU has failed and should be replaced.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Resolved

A SES alert for an expander in the indicated enclosure has been resolved.

### Recommended Actions

- No action is required.

## 556

## Error

An alert condition was detected on an root expander or drawer expander element.

### Recommended Actions

- Replace the module that contains the indicated expander. This could be an IOM, sideplane or a drawer.

Contact technical support for replacement of the module containing the drawer expander.

- If this does not resolve the issue, contact technical support. The enclosure must be replaced.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Warning

An alert condition was detected on an root expander or drawer expander element.

### Recommended Actions

- If uninstalled, the expander associated with the sideplane or drawer will have to be installed. Contact technical support. Otherwise, replace the module that contains the indicated expander. This could be a sideplane or a drawer.

Contact technical support for replacement of the module containing the drawer expander.

- If this does not resolve the issue, contact technical support. The enclosure must be replaced.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Resolved

A SES alert for an expander in the indicated enclosure has been resolved.

### Recommended Actions

- No action is required.

## 557

## Error

An Enclosure Management Processor (EMP) reported an alert condition on a current sensor.

- The EMP is unable to communicate with the indicated current sensor.
- The current sensor is outside critical threshold values.

### Recommended Actions

- If the EMP is unable to communicate with the indicated current sensor.
  - Wait for at least 10 minutes and check if the error resolves.
  - If the error persists, check that all modules in the enclosure are fully seated in their slots and that their latches, if any, are locked.
  - If this does not resolve the issue, note down the FRU. Ensure the partner FRU is not degraded. If the partner FRU is degraded, contact technical support.
  - For all FRU types except the enclosure, if the partner FRU is not degraded, remove and reinsert the indicated FRU.
  - If the indicated FRU is the enclosure, set up a preventive maintenance window and power cycle the enclosure at that time.
  - If this does not resolve the issue, the indicated FRU has probably failed and should be replaced.
- If the current sensor is outside critical threshold values.
  - Check that all modules in the enclosure are fully seated in their slots and that their latches, if any, are locked.
  - If this does not resolve the issue, the indicated FRU has probably failed and should be replaced.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Warning

An Enclosure Management Processor (EMP) reported an alert condition on a current sensor.

- A current sensor is outside the defined warning threshold values.
- A current sensor has been uninstalled.

### Recommended Actions

- If a current sensor has exceeded the defined warning threshold values.
  - Check that all modules in the enclosure are fully seated in their slots and that their latches, if any are locked.
  - If this does not resolve the issue, the indicated FRU has probably failed and should be replaced.
- If a current sensor has been uninstalled.

- Check that the indicated FRU is in the indicated enclosure.
- If the FRU is not in the enclosure, install the FRU immediately.
- If the FRU is in the enclosure, ensure that the FRU is fully seated in its slot and that its latches, if any, are locked.
- If this does not resolve the issue, the indicated FRU has probably failed and should be replaced.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Resolved

A SES alert for a current sensor in the indicated enclosure has been resolved.

### Recommended Actions

- No action is required.

---

## Event Descriptions 560-579

### 562

#### Info

Virtual pool statistics have been reset.

### Recommended Actions

- No action is required.

### 563

#### Info

A disk has been restarted.

### Recommended Actions

- No action is required.

### 565

#### Warning

One of the PCIe buses is running at less than optimal speed.

This event is the result of a hardware problem that has caused the controller to run slower than expected. The system works, but I/O performance is degraded.

### Recommended Actions

- Replace the controller module that logged this event.
- If the problem persists, replace the controller module.

## 566

### Info

One of the DDR ports has been busy for at least 5 minutes.

This event is the result of a speed compensation while handling short data blocks. The system is operational but I/O performance is degraded.

#### Recommended Actions

- No action is required.

## 568

### Info

A disk group has mixed physical sector size disks (for example 512n and 512e disks in the same disk group).

This event is the result of the user selecting disks with sector formats that do not match or a global spare replacement with a different sector format than the disk group. This could result in degraded performance for some workloads.

#### Recommended Actions

- No action is required.

## 571

### Error

Allocated snapshot space exceeded the configured percentage limit of the virtual pool.

