PowerScale OneFS

Event Reference Guide

April 2022

DELL Technologies

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Contents

1

Introduction to this guide

This section contains the following topics:

Topics:

About this guide

This guide describes the OneFS event notification system, provides a list of all event IDs, and explains how to respond to events that might impact the overall health or performance of the cluster.

Your suggestions help us to improve the accuracy, organization, and overall quality of the documentation. Send your feedback to http://bit.ly/isilon-docfeedback. If you cannot provide feedback through the URL, send an email message to docfeedback@dell.com.

Scale-out NAS overview

The scale-out NAS storage platform combines modular hardware with unified software to harness unstructured data. Powered by the OneFS operating system, a cluster to deliver a scalable pool of storage with a global namespace.

The unified software platform provides centralized web-based and command-line administration to manage the following features:

- A cluster that runs a distributed file system
- Scale-out nodes that add capacity and performance
- Storage options that manage files and tiering
- Flexible data protection and high availability
- Software modules that control costs and optimize resources.

Where to get help

The Dell Technologies Support site (https://www.dell.com/support) contains important information about products and services including drivers, installation packages, product documentation, knowledge base articles, and advisories.

A valid support contract and account might be required to access all the available information about a specific Dell Technologies product or service.

Additional options for getting help

This section contains resources for getting answers to questions about PowerScale products.

Dell Technologies support	• https://www.dell.com/support/incidents-online/en-us/contactus/product/isilon-onefs
Telephone support	• United States: 1-800-SVC-4EMC (1-800-782-4362) • Canada: 1-800-543-4782 • Worldwide: 1-508-497-7901

	• Local phone numbers for a specific country or region are available at https://www.dell.com/support/incidents-online/en-us/contactus/product/isilon-onefs.
PowerScale OneFS Documentation Info Hubs	• https://www.dell.com/support/kbdoc/en-us/000152189/powerscale-onefs-info-hubs
Dell Community Board for self-help	• https://www.dell.com/community

2

Introduction to system events

This section contains the following topics:

Topics:

Events overview

Events are individual occurrences or conditions related to the data workflow, maintenance operations, and hardware components of your cluster.

Throughout OneFS there are processes that are constantly monitoring and collecting information on cluster operations.

When the status of a component or operation changes, the change is captured as an event and placed into a priority queue at the kernel level.

Every event has two ID numbers that help to establish the context of the event:

- The event type ID identifies the type of event that has occurred.
- The event instance ID is a unique number that is specific to a particular occurrence of an event type. When an event is submitted to the kernel queue, an event instance ID is assigned. You can reference the instance ID to determine the exact time that an event occurred.

You can view individual events. However, you manage events and alerts at the event group level.

Event groups overview

Event groups are collections of individual events that are related symptoms of a single situation on your cluster. Event groups provide a single point of management for multiple event instances that are generated in response to a situation on your cluster.

For example, if a chassis fan fails in a node, OneFS might capture multiple events related both to the failed fan itself, and to exceeded temperature thresholds within the node. All events related to the fan will be represented in a single event group. Because there is a single point of contact, you do not need to manage numerous individual events. You can handle the situation as a single, coherent issue.

All management of events is performed at the event group level. You can mark an event group as resolved or ignored. You can also configure how and when alerts are distributed for an event group.

Alerts overview

An alert is a message that describes a change that has occurred in an event group.

At any point in time, you can view event groups to track situations occurring on your cluster. You can also create alerts to proactively notify you when there is a change in an event group. For example, you can generate an alert when a new event is added to an event group, when an event group is resolved, or when the severity of an event group changes.

You can adjust the thresholds at which certain events raise alerts. For example, by default, OneFS generates an alert when a disk pool is 95% full. You can adjust that threshold to a lower percentage.

You can configure your cluster to generate alerts only for specific event groups, conditions, severity, or during limited time periods.

Alerts are delivered through channels. You can configure a channel to determine who will receive the alert and when.

Alert channel overview

Alert channels are pathways by which event groups send alerts.

When an alert is generated, the channel that is associated with the alert determines how the alert is distributed and who receives the alert.

You can configure an alert channel to deliver alerts with one of the following mechanisms: SMTP, SNMP, or Connect Home. You can also specify the required routing and labeling information for the delivery mechanism.

Viewing and modifying event groups

You can view event and modify the status of event groups.

View an event group

You can view the details of an event group.

1. Click **Cluster Management** > **Events and Alerts**.
   The **Event group history** tab summarizes the list of all the event groups, and you can customize the list as needed.
   - You can filter the data by date range, event group status, and event group severity.
   - You can search for relevant event groups by entering the search string in the search box.
2. In the **Actions** column of the event group you want to view, click **View event details**.
   You can view details of each event group in a separate window.

Change the status of an event group

You can ignore or resolve an event group.

After you resolve an event group, you cannot reverse that action. Any new events that would have been added to the resolved event group will be added to a new event group.

1. Click **Cluster Management** > **Events and Alerts**.
   The **Event group history** tab summarizes the list of all the event groups, and you can customize the list as needed.
   - You can filter the data by date range, event group status, and event group severity.
   - You can search for relevant event groups by entering the search string in the search box.
2. In the **Actions** column of the event group you want to change, click **Actions**.
3. In the menu that appears, click **Resolve event** to resolve the event group or **Ignore event** to ignore the event group.

   NOTE: You can perform an action on multiple event groups by selecting the check box next to the Event group description of the events that you want to change, then selecting an action from the **Select a bulk action** list.
4. Click **Mark Resolved** or **Ignore** to confirm the action.

View an event

You can view the details of a specific event.

1. Click **Cluster Management** > **Events and Alerts**.
   The **Event group history** tab summarizes the list of all the event groups, and you can customize the list as needed.
   - You can filter the data by date range, event group status, and event group severity.

- You can search for relevant event groups by entering the search string in the search box.

2. In the **Actions** column of the event group that contains the event you want to view, click **View event details**.
3. In the new window, click **+See event instance details** to expand the list of events.
4. Click **View details** next to the event whose details you want to view.

Managing alerts

You can view, create, modify, or delete alerts to determine the information you deliver about event groups.

View maintenance mode

You can view the current maintenance mode status.

Click **Cluster Management** > **Events and Alerts** > **Alert Management**.

Enable maintenance mode

You can enable the CELOG maintenance mode.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. Click **Enable CELOG maintenance mode**.
   A warning message similar to the following appears:

   ```
   Are you sure you want to enable CELOG maintenance mode?
   ```

3. Click **Enable CELOG maintenance mode**.
   The CELOG maintenance mode is enabled and the maintenance window duration to date appears.

Disable maintenance mode

You can disable the CELOG maintenance mode. While disabling, you can view all the events that have occurred during the maintenance mode and clear the details, if needed.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. Click **Disable CELOG maintenance mode**.
   The **Disable CELOG maintenance mode** dialog box with the following details appear:
   - CELOG maintenance window start date and time
   - Duration to date of the maintenance window
   - Details of events that occurred during the maintenance mode
3. Click **Disable CELOG maintenance mode**.
   The CELOG maintenance mode is disabled.

View maintenance history

You can view the maintenance history. Maintenance window history is controlled by the event retention policy for the cluster, which is controlled in the **Settings** tab.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. In the **Maintenance window history** area, view the details.

View alerts by event type id

You can view a list of alerts with their event type IDs with description, category, and associated alert rules and alert channels.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. In the **CELOG alerting** area, view the details of all alerts.

Modify alerts by event type id

You can suppress or un-suppress one or more event type ID depending on its current state.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. In the Actions column of the event type ID you want to modify, click **Suppress** or **Un-suppress**.

   You can perform an action on multiple event type IDs by selecting the check box next to the event type ID of the alerts you want to change, then selecting an action from the **Select a bulk action** list.

View alerting rules

You can view a list of all the alerting rules.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. In the **CELOG alerting** area, click the **Alerting rule** tab.
   You can view the list of alerting rules.

Create an alerting rule

You can create an alerting rule.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. In the **CELOG alerting** area, click the **Alerting rule** tab.
3. Click **Create alert rule**.
   The **Create alert rule** window appears.
4. Enter the following details:
   - Rule name: Enter a name for the new alerting rule.
   - Rule condition: Select a condition (New, New Events, Ongoing, Severity Increase, Severity decrease, Resolved) from the drop-down list.
   - Send an alert only if the event lasts longer than: Enter the numerical value in the text box and select the unit of time from the drop-down list.
   - Applies to: Select the check box next to the relevant alert category.
   - Add event group ID: Click **Add event group ID** to add an event group.
   - Select alert channel for this rule: Select the check box next to the relevant channel name.
5. Click **Create rule**.

Modify an alerting rule

You can modify an existing alerting rule.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. In the **CELOG alerting** area, click the **Alerting rule** tab.
3. In the Actions column of the alerting rule you want to modify, click **Edit rule**.
   The **Edit alert rule** window appears.
4. Modify the following details:
   - Rule name: You cannot modify the name for an existing alerting rule.
   - Rule condition: Select a condition (New, New Events, Ongoing, Severity Increase, Severity decrease, Resolved) from the drop-down list.
   - Send an alert only if the event lasts longer than: Modify the numerical value in the text box and select the unit of time from the drop-down list.
   - Applies to: Select the check box next to the relevant alert category.
   - Add event group ID: Click **Add event group ID** to add a new event group.
   - Select alert channel for this rule: Select the check box next to the relevant channel name.
5. Click **Save Changes**.

Delete an alerting rule

You can delete an existing alerting rule.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. In the **CELOG alerting** area, click the **Alerting rule** tab.
3. In the Actions column of the alerting rule, click **Edit rule**.
   The **Edit alert rule** window appears.
4. Click **Delete** at the bottom of the window.
   The **Confirm delete alert rule** dialog box appears with the following message:

   ```
   This action can not be undone. Are you sure you want to delete this alert rule?
   ```

5. Click **Confirm**.

Managing channels

You can view, create, modify, or delete channels to determine how you deliver information about event groups.

View alert channels

You can view the list of all the alert channels.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. In the **CELOG alerting** area, click the **Alert channel** tab.
3. View the list of alert channels.

Create a channel

You can create and configure new channels to send out alert information.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. In the **CELOG alerting** area, click the **Alert channel** tab.
3. In the **Alert Channel** area, click **Create channel**.
4. Select the **Enable channel** check box to enable or disable the channel.
5. In the **Channel name** field, type the channel name.
6. Select the delivery mechanism for the channel from the **Channel type** list.

   NOTE: Depending on the delivery mechanism you select, different settings appear.

7. If you are creating an SMTP channel, you can configure the following settings:
   a. In the **Send to** field, enter an email address that you want to receive alerts on this channel.
      To add another email address to the channel, click **Add another email address**.
   b. To manually configure the SMTP server settings, select the **Manually configured SMTP server settings** radio button and configure the following fields.
   c. In the **Send from** field, enter the email address that you want to appear in the from field of the alert email messages.
   d. In the **Subject** field, enter the text that you want to appear on the subject line of the alert email messages.
   e. In the **SMTP host or relay address** field, enter your SMTP host or relay address.
   f. In the **SMTP relay port** field, enter the number of your SMTP relay port.
   g. Select the **Use SMTP authentication** check box to specify a username and password for your SMTP server.
   h. Specify your connection security between **NONE** or **STARTTLS**.
   i. From the **Notification batch mode** list, select whether alerts will be batched together, by severity, or by category.
   j. From the **Notification email template** list, select whether email messages will be created from a default or custom email template.
      If you specify a custom template, enter the location of the template on your cluster in the **Custom Template Location** field.

k. In the **Allowed nodes** field, type the node number of a node in the cluster that is allowed to send alerts through this channel.

   To add another allowed node to the channel, click **Add another Node**. If you do not specify any nodes, all nodes in the cluster are considered as allowed nodes.

l. In the **Excluded nodes** field, type the node number of a node in the cluster that is not allowed to send alerts through this channel.

   To add another excluded node to the channel, click **Exclude another node**.

8. If you are creating a CONNECTEMC channel, you can configure the following settings:
   a. In the **Allowed nodes** field, type the node number of a node in the cluster that is allowed to send alerts through this channel.

      To add another allowed node to the channel, click **Add another node**. If you do not specify any nodes, all nodes in the cluster are considered as allowed nodes.

   b. In the **Excluded nodes** field, type the node number of a node in the cluster that is not allowed to send alerts through this channel.

      To add another excluded node to the channel, click **Exclude another node**.

9. If you are creating an SNMP channel, you can configure the following settings:
   a. In the **Community** field, enter your SNMP community string.
   b. In the **Host** field, enter your SNMP hostname or address.
   c. In the **Allowed nodes** field, type the node number of a node in the cluster that is allowed to send alerts through this channel.

      To add another allowed node to the channel, click **Add another node**. If you do not specify any nodes, all nodes in the cluster are considered as allowed nodes.

   d. In the **Excluded nodes** field, type the node number of a node in the cluster that is not allowed to send alerts through this channel.

      To add another excluded node to the channel, click **Exclude another node**.

10. Click **Create channel**.

Modify a channel

You can modify a channel that you have created.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. In the **CELOG alerting** area, click the **Alert channel** tab.
3. In the **Actions** column of the channel you want to modify, click **Edit channel**.
   The **Edit alert channel** window appears.
4. Select the **Enable channel** check box to enable or disable the channel.
5. Select the delivery mechanism for the channel from the **Channel type** list.

   NOTE: Depending on the delivery mechanism you select, different settings appear.

6. If you are modifying an SMTP channel, you can change the following settings:
   a. In the **Send to** field, enter an email address that you want to receive alerts on this channel.

      To add another email address to the channel, click **Add another email address**.

   b. To manually configure the SMTP server settings, select the **Manually configured SMTP server settings** radio button and configure the following fields.
   c. In the **Subject** field, enter the text that you want to appear on the subject line of the alert email messages.
   d. In the **SMTP host or relay address** field, enter your SMTP host or relay address.
   e. In the **SMTP relay port** field, enter the number of your SMTP relay port.
   f. Select the **Use SMTP authentication** check box to specify a username and password for your SMTP server.
   g. Specify your connection security between **NONE** or **STARTTLS**.
   h. From the **Notification batch mode** list, select whether alerts will be batched together, by severity, or by category.
   i. From the **Notification email template** list, select whether email messages will be created from a standard or custom email template.

      If you specify a custom template, enter the location of the template on your cluster in the **Custom template location** field.

j. In the **Allowed nodes** field, type the node number of a node in the cluster that is allowed to send alerts through this channel.

   To add another allowed node to the channel, click **Add another node**. If you do not specify any nodes, all nodes in the cluster are considered as allowed nodes.

k. In the **Excluded nodes** field, type the node number of a node in the cluster that is not allowed to send alerts through this channel.

   To add another excluded node to the channel, click **Exclude another node**.

7. If you are modifying a CONNECTEMC channel, you can change the following settings:
   a. In the **Allowed nodes** field, type the node number of a node in the cluster that is allowed to send alerts through this channel.

      To add another allowed node to the channel, click **Add another node**. If you do not specify any nodes, all nodes in the cluster are considered as allowed nodes.

   b. In the **Excluded nodes** field, type the node number of a node in the cluster that is not allowed to send alerts through this channel.

      To add another excluded node to the channel, click **Exclude another node**.

8. If you are modifying an SNMP channel, you can change the following settings:
   a. In the **Community** field, enter your SNMP community string.
   b. In the **Host** field, enter your SNMP hostname or address.
   c. In the **Allowed nodes** field, type the node number of a node in the cluster that is allowed to send alerts through this channel.

      To add another allowed node to the channel, click **Add another node**. If you do not specify any nodes, all nodes in the cluster are considered as allowed nodes.

   d. In the **Excluded Nodes** field, type the node number of a node in the cluster that is not allowed to send alerts through this channel.

      To add another excluded node to the channel, click **Exclude another node**.

9. Click **Save Changes**.

Delete a channel

You can delete channels that you have created.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. In the **CELOG alerting** area, click the **Alert channel** tab.
3. In the **Alert Channel** area, locate the channel you want to delete.
4. In the **Actions** column of the channel you want to delete, click **Edit channel**.
   The **Edit alert channel** window appears
5. Click **Delete** to confirm the action.

Managing event thresholds

You can list, modify, reset, and view alert thresholds for events that use percentage-based statistics to generate alerts.

View events with configurable thresholds and adjust the threshold values

You can view the events with configurable alert thresholds and adjust the thresholds.

1. Click **Cluster Management** > **Events and Alerts** > **Thresholds**.
2. In the Actions column of the event that you want to adjust, click **Edit thresholds**.
3. In the Threshold value column for each threshold you want to adjust, enter an integer in the range 0-100 for each threshold.
4. Click **Apply changes**.

Maintenance and testing

You can modify event settings to specify retention and storage limits for event data, schedule maintenance history windows, and send test events.

Event data retention and storage limits

You can modify settings to determine how event data is handled on your cluster.

By default, data related to resolved event groups is retained indefinitely. You can set a retention limit to make the system automatically delete resolved event group data after a certain number of days.

You can also limit the amount of memory that event data can occupy on your cluster. By default, the limit is 1 megabyte of memory for every 1 terabyte of total memory on the cluster. You can adjust this limit to be between 1 and 100 megabytes of memory. For smaller clusters, the minimum amount of memory that will be set aside is 1 gigabyte.

When your cluster reaches a storage limit, the system will begin deleting the oldest event group data to accommodate new data.

View event storage settings

You can view your storage and maintenance settings.

Click **Cluster Management** > **Events and Alerts** > **Settings**.

Modify event storage settings

You can modify your storage and maintenance settings.

1. Click **Cluster Management** > **Events and Alerts** > **Settings**.
2. In the **Retain event group and maintenance window history** field, enter the number of days you want resolved event groups and maintenance window history to be stored before they are deleted.
3. In the **Event log storage limit** field, enter the limit for the amount of storage you want to set aside for event data.

   The value in this field represents how many megabytes of data can be stored per terabyte of total cluster storage.
4. Click **Submit**.

Test events and alerts

Test events called heartbeat events are automatically generated. You can also manually generate test alerts.

In order to confirm that the system is operating correctly, test events are automatically sent every day, one event from each node in your cluster. These are referred to as heartbeat events and are reported to an event group named Heartbeat Event.

To test the configuration of channels, you can manually send a test alert through the system.

Send a test alert

You can manually generate a test alert.

1. Click **Cluster Management** > **Events and Alerts** > **Alert Management**.
2. In the **CELOG alerting** area, click the **Alert channel** tab.
3. In the **Alert Channel** area, click **Create channel**.
4. Locate the **Send test alert** area.
5. In the **Test message** field, enter the message that you want to send.
6. Click **Send.**

Gathering cluster logs

You can gather cluster logs and send the logs to PowerScale Technical Support for analysis. Cluster logs can be sent automatically or manually through the cluster command-line and web administration interfaces.

NOTE: Your cluster must be connected to the internet to be able to send log files directly. In newer versions of OneFS, you must also have remote support and SRS enabled. If your cluster does not have an internet connection or if your upload has failed, you can copy the log file from the cluster and upload the log file with an FTP client to an FTP server. For more information about how to gather and send cluster log files, including the command parameters for configuring how logs are uploaded to PowerScale Technical Support, refer to the administration guide for your version of OneFS.

Gather cluster logs through the web administration interface

You can gather and send cluster logs from each node in the cluster to technical support.

1. Click **Cluster Management > Diagnostics**
2. Click **Start Gather**

The compressed log file is listed in the **Gather output files** section.

Gather cluster logs - CLI

You must have root access to run log gathering commands.

To gather the log files in OneFS 9.0.0.0 and earlier, run the following command:

```
isi_gather_info
```

To gather the log files in OneFS 9.1.0.0 and later, run the following command:

```
isi diagnostics gather start
```

The files generated during the log gathering process are stored on the cluster in the `/ifs/data/Isilon_Support/pkg` directory.

Manually send cluster logs through the command-line interface

You can manually send the cluster logs to technical support.

