



Replace drives for HPE DL360

Element Software

amitha

March 09, 2021

Table of Contents

Replace drives for HPE DL360 1
Find more information 8

Replace drives for HPE DL360

Choose from the procedures listed here to replace a drive proactively, replace a drive after it has failed, and replace a cache drive. replace a metadata drive or a block drive in your SolidFire eSDS cluster. The Element UI **Cluster > Drives** page shows the drive wear information.

- [Replace a drive proactively](#)
- [Replace a faulty drive](#)
- [Replace a cache drive](#)

Replace a drive proactively

Perform this procedure if you want to proactively replace a metadata drive or a block drive in your SolidFire eSDS cluster. The Element UI **Cluster > Drives** page shows the drive wear information.

What you'll need

- From the NetApp Element software UI, ensure that your cluster is in good health and there are no warnings or cluster faults. You can access the Element UI by using the management virtual IP (MVIP) address of the primary cluster node.
- Ensure that there are no active jobs running on the cluster.
- Ensure that you have familiarized yourself with all the steps.
- Ensure that you take necessary precautions to prevent electrostatic discharge (ESD) while handling drives.

Steps

1. Perform the following steps in the Element UI:
 - a. In the Element UI, select **Cluster > Drives > Active**.
 - b. Select the drive that you want to replace.
 - c. Make a note of the serial number of the drive. This will help you locate the corresponding BayID in the IPMI interface of the node (HPE Integrated Lights-Out or iLO, in this case).
 - d. Select **Bulk Actions > Remove**. After you remove the drive, the drive goes into the **Removing** state. It stays in the **Removing** state for a while, waiting for the data on the drive to be synced or redistributed to the remaining drives in the cluster. After the remove is complete, the drive moves to the **Available** state.
2. Perform the following steps to locate the drive slot of the drive that you are replacing:
 - a. Log in to the IPMI interface of the node (iLO in this case).
 - b. Select **System Information** from the left-hand navigation, and then select **Storage**.
 - c. Match the serial number you made a note of in the previous step with what you see on the screen.
 - d. Look for the slot number listed against the serial number. This is the physical slot from which you must remove the drive.
3. Now that you have identified the drive, physically remove it as follows:
 - a. Identify the drive bay.

The following image shows the front of the server with the drive bay numbering shown on the left side of the image:



- b. Press the power button on the drive that you want to replace. The LED blinks for 5-10 seconds and stops.
- c. After the LED stops blinking and the drive is powered off, remove it from the server by pressing the red button and pulling the latch.



Ensure that you handle drives very carefully.

After you physically remove the drive, the drive state changes to **Failed** in the Element UI.

4. In the Element UI, select **Cluster > Drives > Failed**.

5. Select the icon under **Actions**, and then select **Remove**.

Now you can go ahead and install the new drive in the node.

6. Make a note of the serial number of the new drive.
7. Insert the replacement drive by carefully pushing the drive into the bay using the latch and closing the latch. The drive powers on when inserted correctly.
8. Perform the following steps to verify the new drive details in iLO:
 - a. Log in to iLO.
 - b. Select **Information > Integrated Management Log**. You will see an event logged for the drive that you added.
 - c. Select **System Information** from the left-hand navigation, and then select **Storage**.
 - d. Scroll till you find information about the bay that you replaced the drive in.
 - e. Verify that the serial number on your screen matches the serial number of the new drive that you replaced.
9. Add the new drive information in the `sf_sds_config.yaml` file for the node in which you replaced the drive.

The `sf_sds_config.yaml` file is stored in `/opt/sf/`. This file includes all the information about the drives in the node. Every time you replace a drive, you must enter the replacement drive information in this file. For more information about this file, see [Contents of the sf_sds_config.yaml file](#).