If the snapshot space limit policy is set to delete snapshots, the system will begin to delete snapshots according to the snapshot retention priority setting until the snapshot space usage drops below the configured limit. Otherwise, the system will begin to use general pool space for snapshots until snapshots are manually deleted. If the storage usage drops below a threshold, event 572 is logged.

#### Recommended Actions

- If the snapshot space limit policy is set to notify only, you should immediately take steps to reduce snapshot space usage or add storage capacity.
- If the snapshot space policy is set to delete, the system will reduce snapshot space automatically, or log event 573 if no snapshots can be deleted.

### Warning

Allocated snapshot space exceeded the high snapshot space threshold.

The high threshold setting indicates that the pool is nearly out of snapshot space. The threshold settings are intended to indicate that the pool is using a significant portion of configured snapshot space and should be monitored. If the storage usage drops below any threshold, event 572 is logged.

#### Recommended Actions

- Reduce the snapshot space usage by deleting snapshots that are no longer needed.

## Info

Allocated snapshot space exceeded either the low or middle snapshot space threshold.

The threshold settings are intended to indicate that the pool is using a significant portion of configured snapshot space and should be monitored. If the storage usage drops below any threshold, event 572 is logged.

### Recommended Actions

- Reduce the snapshot space usage by deleting snapshots that are no longer needed.

## 572

## Info

The indicated virtual pool has dropped below one of its snapshot space thresholds.

This event indicates that a condition reported by event 571 is no longer applicable.

### Recommended Actions

- No action is required.

## 573

## Warning

Allocated snapshot space for a virtual pool cannot be reduced because no snapshots are deletable.

Allocated snapshots cannot be automatically deleted if their retention priority is set to never-delete. Snapshots must also be at the leaf end of a snapshot tree in order to be considered deletable. This event is logged when no snapshots in the pool pass these constraints.

### Recommended Actions

- Manually delete snapshots to reduce allocated snapshot space.

## 574

## Info

A peer connection was created.

### Recommended Actions

- No action is required.

## 575

## Info

A peer connection was deleted.

### Recommended Actions

- No action is required.

## 576

### Info

A replication set was created, or a replication set failed to be created.

#### Recommended Actions

- No action is required.

## 577

### Error

A replication set failed to be deleted.

#### Recommended Actions

- No action is required.

### Info

A replication set was deleted.

#### Recommended Actions

- No action is required.

## 578

### Error

A replication failed to start.

The replication was unsuccessful due to the condition specified within the event. Reasons for replication failure include but are not limited to shutdown of the secondary system, a loss of communication across the peer connection (which may be due to CHAP configuration changes), or a pool out-of-space condition.

#### Recommended Actions

- Resolve the issue specified by the error message included with this event.

### Info

A replication was started.

#### Recommended Actions

- No action is required.

## 579

### Warning

A replication completed with failure.

The replication was unsuccessful due to the condition specified within the event. Reasons for replication failure include but are not limited to shutdown of the secondary system, a loss of communication across the peer connection (which may be due to CHAP configuration changes), or a pool out of space condition.

#### Recommended Actions

- Resolve the issue specified by the error message included with this event.

### Info

A replication completed successfully.

#### Recommended Actions

- No action is required.

---

## Event Descriptions 580-599

## 580

### Info

A replication was aborted.

#### Recommended Actions

- No action is required.

## 581

### Warning

A replication was suspended internally by the system.

The system will suspend the replication internally if it detects an error condition in the replication set and replications cannot continue for any reason. This includes but is not limited to shutdown of the secondary system, a loss of communication across the peer connection (which may be due to CHAP configuration changes), or a pool out-of-space condition.

#### Recommended Actions

- The replication will automatically resume once the condition described in this event is cleared.

### Info

A replication was suspended by the user.

#### Recommended Actions

- No action is required.

## 582

### Info

A replication has queued behind the active replication.

#### Recommended Actions

- No action is required.

## 584

### Info

A peer connection was modified.

#### Recommended Actions

- No action is required.

## 585

### Info

A replication set was modified.

#### Recommended Actions

- No action is required.