Logs generated by this command are stored by default in the `/ifs/data/Isilon_Support/pkg` directory.

1. Run the following command to list all of the files in the temporary directory.

   ```
   ls -l /ifs/data/Isilon_Support/pkg
   ```

2. Note the name of the log file that you want to upload.
3. Run the following command, where <*file-name*> is the name of the log file.

   ```
   isi_gather_info --re-upload /ifs/data/Isilon_Support/pkg/<file-name>
   ```

Download cluster logs and send through FTP

If your cluster is not connected to the Internet, then you can download and send the cluster logs to PowerScale Technical Support through an FTP client.

1. Click **Cluster Management > Diagnostics**.
2. In the **Gather output files** section, click **Download** for the log file that you want to download.

   Follow the download process that is specific to your browser.

3. Open an FTP client.
4. Enter the following settings to connect to the FTP server:
   Host: ftp.emc.com
   User name: anonymous
   Password: your email address
5. Change the destination directory to `incoming`.
6. Upload the log file.

3

Notifications

This section contains the following topics:

Topics:

Event notifications

Event notifications enable you to determine which system events are sent to you.

By default, OneFS is configured to log and send all critical and emergency events to PowerScale Technical Support. Additionally, you can configure event notification rules and receive notification when an event is logged on your cluster.

Events are logged and reported according to the following severity levels:

- Informational (Info): The event is informational and does not require any administrator action.
- Warning (Warn): The event might be a cause for concern and likely requires some administrator action.
- Critical (Crit): The event should be reviewed and an administrator action is required.
- Emergency (Emerg): The event should be reviewed and might be time-sensitive. An administrator action is required.

You can configure the cluster to notify you only when an event of a particular severity has occurred for a particular feature area. For example, you can configure an event notification rule to send you an email if a critical file system event occurs that might potentially cause the cluster to be unwritable.

Event notification methods

You can define the method by which OneFS delivers notifications.

Email	You can send email messages to distribution lists and apply email templates to notifications. You can also specify SMTP, authorization, and security settings.
SNMP trap	You can send SNMP traps to one or more network monitoring stations or trap receivers. Each event can generate one or more SNMP traps. You can download management information base files (MIBs) from the cluster at `/usr/local/share/snmp/mibs/`. The `ISILON-TRAP-MIB.txt` file describes the traps that the cluster can generate, and the `ISILON-MIB.txt` file describes the associated varbinds that accompany the traps.
SRS	You can receive alerts from Secure Remote Services (SRS), which is a secure, IP-based customer service support system. SRS can: ● Send alerts regarding the health of devices. ● Enable support personnel to gather data from devices. ● Allow support personnel to establish remote access to troubleshoot the cluster.

Managing event notification rules

You can create, modify, or delete event notification rules to determine when and how you receive information about specific system events.

Create an event notification rule

You can configure event notification rules based on specified events and event types.

You can configure email notification and SNMP trap generation for a specific event.

1. Click **Cluster management** > **Events and alerts**.
2. In the Alerting rule tab on the Alert management page, click **Create alert rule**.
3. In the **Rule name** field on the Create alert rule window, type a name for the rule.
4. In the **Rule condition** drop-down, select a condition.

   NOTE: Depending on the condition you select, different settings appear.

5. If you are creating a rule with the condition of NEW, SEVERITY INCREASE, SEVERITY DECREASE, or RESOLVED, you can configure the length of time of the event in seconds.
6. If you are creating a rule with the condition of NEW EVENTS, you can configure the Maximum alert limit and the length of time of the event in seconds.
7. If you are creating a rule with the condition ONGOING, you can configure the Interval in seconds and the length of time of the event in seconds.
8. Select the check boxes for the Alert categories that this alert rule applies to.
9. To add an EventGroup ID, click **Add event group ID** and enter the ID.
10. Select the alert channel for this rule.
11. Click **Create rule**.

View event notification rules

You can view a list of event notification rules and details about specific rules.

1. Click **Cluster management** > **Events and alerts** > **Alert management**.
2. On the **Alerting rule** tab, in the **Actions** column, click **Edit rule** of the rule whose settings you want to view.
3. When you have finished viewing the rule details, click **Cancel**.

Modify an event notification rule

You can modify event notification rules that you created. System event notification rules cannot be modified.

1. Click **Cluster management** > **Events and alerts** > **Alert management** > **Alerting rule**.
2. In the **Actions** column for the rule that you want to modify, click **Edit rule**.
3. Modify the event notification rule settings as needed.
4. Click **Save**.

Delete an event notification rule

You can delete event notification rules that you created, but system event notification rules cannot be deleted.

1. Click **Cluster management** > **Events and alerts** > **Alert management** > **Alerting rule**.
2. In the **Actions** column for the rule that you want to delete, click **Edit rule**.
3. In the **Edit alert rule** window, click **Delete**.
4. Click **Delete** to confirm the action.

Send a test event notification

You can generate a test event notification to confirm that event notifications are working as you intend.

1. Click **Cluster management** > **Events and alerts** > **Alert management** > **Alert channel**.
2. In the **Send test alert** area on the Edit alert channel window, enter the test message content and click **Send**.
3. On the Event group history tab, a corresponding test event appears in the Event group list.

Managing event group notification settings

You can view and modify event group notification settings and configure batch notifications.

Event group notification settings

You can specify whether you want to receive event notifications as aggregated batches or as individual notifications for each event. Batch notifications are sent every 10 seconds.

The batch options that are described in this table affect both the content and the subject line of notification emails that are sent in response to system events. You can specify event notification batch options when you configure SMTP email settings.

Setting	Option	Description
Notification batch mode	Batch all	Generates a single email for each event notification.
	Batch by severity	Generates an email that contains aggregated notifications for each event of the same severity, regardless of event category.
	Batch by category	Generates an email that contains aggregated notifications for event of the same category, regardless of severity.
	No batching	Generates one email per event.
Custom notification template	No custom notification template is set	Sends the email notification in the default OneFS notification template format.
	Set custom notification template	Sends the email notifications in the format that you defined in your custom template file.

View event notification settings

You can view email and contact information for event notifications.

Click **Cluster management** > **General settings** > **Email settings**.

Modify event notification settings

You can modify email and contact settings for event notifications.

1. Click **Cluster management** > **General settings** > **Email settings**.
2. In the **Event group notification settings**, edit the settings that you want to change.
3. Click **Save changes**.

Specify event-notification batch mode or template settings

You can choose an event-notification batch option to specify whether you want to receive notifications individually or as an aggregate. You also can specify a custom notification template for email notifications.

You must first create a custom notification template and then upload it to a directory at the same level or below `/ifs`; for example, `/ifs/templates`.

1. Click **Cluster Management** > **General Settings** > **Email Settings**.
2. In the **Event group notification settings** area, select a **Notification batch mode** option. Notifications can be batched all together, by severity, or by category.
3. In the **Notification email template** drop-down, select **Use custom template**.
4. In the **Custom template location** field, click **Browse**, navigate to and select the template file that you want to use, and then click **Select**.
5. Click **Save changes**.

4

Software events

This section contains the following topics:

Topics:

- 100010050
- 100010051
- 100010052
- 100010053
- 100010054
- 100010055
- 100010056
- 100010057
- 100010058
- 100010059
- 100010060
- 100010061
- 100010062
- 100020060
- 100020061
- 100020062
- 100020063
- 100030001
- 200010001
- 200010002
- 200010003
- 200010006
- 200010007
- 200010008
- 200010009
- 200020001
- 200020002
- 200020003
- 200020004
- 200020005
- 200020006
- 200020007
- 200020008
- 200020009
- 200020010
- 200020011
- 200020012
- 200020013
- 200020014
- 200020015
- 200020020
- 200020021
- 200020022
- 200020023
- 200020024
- 200020025
- 200020026
- 200030001
- 200030002
- 300010001
- 300010002
- 300010003
- 300020001
- 300020002
- 300020003
- 400020001

- 400030001
- 400030002
- 400040001
- 400040002
- 400040003
- 400040004
- 400040007
- 400040009
- 400040010
- 400040011
- 400040012
- 400040014
- 400040015
- 400040017
- 400040018
- 400040019
- 400040020
- 400040021
- 400040022
- 400040023
- 400040024
- 400040025
- 400040026
- 400050001
- 400050002
- 400050004
- 400060001
- 400060002
- 400060004
- 400060101
- 400060102
- 400060103
- 400060104
- 400060105
- 400060106
- 400060107
- 400060108
- 400060109
- 400060110
- 400060111
- 400060112
- 400060113
- 400070004
- 400070005
- 400070006
- 400070007
- 400080001
- 400090001
- 400090002
- 400090003
- 400090004
- 400100001
- 400100002
- 400100003
- 400100004
- 400100005

- 400100006
- 400100007
- 400100008
- 400100009
- 400100010
- 400100011
- 400110001
- 400120001
- 400130001
- 400130002
- 400140001
- 400140002
- 400140003
- 400150001
- 400150002
- 400150003
- 400150004
- 400150005
- 400150006
- 400150007
- 400150008
- 400150009
- 400150010
- 400150011
- 400150012
- 400151001
- 400160001
- 400160002
- 400160005
- 400170001
- 400170002
- 400180001
- 400180002
- 400180003
- 400180004
- 400180005
- 400190001
- 400200001
- 400200002
- 400210001
- 400210002
- 400210003
- 400210004
- 400210005
- 400210006
- 400210007
- 400210008
- 400220000
- 400230001
- 400240000
- 400240001
- 400240002
- 400240003
- 400240004
- 400240005
- 400250000

- 400260000
- 500010001
- 500010002
- 500010003
- 500010004
- 500010005
- 600010001
- 600010002
- 600010003
- 600010004
- 600010005
- 700010001
- 700010003
- 700010004
- 700010005
- 700020001
- 700020002
- 700020003
- 700030001
- 700030002
- 700030003
- 700030004
- 700030005
- 700030006
- 700040001
- 700050001
- 700100001
- 800010002
- 800010003
- 800010004
- 800010005
- 800010006
- 800010007
- 800010008
- 800010009
- 800010010
- 1100000001
- 1100000002
- 1100000003
- 1100000004
- 1100000005
- 1100000006
- 1100000007
- 1100000008
- 1100000009

Software events overview

Software events provide information about OneFS and related application software status, such as SyncIQ policy issues and errors.

100010001

The /var partition on the node is at or near capacity.

Description

Allowing nodes to run with a full `/var` partition could lead to system stability issues. This issue is often temporary and can be caused by a problem with log rotation or system files that are eventually deleted automatically. This event also commonly occurs when a file is mistakenly written into the `/var` partition by an administrator working in the node console.

If the `/var` partition returns to a normal usage level of less than 75 percent, and the issue does not recur, you can disregard the event.

Administrator action

Follow the instructions in *OneFS: Event notification: The /var partition is near capacity (95% used) - Event ID: 100010001*, article 000169344.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010002

The /var/crash partition on a node is at or near capacity.

Description

The purpose of the `/var/crash` partition is to preserve data about failed processes and unplanned restarts to enable analysis of those events. This event is usually the result of a process or service stopping unexpectedly and producing a core file, which is a type of log file. Core files record all of the system events when the system terminates unexpectedly. Depending on the state of the process when the problem occurred, these core files can be very large.

Administrator action

- If the cluster has generated a number of core files, you can redirect those core files to another directory with a larger capacity by following the instructions in *Isilon: OneFS-How to redirect core files from the /var/crash directory to a larger capacity directory*, article 00018985 on the Dell EMC Online Support site.
- If you are storing temporary files in the /var/crash partition, move them to a larger capacity directory. If you need to keep the files in the /var/crash partition for some time, quiet the event.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

10010003

The root file system on one or more nodes in the cluster is nearing capacity.

Description

Allowing nodes to run with a full root partition could lead to system stability issues.

This issue is often temporary and can be caused by system files that are eventually deleted automatically. This event also commonly occurs when a file is mistakenly written into the root partition by an administrator working in the node console.

Administrator action

Follow the instructions in *OneFS: Event notification: Node reached 95% or greater used capacity on the root filesystem*, article 000016965 on theDell EMC Online Support site.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

10010004

The /ifs partition on the cluster is near capacity.

Description

By default, this event first appears when the amount of data on the partition reaches the warning threshold of 88% of the partition capacity. The message appears again when the amount of data reaches the critical threshold of 95% of the partition capacity.

Administrator action

Reduce the amount of data that is stored on the cluster or contact your sales representative to discuss your capacity needs.

10010005

The serially-attached SCSI (SAS) PHY monitor detected an error or a change in the disk subsystem.

Description

A drive replacement or upgrade of the node firmware of an Isilon node may generate an alert describing a problem or change in the SAS PHY topology.

Administrator action

If the drive was recently replaced, you can safely ignore this event.

If a drive was not recently replaced, make sure that the drive firmware is updated to the latest version. Drive firmware is available from the Dell EMC Online Support site.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010006

A drive logged an error or a change in the disk subsystem.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010007

The serially-attached SCSI (SAS) PHY monitor detected an excessive bit error rate in the SAS cable traffic.

Description

One or more nodes report a Bit Rate Error (BER) event message.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010008

The serially-attached SCSI (SAS) PHY monitor detected an excessive bit error rate and disabled traffic on the SAS cables.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010009

Disk repair was initiated.

Administrator action

This message is informational. No action is required.

100010010

Disk repair is complete.

Description

OneFS has reprotected all of the data on the drive after a smartfail operation, and the drive has been removed from the group.

Administrator action

Replace the drive by following the instructions in the drive replacement guide for your node type, which can be downloaded from the Dell EMC Online Support site.

100010011

One or more failed drives in this node are ready to be replaced.

Administrator action

Review the messages that were generated during the most recent FlexProtect job in the `/var/log/messages` file.

- If the Flexprotect job completed successfully, replace the drive by following the instructions in the drive replacement guide for your node type, which can be downloaded from the Dell EMC Online Support site.
- If the Flexprotect job failed, contact Dell EMC PowerScale Technical Support for additional troubleshooting.

100010012

The disk has stalled and the disk health is being evaluated.

Administrator action

This message is informational. No action is required.

100010013

There is an error in a disk sector.

Description

If enough disk errors occur, the drive is automatically smartfailed. If a disk smartfails and should be replaced, another event is generated.

Administrator action

This message is informational. No action is required.

100010014

The disk ECC list is full.

Description

The ECC list contains information about disk sector errors. If sufficient errors are recorded for a particular disk, the drive can become unusable and will be smartfailed by the system.

This event also might occur if the ECC list configuration has changed.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010015

One of the disk pools on your cluster is nearing, or has reached, maximum capacity.

Description

If the cluster is too close to maximum capacity there might be insufficient space to restripe data in the event of a hardware failure, which could put your data at risk.

In addition, should the cluster be allowed to approach 100% used capacity, important system processes will cease to function properly until disk usage is reduced.

Administrator action

To reduce the capacity to below 90% you can:

- Modify the affected disk pool
- Remove extraneous data
- Add capacity to your cluster

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010016

The diskpool metadata has been written to more SSDs than is allotted in the layout preferences.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010017

A drive does not match the node configuration information.

Administrator action

Determine the types of drives that are on the node by running the following command:

```
isi devices drive list
```

Any drives that are not functioning correctly will be flagged.

- If the drives are incorrect, replace the drives with the correct types. You can download the latest version of drive replacement guide from the Online Support site.
- If the drives are correct, or if you need assistance obtaining the correct drive type, contact Technical Support.

100010018

The storage capacity of your SSDs in the cluster is approaching capacity.

Administrator action

Reduce the amount of data that is stored on the SSD or contact your sales representative to discuss your capacity needs.

100010019

A serially-attached SCSI (SAS) controller logged an error or a change in the disk subsystem.

Description

When disk drives go into error recovery, they hold the PHY connection open. The LSI 2008 SAS controllers can time out on that open connection. When the LSI controller reaches its timeout threshold, the SAS connection is reset, which causes the SAS BER error messages and event notifications.

Administrator action

Verify that the node firmware and drive firmware are up to date. The latest firmware packages can be downloaded from the Dell EMC Online Support site.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010020

A serially-attached SCSI (SAS) controller logged an error or a change in the disk subsystem.

Description

When disk drives go into error recovery, they hold the PHY connection open. The LSI 2008 SAS controllers can time out on that open connection. When the LSI controller reaches its timeout threshold, the SAS connection is reset, which causes the SAS BER error messages and event notifications.

Administrator action

Verify that the node firmware and drive firmware are up to date. The latest firmware packages can be downloaded from the Dell EMC Online Support site.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

10010021

A serially-attached SCSI (SAS) link has exceeded the maximum Bit Error Rate (BER).

Description

When disk drives go into error recovery, they hold the PHY connection open. The LSI 2008 SAS controllers can time out on that open connection. When the LSI controller reaches its timeout threshold, the SAS connection is reset, which causes the SAS BER error messages and event notifications.

Administrator action

Verify that the node firmware and drive firmware are up to date. The latest firmware packages can be downloaded from the Dell EMC Online Support site.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

10010022

A serially-attached SCSI (SAS) link was disabled for exceeding the maximum Bit Error Rate (BER).

Description

When disk drives go into error recovery, they hold the PHY connection open. The LSI 2008 SAS controllers can time out on that open connection. When the LSI controller reaches its timeout threshold, the SAS connection is reset, which causes the SAS BER error messages and event notifications.

Administrator action

Verify that the node firmware and drive firmware are up to date. The latest firmware packages can be downloaded from the Dell EMC Online Support site.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

10010023

A drive bay error counter has exceeded a configured threshold.

Administrator action

Troubleshooting is required to determine if a hardware component must be replaced.

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010024

A configuration file is missing for the identified drive model.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010025

A SMART status threshold for the identified bay has been exceeded.

Administrator action

Troubleshooting is required to determine if a hardware component must be replaced.

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010026

An encrypted drive is in an insecure state

Administrator action

1. Smartfail the specified drive. Do not remove the drive from the node.
2. Reinstall the smartfailed drive by running the following command:

```
isi devices drive format <bay> --node-lnn <integer>
```

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010027

The drive subsystem determined that the drive firmware version on a drive is incorrect.

Administrator action

Verify that the drive support package and drive firmware are up to date.

You can download the latest drive support packages and drive firmware from the Dell EMC Online Support site.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010028

The drive subsystem does not recognize the drive firmware version of a drive.

Administrator action

Verify that the drive support package and drive firmware are up to date.

You can download the latest drive support packages and drive firmware from the Dell EMC Online Support site.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010029

The drive subsystem does not recognize the model of a drive.

Administrator action

Verify that the drive support package and drive firmware are up to date.

You can download the latest drive support packages and drive firmware from the Dell EMC Online Support site.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010030

An unsupported drive was installed in the node.

Description

This event indicates that an unsupported drive was installed in the node.

Administrator action

Remove the unsupported drive from the affected node and contact Dell EMC PowerScale Technical Support. You do not need to smartfail the drive as the drive was unable to format.

100010031

An encrypted drive was not erased when it was removed from a node.

Description

The data on the drive remains encrypted, but the drive was not erased. The data on the drive still cannot be accessed without the drive encryption key stored on the node. If necessary, take additional steps to destroy the data on the drive.

It is only possible to access the data on the drive if an individual possesses both the drive and the encryption key stored on the node.

Administrator action

We recommend that you destroy the drive data according to your company's security regulations.

100010032

A used drive from another cluster was inserted as a replacement.

Description

A drive that was installed as a replacement was used previously in a different cluster. The event message provides you with the current node, bay, location, drive type, and Logical Number (LNUM) of the drive.

Administrator action

Take one of the following actions to resolve the issue:

- If the drive was inserted in error, replace the drive with a new drive according to the instructions in the *PowerScale Drive Replacement Guide* for your platform.
- If you want to use the drive that was inserted, you need to re-format the drive to erase it.

100010033

A used drive from another node in the cluster was inserted as a replacement.

Description

A drive that was installed as a replacement was used previously in the cluster. The event message provides you with the current node, bay, location, drive type, and Logical Number (LNUM) of the drive.