- a. Establish an SSH connection to the node by using PuTTY.
- b. In the PuTTY configuration window, enter node MIP in the **Host Name (or IP address)** field.
- c. Select **Open**.
- d. In the terminal window that opens, log in with your username and password.
- e. Run the `# cat /opt/sf/sf_sds_config.yaml` command to list the contents of the file.
- f. Replace the entries in the `dataDevices` or `cacheDevices` lists for the drive you replaced with the new drive information.
- g. Run `# systemctl start solidfire-update-drives`.

You see the Bash prompt after this command runs. You should go to the Element UI after this to add

the drive to the cluster. The Element UI shows an alert for a new drive that is available.

10. Select **Cluster > Drives > Available**.

You see the serial number of the new drive that you installed.

11. Select the icon under **Actions**, and then select **Add**.

12. Refresh the Element UI after the block sync job completes. You see that the alert about the drive available has cleared if you access the **Running Tasks** page from the **Reporting** tab of the Element UI.

Replace a faulty drive

If your SolidFire eSDS cluster has a faulty drive, the Element UI displays an alert. Before you remove the drive from the cluster, verify the reason for failure by looking at the information in the IPMI interface for your node/server. These steps are applicable if you are replacing a block drive or a metadata drive.

What you'll need

- From the NetApp Element software UI, verify that the drive has failed. Element displays an alert when a drive fails. You can access the Element UI by using the management virtual IP (MVIP) address of the primary cluster node.
- Ensure that you have familiarized yourself with all the steps.
- Ensure that you take necessary precautions to prevent electrostatic discharge (ESD) while handling drives.

Steps

1. Remove the failed drive from the cluster as follows using the Element UI:
 - a. Select **Cluster > Drives > Failed**.
 - b. Note the node name and serial number associated with the failed drive.
 - c. Select the icon under **Actions**, and then select **Remove**. If you see warnings of the service associated with the drive, wait until bin sync completes, and then remove the drive.
2. Perform the following steps to verify the drive failure and view the events logged that are associated with the drive failure:
 - a. Log in to the IPMI interface of the node (iLO in this case).
 - b. Select **Information > Integrated Management Log**. The reason for the drive failure (for example, SSDWearOut) and the location is listed here. You can also see an event stating that the status of the drive is degraded.
 - c. Select **System Information** from the left-hand navigation, and then select **Storage**.
 - d. Verify the information available about the failed drive. The status of the failed drive will say **Degraded**.
3. Remove the drive physically as follows:
 - a. Identify the drive slot number in the chassis.

The following image shows the front of the server with the drive bay numbering shown on the left side of the image:



- b. Press the power button on the drive that you want to replace. The LED blinks for 5-10 seconds and

stops.

- c. After the LED stops blinking and the drive is powered off, remove it from the server by pressing the red button and pulling the latch.



Ensure that you handle drives very carefully.

4. Insert the replacement drive by carefully pushing the drive into the bay using the latch and closing the latch. The drive powers on when inserted correctly.
5. Verify the new drive details in iLO:
 - a. Select **Information > Integrated Management Log**. You see an event logged for the drive that you added.
 - b. Refresh the page to see the events logged for the new drive that you added.
6. Verify the health of your storage system in iLO:
 - a. Select **System Information** from the left-hand navigation, and then select **Storage**.
 - b. Scroll till you find information about the bay in which you installed the new drive.
 - c. Make a note of the serial number.
7. Add the new drive information in the `sf_sds_config.yaml` file for the node in which you replaced the drive.

The `sf_sds_config.yaml` file is stored in `/opt/sf/`. This file includes all the information about the drives in the node. Every time you replace a drive, you must enter the replacement drive information in this file. For more information about this file, see [Contents of the sf_sds_config.yaml file](#).

- a. Establish an SSH connection to the node by using PuTTY.
- b. In the PuTTY configuration window, enter node MIP in the **Host Name (or IP address)** field.
- c. Select **Open**.
- d. In the terminal window that opens, log in with your username and password.
- e. Run the `# cat /opt/sf/sf_sds_config.yaml` command to list the contents of the file.
- f. Replace the entries in the `dataDevices` or `cacheDevices` lists for the drive you replaced with the new drive information.
- g. Run `# systemctl start solidfire-update-drives`.