## 586

### Error

Resuming the replication was unsuccessful due to the condition specified within the event. Reasons for replication failure include but are not limited to shutdown of the secondary system, a loss of communication across the peer connection (which may be due to CHAP configuration changes), or a pool out-of-space condition.

#### Recommended Actions

- Resolve the issue specified by the error message included with this event.

### Info

A replication was resumed.

#### Recommended Actions

- No action is required.

## 587

### Info

A pending replication was removed from the queue.

#### Recommended Actions

- No action is required.

## 590

### Error

A disk group has been quarantined.

This condition resulted from a controller flush/restore failure.

#### Recommended Actions

- To restore the disk group, use the CLI dequarantine command to dequarantine the disk group. If more than one disk group is quarantined you must individually dequarantine each disk group, whether it is fault tolerant or not. When dequarantine is complete, the disk group will return to the state it was in before being quarantined. For example, if the disk group was reconstructing before being quarantined, the disk group will resume reconstructing where it stopped.
- For a linear disk group, if you want to find where parity is incorrect, use the CLI scrub vdisk command with the fix parameter disabled. This step is optional and not required to fix data integrity issues.
- For a fault tolerant disk group, run either scrub disk-groups for a virtual disk group or scrub vdisk with the fix parameter enabled for a linear disk group. This step will make the parity consistent with the existing user data, and is required to fix data integrity issues.
- For a reconstructing disk group, let reconstruction finish, then run either scrub disk-groups for a virtual disk group or scrub vdisk with the fix parameter enabled for a linear disk group. This step will make the parity consistent with the existing user data, and is required to fix data integrity issues.
- Restore the data to the disk group from a backup copy.

## 593

### Info

A PCIe bus has transitioned to a different speed.

#### Recommended Actions

- No action is required.

## 594

### Info

The indicated disk in the indicated disk group is missing and the disk group is quarantined. While the disk group is quarantined, in linear storage any attempt to access its volumes from a host will fail. In virtual storage, all volumes in the pool will be forced read-only. If all of the disks become accessible, the disk group will be dequarantined automatically with a resulting status of FTOL. If not all of the

disks become accessible but enough become accessible to allow reading from and writing to the disk group, it will be dequarantined automatically with a resulting status of FTDN or CRIT. 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 dequarantine, see the WBI or CLI documentation.

---

**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.

---

---

**CAUTION:** 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.

## 595

### Info

This event reports the serial number of each controller module in this system.

### Recommended Actions

- No action is required.

## 598

### Info

Drive has failed a performance measurement.

#### Recommended Actions

- Monitor the disk.

### Warning

Drive has failed a performance measurement.

#### Recommended Actions

- Monitor the disk.

## 599

### Error

The firmware has yet to retrieve Enclosure Power control status.

The Enclosure Power element provides enclosure level power control. This could occur shortly after a reboot or module insertion. It should only be treated as an error if it persists for more than 30 seconds after a reset.

#### Recommended Actions

- Contact technical support.
- When the problem is resolved, an event with the same code will be logged with Resolved severity.

### Warning

The firmware has yet to retrieve Enclosure Power control status.

The Enclosure Power element provides enclosure level power control.

#### Recommended Actions

- No action is required.

### Resolved

The firmware has yet to retrieve Enclosure Power control status.

The Enclosure Power element provides enclosure level power control.

#### Recommended Actions

- No action is required.



This chapter provides the following information:

- [Event Descriptions 600-619](#)
- [Event Descriptions 620-699](#)

---

## Event Descriptions 600-619

### 602

#### **Error, Warning**

An alert condition was detected on a Midplane Interconnect element.

The Midplane Interconnect element reports status associated with the interface between the SBB I/O module and the midplane. This is typically some form of communication problem on the midplane interconnect.

#### **Recommended Actions**

- Contact technical support. Provide logs to technical support personnel for analysis.
- When the problem is resolved, an event with the same code will be logged with Resolved severity.

### 603

#### **Error, Warning**

An alert condition for a SAS Connector element has been detected.

The SAS Connector element report status information for both external and internal SAS port connectors.

#### **Recommended Actions**

- Contact technical support.
- When the problem is resolved, an event with the same code will be logged with Resolved severity.

## Info

An alert condition for a SAS Connector element has been detected.