Administrator action

Take one of the following actions to resolve the issue:

- If the drive was inserted in error, replace the drive with a new drive according to the instructions in the *PowerScale Drive Replacement Guide* for your platform.
- If you want to use the drive that was inserted, you need to re-format the drive to erase it.

100010034

A FlexProtect job is in progress for a drive.

Description

Flexprotect is currently preparing a drive to be added to the filesystem. The event message provides you with the current node, bay, location, drive type, and Logical Number (LNUM) of the drive.

Administrator action

This message is informational. No action is required.

100010035

A drive that previously failed was inserted as a replacement.

Description

A drive that previously failed was installed as a replacement. The event message provides you with the current node, bay, location, drive type, and Logical Number (LNUM) of the drive.

Administrator action

Replace the failed drive with a new drive according to the instructions in the *PowerScale Drive Replacement Guide* for your platform.

100010036

The disk repair process for a drive is complete.

Description

The smartfail process has completed and the specified drive is ready to be replaced. The event message provides you with the current node, bay, location, drive type, and Logical Number (LNUM) of the drive.

Administrator action

Replace the drive with a new drive according to the instructions in the *PowerScale Drive Replacement Guide* for your platform.

100010037

A new drive was not formatted correctly, and as a result, the drive was not added to the filesystem.

Description

A drive that was installed in a node was not formatted correctly. The event message provides you with the current node, bay, location, drive type, and Logical Number (LNUM) of the drive.

Administrator action

Take one of the following actions to resolve the issue:

- You can try to format the drive again to see if that resolves the issue.
- Replace the drive with a new drive according to the instructions in the *PowerScale Drive Replacement Guide* for your platform.

100010038

The following drive was inserted into a bay that contains another drive that is in the process of smartfailing. Reinsert the drive that is smartfailing, or wait for the FlexProtect job to complete before inserting a new drive.

Description

A drive was installed as a replacement in a drive bay that is still in the process of Smartfailing the failed drive that was removed. The event message provides you with the current node, bay, location, drive type, and Logical Number (LNUM) of the drive.

Administrator action

You must wait for the FlexProtect job to complete before you insert a replacement drive in the drive bay.

You can replace the failed drive in the drive bay and wait for the drive to finish Smartfailing, or leave the bay empty until the FlexProtect job finishes.

100010039

Unprovisionable drive(s): {unprovisionable}

Description

One or more drives in the cluster cannot be provisioned.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010041

Use of a non-supported boot flash drive has been attempted.

Description

The non-supported boot flash drive must be replaced. Boot flash drives are not customer-replaceable parts. If you are unsure whether a potential replacement boot flash drive is supported, contact Dell EMC PowerScale Technical Support.

Administrator action

Contact Dell EMC PowerScale Technical Support to diagnose and resolve the issue.

100010042

An error occurred during the SmartPools upgrade, and as a result the upgrade did not complete.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010043

A node boot flash drive must be replaced.

Description

Boot flash drives are not customer-replaceable parts. Replacement of this device will require turning off power to the node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010044

A node boot flash drive must be replaced.

Description

Boot flash drives are not customer-replaceable parts. Replacement of this device will require turning off power to the node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010045

A node boot flash drive is receiving excessive writes.

Description

This event is only generated during troubleshooting by Dell EMC PowerScale Technical Support.

Administrator action

Contact Dell EMC PowerScale Technical Support to diagnose and resolve the issue.

100010046

A node pool has a node whose SSD count does not match the SSD counts of other nodes in the pool.

Description

A mismatch of the number of SSDs in nodes in a pool can be caused by an SSD failure in one of the nodes in the pool or, in the case of Generation 6 hardware, nodes do not contain the same number of cache SSDs and L3 cache is not enabled for the node pool.

Administrator action

In the case of a failed SSD, replace the failed drive. In the case of Generation 6 hardware, where the cache SSD counts of nodes in the chassis do not match, either turn L3 cache on for the node pool, or add cache SSDs to ensure that every node contains the same number of cache SSDs.

100010050

The smartfail process completed on a drive.

Description

A drive has been smartfailed from the cluster. The event message provides you with the chassis serial number, node, sled, the drive slot number within the sled, drive type, and Logical Number (LNUM) of the drive.

Administrator action

Replace the smartfailed drive with a new drive according to the instructions in the *PowerScale Drive Replacement Guide* for your platform.

100010051

The smartfail process completed on a drive.

Description

A drive has been smartfailed from the cluster. The event message provides you with the chassis serial number, node, sled, the drive slot number within the sled, drive type, and Logical Number (LNUM) of the drive.

Administrator action

Replace the smartfailed drive with a new drive according to the instructions in the *PowerScale Drive Replacement Guide* for your platform.

100010052

A drive is no longer appearing as part of the cluster and is being smartfailed.

Description

A drive appears to be missing from the cluster and the smartfail process has been initiated to officially remove the drive from the cluster. The event message provides you with the chassis serial number, node, sled, the drive slot number within the sled, drive type, and Logical Number (LNUM) of the drive.

Administrator action

Wait for the smartfail process to complete, then replace the drive with a new drive according to the instructions in the *PowerScale Drive Replacement Guide* for your platform.

100010053

A drive that is write-cache enabled was installed in a Generation 6 platform. Write-cache enabled drives are not compatible with Generation 6 nodes.

Description

A write-cache enabled drive was installed in a 6th Generation node and is not compatible with the node. The event message provides you with the chassis serial number, node, sled, the drive slot number within the sled, drive type, and Logical Number (LNUM) of the drive.

Administrator action

Replace the drive with a compatible drive according to the instructions in the *PowerScale Drive Replacement Guide* for your platform.

100010054

A drive was inserted in a bay that is disabled.

Description

A drive was installed in a sled that is in a bay that is disabled. While, OneFS recognizes service requests for devices in disabled bays, you must resolve the issue with the disabled bay before you can service a drive. To ensure proper functionality, remove the drive from the chassis immediately. The event message provides you with the chassis serial number, node, sled, the drive slot number within the sled, drive type, and Logical Number (LNUM) of the drive.

Administrator action

Remove the drive from the sled that is in the disabled bay. After you have resolved the issue with the disabled bay, re-install the drive according to the instructions in the *PowerScale Drive Replacement Guide* for your platform.

100010055

A drive that is write-cache enabled was installed in a Generation 6 platform. Write-cache enabled drives are not compatible with Generation 6 nodes and the drive has been smartfailed.

Description

A write-cache enabled drive was installed in a 6th Generation node and is not compatible with the node. The drive was smartfailed. The event message provides you with the chassis serial number, node, sled, the drive slot number within the sled, drive type, and Logical Number (LNUM) of the drive.

Administrator action

Replace the drive with a compatible drive according to the instructions in the *PowerScale Drive Replacement Guide* for your platform.

100010056

The write-cache is enabled for a drive in a Generation 6 platform. Write-cache enabled drives are not compatible with Generation 6 nodes.

Description

A write-cache enabled drive was installed in a 6th Generation node and is not compatible with the node. The event message provides you with the chassis serial number, node, sled, the drive slot number within the sled, drive type, Logical Number (LNUM), and serial number of the drive.

Administrator action

Replace the drive with a compatible drive according to the instructions in the *PowerScale Drive Replacement Guide* for your platform.

100010057

There is a missing m.2 card on Gen6.

Description

There is a missing m.2 card in a Generation 6 node.

Administrator action

Replace the m.2 card.

100010058

A node pool does not meet the minimum storage space requirement for large files: {node_pool_name} (id={node_pool_id})

Description

The event indicates that the Large File feature was previously enabled, but the node pool does not meet the minimum storage space capacity requirement.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010059

Node pool {nodepool_name} (node pool ID: {nodepool_id}) is at, or over capacity for large files.

Description

The event indicates that the Large File feature is enabled and a node pool has passed the space in-use in threshold. System performance can be affected.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100010060

Error detected in PCI drive: Location {location}, Type {media_type}, LNUM {disk}: {aer}.

Description

PCI advanced error reporting is reporting an NVMe drive error.

Administrator action

Replace the NVMe drive.

100010061

Error detected in NVMe drive: Location {location}, Type {media_type}, LNUM {disk}: {detail}.

Description

NVMe error log has reported a new error. The event checks the error log page of an NVMe drive. The error log page contains the last command that failed IIRC.

Administrator action

Replace the failed drive.

100010062

A PCI link error is detected with NVMe drive connectivity.

Description

A PCI switch or bay link for NVMe drive connectivity is unhealthy and may require maintenance or replacement.

Administrator action

Determine which PCI link is unhealthy. The PCI link location will be specified in the event. Some errors may be recoverable by reseating cards or cables, but other errors may require replacement of hardware.

100020060

Drive sled was removed from the a chassis and the sled service timeout limit has expired.

Description

A drive sled was removed from a chassis and wasn't replaced before the sled service timeout limit was exceeded. As a result, the drive sled was smartfailed. The event message provides you with the node, sled, and bay, as well as the drive type, and Logical Numbers (LNUMs) of the drives in the sled.

Administrator action

Replace the drive sled.

100020061

A drive sled was unexpectedly removed from a chassis. All drives in the sled were suspended.

Description

A drive sled was removed without pressing the circular gray button on the lower left corner of the sled face and waiting for the Do Not Remove LED to turn off. As a result, all drives in the sled were suspended.

The event message provides you with the node, sled, and bay, as well as the drive type, and Logical Numbers (LNUMs) of the drives in the sled.

Administrator action

Replace the drive sled.

100020062

A fault was detected in a drive sled.

Description

A drive sled has failed and must be replaced. The event message provides you with the failed drive sled and the node that the sled is a part of.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

100020063

A drive sled has been removed from the chassis for longer than the timeout limit.

Description

A drive sled was removed from a chassis and wasn't replaced before the sled service timeout limit was exceeded. As a result, the drive sled was smartfailed. The event message provides you with the node and drive sled that was removed.

Administrator action

Replace the drive sled.

100030001

Drives have been marked as draining but usage is low.

Description

Drives have been marked as draining but usage is low. Please remove the draining flag from the drive.

Administrator action

Remove the draining flag from the drive.

200010001

One or more nodes in your cluster is offline or unreachable.

One or more nodes are offline due to one of the following conditions:

- A node was intentionally shut down for maintenance.
- A node lacks internal network connectivity. Internal connectivity is how a node communicates with other nodes on the cluster.
- A node cannot join the group.

Administrator action

If the Cluster Status page in the OneFS web administration interface indicates that a node is down, complete the following steps.

1. Determine whether the node is turned on. Visually inspect the node to verify that the power light is on.
2. If the node is turned off, attempt to turn the node on.
   - If the node turns on, view the Cluster Status page to determine whether the node has rejoined the cluster. If the node does not rejoin the cluster, proceed to step 3.
   - If the node rejoins the cluster, and the cluster is operational, the event will cancel.
   - If the node does not turn on, make sure that any circuit breakers in the power path are closed and that the power outlets for the node are active. If the node is not receiving power, resolve the power supply issue. If the node is receiving power, contact Technical Support.
3. If the node is on but did not rejoin the cluster, attempt to establish remote access through a secure shell (SSH) session. If the SSH session fails, attempt to establish remote access through the serial console.
4. If neither the SSH session nor the serial console is responsive, press CTRL+T in the SSH session or in the serial console.
   - If pressing CTRL+T produces output, record the output, and then contact Technical Support for failure analysis.
   - If the node is unresponsive, turn the power off and then on again.

200010002

A node that was previously offline has rejoined the group.

Administrator action

This message is informational. No action is required.

200010003

One or more nodes are offline.

One or more nodes are offline due to one of the following conditions:

- A node was intentionally shut down for maintenance.
- A node lacks internal network connectivity. Internal connectivity is how a node communicates with other nodes on the cluster.
- A node cannot join the group.

Administrator action

If the Cluster Status page in the OneFS web administration interface indicates that a node is down, complete the following steps.

1. Determine whether the node is turned on. Visually inspect the node to verify that the power light is on.
2. If the node is turned off, attempt to turn the node on.
   - If the node turns on, view the Cluster Status page to determine whether the node has rejoined the cluster. If the node does not rejoin the cluster, proceed to step 3.
   - If the node rejoins the cluster, and the cluster is operational, the event will cancel.
   - If the node does not turn on, make sure that any circuit breakers in the power path are closed and that the power outlets for the node are active. If the node is not receiving power, resolve the power supply issue. If the node is receiving power, contact Technical Support.
3. If the node is on but did not rejoin the cluster, attempt to establish remote access through a secure shell (SSH) session. If the SSH session fails, attempt to establish remote access through the serial console.
4. If neither the SSH session nor the serial console is responsive, press CTRL+T in the SSH session or in the serial console.
   - If pressing CTRL+T produces output, record the output, and then contact Technical Support for failure analysis.
   - If the node is unresponsive, turn the power off and then on again.

200010006

The identified node group is underprotected from data loss.

Description

The OneFS protection system relies on a particular number of drives and nodes being available for the pool, depending on your settings. If this requirement is not met, your data could be at risk.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

200010007

The identified node is not provisioned.

Description

When a node is not provisioned, that means that it is not associated with a node pool. You cannot write to a node that is not provisioned.

Administrator action

1. If the unprovisioned node recently received a hardware upgrade, confirm that all nodes in the node pool received the same upgrade and that the nodes are still equivalent.
2. Attempt to provision the node into a node pool by running the following command:

   ```
   isi_evaluate_provision_drive
   ```

3. If the issue persists, contact Support.

   Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

200010008

The identified node pool is underprovisioned.

Description

When a node pool is underprovisioned, that means that there only two nodes remaining in the node pool.

In most cases, this event occurs when a node that was previously associated with a node pool has become unprovisioned, or has been smartfailed.

When a node pool is underprovisioned, you can still write to the remaining two nodes in the pool. However, in order to retain appropriate protection levels, it is important to locate the node or nodes that have dropped from the pool and restore them in the node pool.

Administrator action

1. Locate the nodes that are no longer associated with the underprovisioned node pool.
   a. If a node became unprovisioned, you will see that node identified in a separate event, 200010007. You can attempt to provision the node back into the node pool by running the following command:

      ```
      isi_evaluate_provision_drive
      ```

   b. If a node was smartfailed, you can identify the node by running the following command:

      ```
      isi status
      ```

      Address the cause of the smartfail and add the node back to the cluster. The node will be automatically provisioned to its original node pool.
2. If the issue persists, contact Support.

   Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

200010009

Node has recovered from a panic.

Administrator action

A service request (SR) has been opened on your behalf.

200020001

An Ethernet link is not operating at maximum throughput.

Administrator action

Determine whether the issue is related to the cable or the node. Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Verify the following connections:
   - All of the cables in the rack and in adjacent racks are connected securely to the node and neither the cable nor the connector is damaged.
   - The cable is rated for the appropriate Ethernet speed.
   - The switch port speed is set to the same or higher speed as the Ethernet cable.
   - The switch port is set to the Full Duplex setting.
2. If the cable is connected securely, plug the cable into a different node that has a network port functioning at full speed and that has an identical network configuration. When you plug the cable into the other node, leave the other end of the cable plugged into the same switch port.
   - If the issue persists after plugging the cable into a different node, replace the cable.
   - If the issue persists after replacing the cable, move the cable to another port on the switch.
   - If the issue persists after moving the cable to another port on the switch, review the switch logs and consult your switch user manual.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

200020002

The 10 GigE interfaces on one or more nodes have experienced network connectivity issues.

Administrator action

Determine whether the issue is related to the cable or the node. Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Verify the following connections:
   - All of the cables in the rack and in adjacent racks are connected securely to the node and neither the cable nor the connector is damaged.
   - The cable is rated for the appropriate Ethernet speed.
   - The switch port speed is set to the same or higher speed as the Ethernet cable.
   - The switch port is set to the Full Duplex setting.
2. If the cable is connected securely, plug the cable into a different node that has a network port functioning at full speed and that has an identical network configuration. When you plug the cable into the other node, leave the other end of the cable plugged into the same switch port.
   - If the issue persists after plugging the cable into a different node, replace the cable.
   - If the issue persists after replacing the cable, move the cable to another port on the switch.
   - If the issue persists after moving the cable to another port on the switch, review the switch logs and consult your switch user manual.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

200020003

Multiple internal network issues were detected.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. If internal network issues are not ongoing and this issue is not a recurring problem, attempt the following steps:
   - Make sure that all InfiniBand (IB) cables in the rack and in adjacent racks are connected securely to the node and are undamaged.
   - Make sure there are no kinks or sharp bends in the cables. InfiniBand (IB) cables have an approximate maximum bend radius of 1.5 inches (38mm).
2. If internal network issues are ongoing but do not continually recur, attempt the following steps.
   - Determine whether the issue is related to the cable or the node. Plug the affected cable into another node that has a functioning network port and an identical network configuration. Be sure to plug the cable into the same network port on the different node.
   - If the issue persists after plugging the cable into a different node, replace the cable.
   - If the issue persists after replacing the cable, move the cable to another port on the switch.
   - Confirm that multiple clusters are not connected to the same IB switch. This configuration is unsupported and can cause this issue. Each cluster must be connected to a dedicated IB switch.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

200020004

One or more nodes have experienced network connectivity issues on their aggregated network interfaces.

Administrator action

Determine whether the issue is related to the cable or the node. Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Verify the following connections:
   - All of the cables in the rack and in adjacent racks are connected securely to the node and neither the cable nor the connector is damaged.
   - The cable is rated for the appropriate Ethernet speed.
   - The switch port speed is set to the same or higher speed as the Ethernet cable.
   - The switch port is set to the Full Duplex setting.
2. If the cable is connected securely, plug the cable into a different node that has a network port functioning at full speed and that has an identical network configuration. When you plug the cable into the other node, leave the other end of the cable plugged into the same switch port.
   - If the issue persists after plugging the cable into a different node, replace the cable.
   - If the issue persists after replacing the cable, move the cable to another port on the switch.
   - If the issue persists after moving the cable to another port on the switch, review the switch logs and consult your switch user manual.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

200020005

One of the nodes in your cluster has lost network connectivity on one or both of its external interfaces.

Administrator action

Determine whether the issue is related to the cable or the node. Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Verify the following connections:
   - All of the cables in the rack and in adjacent racks are connected securely to the node and neither the cable nor the connector is damaged.
   - The cable is rated for the appropriate Ethernet speed.
   - The switch port speed is set to the same or higher speed as the Ethernet cable.
   - The switch port is set to the Full Duplex setting.
2. If the cable is connected securely, plug the cable into a different node that has a network port functioning at full speed and that has an identical network configuration. When you plug the cable into the other node, leave the other end of the cable plugged into the same switch port.
   - If the issue persists after plugging the cable into a different node, replace the cable.
   - If the issue persists after replacing the cable, move the cable to another port on the switch.
   - If the issue persists after moving the cable to another port on the switch, review the switch logs and consult your switch user manual.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

200020006

The link status of the identified InfiniBand interface is changing rapidly and repeatedly.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. If internal network issues are not ongoing and this issue is not a recurring problem, attempt the following steps:
   - Make sure that all InfiniBand (IB) cables in the rack and in adjacent racks are connected securely to the node and are undamaged.
   - Make sure there are no kinks or sharp bends in the cables. InfiniBand (IB) cables have an approximate maximum bend radius of 1.5 inches (38mm).
2. If internal network issues are ongoing but do not continually recur, attempt the following steps.
   - Determine whether the issue is related to the cable or the node. Plug the affected cable into another node that has a functioning network port and an identical network configuration. Be sure to plug the cable into the same network port on the different node.
   - If the issue persists after plugging the cable into a different node, replace the cable.
   - If the issue persists after replacing the cable, move the cable to another port on the switch.
   - Confirm that multiple clusters are not connected to the same IB switch. This configuration is unsupported and can cause this issue. Each cluster must be connected to a dedicated IB switch.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

200020007

The internal network switch and the SNMP server are not communicating.

Description

The connection between the internal network switch and the SNMP server has failed.