You see the Bash prompt after this command runs. You should go to the Element UI after this to add the drive to the cluster. The Element UI shows an alert for a new drive that is available.

8. Select **Cluster > Drives > Available**.

You see the serial number of the new drive that you installed.

9. Select the icon under **Actions**, and then select **Add**.
10. Refresh the Element UI after the block sync job completes. You see that the alert about the drive available has cleared if you access the **Running Tasks** page from the **Reporting** tab of the Element UI.

Replace a cache drive

Perform this procedure if you want to replace the cache drive in your SolidFire eSDS cluster. The cache drive

is associated with metadata services. The Element UI **Cluster > Drives** page shows the drive wear information.

What you'll need

- From the NetApp Element software UI, ensure that your cluster is in good health and there are no warnings or cluster faults. You can access the Element UI by using the management virtual IP (MVIP) address of the primary cluster node.
- Ensure that there are no active jobs running on the cluster.
- Ensure that you have familiarized yourself with all the steps.
- Ensure that you remove metadata services from the Element UI.
- Ensure that you take necessary precautions to prevent electrostatic discharge (ESD) while handling drives.

Steps

1. Perform the following steps in the Element UI:

- a. In the Element UI, select **Cluster > Nodes > Active**.
- b. Make a note of the node ID and management IP address of the node in which you are replacing the cache drive.
- c. If the cache drive is healthy and you are proactively replacing it, select **Active Drives**, locate the metadata drive, and remove it from the UI.

After you remove it, the metadata drive goes to **Removing** state first, and then to **Available**.

- d. If you are performing replacement after the cache drive failed, the metadata drive will be in **Available** state, and listed under **Cluster > Drives > Available**.
- e. In the Element UI, select **Cluster > Drives > Active**.
- f. Select the metadata drive associated with the NodeName, where you want to do the cache drive replacement.
- g. Select **Bulk Actions > Remove**. After you remove the drive, the drive goes into the **Removing** state. It stays in the **Removing** state for a while, waiting for the data on the drive to be synced or redistributed to the remaining drives in the cluster. After the remove is complete, the drive moves to the **Available** state.

2. Perform the following steps to locate the drive slot of the cache drive that you are replacing:

- a. Log in to the IPMI interface of the node (iLO in this case).
- b. Select **System Information** from the left-hand navigation, and then select **Storage**.
- c. Locate the cache drive.



Cache drives are of lesser capacity than storage drives.

- d. Look for the slot number listed for cache drive. This is the physical slot from which you must remove the drive.

3. Now that you have identified the drive, physically remove it as follows:

- a. Identify the drive bay.

The following image shows the front of the server with the drive bay numbering shown on the left side of the image:



- b. Press the power button on the drive that you want to replace. The LED blinks for 5-10 seconds and stops.
- c. After the LED stops blinking and the drive is powered off, remove it from the server by pressing the red button and pulling the latch.



Ensure that you handle drives very carefully.

After you physically remove the drive, the drive state changes to **Failed** in the Element UI.

4. Make a note of the HPE model number and the ISN (serial number) of the new cache drive.
5. Insert the replacement drive by carefully pushing the drive into the bay using the latch and closing the latch. The drive powers on when inserted correctly.
6. Perform the following steps to verify the new drive details in iLO:
 - a. Log in to iLO.
 - b. Select **Information > Integrated Management Log**. You see an event logged for the drive that you added.
 - c. Select **System Information** from the left-hand navigation, and then select **Storage**.
 - d. Scroll till you find information about the bay that you replaced the drive in.
 - e. Verify that the serial number on your screen matches the serial number of the new drive that you installed.
7. Add the new cache drive information in the `sf_sds_config.yaml` file for the node in which you replaced the drive.

The `sf_sds_config.yaml` file is stored in `/opt/sf/`. This file includes all the information about the drives in the node. Every time you replace a drive, you should enter the replacement drive information in this file. For more information about this file, see [Contents of the sf_sds_config.yaml file](#).