The SAS Connector element report status information for both external and internal SAS port connectors.

### Recommended Actions

- No action is required.

## Resolved

An alert condition for a SAS Connector element has been detected.

The SAS Connector element report status information for both external and internal SAS port connectors.

### Recommended Actions

- No action is required.

## 604

### Warning

A replication snapshot was attempted and failed.

A replication-set has been configured to retain snapshots of the volume. An error is possible if the snapshot fails.

### Recommended Actions

- Monitor the health of the local system, the replication-set, the volume, and the peer-connection. A full storage pool may be the cause of this fault.
  - Check the peer-connection system health and state.
  - Ensure that the Maximum Licensable Snapshots limit (shown by the CLI show license command) was not exceeded.

## 605

### Warning

Inactive processing core.

The controller module has multiple processing cores. The system has enough active cores to operate but performance is degraded.

### Recommended Actions

- Attempt to restart all the processing cores as follows:
  - Shut down the controller module that logged this event.
  - Remove the controller module, wait 30 seconds, and then reinsert the controller module.
- If this event is logged again, contact technical support.

## 606

### Error

A controller contains unwritten cache data for a volume, and its supercapacitor has failed to charge. Due to the supercapacitor failure, if the controller loses power, it will not have backup power to flush the unwritten data from cache to CompactFlash.

#### Recommended Actions

- Verify that the cache-write policy is write-through for all volumes.
- Contact technical support for information about replacing the controller module.

## 607

### Warning

The local controller is rebooting the other controller.

#### Recommended Actions

- No action is required.

## 608

### Error

A back-end cabling error was detected.

#### Recommended Actions

- If the message says both controllers are connected with a undefined error type, one of the cables is incorrectly connected to a controller egress port forming a loop in the SAS topology. Check back-end cabling from each controller egress port to determine the incorrect connection.
- If the message says controller egress ports are connected to each other, one of the cables is incorrectly connected to a controller egress port forming a loop in the SAS topology. Check back-end cabling and make sure that SAS cables are connected to the correct ports for the port specified.
- If the message says an EBOD loop has been created, one of the cables is incorrectly connected to an expansion enclosure egress port forming a loop in the SAS topology. Check back-end cabling and make sure that SAS cables are connected to the correct ports for the port specified.
- If the message says a cable is connected to the middle port but that port is not supported, check back-end cabling and make sure that SAS cables are connected to the correct ports for the port specified. Move the cable from the middle port of the IOM to the left or right port, as appropriate.

## 609

### Error

An alert condition was detected on a door lock element. The door lock element reports status associated with the enclosure drawer. The drawer has been reporting as open for a long period of time. This may reduce cooling, potentially causing the enclosure to overheat.

**Recommended Actions**

- Check that the drawer is fully closed and latched.

When the problem is resolved, an event with the same code will be logged with Resolved severity.

**Info**

An alert condition was detected on a door lock element. The door lock element reports status associated with the enclosure drawer. The drawer sensor is reporting uninstalled.

**Recommended Actions**

- No action is required.

**Resolved**

A previous Informational or Error condition for the door lock element has been resolved.

**Recommended Actions**

- No action is required.

**610****Error**

An alert condition was detected on a sideplane element.

**Recommended Actions**

- Check the drawer that the indicated sideplane is in fully closed and latched.
- If this does not resolve the issue, contact technical support. The enclosure must be replaced.

**Warning**

An alert condition was detected on a sideplane element.

**Recommended Actions**

- The sideplane associated with the drawer must be installed. Contact technical support.

**Resolved**

A previous Warning or Error condition for the sideplane element has been resolved.

**Recommended Actions**

- No action is required.

**611****Error**

Email notification failed due to either:

- An unreachable SMTP server or a difference between the sender and SMTP server domains.
- Improper configuration.

**Recommended Actions**

- Please verify the configured parameters and ask the recipients to confirm that they received the message.

**Info**

Email notification sent successfully. Please ask the recipients to confirm that they received the message.

**Recommended Actions**

- Please verify the configured parameters and ask the recipients to confirm that they received the message.

**612****Info**

An alert condition was detected on an internal chassis SAS connector.

The event message specifies the location of the internal SAS connector in the chassis.