Administrator action

Confirm that the internal network switch is configured correctly.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

200020008

A fan has failed in an internal network switch.

Description

The fan is not a replaceable part. You will need to schedule a maintenance window to replace the switch.

The event message provides you with the switch serial number and the network fabric supported by the switch.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

200020009

A power supply has failed in an internal network switch.

Description

The power supply is not a replaceable part. You will need to schedule a maintenance window to replace the switch.

The event message provides you with the switch serial number and the network fabric supported by the switch.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

200020010

An internal network switch has failed.

Description

You will need to schedule a maintenance window to replace the switch.

The event message provides you with the switch serial number and the network fabric supported by the switch.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

200020011

A 40-Gigabit Ethernet link is not operating at maximum throughput.

Administrator action

Determine whether the issue is related to the cable or the node. Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Verify the following connections:
   - All of the cables in the rack and in adjacent racks are connected securely to the node and neither the cable nor the connector is damaged.
   - The cable is rated for the appropriate Ethernet speed.
   - The switch port speed is set to the same or higher speed as the Ethernet cable.
   - The switch port is set to the Full Duplex setting.

2. If the cable is connected securely, plug the cable into a different node that has a network port functioning at full speed and that has an identical network configuration. When you plug the cable into the other node, leave the other end of the cable plugged into the same switch port.
   - If the issue persists after plugging the cable into a different node, replace the cable.
   - If the issue persists after replacing the cable, move the cable to another port on the switch.
   - If the issue persists after moving the cable to another port on the switch, review the switch logs and consult your switch user manual.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

20002012

A management Ethernet link is not operating at maximum throughput.

Administrator action

Determine whether the issue is related to the cable or the node. Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Verify the following connections:
   - All of the cables in the rack and in adjacent racks are connected securely to the node and neither the cable nor the connector is damaged.
   - The cable is rated for the appropriate Ethernet speed.
   - The switch port speed is set to the same or higher speed as the Ethernet cable.
   - The switch port is set to the Full Duplex setting.
2. If the cable is connected securely, plug the cable into a different node that has a network port functioning at full speed and that has an identical network configuration. When you plug the cable into the other node, leave the other end of the cable plugged into the same switch port.
   - If the issue persists after plugging the cable into a different node, replace the cable.
   - If the issue persists after replacing the cable, move the cable to another port on the switch.
   - If the issue persists after moving the cable to another port on the switch, review the switch logs and consult your switch user manual.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

20002013

An internal network Ethernet link is not operating at maximum throughput.

Administrator action

Determine whether the issue is related to the cable or the node. Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Verify the following connections:
   - All of the cables in the rack and in adjacent racks are connected securely to the node and neither the cable nor the connector is damaged.
   - The cable is rated for the appropriate Ethernet speed.
   - The switch port speed is set to the same or higher speed as the Ethernet cable.
   - The switch port is set to the Full Duplex setting.
2. If the cable is connected securely, plug the cable into a different node that has a network port functioning at full speed and that has an identical network configuration. When you plug the cable into the other node, leave the other end of the cable plugged into the same switch port.
   - If the issue persists after plugging the cable into a different node, replace the cable.
   - If the issue persists after replacing the cable, move the cable to another port on the switch.

- If the issue persists after moving the cable to another port on the switch, review the switch logs and consult your switch user manual.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

20002014

10-Gigabit Ethernet link {ifname} running below capacity.

Description

10, or 25-Gigabit Ethernet link {*ifname*} running below capacity.

Administrator action

1. Validate that the ethernet switch is capable and configured for operation at the expected speed.
2. Validate that cables and modules used support the expected speed. For details on cables and modules, see *FAQ: Optics and cables*, article 000134129.

20002015

100-Gigabit Ethernet link {ifname} running below capacity.

Description

100 Gigabit Ethernet link {*ifname*} running below capacity.

Administrator action

1. Validate that the ethernet switch is capable and configured for operation at the expected speed.
2. Validate that cables and modules used support the expected speed. For details on cables and modules, see *FAQ: Optics and cables*, article 000134129.

20002020

One of the nodes in your cluster has lost external network connectivity.

Administrator action

Determine whether the issue is related to the cable or the node. Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Verify the following connections:
   - All of the cables in the rack and in adjacent racks are connected securely to the node and neither the cable nor the connector is damaged.
   - The cable is rated for the appropriate Ethernet speed.
   - The switch port speed is set to the same or higher speed as the Ethernet cable.
   - The switch port is set to the Full Duplex setting.
2. If the cable is connected securely, plug the cable into a different node that has a network port functioning at full speed and that has an identical network configuration. When you plug the cable into the other node, leave the other end of the cable plugged into the same switch port.

- If the issue persists after plugging the cable into a different node, replace the cable.
- If the issue persists after replacing the cable, move the cable to another port on the switch.
- If the issue persists after moving the cable to another port on the switch, review the switch logs and consult your switch user manual.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

20020021

The Dell switch has a cabling issue.

Description

The alert dynamically reports that the Dell switch is mis-cabled, or has another cabling issue.

Administrator action

1. Run the `isi event view <ID>` command to determine the problem and the specific cabling issues based on the output.
2. Correct the cabling issue.

The alert automatically ends when the cabling issue is corrected.

20020022

The back-end fabric is unable to contact the back-end Dell master switch.

Description

The management service is down and is critical for:

- Leaf spine failover operation to function correctly.
- Dell Switch operating system communications between OneFS and Dell Ethernet switches.
- SNMP in Arista switches

The event message provides the switch serial number and the supported switch network fabric.

Administrator action

1. Run the `isi event view <ID>` command to determine the problem and the specific switch fabric issues that are based on the output.
2. Correct the fabric or cabling issue.

The alert automatically ends when the issue is corrected.

20020023

A Leaf and Spine switch bandwidth differs.

Description

In a leaf-spine configuration, verify that uplink bandwidth is equal to downlink bandwidth on a leaf switch:

uplink bandwidth = total bandwidth between leaf and each spine.

downlink bandwidth = total bandwidth between leaf and all Isilon nodes.

This indicates the uplink bandwidth does not equal the downlink bandwidth, causing an imbalance in the fabric and subsequent bottleneck.

Administrator action

Correct wiring to ensure uplink bandwidth is equal to downlink fabric for a single fabric (int-a and/or int-b as specified).

200020024

Fabric bandwidth incongruence.

Description

In a leaf-spine configuration, verify that uplink bandwidth is equal on int-a and int-b fabric.

This indiacates that the primary and failover bandwidths differ. When a failover occurs the performance of the cluster may be impacted.

Administrator action

Correct wiring to ensure primary and secondary fabric are in the same configuration so their bandwidths are identical.

200020025

Back-end network non-connectivity.

Description

Back-end network non-connectivity detected: No connectivity between nodes 215 and 252 on the Int-a network.

Administrator action

Obtain a full list of nodes with no connectivity between them, by running the `isi_check_be` command. Restore connectivity.

200020026

A network interface card is unhealthy.

Description

A network interface card is unhealthy and may require maintenance or replacement depending on the card.

Administrator action

Determine which network interface card is unhealthy. The interface name is specified in the event. Some errors are self-recoverable, but other errors might require replacement of the network card.

200030001

The cluster does not have up-to-date firmware.

Description

The installed node firmware package has not been applied to all nodes in the cluster. The cluster requires a firmware upgrade.

Administrator action

To determine the recommended update schedule, run the `isi upgrade firmware assess` command.

200030002

The cluster does not have the available firmware packages.

Description

To ensure the cluster is running qualified firmware, firmware packages must be present. Packages missing: Node Firmware Package, Drive Firmware Package.

Administrator action

Download and install the latest node firmware and drive support packages.

- To install the drive firmware, download the latest Drive Support Package, and run the `isi_dsp_install` command.
- To install the node firmware, download the latest Node Firmware Package, and run the `isi upgrade firmware assess --fw-pkg=<path to NFP>` command.

300010001

The node is being rebooted for maintenance purposes.

Administrator action

This message is informational and occurs only when an administrator runs a command to reboot the node. No action is required.

300010002

The node is being shut down for maintenance purposes.

Administrator action

This message is informational and occurs only when an administrator runs a command to shut down the node. No action is required.

300010003

The node has failed to reboot within the specified time period.

One or more nodes are offline due to one of the following conditions:

- A node was intentionally shut down for maintenance.
- A node lacks internal network connectivity. Internal connectivity is how a node communicates with other nodes on the cluster.
- A node cannot join the group.

Administrator action

If the Cluster Status page in the OneFS web administration interface indicates that a node is down, complete the following steps.

1. Determine whether the node is turned on. Visually inspect the node to verify that the power light is on.
2. If the node is turned off, attempt to turn the node on.
   - If the node turns on, view the Cluster Status page to determine whether the node has rejoined the cluster. If the node does not rejoin the cluster, proceed to step 3.
   - If the node rejoins the cluster, and the cluster is operational, the event will cancel.
   - If the node does not turn on, make sure that any circuit breakers in the power path are closed and that the power outlets for the node are active. If the node is not receiving power, resolve the power supply issue. If the node is receiving power, contact Technical Support.
3. If the node is on but did not rejoin the cluster, attempt to establish remote access through a secure shell (SSH) session. If the SSH session fails, attempt to establish remote access through the serial console.
4. If neither the SSH session nor the serial console is responsive, press CTRL+T in the SSH session or in the serial console.
   - If pressing CTRL+T produces output, record the output, and then contact Technical Support for failure analysis.
   - If the node is unresponsive, turn the power off and then on again.

300020001

A read-only transition failed on the node.

Description

A read-only transition occurs when a node or drive changes from a writeable state to a read-only state, such as during a node shutdown. This event is generated if this transition fails for any reason.

Administrator action

If the node or drive changes to a read-only state, and this event is temporary, no action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

300020002

Validation of a node journal backup failed.

Administrator action

This message is informational. No action is required.

300020003

The node encountered an error performing final shutdown.

Administrator action

Attempt to shut down the node again.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400020001

LWIO is throttling due to current memory threshold settings.

Description

This event is generated when the SMB LWIO process has exceeded the memory usage threshold and new connections are being denied.

Administrator action

Examine SMB usage on the affected node, particularly any active sessions or open files.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400030001

A process failed to restart, despite several attempts to start it.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400030002

The master control program (MCP) stopped a process.

Administrator action

This message is informational. No action is required.

400040001

A SyncIQ policy issue was detected.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400040002

A SyncIQ policy failed.

Administrator action

This event provides information about the specific policy that has failed and information about possible causes. Attempt to resolve the issue identified in the event message.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400040003

A SyncIQ policy cannot start.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400040004

The target cluster for a SyncIQ job cannot create a requested snapshot.

Description

If the cluster is split, this event might appear for a node in the minority group, the group that has fewer than half of the nodes. If the event is for a node on the minority group, you can safely ignore and quiet the event. This message might also appear if the target cluster is limited by a configured quota.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Determine whether the message appeared on the minority or majority group by running the following command from the node that is reporting the event:

   ```
   sysctl efs.gmp.has_quorum
   ```

   - If the command returns 1, the node is in the majority group. Proceed to step 2.
   - If the command returns 0, the node is in the minority group, and you can safely quiet the event.

2. Confirm that the number of snapshots does not exceed the system-wide and directory limits. (The system-wide limit is 20,000 and the directory limit is 1,000.)

   If the number of snapshots is in excess of limitations, delete extraneous snapshots according to the instructions in the OneFS CLI Administration Guide and OneFS Web Administration Guide.

3. Remove any quotas that are limiting capacity for snapshots on the SyncIQ target directory.
4. If the cluster space is approaching capacity, delete files or add capacity to the cluster. You can determine the available capacity on your cluster by running the following command:

```
df -h /ifs
```

400040007

Files have been modified on the target cluster. SyncIQ is overwriting those modified files.

Description

By design, SyncIQ automatically overwrites all of the files on the target cluster. Therefore, any files that are manually saved on the target cluster are overwritten and are not preserved or replicated on the SyncIQ source cluster. SyncIQ is not designed to support bidirectional synchronization of data.

This event occurs if files were modified on the target cluster and SyncIQ is overwriting those modified files.

Administrator action

Make sure that your workflow does not require files that are written manually to a SyncIQ target be preserved on or replicated to the SyncIQ source.

400040009

The SyncIQ scheduler is unable to start the scheduled policy.

Description

This event might be caused by a connection failure between source clusters and target clusters.

Administrator action

1. Verify that the node pool specified in the event message includes one or more nodes that are able to reach the target cluster.
2. If no nodes from the source cluster are able to reach the target cluster, attempt to resolve the connectivity issue.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400040010

A SyncIQ policy configuration error occurred.

Description

This event varies based on the following possible errors that are included in the event message:

- The SyncIQ module detected a problem with the configuration of the policy.
- The SyncIQ policy target path overlaps the target path of another policy.
- The SyncIQ policy target path overlaps the source path when syncing to the same cluster.

Administrator action

Review the SyncIQ policy to resolve any overlapping paths or other configuration errors.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400040011

SyncIQ is attempting to sync to an incompatible target version.

Administrator action

SyncIQ only supports syncing to the same or newer target version. Consider upgrading the target cluster to the same version of OneFS that is on the source cluster.

400040012

A SyncIQ configuration error occurred.

Description

An error occurred when parsing the following SyncIQ global configuration file: `/ifs/.ifsvar/modules/tsm/config/siq-conf.gc`.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400040014

SyncIQ failed to contact the target cluster.

Description

This event occurs under the following conditions:

- SyncIQ is unable to connect to a resource on the target cluster. SyncIQ is unable to connect to a local resource.
- SyncIQ is unable to connect to a daemon (bandwidth, throttle, pworker, or scheduler).

Administrator action

1. If SyncIQ is unable to connect to a resource on the target cluster:
   - Ping the target cluster to determine if the cluster is reachable.
   - Verify that a SyncIQ license is activated and that SyncIQ is enabled on the target cluster.
   - Verify that the SyncIQ daemon process is running.
2. If SyncIQ is unable to connect to a local source, the node running the SyncIQ coordinator was unable to obtain a list of local IP addresses. In this case, check whether the node has split from the group.
3. If SyncIQ is unable to connect to one or more of the daemons mentioned above, verify that the SyncIQ daemons are running on every node by running the following command:

   ```
   ps auwx | grep isi_migr
   ```

   The output should display:

   ```
   isi_migr_bandwidth
   isi_migr_pworker
   isi_migr_sworker
   isi_migr_sched
   ```

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400040015

SyncIQ failed to take a snapshot for a policy.

Description

This event occurs under the following conditions:

- SyncIQ failed to take a snapshot on the source cluster.
- SyncIQ failed to take a snapshot on the target cluster.

Administrator action

1. For a source snapshot error, the appended error string will indicate the problem.
2. For a target snapshot failure, check whether a SnapshotIQ license is active on both source and target clusters.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400040017

There was an error with a policy on the SyncIQ file system.

Description

This event occurs under the following conditions:

- SyncIQ encountered a file system error.

- SyncIQ encountered a file system error on the source cluster.
- SyncIQ encountered a file system error on the target cluster.

NOTE: If you are also experiencing events about hardware failure, investigate the cause of those events before attempting the following actions.

Administrator action

1. Determine if either the source cluster or the target cluster is in read-only mode by running the following command:

   ```
   isi readonly
   ```

   The output for that command displays the mode for each node. If any node is set to read-only mode, set all nodes to read-write mode by running the following command:

   ```
   isi_for_array isi readonly off
   ```

2. Confirm that the file referenced in the error message exists in the synchronization dataset on the source or target cluster that encountered the failure and that the files are accessible.
3. Check the SyncIQ logs to see if any error information is available.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400040018

The SyncIQ policy failed to upgrade.

Description

This event occurs whenever the upgrade sync run for a policy has failed while in progress, or when the source record for the policy is missing.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400040019

SyncIQ has encountered a problem connecting to a target cluster that is configured in a SyncIQ policy.

Administrator action

1. Confirm that the source cluster can ping the target cluster.
2. Confirm that a traceroute can be completed to the target cluster.
3. Test to confirm that TCP ports 5666, 5667, 2097, 2098, 3147, and 3148 are open on your network by running the following command from the OneFS command-line interface on either the source or the target cluster:

   ```
   telnet <IP_address> <port_number>
   ```

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400040020

The Recovery Point Objective (RPO) was exceeded for a SyncIQ policy.

Description

A replication job failed to complete within the time period specified by the SyncIQ policy.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- Verify that the SyncIQ policy is running.
- Adjust the policy's RPO alert interval to accommodate a longer policy run time.
- Review the policy's RPO and identify policy changes that might allow the replication to meet the configured RPO. For example, you can adjust schedules, workers, and bandwidth restrictions to reduce replication time.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400040021

A SyncIQ SnapRevert job resolved conflicts between WORM committed files.

Description

If WORM committed file conflicts are resolved as part of a SnapRevert job, OneFS will generate a report that provides details about the file conflicts. This event will specify the replication policy associated with the SnapRevert job, and provide the location of the conflict report.

Administrator action

Review the conflict report for additional information about the file conflicts.

This event will not resolve automatically. You must manually resolve the event group.

400040022

A SyncIQ policy failed to establish an encrypted connection with the target.

Administrator action

Review the report for additional information about the encryption errors.

The event will not resolve automatically. Manually establish an encrypted connection with the target.

400040023

SyncIQ encountered an error during service export.

Administrator action

Review the report for additional information about the service export errors.

400040024

A SyncIQ policy detected unsupported WORM settings on the target.

Description

If there are unsupported WORM settings on the target, OneFS generates a report that provides details about the policy settings.

Administrator action

Review the report for additional information about WORM settings.

The event will not resolve automatically. Manually adjust the WORM settings.

400040025

A SyncIQ policy is waiting for the Cloudpools preparation of a stubbed LIN.

Administrator action

The event will not resolve automatically.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400040026

Maximum file name length support differs between SyncIQ source and target cluster.

Description

When syncing from a cluster with filenames longer than 255 bytes, to a cluster where support for longer names has not yet been enabled, the sync will fail.

Administrator action

Both clusters must have long-name support enabled for SyncIQ to continue to function normally. This may require an upgrade of the cluster that does not have long-name support enabled.

400050001

This event was generated as a test.

Description

A user requested that a test event be generated, either through the OneFS web administration interface or by running the following command:

```
isi event test create
```

Administrator action

This message is informational. No action is required.

400050002

This event was generated as a test.

Description

A user requested that a test event be generated, either through the OneFS web administration interface or by running the following command:

```
isi event test create
```

Administrator action

This message is informational. No action is required.

400050004

This is a heartbeat event that confirms that the event system is healthy.

Description

In order to confirm that the system is operating correctly, test events are automatically sent every day, one event from each node in your cluster. These are referred to as heartbeat events and are reported to an event group named Heartbeat Event.

Administrator action

This message is informational. No action is required.

400060001

The AVScan service is enabled, but a URL to an antivirus ICAP server has not been entered.

Administrator action

1. Perform one of the following tasks:
   - If you do not intend to configure a supported ICAP antivirus server, disable the antivirus service through the OneFS web administration interface.
   - Add an ICAP antivirus server through the web administration interface. For more information about the ICAP server, see the OneFS Web Administration Guide.
2. Cancel the existing event.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400060002

The cluster cannot reach any antivirus ICAP servers or the ICAP server is unresponsive.

Administrator action

1. Attempt the following possible solutions in the order listed. If a solution resolves the issue, there is no need to perform the subsequent solutions.
   - Make sure the cluster can reach the ICAP server through the network connection.
   - If the network connection is working as expected, determine if the ICAP server is responsive.
2. Cancel the existing event.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400060004

The virus-scanning software has identified a file that is infected by a virus.

Administrator action

Review the antivirus scan report and address the quarantined files as needed.

400060101

No configured CEE/CAVA servers.

Description

The anti-virus service on the cluster has stopped because no external anti-virus servers are configured.

Administrator action

Configure one or more external anti-virus servers, and then check the anti-virus service.

400060102

All CEE/CAVA anti-virus servers disabled.

