



# Replace drives for Dell R640

## Element Software

amitha

March 09, 2021

# Table of Contents

Replace drives for Dell R640 ..... 1  
Find more information ..... 8

# Replace drives for Dell R640

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 slot number of the drive in the Integrated Dell Remote Access Controller (iDRAC).
  - 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 (iDRAC in this case).
  - b. Select **Storage** from the menu, and then select **Physical Disks**.
  - c. To find the serial number of the drive, select the **+** icon next to each PCIe SSD.
  - d. Match the serial number you see on the screen with what you noted in the Element UI.
  - e. Look for the slot number listed against the serial number.

This is the physical slot from which you should remove the drive.

3. Now that you have identified the drive, physically remove it as follows:
  - a. Identify the drive slot number in the chassis.

The following image shows the front of the server with the slot numbering shown for each drive:



- b. Press the button on the drive.

The latch comes out.

- c. Physically pull the drive out from the slot.



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 iDRAC:
  - a. Log in to iDRAC.
  - b. Select **Maintenance > System Event Log**.

You see an event logged for the drive that you added.

- c. Select **Storage** from the menu, and then select **Physical Disks**.
  - d. Verify that the new drive you inserted shows up in the corresponding slot in the UI.
  - e. To find the serial number of the drive, select the **+** icon next to each PCIe SSD.
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 (iDRAC in this case).
  - b. Select **Maintenance > System Event Log** to see the reason for the drive failure (for example, SSDWearOut or drive not inserted properly).

You can also see an event showing the status of the drive.

- c. Select **Storage** from the menu, and then select **Physical Disks**.
  - d. Find the slot number of the failed drive using the serial number that you noted in the Element UI.
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 slot numbering shown for each drive:



b. Press the button on the drive.

The latch comes out.

c. Physically pull the drive out from the slot.



Ensure that you handle drives very carefully.

4. Insert the replacement drive by carefully pushing the drive into the slot using the latch and closing the latch.

The drive powers on when inserted correctly.

5. Verify the new drive details in iDRAC:

a. Select **Maintenance > System Event Log**. You see an event logged for the drive that you added.

b. Select **Storage** from the menu, and then select **Physical Disks**.

c. Verify that the new drive you inserted shows up in the corresponding slot in the UI.

d. To find the serial number of the drive, select the **+** icon next to each PCIe SSD.

6. 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.

7. Select **Cluster > Drives > Available**.

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

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

9. 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 (iDRAC in this case).
  - b. Select **Storage** from the menu, and then select **Physical Disks**.
  - c. Locate the cache drive.



Cache drives are of lesser capacity (375 GB) than storage drives, and are PCIe SSDs.

- d. Look for the slot number listed for cache drive.

This is the physical slot from which you should remove the drive.

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

- a. Identify the drive slot number in the chassis.

The following image shows the front of the server with the slot numbering shown for each drive:



- b. Press the button on the drive.

The latch comes out.

- c. Physically pull the drive out from the slot.



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 model number and the ISN (serial number) of the new cache drive.
5. Insert the replacement drive by carefully pushing the drive into the slot 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 iDRAC:
  - a. Select **Maintenance > System Event Log**. You see an event logged for the drive that you added.
  - b. Select **Storage** from the menu, and then select **Physical Disks**.
  - c. Verify that the new drive you inserted shows up in the corresponding slot in the UI.
  - d. To find the serial number of the drive, select the **+** icon next to each PCIe SSD.
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:



Node	SN	Model	Namespace	Usage	Format	FW Rev
/dev/nvme0n1	PHLJ029506A54P0DGN	INTEL SSDPE2KX040T8	1	4.00 TB / 4.00 TB	512 B + 0 B	V0V10131
/dev/nvme1n1	PHKE91400006375AGN	INTEL SSDPE21K375GA	1	375.08 GB / 375.08 GB	512 B + 0 B	E2010435
/dev/nvme2n1	PHLJ030004VJ4P0DGN	INTEL SSDPE2KX040T8	1	4.00 TB / 4.00 TB	512 B + 0 B	V0V10131
/dev/nvme3n1	PHLJ029507N84P0DGN	INTEL SSDPE2KX040T8	1	4.00 TB / 4.00 TB	512 B + 0 B	V0V10131
/dev/nvme4n1	PHLJ030004W04P0DGN	INTEL SSDPE2KX040T8	1	4.00 TB / 4.00 TB	512 B + 0 B	V0V10131
/dev/nvme5n1	PHLJ030101R54P0DGN	INTEL SSDPE2KX040T8	1	4.00 TB / 4.00 TB	512 B + 0 B	V0V10131
/dev/nvme6n1	PHLJ029509QX4P0DGN	INTEL SSDPE2KX040T8	1	4.00 TB / 4.00 TB	512 B + 0 B	V0V10131
/dev/nvme7n1	PHLJ030101S44P0DGN	INTEL SSDPE2KX040T8	1	4.00 TB / 4.00 TB	512 B + 0 B	V0V10131
/dev/nvme8n1	PHLJ029509QZ4P0DGN	INTEL SSDPE2KX040T8	1	4.00 TB / 4.00 TB	512 B + 0 B	V0V10131
/dev/nvme9n1	PHLJ030101R44P0DGN	INTEL SSDPE2KX040T8	1	4.00 TB / 4.00 TB	512 B + 0 B	V0V10131

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: "team0"
  storageInterface: "team1"
dataDrives:
  - "/dev/nvme0n1"
  - "/dev/nvme2n1"
  - "/dev/nvme3n1"
  - "/dev/nvme4n1"
  - "/dev/nvme5n1"
  - "/dev/nvme6n1"
  - "/dev/nvme7n1"
  - "/dev/nvme8n1"
  - "/dev/nvme9n1"
cacheDevices:
  - "/dev/disk/by-id/nvme-INTEL SSDPE21K375GA PHKE913200Z3375AGN"

```

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"> <li>1. Run <code># systemctl restart solidfire</code>. This takes around three minutes.</li> <li>2. Check the <code>solidfire</code> status by running <code>system status solidfire</code>.</li> <li>3. Go to step 9.</li> </ol>
The new inserted cache drive does not show up after you run the <code>nvme list</code> command	<ol style="list-style-type: none"> <li>1. Reboot the node.</li> <li>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.</li> <li>3. Go to step 9.</li> </ol>



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.