**Recommended Actions**

- No action is required.

**613****Error**

An alert condition was detected on an IOM.

**Recommended Actions**

- Either install the indicated IOM or attempt to reseal it.
- If this does not resolve the issue, replace the IOM.

**Warning**

An alert condition was detected on an IOM.

**Recommended Actions**

- If uninstalled, install the indicated IOM otherwise attempt to reseal it.
- If this does not resolve the issue, replace the IOM.

**Info**

An IOM was uninstalled.

**Recommended Actions**

- No action is required.

## Resolved

A previous Warning or Error condition for the IOM has been resolved.

### Recommended Actions

- No action is required.

## 615

### Info

A rebalance operation for an ADAPT disk group has started.

### Recommended Actions

- No action is required.

## 616

### Warning

A rebalance operation for an ADAPT disk group has completed partially (because of drive failure).

### Recommended Actions

- No action is required.

### Info

A rebalance operation for an ADAPT disk group has completed.

### Recommended Actions

- No action is required.

## 617

### Warning

The spare capacity goal is not met.

This event indicates that the available space in the system is insufficient to provide the level of full fault tolerance that is specified by the target spare capacity. Spare capacity availability can be influenced by operations that require available space in the system, such as reconstructing data from a failed disk.

### Recommended Actions

- Add disks to the disk group, or replace any disks that may have failed. The system will automatically increase the spare capacity to meet the requirements placed on the system by the target spare capacity.

## 618

### Resolved

The spare capacity goal is met.

#### Recommended Actions

- No action is required.

## 619

### Error

Disk access will change depending on the port used to connect to the expander.

#### Recommended Actions

- Load a valid firmware bundle to disable zoning.

### Info

The controller has been injected with a fault to introduce a broadcast receiver (BR) link error.

#### Recommended Actions

- No action is required.

### Resolved

Expander zoning has been disabled for the indicated enclosure.

#### Recommended Actions

- No action is required.

## 626

### Info

Detected an unsupported TPID (midplane Type ID).

#### Recommended Actions

- No action is required.

## 627

### Info

Detected an unknown TPID (midplane Type ID).

#### Recommended Actions

No action is required.

## 628

### Error

A firmware mismatch has been identified for an expansion enclosure. A firmware mismatch could result from attaching an enclosure configured as a JBOD (instead of an EBOD) or installing a new IOMFRU with incompatible firmware.

#### Recommended Actions

Upgrade the firmware to the appropriate level for connecting your expansion enclosures to the controller enclosure. If you receive this event when no new enclosures or IOMs have been added, please contact support.

---

## Event Descriptions 620-699

Currently there are no event descriptions for 620-699.



## Chapter 8

# Additional Information

---

This chapter provides the following information:

- [Deprecated Events](#)
- [Removed Events](#)
- [Using the trust Command](#)
- [Events Sent as Indications to SMI-S Clients](#)

---

## Deprecated Events

No events are deprecated.

---

## Removed Events

The following table lists events that have been removed and specifies other events that the system will report instead, if any. If you have scripts that refer to removed events, update the scripts.

**Table 3** Removed Events

Removed event	Replacement event
168	551-556
169	551-556 with Resolved Severity

---

## 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/drive is in a leftover or failed condition, the trust command should never be used. The trust command can only be used if the disk group is in an **Offline** state.

## Single Disk/Drive Failed

If a single disk/drive in a disk group has failed or been placed into a Leftover state due to errors, reintegrating the disk/drive into the same or a different disk group has the potential to cause data loss. A disk/drive that has failed or been placed into a Leftover state due to multiple errors should be replaced with a new disk/drive. Assign the new disk/drive 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

## Resynchronize Leftover Disks/Drives

The trust command attempts to resynchronize leftover disks/drives in order to make any leftover disk/drive 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/drive 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. If you are uncertain whether to use this command, contact technical support for assistance.

---

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.

---

## 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 *QXS 12G Disk Management Utility User Guide*.

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

**Table 4** 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

**Table 4** Events and Corresponding SMI-S Indications

<b>FRU/Event category</b>	<b>Corresponding SMI-S class</b>	<b>Operation status values that would trigger alert conditions</b>
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