Description

The anti-virus service on the cluster has stopped because no external anti-virus servers have been enabled.

Administrator action

Enable one or more external anti-virus servers, and then check the anti-virus service.

400060103

Not enough CEE/CAVA Servers for cluster size.

Description

Some nodes in the cluster cannot connect to the anti-virus service because the maximum number of available connections has been exceeded. Anti-virus servers can support a maximum of 20 connections.

Administrator action

Add additional anti-virus servers.

400060104

All CEE/CAVA servers are offline.

All anti-virus servers on a node are currently reporting as offline.

Administrator action

Verify that the node is connected to the anti-virus servers and that the servers are operational.

400060105

The anti-virus software on the CEE/CAVA server has errors, or is working improperly.

Administrator action

Check the CEE installation guide and the anti-virus software vendor documentation for proper setup.

400060106

The CEE/CAVA server is offline.

Description

One node reports that anti-virus server is offline.

Administrator action

Verify that the anti-virus server is connected to the node and that the server is operational.

400060107

All access zones have CEE or CAVA anti-virus disabled.

Description

The anti-virus service is disabled on all access zones in the cluster.

Administrator action

This event is informational and does not require any action.

400060108

The anti-virus service found an infected file.

Description

A file was found to be infected by the anti-virus server.

Administrator action

This event is informational and does not require any action.

400060109

The anti-virus access zone is missing.

Description

The required AvVendor access zone is missing. Anti-virus scanning cannot occur without the access zone.

Administrator action

To create the AvVendor access zone, disable, and then enable CAVA anti-virus.

400060110

The anti-virus IP Pool is missing or is misconfigured.

Description

The anti-virus IP Pool is misconfigured or missing. Anti-virus scanning cannot occur until the IP Pool is properly configured.

Administrator action

Ensure that the CAVA pool is properly configured, and that it displays in the `isi antivirus cava settings view` command output.

400060111

The CAVA agent that is installed on the Windows server is the wrong version.

Administrator action

Update the CEE or CAVA software on the anti-virus server to the correct version.

400060112

The required SMB service is unavailable on a node.

Administrator action

Check the SMB service on the node and resolve any issues.

400060113

The CAVA Filter Driver is offline.

Description

The CAVA Filter Driver on a node is not responding.

Administrator action

Verify that the filter driver is running by using the `/usr/likewise/bin/lwsm status flt_avscan` command, and resolve any issues.

400070004

An evaluation license for a OneFS software module is scheduled to expire soon.

Administrator action

To purchase the software module before the evaluation license expires, contact your sales representative.

400070005

An evaluation license for a OneFS software module has expired.

Administrator action

To purchase the software module, contact your sales representative.

400070006

Activation of a license was not completed.

Description

It has been over 90 days since the cluster was upgraded to a version of OneFS that requires a software license.

Administrator action

You must complete the license activation process. For more information on how to obtain a signed license file from Software Licensing Central (SLC), refer to the Licensing section of the OneFS Web Administration Guide or OneFS CLI Administration Guide.

400070007

The cluster is using software that is not licensed.

Description

The capacity of the cluster was recently upgraded. An updated license file is required.

Administrator action

You must generate a new license activation file that contains all of your current entitlements and submit it to Software Licensing Central (SLC). For more information on how to generate a new license activation file and obtain a signed license file, refer to the Licensing section of the *OneFS Web Administration Guide* or *OneFS CLI Administration Guide*.

400080001

A firmware upgrade has failed: {msg}

Administrator action

Attempt to reapply the firmware by running the following command from the node that reported the error:

```
isi upgrade firmware start --nodes-to-upgrade <local-node-lnn>
```

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400090001

This event is generated once each month to provide general cluster information.

Administrator action

This message is informational. No action is required.

400090002

The cluster contains a mix of encrypting nodes and non-encrypting nodes.

Description

You are currently migrating from a non-encrypted cluster to an encrypted cluster. If data is written to the cluster while this event is active, the new data might not be encrypted.

Administrator action

Confirm that the event has cleared after Professional Services finishes the cluster migration.

400090003

Secure Remote Support (SRS) is not configured.

Description

This is a recurring event that notifies you if SRS is not configured.

Administrator action

Enable SRS according to the instructions in the OneFS CLI Administration Guide and the OneFS Web Administration Guide.

If you do not want to enable SRS, you can disable this recurring SRS event by running the following command:

```
isi_monthly_esrs_disabled_alert --disable-event
```

400090004

The cluster lost connection to the Secure Remote Support (SRS) gateway server.

Description

In order for the cluster to activate SRS services, OneFS must communicate with the SRS gateway server. This event might be the result of the following issues:

- There is an SRS configuration error in OneFS.
- The cluster is unable to communicate with the SRS gateway server.

Administrator action

Verify your SRS configuration in OneFS. For information regarding SRS configuration, refer to the OneFS Web Administration Guide or the OneFS CLI Administration Guide.

Confirm that your SRS gateway server is powered up and connected to the cluster's external network.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400100001

The job state changed.

Administrator action

This message is informational. No action is required.

400100002

Job engine phase begin.

Administrator action

This message is informational. No action is required.

400100003

A job phase ended.

Administrator action

This message is informational. No action is required.

400100004

The job failed.

Administrator action

This message is informational. No action is required.

400100005

A job policy event occurred.

Administrator action

This message is informational. No action is required.

400100006

Job {job_type} failed to start as scheduled.

Administrator action

Another instance of the job is still running. Investigate why the previous instance of the job failed, and then run the job manually.

400100007

A job engine event occurred. The cluster is full and data can no longer be written.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400100008

A job engine event occurred. A file write operation is stalled, or writing very slowly to the cluster.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400100009

One or more nodes have been excluded from participating in this job.

Description

The job engine was configured to exclude one or more nodes from participating in this job.

Administrator action

This message is informational. No action is required.

400100010

One or more nodes that do not exist have been excluded from participating in this job.

Description

The job engine was configured to exclude nodes from this job, but one or more does not exist.

Administrator action

Examine job engine configuration and ensure that excluded nodes are configured properly

400100011

The cluster must be restriped, but FlexProtect is not running.

Description

The cluster contains one or more failed devices, and a FlexProtect operation needs to be run. FlexProtect will start automatically in the case of drive failures; however, in the case of node failures, FlexProtect requires user intervention in order to start.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400110001

The system is running low on memory, and the specified process was stopped to free memory.

Administrator action

1. Identify any new workflows or determine whether an increased number of clients are accessing the cluster. Make sure that the load is spread equally across all nodes through SmartConnect or an external load balancer.
2. If any changes have been made to the HTTPD configuration to increase the supported number of connections, revert those changes.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400120001

One of the boot disks is unhealthy, and the boot data is no longer being mirrored across the two boot disks. Loss of the remaining boot disk will lead to node failure.

Administrator action

Boot disks are not customer-replaceable components. Contact Dell EMC PowerScale Technical Support.

400130001

The NFS export rules are configured in such a way that the client cannot mount the path.

Administrator action

If you want the client to have permission to mount the specified directory, modify your NFS export rules to permit access. For more information, see the NFS section of the OneFS Web Administration Guide.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400130002

When processing NFS export rules, an attempt to look up the DNS name for the specified host failed.

Administrator action

1. Make sure that the DNS server includes an entry for the reported host name, and that non-PowerScale hosts can resolve the host name.
2. Check for misspellings and typographical errors in the cluster NFS configuration. Where possible, use a fully qualified domain name (FQDN) for every host in the NFS export rules. If you are not using a FQDN, make sure that each full domain name for the host is part of the DNS search path in the cluster configuration.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400140001

The NFSv4 server could not look up the user or group name to map to a user id (UID) or group id (GID).

Description

NFSv4 uses the string 'user@domain' instead of numeric UIDs and GIDs. This event occurs if there is a mismatch between the client domain name and the NFSv4 server domain name, or if the user database configuration is incorrect.

For example, the domain name stored in the `idmap.conf` file on the client might not match the NFSv4 domain name of the cluster (vfs.nfsrv.nfsv4.domainname).

Administrator action

1. Verify that the `idmap.conf` file is not missing or corrupted on the client.
2. Compare the value for *domainstring* in the `idmap.conf` file and the value for the cluster NFSv4 domain name that is returned from the following cluster sysctl command:

   ```
   isi_for_array -s sysctl vfs.nfsrv.nfsv4.idmap_replacedomain
   ```

3. If these values are different, update the `idmap.conf` file to match the domain of the cluster. Alternatively, you can match the *domainstring* and the value for the NFSv4 domain name of the cluster in the two files by running the following sysctl command:

   ```
   isi_sysctl_cluster vfs.nfsrv.nfsv4.idmap_replacedomain=1
   ```

4. Verify that Active Directory, NIS, and LDAP servers are configured correctly.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400140002

NFS could not translate a 64-bit cookie to a 32-bit cookie.

Description

OneFS readdir cannot translate 64-bit cookies to 32-bit for a directory.

Administrator action

Reduce the number of entries in the directory that the readdir is targeting. This event might be the result of a directory with a large number of entries.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400140003

To use the NFSv3-over-RDMA feature, the cluster must have an RDMA-capable front-end Network Interface Card.

Administrator action

Please add an RDMA-capable node, a node with RDMA-capable front-end Network Interface Card, or disable the NFSv3-over-RDMA feature by running the command:

```
isi nfs settings global modify --nfsv3-rdma-enabled=false
```

400150001

A OneFS upgrade started.

Administrator action

This message is informational. No action is required.

400150002

A OneFS upgrade finished successfully.

Description

This event appears when a OneFS upgrade or rollback finishes successfully.

Administrator action

This message is informational. No action is required.

400150003

A OneFS upgrade is in progress.

Description

A OneFS upgrade process has been running for over a week.

Administrator action

Confirm that the upgrade is making progress. If you feel that the upgrade process stopped, contact Dell EMC PowerScale Technical Support.

400150004

A step in the OneFS upgrade process is taking longer than expected.

Administrator action

Confirm that the upgrade is still making progress. If you feel that the upgrade process has discontinued, contact Dell EMC PowerScale Technical Support.

400150005

The rollback of a OneFS upgrade started.

Administrator action

This message is informational. No action is required.

400150006

Agent not ready.

Description

The Agent is commanded by the Supervisor to execute a hook or command, but if the node does not have superblock quorum and quorum, the Agent will not start the hook. If the hook does not start, the upgrade Supervisor re-sends the command at periodic intervals indefinitely.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400150007

Upgrade Hang - unable to communicate with Upgrade Agent on devids: {devids}

Description

By design, the upgrade framework design prevents any hooks or commands from starting if there are unresponsive nodes. Specifically, if the Agent on any node does not reply to status commands from the Supervisor. As long as this condition persists, all upgrade process stops.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400150008

Hook/command running too long.

Description

The upgrade hook/command takes much too long to complete.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400150009

A parallel upgrade is stalled at the `PendingReboot` hook.

Description

Upgrade Alert - disk pool db and LKF failure domain mismatch - unable to reboot nodes without potential client disruption.

Administrator action

If a parallel OneFS upgrade is stalled, and the nodes are stuck at the `PendingReboot` hook, contact Dell EMC PowerScale Technical Support. If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400150010

Upgrade stalled on node - unable to initiate a reboot without risking Data Unavailability (DU).

Description

There are node(s) or drive(s) in a degraded or down state that are preventing reboot.

Administrator action

Identify which nodes or drives are down or degraded by running: `isi_group_info` and either repair or smartfail the down or degraded nodes out of the cluster.

If temporary DU is permissible, run `isi upgrade unblock` to allow nodes to reboot when ready.

400150011

Upgrade Drain Alert - unable to reboot node without potential client disruption

Description

SMB clients are connected to a 'draining' node, which is stalling the upgrade by preventing a node from rebooting.

Administrator action

Identify if any clients are still connected to the node by reviewing the Upgrade Status page on the WebUI, or by running `isi smb sessions list` on the draining node. Alternatively, you can reboot the node.

400150012

Error installing HEALTHCHECK patch

Description

There was an error while installing the healthcheck patch.

Administrator action

Determine if you are at the most recent patch level and upgrade if necessary.

400151001

No secure image was found for use in expanding the cluster.

Description

An install image that matches the current committed version is required in order for nodes to be joined to the cluster.

Administrator action

Import the correct image by running the following command, where `<example.isi>` is the file name of the image you want to import: `isi upgrade catalog import <example.isi>`.

400160001

The cluster cannot reach an external Common Event Enabler (CEE) server, or the CEE server is unresponsive.

Administrator action

1. Ping the CEE server.
   - If the ping operation fails, confirm that network connectivity exists between the cluster and the CEE server.
     - If you can establish contact between the cluster and the CEE server, attempt to ping the CEE server again to see if the issue has resolved.
     - If you cannot establish contact between the cluster and the CEE server, attempt to resolve any network connectivity issues.
   - If the ping operation succeeds, verify that your CEE server is online and functional.
2. Cancel the existing event.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400160002

Audit System cannot provide service.

Description

The audit system cannot provide service.

Administrator action

Verify that the audit services is enabled by running: `isi services -a isi_audit_d enable`.

400160005

Audit daemon failed to persist one or more events.

Description

The audit daemon failed to save one or more events. The daemon will continue to retry writing the events to storage.

Administrator action

Verify that the `/ifs/.ifsvar/audit` directory allows for writes.

400170001

A periodic check against the store finds expiring certificates.

Administrator action

1. To refresh or renew SSL certificates, follow the instructions in *Isilon OneFS: How to replace or renew the SSL certificate used for the Isilon web administration interface*, article 000157711.
2. Cancel the existing event.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400170002

A periodic check against the store finds expired certificates.

Administrator action

1. To refresh or renew SSL certificates, follow the instructions in *Isilon OneFS: How to replace or renew the SSL certificate used for the Isilon web administration interface*, article 000157711.
2. Cancel the existing event.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400180001

Inline dedupe allocation failed on node {lnn}, occurrence {occurrence}

Description

Memory allocation failed for the inline dedupe index.

Administrator action

1. Disable inline dedupe.
2. Free up memory by running the `isi_flush` command.
3. Enable inline dedupe.
4. If the issue does not resolve, restart the node.

400180002

Inline dedupe allocation in progress on node {lnn}, occurrence {occurrence}

Description

A request for inline dedupe to be enabled is pending allocation of the index.

Administrator action

Wait until the event clears. The dedupe process succeeds or fails within a few minutes.

400180003

Inline dedupe allocation not supported on node {lnn}, occurrence {occurrence}

Description

The node (must be an F810) is not permitted to enable inline dedupe.

Administrator action

This is an informational event only, and no action is required.

400180004

Inline dedupe running degraded with smaller index on node {lnn}, occurrence {occurrence}.

Description

Inline dedupe was unable to acquire all the memory for the preferred index size so has reduced the size of the index.

Administrator action

The inline dedupe index issue is not critical, and can be safely ignored. Inline dedupe is enabled and fully effective but the index is fragmented. The issue might be resolved by completing the following steps:

1. Disable inline dedupe.
2. Free up memory by running the `isi_flush` command.
3. Enable inline dedupe.
4. If the issue does not resolve, restart the node.

400180005

Inline dedupe index has non-standard layout on node {lnn}, occurrence {occurrence}

Description

Inline dedupe was unable to allocate the preferred structure for the index so has changed the layout to preserve the requested size.

Administrator action

The inline dedupe index issue is not critical, and can be safely ignored. Inline dedupe is enabled and fully effective but the index is fragmented. The issue might be resolved by completing the following steps:

1. Disable inline dedupe.
2. Free up memory by running the `isi_flush` command.
3. Enable inline dedupe.
4. If the issue does not resolve, restart the node.

400190001

Invalid dedupe directory {path}

Description

SmartDedupe has been configured with an invalid path, or a non-existent directory.

Administrator action

Check the paths in the SmartDedupe configuration and verify that they exist:

- If the directory exists but is not accessible then raise contact Dell EMC Technical support. If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.
- If the directory does not exist then remove it from the dedupe configuration.

400200001

Security verification check failed

Administrator action

This message is informational. No action is required.

400200002

Security verification check ran successfully

Administrator action

This message is informational. No action is required.

400210001

The encryption key manager for self-encrypting drives (SED) is unable to start on the indicated node.

Administrator action

Do not reboot the node. There are several possible causes for this error.

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400210002

The encryption key manager for Cloudpools is unable to start on the indicated node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

400210003

Key manager control key is unavailable in hardware.

Administrator action

A service request (SR) has been opened on your behalf. Technical Support will contact you.

400210004

KMIP Server returned an error.

Administrator action

Check KMIP configurations and logs on the KMIP server.

400210005

Network error occurred when reaching KMIP server.

Administrator action

Check the network configuration and verify that all nodes can connect to the key management interoperability protocol (KMIP) server.

400210006

KMIP Key Migration failed.

Administrator action

Check that the KMIP configurations are valid and that all nodes can contact the KMIP server.

400210007

A certificate for a KMIP server is about to expire.

Administrator action

Renew or replace the expired certificate.

400210008

A certificate for a KMIP server has expired.

Administrator action

Renew or replace the expired certificate.

400220000

PDM degraded, too many operations

Description

Multiple Policy Domain Manager (PDM) operations are pending. PDM operations run in the background. These are completed by the `DomainTag` job and `isi_pdm_d` daemon. This event can happen if operations are being created faster than expected, or if the operations are not completing. This can lead to performance degradation affecting writes within directories that are governed by snapshots.

Administrator action

- Warning: Enable the `DomainTag` job if it has been disabled. The event is likely to resolve without further action.
- Critical: Enable `DomainTag` job if it has been disabled. Open a Service Request and attach any information received in email regarding the event notification.

400230001

Invalid configuration changes that are made to SSHD by the user.

Description

Last Attempt on Modifying SSHD settings failed due to: [*/etc/ssh/sshd_config line 30: Directive 'Subsystem' is not allowed within a Match block*].

Administrator action

Correct the configuration changes as per the reported error.

400240000

S3 Service failed to start.

Administrator action

Check cluster status. If the cluster is healthy, but S3 is failing, contact Dell EMC technical support.

400240001

Identity query failed user=1000 to name status=STATUS_ACCESS_DENIED.

Description

S3 failed to resolve name from ID.

Administrator action

Verify that the naming service is operating correctly.

400240002

S3 name query failed user=alice to id status=STATUS_ACCESS_DENIED.

Description

S3 failed to resolve ID from name.

Administrator action

Verify that the naming service is operating correctly.

400240003

S3 could not parse mpu info for bucket id : 123456. Upload Id 987654. SBT may be broken.

Description

Could not parse mpu info for bucket id.

Administrator action

Contact Dell EMC technical support to open a Service Request (SR).

400240004

S3 key in SBT is invalid. SBT may be broken. Current Basekey = a/b/c.

Description

S3 key in SBT is invalid. SBT may be broken.

Administrator action

Contact Dell EMC technical support to open a Service Request (SR).

400240005

S3 key in SBT has maxed out. SBT may be full for bucket - 123456.

Description

S3 key in SBT has maxed out. SBT may be full for bucket.

Administrator action

Contact Dell EMC technical support to open a Service Request (SR).

400250000

Noncompatible, user-specified patches were found, and ignored.

Description

The specified patches have conflicts with the embedded patches in the prepatched image. The conflicts were ignored, and the upgrade continued.

Administrator action

This event is informational, no action is required.

400260000

PW account was updated.

Description

Raise a CELOG event to notify system administrators when the pw account is updated.

Administrator action

This event is informational, no action is required.

500010001

The SmartQuotas module has notified a user of a quota violation.

Description

You can disable notifications for this event or modify the SmartQuotas rules. For information about configuring SmartQuotas rules, see the OneFS Web Administration Guide.

Administrator action

This message is informational. No action is required.

500010002

The SmartQuotas notification functionality failed.

Description

This error typically occurs as a result of one or more of the following conditions:

- The mail server is configured incorrectly.

- The mail server or the authentication server is down.
- A quota address mapping rule is configured incorrectly.