- a. Establish an SSH connection to the node by using PuTTY.
- b. In the PuTTY configuration window, enter node MIP address (that you made a note of from the Element UI earlier) in the **Host Name (or IP address)** field.
- c. Select **Open**.
- d. In the terminal window that opens, log in with your username and password.
- e. Run the `nvme list` command to list the NVMe devices.

You can see the model number and serial number of the new cache drive. See the following sample output:


```
[root@NLABICT062226 ~]# nvme list
Node          SN              Model          Namespace Usage
-----
/dev/nvme0n1  KI9AN0136I02AAU17  VK003840KWWFP  1          3.84 TB / 3.84 TB
/dev/nvme1n1  PHKE913200XM375AGN  E0000375KWJUC  1          375.08 GB / 375.08 GB
/dev/nvme1n1  KI9AN0136I02AAU00  VK003840KWWFP  1          3.84 TB / 3.84 TB
/dev/nvme2n1  KI05T0003I1205C14  VK003840KWWFP  1          3.84 TB / 3.84 TB
/dev/nvme3n1  KI05T0003I1205C0W  VK003840KWWFP  1          3.84 TB / 3.84 TB
/dev/nvme4n1  KI05T0003I1205C10  VK003840KWWFP  1          3.84 TB / 3.84 TB
/dev/nvme5n1  KI05T0003I1205C1P  VK003840KWWFP  1          3.84 TB / 3.84 TB
/dev/nvme7n1  KI05T0003I1205C1L  VK003840KWWFP  1          3.84 TB / 3.84 TB
/dev/nvme8n1  KI05T0003I1205C13  VK003840KWWFP  1          3.84 TB / 3.84 TB
/dev/nvme9n1  KI9AN0136I02AAU1Q  VK003840KWWFP  1          3.84 TB / 3.84 TB
[root@NLABICT062226 ~]#
```

f. Add the new cache drive information in `/opt/sf/sf_sds_config.yaml`.

You should replace the existing cache drive model number and serial number with the corresponding information for the new cache drive. See the following example:

```
schemaVersion: "2.0"

network:
  managementInterface: "team1G"
  storageInterface: "team10G"
dataDrives:
  - "/dev/nvme0n1"
  - "/dev/nvme1n1"
  - "/dev/nvme2n1"
  - "/dev/nvme3n1"
  - "/dev/nvme4n1"
  - "/dev/nvme5n1"
  - "/dev/nvme7n1"
  - "/dev/nvme8n1"
  - "/dev/nvme9n1"
cacheDevices:
  - "/dev/disk/by-id/nvme-E0000375KWJUC PHKE913200XM375AGN"
```

g. Save the `/opt/sf/sf_sds_config.yaml` file.

8. Perform the steps for the scenario that is applicable to you:

Scenario	Steps
The new inserted cache drive shows up after you run the <code>nvme list</code> command	<ol style="list-style-type: none"> 1. Run <code># systemctl restart solidfire</code>. This takes around three minutes. 2. Check the <code>solidfire</code> status by running <code>system status solidfire</code>. 3. Go to step 9.
The new inserted cache drive does not show up after you run the <code>nvme list</code> command	<ol style="list-style-type: none"> 1. Reboot the node. 2. After the node reboots, verify that the <code>solidfire</code> services are running by logging in to the node (using PuTTY), and running the <code>system status solidfire</code> command. 3. Go to step 9.



Restarting `solidfire` or rebooting the node causes some cluster faults, which eventually clear in about five minutes.

9. In the Element UI, add back the metadata drive that you removed:
 - a. Select **Cluster > Drives > Available**.
 - b. Select the icon under Actions, and select **Add**.
10. Refresh your Element UI after the block sync job completes.

You can see that the alert about the drive available has cleared along with other cluster faults.

Find more information

- [NetApp SolidFire Resources Page](#)
- [Documentation for earlier versions of NetApp SolidFire and Element products](#)

Copyright Information

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system- without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.