Administrator action

Review the settings for quota notification rules and correct any apparent errors. For information about configuring SmartQuotas, see the OneFS Web Administration Guide.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

500010003

The SmartQuotas configuration file is corrupt or invalid.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

500010004

The SmartQuotas configuration file is corrupt or invalid.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

500010005

The SmartQuotas module failed to generate a requested quota report.

Administrator action

Review your SmartQuotas report settings and make sure that the settings are configured correctly. For information about SmartQuotas, see the OneFS Web Administration Guide.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

600010001

The snapshot daemon failed to create the scheduled snapshot.

Description

If the cluster is split, this message might appear on the minority group, the group that has fewer than half of the nodes. In this case, the error persists until the cluster is healthy, and you can safely ignore the error.

Administrator action

1. Determine whether a node is in the minority or majority group, by running the following command from the node that is reporting the error:

   ```
   sysctl efs.gmp.has_quorum
   ```

   - If the command returns `0`, the error occurred on the minority group. The message might continue until the cluster is healthy. No further action is required.
   - If the command returns `1`, the node is in the majority group. Proceed to step 2.
2. Determine whether the number of snapshots exceeds the system-wide and directory limits. (The system-wide limit is 20,000, and the directory limit is 1,000.)
3. If the number of snapshots is at or exceeds the limits, you can delete the extraneous snapshots. If the snapshots are within system limits, the event will automatically clear.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

600010002

The snapshot daemon failed to delete an expired snapshot.

Description

The system cannot remove an expired snapshot lock. This error can occur when a disk is unwritable.

If the cluster is split, this error might occur on the minority group, or the group that contains fewer than half of the nodes. In this case, the message persists until the cluster is healthy, and you can safely ignore the error.

Administrator action

1. Determine whether a node is in the minority or majority group, by running the following command from the node that is reporting the error:

   ```
   sysctl efs.gmp.has_quorum
   ```

   - If the command returns `0`, the error occurred on the minority group. The message might continue until the cluster is healthy. No further action is required.
   - If the command returns `1`, the node is in the majority group. Proceed to step 2.
2. If the error occurred on the majority group, perform the following tasks:
   - Confirm that the cluster contains free disk space. If the cluster is more than 99 percent full, delete files or add storage capacity to the cluster. You can view the percentage of available disk space on the cluster by running the following command:

     ```
     isi status -q
     ```
   - Verify that the `isi_job_d` process is running on all nodes by first logging in to any node through a secure shell (SSH) connection or the serial console and then running the following command:

     ```
     isi_for_array -s 'pgrep isi_job_d|wc -l'|grep '[^0-9]0$'
     ```

     Nodes that are listed in the output do not have the isi_job_d process running, and cannot run any system jobs.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

600010003

The snapshot daemon failed to remove a snapshot lock.

Description

The system cannot remove an expired snapshot lock. This error can occur when a disk is unwritable.

If the cluster is split, this error might occur on the minority group, or the group that contains fewer than half of the nodes. In this case, the message persists until the cluster is healthy, and you can safely ignore the error.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

600010004

The snapshot_schedule.xml file is corrupt or unreadable.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

600010005

The amount of data stored on the cluster is approaching or has exceeded the snapshot reserve space.

Description

Exceeding the snapshot reserve space does not result in a failure to write snapshots to the cluster. The system can write snapshots to any available disk space, and snapshots can exceed the snapshot reserve space. However, problems occur when the available space in the cluster is less than the snapshot reserve space. If the cluster exceeds the snapshot reserve space, all attempts to write non-snapshot data to the cluster fail.

Administrator action

Review cluster and snapshot usage data, and then perform either of the following tasks:

- Delete some snapshots to reduce the amount of snapshot reserve space in use.
- Disable the snapshot reserve space. On PowerScale clusters, snapshot reserve space is not required to write snapshots to disk.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700010001

The cluster time differs from the Windows Active Directory server.

Description

The timestamp between one or more nodes and the Active Directory host differs by at least five minutes. The time discrepancy can result in authentication failures due to mismatches with the Kerberos ticket timestamp.

Mismatched clock settings might be a result of one of the following:

- The cluster synchronized time with a domain controller whose clock is set incorrectly.
- The cluster is not synchronizing time with the domain controller.

Administrator action

1. Check the time on the Active Directory server, and then perform one of the following steps:
   - If the Active Directory time is incorrect, adjust the time on the Active Directory server.
   - If the Active Directory time is correct, adjust the cluster time to match the time on the Active Directory server through the OneFS web administration interface.
2. If enabled, disable the Network Time Protocol (NTP). You cannot synchronize time through both an Active Directory server and NTP on the same cluster.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700010003

The Windows time server could not be contacted. The cluster time is not synchronized.

Administrator action

Check the time on the Windows time server.

- If the Windows time server is incorrect, adjust the time on the Windows time server.
- If the Windows time server is correct, adjust the cluster time to match the Windows time server time through the OneFS web administration interface.

700010004

An SMB upgrade error occurred, which might affect the behavior of the SMB service.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700010005

An authentication upgrade failure has occurred.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700020001

The Windows UID map range is full. Authentication might fail until the range is increased.

Description

The user ID (UID) range for mapping Microsoft Active Directory groups has run out of IDs and must be expanded. This condition might occur if there is a large Active Directory Services hierarchy or an unusually small UID range.

Administrator action

1. Verify the current UID range by running the following command:

   ```
   isi auth settings mapping view --zone=<zone_name>
   ```

2. Expand the UID range according to the instructions in the OneFS CLI Administration Guide and the OneFS Web Administration Guide.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700020002

The Windows GID map range is full. Authentication might fail until the range is increased.

Description

The group ID (GID) range for mapping Microsoft Active Directory groups has run out of IDs and must be expanded. This condition might occur if there is a large Active Directory Services hierarchy or an unusually small GID range.

Administrator action

1. Verify the current GID range by running the following command:

   ```
   isi auth settings mapping view --zone=<zone_name>
   ```

2. Expand the UID range according to the instructions in the OneFS CLI Administration Guide and the OneFS Web Administration Guide.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700020003

The Windows networking service failed to parse idmap rules.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700030001

The Active Directory account data that is stored on the cluster was deleted or damaged.

Administrator action

Unjoin and rejoin the Active Directory domain.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700030002

The Active Directory server is offline. Authentication services might be interrupted.

Description

The node cannot contact an authentication server for the specified domain. By design, the node periodically attempts to re-establish communication with the domain. If the node successfully connects to the domain, the event clears itself.

Administrator action

Confirm that TCP port 389 is open on your network.

If the event does not clear itself within five minutes or if the event recurs, perform the following steps on the node on which the issue occurred:

1. Ping the authentication server.
   - If the ping operation fails, confirm that network connectivity exists between the cluster and the authentication server.
     - If you can establish contact between the cluster and the authentication server, attempt to ping the authentication server again to see if the issue has resolved.
     - If you cannot establish contact between the cluster and the authentication server, attempt to resolve any network connectivity issues.
   - If the ping operation succeeds, verify that your authentication server is online and functional.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700030003

The node cannot perform read or write operations on the authentication database files.

Description

This event typically appears when a node does not have quorum or otherwise cannot access the authentication database files.

This event also sometimes appears when a node starts up. In that case, the event frequently resolves itself within five minutes. If the event resolves itself, no action is required.

Administrator action

If the event does not resolve itself within five minutes, review the node status through the PowerScale web administration interface and resolve any apparent issues.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700030004

The authentication service is unavailable.

Description

The node cannot contact an authentication server for the specified domain. By design, the node periodically attempts to re-establish communication with the domain. If the node successfully connects to the domain, the event clears itself.

Administrator action

If the event does not clear itself within five minutes or if it recurs, perform the following steps on the node on which the issue occurred:

1. Ping the authentication server.
   - If the ping operation fails, confirm that network connectivity exists between the cluster and the authentication server.
     - If you can establish contact between the cluster and the authentication server, attempt to ping the authentication server again to see if the issue has resolved.
     - If you cannot establish contact between the cluster and the authentication server, attempt to resolve any network connectivity issues.
   - If the ping operation succeeds, verify that your authentication server is online and functional.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700030005

An Active Directory service provider is missing a required SPN.

Administrator action

1. View the configured SPNs for the Active directory service provider by running the following command:

```
isi auth ads spn check <provider-name>
```

2. Repair any missing SPNs by running the following command:

```
isi auth ads spn fix <provider-name>
```

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700030006

The Active Directory machine is invalid.

Administrator action

Verify the password for the Active Directory machine username and attempt to rejoin to the Active Directory domain.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700040001

LDAP servers are offline. Authentication services might be interrupted.

Description

The node cannot contact an authentication server for the specified domain. By design, the node periodically attempts to re-establish communication with the domain. If the node successfully connects to the domain, the event clears itself.

Administrator action

If the event does not clear itself within five minutes or if it recurs, perform the following steps on the node on which the issue occurred:

1. Ping the authentication server.
   - If the ping operation fails, confirm that network connectivity exists between the cluster and the authentication server.
     - If you can establish contact between the cluster and the authentication server, attempt to ping the authentication server again to see if the issue has resolved.
     - If you cannot establish contact between the cluster and the authentication server, attempt to resolve any network connectivity issues.
   - If the ping operation succeeds, verify that your authentication server is online and functional.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700050001

NIS servers are offline. Authentication services might be interrupted.

Description

The node cannot contact an authentication server for the specified domain. By design, the node periodically attempts to re-establish communication with the domain. If the node successfully connects to the domain, the event clears itself.

Administrator action

If the event does not clear itself within five minutes or if it recurs, perform the following steps on the node on which the issue occurred:

1. Ping the authentication server.
   - If the ping operation fails, confirm that network connectivity exists between the cluster and the authentication server.
     - If you can establish contact between the cluster and the authentication server, attempt to ping the authentication server again to see if the issue has resolved.
     - If you cannot establish contact between the cluster and the authentication server, attempt to resolve any network connectivity issues.
   - If the ping operation succeeds, verify that your authentication server is online and functional.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

700100001

An LWIO parameter is invalid.

Administrator action

This message is informational. No action is required.

800010002

The system detected a metadata referential integrity error that requires manual intervention to resolve.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

800010003

An Isilon Data Integrity (IDI) failure was detected.

Description

The system cannot verify data integrity. The system will attempt to resolve this issue automatically through the Dynamic Sector Recovery (DSR) process. If the DSR process is unsuccessful, manual intervention is required.

When this event occurs, you will see two separate events with the same event ID of 800010003.

- The first event states that an IDI was detected and that the system is attempting to resolve the issue through the DSR process. This event is sent with a `Critical` severity level.
- The second event provides information that will assist Dell EMC PowerScale Technical Support with the debugging process. This event is sent with an `Info` severity level.

Administrator action

For information about this event, see the following articles on Dell EMC Online Support:

- *Isilon Event notification: Detected IDI failure, attempting DSR - Event ID: 800010003*, article 000041456

- *Isilon Dynamic sector recovery (DSR) failures - Event ID: 800010005*, article 000041433

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

800010004

The cluster has encountered a file system error.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

800010005

The Dynamic Sector Repair (DSR) process failed to resolve a data verification error.

Description

The system has failed to resolve this issue automatically. Manual intervention is required.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

800010006

A node has reported that the number of available file descriptors is approaching the maximum limit.

Administrator action

1. Identify a list of the process types with the largest number of file descriptors by running the following command from the OneFS command-line interface:

```
fstat|awk '{ print $2 }'|sort|uniq -c|sort -rn|head
```

2. Contact Dell EMC PowerScale Technical Support and provide the output from the command.

800010007

An Isilon Data Integrity (IDI) network checksum error was detected.

Description

One or more nodes appear to have faulty back-end networking hardware within the node, or a faulty motherboard.

Administrator action

Contact Dell EMC PowerScale Technical Support to determine if one of the components or motherboard must be replaced on the node or nodes.

800010008

The NVRAM journal is larger than the journal backup partition.

Description

The system has resized the journal partition.

Administrator action

This message is informational. No action is required.

800010009

There was an error calculating the partition size of the NVRAM journal backup.

Description

The system is resetting the partition to the default size of 512 MB.

Administrator action

This message is informational. No action is required.

800010010

A node was unable to verify the backup copy of its journal on its peer node.

Description

There are two possible issues that might result in this event:

- The local copy of a node's journal is not valid.
- There was an error when a node tried to verify the local copy of its journal against the mirror copy of its journal on its peer node.

Administrator action

If the link between peer nodes has been temporarily interrupted then this event will resolve when the link between the two nodes is reestablished.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

1100000001

A CloudPools network connection failed.

Description

The network connection for the specified account failed and CloudPools is unable to access files in the cloud provider.

Administrator action

This event might be the result of internal or external network issues. Check to make sure that internal network connections are healthy, then make sure you are able to reach the cloud provider.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

1100000002

A CloudPools user failed to authenticate.

Description

The cloud provider was not able to authenticate the specified cloud account.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- If the password for the cloud account was changed recently, confirm that the change is reflected in the local CloudPools file.
- Confirm that the cloud account is not attempting to log in with an incorrect username or password.
- Confirm that the cloud account's username or password has not been removed from the system.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

1100000003

A cloud account attempted to access a file that it does not have permissions for.

Description

A cloud account attempted to access a file, but is not authorized to access the file stored in the cloud provider.

Administrator action

Modify the cloud account to allow access to the requested file type.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

1100000004

An Amazon S3 telemetry reporting bucket was not found.

Description

CloudPools attempted to access usage reports from an Amazon S3 cloud provider. The S3 telemetry reporting bucket, where usage reports are stored, could not be found.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- Confirm that the specified bucket was not deleted from the cloud. If the bucket was deleted, determine if the bucket was deleted in error or if a security breach allowed direct access to the cloud provider account.
- Confirm that the URL of the cloud provider has not changed.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

1100000005

A Cloudpool capacity threshold was exceeded.

Descriptions

By default, this event first appears when the amount of data on the cloud provider reaches 70%. The event will notify you when the following capacity thresholds are exceeded:

70%	Informational
80%	Warning
90%	Critical
95%	Emergency

Administrator action

Reduce the amount of data stored on your cloud provider.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

1100000006

CloudPools data is corrupted.

Description

The MD5 hash does not match. Data retrieved from the cloud provider is corrupted.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

1100000007

CloudPools no usable account found.

Description

CloudPools is unable to write data to the cloud storage account because it is disabled.

Administrator action

Confirm that the account is disabled by running the following command:

```
isi cloud accounts view <id>
```

If the account is disabled, enable the account by running the following command:

```
isi cloud accounts modify <id> --enabled yes
```

1100000008

CloudPools could not verify a provider certificate.

Description

The certificate for the specified cloud provider is not properly installed or is not valid.

Administrator action

Contact the cloud provider to obtain a valid certificate.

1100000009

Invalid CloudPools gconfig settings.

Description

CloudPools has invalid gconfig settings and requires modification.

Administrator action

A service request (SR) has been opened on your behalf.

5

Hardware events

This section contains the following topics:

Topics:

- 900020033
- 900020034
- 900020035
- 900060001
- 900060002
- 900060003
- 900060004
- 900060005
- 900060006
- 900060007
- 900060008
- 900060009
- 900060010
- 900060011
- 900060012
- 900060013
- 900060014
- 900060015
- 900060016
- 900060017
- 900060018
- 900060019
- 900060020
- 900060021
- 900060022
- 900060023
- 900060024
- 900060025
- 900060026
- 900060027
- 900060028
- 900060029
- 900060030
- 900060031
- 900060032
- 900060033
- 900060034
- 900060035
- 900060036
- 900060037
- 900060038
- 900060039
- 900060040
- 900080001
- 900080002
- 900080003
- 900080004
- 900080005
- 900080006
- 900080007
- 900080008
- 900080009
- 900080010
- 900080011
- 900080012
- 900080013

- 900080014
- 900080015
- 900080016
- 900080017
- 900080018
- 900080019
- 900080020
- 900080021
- 900080022
- 900080023
- 900080024
- 900080025
- 900080026
- 900080027
- 900080028
- 900080029
- 900080030
- 900080031
- 900080032
- 900080033
- 900080034
- 900080035
- 900080036
- 900080037
- 900100001
- 900100004
- 900100018
- 900100019
- 900100020
- 900100021
- 900100022
- 900100023
- 900100024
- 900100025
- 900100026
- 900100027
- 900100028
- 900100029
- 900100030
- 900100031
- 900100032
- 900110001
- 900110002
- 900110003
- 900110004
- 900110005
- 900120001
- 900120002
- 900120003
- 900120004
- 900120005
- 900130001
- 900130002
- 900130003
- 900130004
- 900130005

- 900130006
- 900130007
- 900130008
- 900130009
- 900130010
- 900130011
- 900130013
- 900130014
- 900130015
- 900140001
- 900140002
- 900140003
- 900140004
- 900140005
- 900150001
- 900160001
- 900160002
- 900160003
- 900160004
- 900160005
- 900160006
- 900160007
- 900160008
- 900160009
- 900160010
- 900160011
- 900160012
- 900160013
- 900160014
- 900160015
- 900160016
- 900160017
- 900160018
- 900160019
- 900160020
- 900160021
- 900160022
- 900160023
- 900160024
- 900160100
- 900160102
- 900160101
- 900170001
- 900170002
- 900180001
- 900180002
- 900180003
- 900180004
- 900180005
- 900180006
- 900180007
- 900180008
- 900180009
- 900180010
- 900180011
- 900180012

- 900180013
- 900180014
- 900180015
- 900180016
- 900180028
- 900180029
- 900180030
- 900180031
- 900180032
- 910100001
- 910100002
- 910100003
- 910100004
- 910100005
- 910100006
- 910100007
- 920100000
- 920100001
- 920100002
- 920100003
- 920100004
- 920100005
- 920100006
- 920100007
- 920100008
- 920100009
- 930100000
- 930100001
- 930100002
- 930100003
- 930100004
- 930100005
- 930100006
- 940100001
- 940100002

Hardware events overview

Hardware events provide information about hardware-specific status, such as voltage, power supply, and fan speed issues.

900010001

There is an error on the node motherboard, such as a faulty clock battery.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900010002

The node needs a battery replacement.

Administrator action

View the status of each battery by running the following command from any node in the cluster:

```
isi_for_array -s 'isi_hw_status -bg | grep Battery'
```

- For legacy hardware follow the instructions to initiate a battery test by following the instructions in *Understanding Isilon node battery testing*, article 000016079.
- For NL410, X210, S210, X410, and HD400 nodes, contact Dell EMC PowerScale Technical Support for troubleshooting.

900010003

There is an issue with the NVRAM card.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900010004

A sensor has detected that the node chassis is open.

Description

This event typically appears when maintenance is being performed on the inside of the node, while the node is powered on. Or, this event might appear if one of the NVRAM battery trays were pulled out.

Administrator action

1. Make sure that the battery tray is properly inserted.
2. If the node chassis is open, close the chassis.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900010005

A memory, PCI, or PCIe bus error has occurred in the node.

Administrator action

Troubleshooting is required to determine if a hardware component must be replaced.

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900010006

A memory, PCI, or PCIe bus error has occurred in the node.

Administrator action

Troubleshooting is required to determine if a hardware component must be replaced.

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900010007

Correctable memory error rate exceeded for <DIMM>. Replace the DIMM as soon as possible.

Administrator action

A service request (SR) has been opened on your behalf. Technical Support will contact you.

900010008

A hardware issue was detected with the I2C bus.

Description

The I2C bus is a controller or bus that carries information from various sensors (for fan speed, power supply voltage, and temperature) in a node chassis.

To resolve this issue, you must shutdown the node, disconnect the power cables, and then press the power button on the node to discharge any remaining stored power in the node.

It is not critical to complete these steps immediately, but this event will continue to appear until the issue has been addressed. While this event is active, the node will not report correct values for the temperature, fans, or power supply health on the node.

Administrator action

1. Connect to the affected node through SSH or serial cable.
2. Shut down the node by running the following command:

```
shutdown -p now
```

3. Wait for the node to shut down, and then disconnect both power supply cables.
4. Press the power button on the node to discharge any remaining stored power.
5. Re-connect the power cables and then start the node.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900010009

The node has the wrong drive ratio.

Description

An HDD or SSD has been replaced by a drive of the other type, which changed the configured ratio of SSD to HDD. This event might be generated by a drive failure, which can also affect the configured ratio.

Administrator action

Determine the types of drives that are on the node by running the following command:

```
isi devices drive list
```

Any drives that are not functioning correctly will be flagged.

- If the drives are incorrect, replace the drives with the correct types. You can download the latest version of drive replacement guide from the Online Support site.
- If the drives are correct, or if you need assistance obtaining the correct drive type, contact Technical Support.

900010010

The node has a 812 (3/4 chassis) SKU.

Description

This event is generated if the model number of the chassis is not properly updated after replacing a node chassis.

Administrator action

Contact Dell EMC PowerScale Technical Support to update the model number of the node.

900010011

The Baseboard Management Controller (BMC) or Chassis Management Controller (CMC) are unresponsive.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900010012

A fuse might have failed in the node.

Description

If a fuse has failed, the suitcase must be replaced. Do not shut down the node until a replacement suitcase arrives.

Administrator action

Contact Dell EMC PowerScale Technical Support for a potential suitcase replacement.

900010013

The firmware update failed for the specified device.

Administrator action

Attempt to update the firmware again by running the following command from the node that reported the event:

```
isi firmware update --local
```

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020001

A sensor in the front panel of a node has exceeded the specified threshold.

Description

This event can occur intermittently without harm to the system.

Administrator action

1. Cancel or quiet the event.
2. If the event recurs, shutdown and restart the node by completing the following steps:
   - Connect to the affected node through SSH or serial cable.
   - Shut down the node by running the following command:

     ```
     shutdown -p now
     ```

   - Wait for the node to shut down, and then disconnect both power supply cables.
   - Press the power button on the node to discharge any remaining stored power.
   - Reconnect the power cables and then start the node.
3. (HD400 only.) Re-seat the front panel connector by checking that the ribbon cable is properly attached and properly seated.
4. (All other nodes.) Re-seat the front panel.
5. Move the front panel from a functioning node to the affected node and see if the event clears.
6. Install the front panel from the affected node on another node to determine if the problem is with the front panel or with the node.

   If the problem follows the front panel, contact Technical Support to request a new front panel.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020002

A power supply fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900020003

A power supply fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900020004

A power supply fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900020005

A power supply fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900020006

A chassis fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900020007

A chassis fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900020008

A chassis fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900020009

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020010

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020011

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020012

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020013

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020014

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020015

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020016

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020017

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020018

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020019

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020020

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020021

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020022

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020023

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020024

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020025

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020026

The internal or ambient temperature around a node has exceeded the allowable threshold for the CPU.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020027

The internal or ambient temperature around a node has exceeded the allowable threshold for the CPU.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020028

The internal or ambient temperature around a node has exceeded the allowable thresholds for a power supply.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020029

The internal or ambient temperature around a node has exceeded the allowable thresholds for a power supply.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.

- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020030

The internal or ambient temperature around a node has exceeded the allowable threshold.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020031

The internal or ambient temperature around the front panel of a node has exceeded the allowable threshold.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020032

The internal or ambient temperature around a node has exceeded the allowable thresholds for the chassis.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020033

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900020034

The node is reporting less than the expected amount of physical memory.

Description

This event typically appears because a DIMM has failed, is poorly seated, or an incorrect type of DIMM is installed.

Administrator action

Contact Technical Support to determine if a DIMM replacement is required.

900020035

The internal or ambient temperature around a node has exceeded the allowable threshold for the CPU.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060001

A sensor in the front panel of a node has exceeded the specified threshold.

Description

This event can occur intermittently without harm to the system.

Administrator action

1. Cancel or quiet the event.
2. If the event recurs, shutdown and restart the node by completing the following steps:
   - Connect to the affected node through SSH or serial cable.
   - Shut down the node by running the following command:

     ```
     shutdown -p now
     ```

   - Wait for the node to shut down, and then disconnect both power supply cables.
   - Press the power button on the node to discharge any remaining stored power.
   - Reconnect the power cables and then start the node.
3. (HD400 only.) Re-seat the front panel connector by checking that the ribbon cable is properly attached and properly seated.
4. (All other nodes.) Re-seat the front panel.
5. Move the front panel from a functioning node to the affected node and see if the event clears.
6. Install the front panel from the affected node on another node to determine if the problem is with the front panel or with the node.

   If the problem follows the front panel, contact Technical Support to request a new front panel.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060002

A power supply fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900060003

A power supply fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900060004

A chassis fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900060005

A chassis fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900060006

A chassis fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900060007

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060008

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060009

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060010

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060011

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060012

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060013

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060014

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060015

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060016

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060017

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060018

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060019

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060020

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060021

The internal or ambient temperature around a node has exceeded the allowable threshold for the CPU.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060022

The internal or ambient temperature around a node has exceeded the allowable threshold for the CPU.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060023

The internal or ambient temperature around a node has exceeded the allowable threshold.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060024

The internal or ambient temperature around the front panel of a node has exceeded the allowable threshold.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060025

The internal or ambient temperature around a node has exceeded the allowable thresholds for the chassis.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060026

The internal or ambient temperature around a node has exceeded the allowable threshold for the CPU.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060027

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060028

The node is reporting less than the expected amount of physical memory.

Description

This event typically appears because a DIMM has failed, is poorly seated, or an incorrect type of DIMM is installed.

Administrator action

Contact Technical Support to determine if a DIMM replacement is required.

900060029

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060030

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060031

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060032

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060033

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060034

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060035

A power supply fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900060036

A power supply fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900060037

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060038

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060039

The internal or ambient temperature around a node has exceeded the allowable thresholds for a power supply.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900060040

The internal or ambient temperature around a node has exceeded the allowable thresholds for a power supply.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.

- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080001

A sensor in the front panel of a node has exceeded the specified threshold.

Description

This event can occur intermittently without harm to the system.

Administrator action

1. Cancel or quiet the event.
2. If the event recurs, shutdown and restart the node by completing the following steps:
   - Connect to the affected node through SSH or serial cable.
   - Shut down the node by running the following command:

     ```
     shutdown -p now
     ```

   - Wait for the node to shut down, and then disconnect both power supply cables.
   - Press the power button on the node to discharge any remaining stored power.
   - Reconnect the power cables and then start the node.
3. (HD400 only.) Re-seat the front panel connector by checking that the ribbon cable is properly attached and properly seated.
4. (All other nodes.) Re-seat the front panel.
5. Move the front panel from a functioning node to the affected node and see if the event clears.
6. Install the front panel from the affected node on another node to determine if the problem is with the front panel or with the node.

   If the problem follows the front panel, contact Technical Support to request a new front panel.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080002

A power supply fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900080003

A power supply fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900080004

A power supply fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900080005

A power supply fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900080006

A chassis fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900080007

A chassis fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900080008

A chassis fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900080009

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080010

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080011

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080012

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080013

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080014

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080015

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080016

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080017

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080018

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080019

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080020

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080021

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080022

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080023

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.

2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080024

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes

LED	Power status	Node type
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080025

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series

LED	Power status	Node type
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080026

The internal or ambient temperature around a node has exceeded the allowable threshold for the CPU.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080027

The internal or ambient temperature around a node has exceeded the allowable threshold for the CPU.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080028

The internal or ambient temperature around a node has exceeded the allowable thresholds for a power supply.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080029

The internal or ambient temperature around a node has exceeded the allowable thresholds for a power supply.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080030

The internal or ambient temperature around a node has exceeded the allowable threshold.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080031

The internal or ambient temperature around the front panel of a node has exceeded the allowable threshold.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.

- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080032

The internal or ambient temperature around a node has exceeded the allowable thresholds for the chassis.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080033

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080034

The node is reporting less than the expected amount of physical memory.

Description

This event typically appears because a DIMM has failed, is poorly seated, or an incorrect type of DIMM is installed.

Administrator action

Contact Technical Support to determine if a DIMM replacement is required.

900080035

The internal or ambient temperature around a node has exceeded the allowable threshold for the CPU.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080036

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900080037

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900100001

The NVRAM in the indicated node experienced a single-bit error. The error was automatically corrected by ECC.

Administrator action

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900100004

The PCIe lane width was not successfully negotiated.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900100018

The NVRAM board failed to respond to an Identify Controller command.

Description

During periods of I/O inactivity, an Identify Controller command is issued to make sure that the NVRAM card is still healthy. If the NVRAM card does not respond, the card might have failed and the system will force the node to reboot.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900100019

A message was sent from the NVRAM card to the host node.

Description

These events are often related to correctable errors that do not require further attention.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900100020

The NVRAM card is not responding and the node has been set to read-only.

Description

Connection between the NVRAM card and the node was lost. In some cases, restarting the node will reload the NVRAM driver and restore the connection.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900100021

The NVRAM card did not respond to an arm vault command and the node has been set to read-only.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900100022

The NVRAM card did not respond to an NVRAM command and the node has been set to read-only.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900100023

The NVRAM card firmware reported a correctable error.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900100024

The NVRAM card firmware reported an uncorrectable error.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900100025

The NVRAM flash vault in the indicated node failed to properly disarm on node startup.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900100026

The NVRAM card did not obtain necessary msi-x resources.

Description

The NVRAM card will continue to function in this state, but performance might be affected.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900100027

The PCI lane speed was not successfully negotiated.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900100028

NVDIMM has lost persistence in the chassis ({chassis}). To protect the journal, setting the node to read-only.

Description

This event happens when NVDIMM detects something in its subsystem that will not allow it to persist DRAM through a power-loss event. This could be related to the DRAM itself or the battery backup unit (BBU).

Administrator action

The node is placed into a read-only state to protect the journal. If the NVDIMM was bad during startup, the node will not be armed. Therefore, to re-gain persistence, the node automatically restarts (the cluster is notified, the cluster waits 60 seconds, and then restarts). If the issue does not clear itself, escalate, and determine if the issue is related to the NVDIMM or the battery.

A service request (SR) has been opened on your behalf. Technical Support will contact you.

900100029

NVDIMM has regained persistence in the chassis ({chassis}). The node will reboot to re-arm NVDIMM.

Description

The node lost persistence but has recovered. The node requires a reboot to re-arm the NVDIMM.

Administrator action

Wait for the node restart (this should happen 60 seconds from when the event occurs), or manually restart the node if the message continues.

900100030

NVDIMM in the DIMM slot has failed in the chassis ({chassis}). Until the NVDIMM is replaced, setting the node to read-only

Description

NVDIMM has failed. Node will transition to read-only mode until the NVDIMM has been replaced.

Potential causes

- NVDIMM experiences similar wear to regular DIMMs.
- The NVDIMM is seated incorrectly, damaged, or not responding.
- End of supported life duration.
- Unable to save the data during the previous system shutdown operation, or power loss. The NVDIMM is placed in write-protect mode.

Administrator action

- Refer to the error message and follow the OEM instructions.
- For an unreachable NVDIMM, reseat the DIMM. If the issue continues, replace the NVDIMM.

- Ensure that the NVDIMM is replaced in the correct slot.
- For an end of life, bad, or degraded NVDIMM replace the NVDIMM.

900100031

NVDIMM in the chassis ({chassis}) is in the wrong DIMM slot.

Description

NVDIMM was replaced in the wrong DIMM slot. Refer to the OEM manual for proper DIMM replacement procedure.

Administrator action

Replace the NVDIMM in the correct slot.

900100032

NVDIMM subsystem health is not being monitored in the chassis ({chassis}). Until the issue is resolved, setting the node to read-only

Description

NVDIMM subsystem health is not being monitored. Node will transition to read-only mode until the issue has been resolved.

Causes: Dell PowerTools is not responding to service requests.

Note: The monitoring subsystem will attempt several times before going read-only and triggering this event. Currently it is set to 100 seconds. This is subject to change as per the monitoring subsystem configuration of this node.

Administrator action

Check services, if issues persist, contact Technical Support.

900110001

The internal or ambient temperature around a node has exceeded the allowable threshold for the CPU.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900110002

A sensor in the front panel of a node has exceeded the specified threshold.

Description

This event can occur intermittently without harm to the system.

Administrator action

1. Cancel or quiet the event.
2. If the event recurs, shutdown and restart the node by completing the following steps:
   - Connect to the affected node through SSH or serial cable.
   - Shut down the node by running the following command:

     ```
     shutdown -p now
     ```

   - Wait for the node to shut down, and then disconnect both power supply cables.
   - Press the power button on the node to discharge any remaining stored power.
   - Reconnect the power cables and then start the node.
3. (HD400 only.) Re-seat the front panel connector by checking that the ribbon cable is properly attached and properly seated.
4. (All other nodes.) Re-seat the front panel.
5. Move the front panel from a functioning node to the affected node and see if the event clears.
6. Install the front panel from the affected node on another node to determine if the problem is with the front panel or with the node.

   If the problem follows the front panel, contact Technical Support to request a new front panel.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900110003

The node is reporting less than the expected amount of physical memory.

Description

This event typically appears because a DIMM has failed, is poorly seated, or an incorrect type of DIMM is installed.

Administrator action

Contact Technical Support to determine if a DIMM replacement is required.

900110004

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900110005

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900120001

The internal or ambient temperature around a node has exceeded the allowable threshold for the CPU.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900120002

A sensor in the front panel of a node has exceeded the specified threshold.

Description

This event can occur intermittently without harm to the system.

Administrator action

1. Cancel or quiet the event.
2. If the event recurs, shutdown and restart the node by completing the following steps:
   - Connect to the affected node through SSH or serial cable.
   - Shut down the node by running the following command:

     ```
     shutdown -p now
     ```

   - Wait for the node to shut down, and then disconnect both power supply cables.
   - Press the power button on the node to discharge any remaining stored power.
   - Reconnect the power cables and then start the node.
3. (HD400 only.) Re-seat the front panel connector by checking that the ribbon cable is properly attached and properly seated.

4. (All other nodes.) Re-seat the front panel.
5. Move the front panel from a functioning node to the affected node and see if the event clears.
6. Install the front panel from the affected node on another node to determine if the problem is with the front panel or with the node.

   If the problem follows the front panel, contact Technical Support to request a new front panel.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900120003

The node is reporting less than the expected amount of physical memory.

Description

This event typically appears because a DIMM has failed, is poorly seated, or an incorrect type of DIMM is installed.

Administrator action

Contact Technical Support to determine if a DIMM replacement is required.

900120004

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

- Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
- If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.

4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900120005

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

- If the issue follows the power supply, the power supply must be replaced.

5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900130001

The internal or ambient temperature around a node has exceeded the allowable threshold for the CPU.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900130002

A sensor in the front panel of a node has exceeded the specified threshold.

Description

This event can occur intermittently without harm to the system.

Administrator action

1. Cancel or quiet the event.
2. If the event recurs, shutdown and restart the node by completing the following steps:
   - Connect to the affected node through SSH or serial cable.
   - Shut down the node by running the following command:

     ```
     shutdown -p now
     ```

   - Wait for the node to shut down, and then disconnect both power supply cables.
   - Press the power button on the node to discharge any remaining stored power.
   - Reconnect the power cables and then start the node.
3. (HD400 only.) Re-seat the front panel connector by checking that the ribbon cable is properly attached and properly seated.
4. (All other nodes.) Re-seat the front panel.
5. Move the front panel from a functioning node to the affected node and see if the event clears.
6. Install the front panel from the affected node on another node to determine if the problem is with the front panel or with the node.

   If the problem follows the front panel, contact Technical Support to request a new front panel.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900130003

The node is reporting less than the expected amount of physical memory.

Description

This event typically appears because a DIMM has failed, is poorly seated, or an incorrect type of DIMM is installed.

Administrator action

Contact Technical Support to determine if a DIMM replacement is required.

900130004

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900130005

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900130006

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900130007

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900130008

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900130009

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900130010

A power supply fan in the might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900130011

A power supply fan in the might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

900130013

The internal or ambient temperature around a node has exceeded the allowable thresholds for a power supply.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900130014

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900130015

One of the power supplies in a node has failed or lost power.

Description

It is possible that a power cable was unplugged during recent maintenance or the circuit supplying power to the affected power supply has failed.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

1. Confirm that both power cables are properly connected to the node.
2. View the LED lights on the power supplies and confirm the status of the power supply:

LED	Power status	Node type
Steady green	Good	All nodes
Blinking green	Good, but the node is currently powered down	36000X, 3600NL, 72000X, 72000N
Steady amber	Good, but the node is currently powered down	X-Series, S-Series
Blinking amber	A power supply failure has occurred	X-Series, S-Series
No light	Insufficient or no A/C power	All nodes

3. If only one node reports the issue, determine the cause of the problem by performing the following steps.

   CAUTION: Do not move the power cable to another power supply in the same node as this will cause the node to lose power.

   - Locate the electrical outlet to which the problematic power supply is connected, and then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
   - If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
4. If the issue persists, take one power supply out of a different working node and attach the power supply to the affected node.

   CAUTION: Do not switch power supplies in the same node as this will cause the node to lose power.

   - If the issue follows the power supply, the power supply must be replaced.
5. If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
   - Power Distribution Unit (PDU) functionality and status of any circuit breakers in the power path
   - Power quality such as voltage, frequency values, and stability
   - Uninterruptible Power Supply (UPS) health
6. If the issue is not constant and is limited to one node, move the power to another circuit. Next, one at a time, move both power supplies.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900140001

A node reported a voltage measurement event group.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900140002

A fan sensor value on the indicated node is outside of the normal range.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900140003

A voltage sensor value on the indicated node is outside of the normal range.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900140004

The internal or ambient temperature around a node has exceeded the allowable threshold.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.

- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900140005

There are multiple power supply issues that might cause this event to occur.

Administrator action

1. If the event message specifies an issue with the power supply temperature, verify that the ambient temperature where the cluster is located has not exceeded the specifications.
2. If the event message specifies any other power supply input issue, verify that your power source is functional.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900150001

The node is reporting less than the expected amount of physical memory.

Description

This event typically appears because a DIMM has failed, is poorly seated, or an incorrect type of DIMM is installed.

Administrator action

Contact Technical Support to determine if a DIMM replacement is required.

900160001

A node is not connected to its peer node.

Description

When a node is disconnected from its peer node, the node journals are not mirrored and data is at risk. The event message provides you with the chassis and node slot where the disconnected node is located.

Administrator action

Contact Dell EMC PowerScale Technical Support to determine if a replacement node is required.

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160002

A node cannot connect to its peer node.

Description

When a node is disconnected from its peer node, the node journals are not mirrored and data is at risk. The event message provides you with the chassis and node slot where the disconnected node is located.

Administrator action

Contact Dell EMC PowerScale Technical Support to determine if a replacement node is required.

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160003

A compute node has failed and might need to be replaced.

Description

The event message provides you with the chassis and node slot where the failed node is located.

The compute node is not a customer-replaceable part.

Administrator action

Contact Dell EMC PowerScale Technical Support to determine if a replacement node is required.

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160004

A DIMM fault was detected.

Description

This event appears because a DIMM has failed, is poorly seated, or an incorrect type of DIMM is installed.

The event message provides you with the chassis and node slot where the DIMM is located.

A DIMM is not a customer-replaceable part.

Administrator action

Contact Dell EMC PowerScale Technical Support to determine if a replacement DIMM is required.

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160005

A node has powered down in response to a thermal issue.

Description

A node powered down to protect it from a thermal event. The elevated temperature might be the result of:

- a hardware fault that has increased the temperature inside the node.
- a high ambient temperature in the data center.

The event message provides you with the chassis and node slot of the node that powered down.

Administrator action

If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages. Verify that air flow within the rack, and through the front and rear panel vents of the chassis, is not obstructed in any way.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160006

Fan fault detected.

Description

A node fan has failed. The fan module must be replaced. The fan module is not a customer-replaceable part.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160007

Power supply fault detected.

Description

A power supply has failed. The power supply must be replaced. A power supply is a customer-replaceable part.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160008

Battery Backup Unit fault detected.

Description

A battery module has failed. The battery module must be replaced. The battery module is not a customer-replaceable part.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160009

Internal M.2 drive fault detected.

Description

An internal M.2 drive has failed. The M.2 drive must be replaced. The M.2 drive is not a customer-replaceable part.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160010

IO module fault detected.

Description

A fault was detected in a network card. Troubleshooting is required to determine if a hardware component must be replaced.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160011

Internal fault detected.

Description

An internal fault was detected. Troubleshooting is required to determine if a hardware component must be replaced.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160012

External fault detected.

Description

An external fault was detected. Troubleshooting is required to determine if a hardware component must be replaced.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160013

Non-transparent bridge fault detected.

Description

A non-transparent bridge fault is an indicator that a compute node has failed. The compute node must be replaced. The compute node is not a customer-replaceable part.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160014

A hardware issue was detected with the I2C bus.

Description

The I2C bus is a controller or bus that carries information from various sensors (for fan speed, power supply voltage, and temperature) in a node chassis.

It is not critical to address this issue immediately, but this event will continue to appear until the issue has been addressed. While this event is active, the node will not report correct values for the temperature, fans, or power supply health on the node.

Troubleshooting is required to determine if a hardware component must be replaced.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160015

Drive interface board fault detected.

Description

A drive interface board fault is an indicator that a compute node has failed. The compute node must be replaced. The compute node is not a customer-replaceable part.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160016

Midplane fault detected.

Description

A midplane fault is an indicator that a chassis has failed. The chassis must be replaced. The chassis is not a customer-replaceable part.

The event message provides you with the affected chassis.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160017

Unable to restore journal after powering up the node.

Description

A node restarted after losing power, but was unable to restore the node journal.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

This message is informational. No action is required at this time.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160018

A power supply is no longer providing power to the system.

Description

One of the power supplies in a node pair is no longer providing power. As a result, there is no longer a redundant power supply available to the node pair.

It is possible that a power cable came loose during recent maintenance. If there is no recent history of maintenance, it is likely that one of the power supplies in the node is not receiving power or has malfunctioned.

Administrator action

Perform the following suggestions in the order presented. If a suggestion solves the issue, there is no need to perform the other steps.

- Confirm that both power cables are properly connected to the node.
- Locate the electrical outlet to which the problematic power supply is connected, then determine if the outlet is functioning properly by plugging the power cable into a different electrical outlet.
- If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.

  CAUTION: Do not move the power cable to the power supply in the same node pair.
- If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.

  CAUTION: Do not switch power supplies in the same node pair as this will cause the node to lose power.

  - If the issue follows the power supply, the power supply needs to be replaced.
  - If the issue stays with the node, please contact Dell EMC PowerScale Technical Support for help with troubleshooting the node.
- If multiple nodes report power supply issues, it is likely that the issue is environmental. Check each of the following items to confirm the health of the power subsystem:
  - PDU functionality and status of any circuit breakers in the power path
  - Power quality such as voltage, frequency values, and stability
  - Uninterruptible Power Supply (UPS) health

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160019

Journal protection in the event of a power failure not enabled.

Description

OneFS failed to enable the subsystem that copies the contents of the local and peer node journals to the M.2 vault card in the event of power loss.

If the node loses power unexpectedly, OneFS will not be able to copy the journals.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

You can attempt to enable the subsystem by running the following command:

```
isi_pmp -a
```

If the subsystem is not enabled, there may be an issue with the M.2 vault card or another component in the node.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160020

A hardware error was corrected.

Description

This event can be an indicator of issues with a compute node. The compute node might need to be replaced. The compute node is not a customer-replaceable part.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160021

A node's Baseboard Management Controller (BMC) is not responding.

Description

The BMC monitors hardware components such as batteries and power supplies. To make sure that these hardware components continue to be monitored, the node must be serviced immediately. Troubleshooting is required to determine if a hardware component must be replaced.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160022

The Baseboard Management Controller (BMC) in a peer node is not responding.

Description

The BMC monitors hardware components such as batteries and power supplies. To make sure that these hardware components continue to be monitored, the peer node must be serviced immediately. Troubleshooting is required to determine if a hardware component must be replaced.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160023

A node journal is in an unprotected state.

Description

When a node is disconnected from its peer node, the node journals are not mirrored and data is at risk. The event message provides you with the chassis and node slot where the disconnected node is located.

Administrator action

1. Confirm that both nodes are cabled correctly and powered up.
2. If the issue persists, contact Dell EMC PowerScale Technical Support to determine if a replacement node is required.

   Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160024

A delayed reboot event occurred and the identified node was set to read-only to protect the journal.

Description

When node firmware identifies certain risk conditions, the firmware places the node in a delayed reboot state.

If the condition is cleared quickly, the delayed reboot state will be cleared and the node will return to read-write mode. If the condition persists, the node will reset.

Administrator action

1. Wait to see if the delayed reboot state clears and the node is restored.
2. If the node resets, monitor the node to make sure it successfully rejoins the cluster.
3. If issues with the node persist, contact Dell EMC PowerScale Technical Support to determine if maintenance is required.

   Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

900160100

A network interface card (NIC) is not available in the specified node.

Administrator action

Troubleshooting is required to determine if the data compression NIC must be replaced.

900160102

A network interface card (NIC) reset occurred in the specified node.

Description

The NIC reset is an information-only event. In the rare instance that a compression NIC resets, this event provides a tracking marker for Dell EMC PowerScale Technical Support.

Administrator action

No action is required.

900160101

A network interface card (NIC) is not operating correctly in the specified node.

Administrator action

Troubleshooting is required to determine if the data compression NIC must be replaced.

900170001

The IP address assigned to the BMC LAN interface by the DHCP server overlaps with a subnet already in use by the external network. The BMC LAN IP address is not tracked and used in validation for the external interfaces. {*msg*}

Administrator action

If DHCP is used for the BMC LAN address range, ensure that the external interface IP allocation does not overlap.

900170002

The system requires the following minimum firmware levels to support remote IPMI management. SSP: Minimum version: {*version*}.

Administrator action

Update the firmware.

900180001

Failed to communicate with the iDRAC management service: {protocol}.

Description

The Integrated Dell Remote Access Controller (iDRAC) is experiencing connectivity problems. The controller monitors hardware components such as batteries, power supplies, PowerTools, the iDRAC HTTPS port, and the USB network. To ensure that the hardware components are monitored, contact Dell EMC support as soon as possible.

Administrator action

The event may clear up automatically. However, if the connectivity problem persists, check the services and contact Dell EMC technical support.

900180002

Failed to communicate with the Internal Dual SD Module (IDSDM).

Description

The Internal Dual SD Module (IDSDM) located in the chassis is not populated, has failed, or is experiencing connectivity failures. Contact Dell EMC Support to diagnose this failure and ensure that this hardware platform is running to full specification.

Administrator action

Replace the IDSDM.

900180003

Node {Inn} requires a reboot for BIOS changes to take effect.

Description

Node BIOS settings were changed programmatically. BIOS changes on PowerEdge servers require a reboot to take effect.

Administrator action

Reboot the node.

900180004

A fan has failed.

Description

There is a bad, degraded, or missing fan.

Administrator action

Replace the fan.

900180005

NVDIMM battery has failed.

Description

There is a bad, degraded, or missing NVDIMM battery.

Administrator action

Replace the NVDIMM battery.

900180006

NVDIMM battery charge is low and is below an acceptable threshold and that a vault might fail.

Description

There is a low charge in the NVDIMM battery. When the event is identified, data loss is prevented by placing the node into read-only mode.

Administrator action

If the battery does not service, or charge on its own, replace the battery.

900180007

NVDIMM battery charge has been low for the exceeded time threshold.

Description

The NVDIMM battery has had a low charge for 86,400 s (24 hours).

Administrator action

Replace the NVDIMM battery by opening the chassis and locating the NVDIMM change LEDs on the black box. The Dell PowerEdge R640 documentation provides complete details on changing the battery.

900180008

DIMM has failed on a node.

Description

There is a bad, damaged, degraded, or missing DIMM in a node, or the correctable memory error rate is exceeded.

Administrator action

Ensure that the DIMM is seated correctly, or replace the DIMM when there is failure or degradation.

Discover the failed DIMM number and location by running using any of the following methods:

- Run the `isi_hwmon -b DIMMHealthMonitoring` command.
- Run `si_hwmon.log`.
- In the iDRAC UI, view the unhealthy DIMM slots.

900180009

A physical security sensor has detected that an intrusion error has occurred.

Description

PowerEdge servers contain physical security sensors to detect a chassis that has been left open or tampered with.

Administrator action

Inspect the physical chassis for intrusion, and close the chassis.

900180010

A physical security sensor is unhealthy and requires maintenance.

Description

There is a damaged, failing to respond, or degraded physical security sensor.

Administrator action

Replace the physical security sensor.

900180011

The system board sensor {sensor_name} has detected that a component is operating {adj} the recommended temperature range.

Description

The data center is too cold or too hot.

Administrator action

- Check the data center thermostat.
- Ensure that enclosed fans are working correctly.
- Ensure that objects are not blocking ventilation.
- See the enclosure documentation.

900180012

The chassis temperature sensor '{sensor_name}' is unhealthy and requires maintenance.

Description

There is a bad, damaged, or degraded temperature sensor.

Administrator action

Replace the temperature sensor.

900180013

A power supply is unhealthy and may require maintenance.

Description

- The power supply unit is not correctly seated.
- The power cable may be disconnected or improperly connected.
- The power supply unit has failed.
- There is an input power failure.

Administrator action

- Reseat the power supply unit.
- Reconnect the power cable to the power supply unit.
- Install a new power supply unit.
- Restore the input power supply.

900180014

Power supply has lost redundancy.

Description

One or more power supplies in a redundancy set are experiencing the following issues:

- The power supply unit is incorrectly seated.
- The power cable may be improperly connected, or disconnected.
- The power supply unit has failed.
- There is an input power failure.

Administrator action

Inspect the specified redundant power supply set and restore redundancy.

900180015

System {desc} sensors have detected degraded or unhealthy components that may require maintenance. Sensors: {sensor_list}.

Description

The event is generic for groups of sensor types. It is useful for voltage and amperage sensors. Only one event of this type is created per node. Only the most current list of specific sensors that are affected is listed to reduce the number of events per cluster.

Administrator action

For the affected sensors, see the Event and Error Message Reference Guide.

900180016

System {desc} sensors are unhealthy and may require maintenance. Sensors: {sensor_list}.

Description

A sensor is damaged, or is unhealthy.

Administrator action

Replace the sensors provided in the list of damaged sensors.

900180028

NVDIMM has lost persistence. Setting the node to read-only to protect the journal.

Description

The NVDIMM detects something in its subsystem that does not allow it to persist DRAM through a power-loss event. The event might be related to the DRAM itself or the BBU (battery backup unit).

Administrator action

The node is placed into a read-only state to protect the journal. If the NVDIMM was bad during startup, the node is not armed. The node automatically restarts (the cluster is notified, the cluster waits 60 s, and then restarts), and regains persistence. If the issue does not clear itself, escalate, and determine if the issue is related to the NVDIMM or the battery.

900180029

NVDIMM has regained persistence. Node reboots itself to rearm NVDIMM.

Description

The node lost persistence but has recovered. The node requires a reboot to rearm the NVDIMM.

Administrator action

Wait for the node restart (60 s from when the event occurs), or manually restart the OneFS node if the message continues.

900180030

NVDIMM has failed. Node transitions to read-only mode until the NVDIMM has been replaced.

Description

- NVDIMM experiences similar wear to regular DIMMs.
- The NVDIMM is seated incorrectly, damaged, or not responding.
- End of supported life duration.
- Unable to save the data during the previous system shutdown operation, or power loss. The NVDIMM is placed in write-protect mode.

Administrator action

- See the error message and follow the OEM instructions.
- For an unreachable NVDIMM, reseat the DIMM. If the issue continues, replace the NVDIMM.
- Ensure that the NVDIMM is replaced in the correct slot.
- For an end of life, bad, or degraded NVDIMM, replace the NVDIMM.

900180031

NVDIMM is in the wrong DIMM slot.

Description

NVDIMM was replaced in the wrong DIMM slot.

Administrator action

See the OEM manual for proper DIMM replacement procedure, and replace the NVDIMM int the correct slot.

900180032

NVDIMM subsystem health is not being monitored. The node transitions to read-only mode until the issue has been resolved.

Description

Dell PowerTools are not responding to service requests.

Administrator action

To determine the problem, run the `isi_hwmon -b IDRACServices` command.

910100001

A fan in the node might have failed.

Description

If the fan speed temporarily falls out of optimal range for less than a minute or so and the event does not repeat, this event might be a false alarm.

Administrator action

Follow the instructions in *Event notification: Fan speed out of spec*, article 000083406 to determine if this event is a false alarm.

If this event is not a false alarm, contact Technical Support.

910100002

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

910100003

The internal or ambient temperature around a node has exceeded the allowable threshold.

Description

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

910100004

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

910100005

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

910100006

A voltage component is out of specification.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

910100007

A sensor in the front panel of a node has exceeded the specified threshold

Description

This event can occur intermittently without harm to the system.

Administrator action

1. Cancel or quiet the event.
2. If the event recurs, shutdown and restart the node by completing the following steps:
   - Connect to the affected node through SSH or serial cable.
   - Shut down the node by running the following command:

     ```
     shutdown -p now
     ```

   - Wait for the node to shut down, and then disconnect both power supply cables.
   - Press the power button on the node to discharge any remaining stored power.
   - Reconnect the power cables and then start the node.
3. (HD400 only.) Re-seat the front panel connector by checking that the ribbon cable is properly attached and properly seated.
4. (All other nodes.) Re-seat the front panel.
5. Move the front panel from a functioning node to the affected node and see if the event clears.
6. Install the front panel from the affected node on another node to determine if the problem is with the front panel or with the node.

   If the problem follows the front panel, contact Technical Support to request a new front panel.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

920100000

There are multiple temperature issues that might cause this event to occur.

Description

The internal or ambient temperature around one or more nodes has exceeded the allowable thresholds.

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the above steps do not resolve the issue, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

920100001

There are multiple battery issues that might cause this event to occur.

Description

A connection with a single battery is lost. The HD400, S210, and X410 nodes will continue to operate when connection with a single battery is lost; other node models will be placed in a read-only state until connection with both batteries is restored.

Administrator action

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

920100002

The Chassis Management Controller (CMC) is not monitoring the specified sensor.

Description

The subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have failed. The event will clear when the CMC starts monitoring the sensor again.

Administrator action

If the event occurs once, you can safely ignore it.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

920100003

An HD400 drive drawer is open and the 5 minute service window timer has started.

Description

Drives will be powered off in five minutes to protect them from overheating.

Administrator action

Close the drive drawer before the service window timer expires. If you are not performing service on the node, make sure that the drive drawer is properly closed by sliding the drawer out and re-closing the drawer firmly and carefully. If the drive drawer is not closed before the service window timer expires, the node reboots, but the node will not rejoin the cluster until temperatures are within acceptable thresholds.

If the event does not clear itself when maintenance is complete, or if maintenance is not being performed on the node and the above steps do not resolve the issue, follow the instructions to gather logs, and contact Dell EMC PowerScale Technical Support.

920100004

There are multiple fan failures. 5 minute drive power down warning.

Administrator action

When multiple fans fail or a fan module is removed for more than two minutes, the node will reboot, and the drives will power down within five minutes to prevent the drives from overheating. The drives will remain powered down until the failed fan modules are replaced. Replace a fan module if the fan has failed or re-insert a fan module if it has been pulled for maintenance.

If the event does not clear itself when maintenance is complete, or if maintenance is not being performed on the node, follow the instructions to gather logs, and contact Dell EMC PowerScale Technical Support.

920100005

A single fan has failed in one of the suitcase fan trays.

Administrator action

- If a fan tray is not fully seated, re-seat the tray and see if the fan resumes operation.
- If the fan does not resume operation, the fan tray might need to be replaced.

Troubleshooting is required to determine if a hardware component must be replaced.

Gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

920100006

A sensor on a node indicates an elevated temperature. Drives are overheating. The node will reboot immediately. Drive power will discontinue in five minutes.

Description

The node will reboot but will not rejoin the cluster until temperatures are within acceptable thresholds. If not cooled within five minutes, the drives will stay powered down until the inlet temperature is at an acceptable level to restart.

Ambient temperature is only measured by front panel sensors. If you receive an event that indicates that the front panel is out of specification, the temperature in your data center might need to be adjusted.

If a node is subjected to high temperatures for an extended period of time, the CPU is throttled and the node goes into read only-mode to help prevent potential data loss due to component failure. If the node temperature reaches critical levels, it is possible that the node will shut down entirely.

Administrator action

Perform the following steps in the order listed. If the issue resolves after a step, there is no need to complete the subsequent steps.

- (HD400 only) Make sure that the drive drawer is properly shut by sliding it out and re-closing it firmly but carefully.
- Review the temperature statistics for the affected sensor, which are included in the event. If the temperature is consistently elevated, the problem is likely a high ambient temperature in the data center. Address any changes in the cluster environment such as air conditioning outages.
- Verify that air flow within the rack, and through the front and rear panel vents of the node, is not obstructed in any way.
- Make sure that the faceplate on the affected node is installed, properly seated, and undamaged. In some cases, removing and re-seating the faceplate will resolve this issue.
- Run the `isi_hw_status` command. Review the output to determine whether there is a slow or failed fan that was not otherwise reported.
- Check for high CPU and disk usage in the node. High usage can contribute to high temperatures within the node.

If the steps above were unsuccessful in clearing this event, the subsystem that monitors the health of the hardware (such as the temperature and fan speeds) might have encountered a problem. This event can occur intermittently without harm to the system and you can safely quiet the event unless the issue persists.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

920100007

All drives in the node are powering down.

Description

This event occurs if the five minute warning has not been cleared from events 920100003, 920100004, or 920100006.

Administrator action

Address the events that occurred before the drives were powered down.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

920100008

One of the drives is overheating.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

920100009

One of the drives is overheating.

Administrator action

Reboot the node. If the event clears and does not recur, no other action is required.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

930100000

A sensor is reporting fan values that are outside expected specifications.

Description

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Monitor your cluster for other events that might be related to this event.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

930100001

A sensor is reporting electrical values that are outside expected specifications.

Description

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Monitor your cluster for other events that might be related to this event.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

930100002

A sensor is reporting temperature values that are outside expected specifications.

Description

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Monitor your cluster for other events that might be related to this event.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

930100003

A sensor is reporting electrical values that are outside expected specifications.

Description

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Monitor your cluster for other events that might be related to this event.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

930100004

A sensor is reporting electrical values that are outside expected specifications.

Description

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Monitor your cluster for other events that might be related to this event.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

930100005

A sensor is reporting values that are outside expected specifications.

Description

This event will tell you which sensor is reporting the unexpected values.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Monitor your cluster for other events that might be related to this event.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

930100006

A sensor is reporting values that are outside expected specifications.

Description

This event will tell you which sensor is reporting the unexpected values.

The event message provides you with the chassis and node slot of the affected node.

Administrator action

Monitor your cluster for other events that might be related to this event.

If the event persists, gather logs, and then contact Technical Support for additional troubleshooting. For instructions on how to gather cluster logs, see Gathering cluster logs.

940100001

OneFS {version} is currently running and is not supported on this hardware. Unsupported OneFS Version.

Description

The OneFS version that is currently running is not supported on this hardware.

Administrator action

Contact Technical Support to obtain the supported software version for this hardware.

940100002

OneFS {version} is currently running on unsupported nodes (devid(s) {devids}). {msg}.

Description

OneFS {version} is currently running on the specified nodes in this cluster.

Administrator action

Contact Technical Support to obtain the supported software version for this hardware.