FRU replacement highlights

The information you should know before replacing FRU

Lenovo

Service considerations

Before taking action, consider the following service factors:

- All DM systems are clustered
  - This means that both controller A and controller B communicate over a dedicated connection using both the SAS back end and two dedicated UTA 10 Gb Ethernet ports (e0a and e0b)
- Any time you need to service the box, you need to first disable the auto-failback feature on the controller or node that is being serviced to prevent it from taking ownership of resources.
- You will also need to ensure that any node or controller that is being serviced is first halted to allow all resources to transfer to the good or non-degraded controller.
- Keep in mind that a DM cluster can be more than just the two nodes in the single system that is being serviced. In such cases, you will need to validate that each node can access the shared quorum resource for the cluster.
- You cannot place a node that owns the quorum resource into service. It will need to be transferred to one of the surviving nodes.
- There is no way to halt both node A and node B at the same time.
- All license keys will also have to be regenerated after the controller replacement.

Before and after FRU replacement

Depending on the type of FRU involved, some actions should be taken before the replacement.

- Identify the machine type and serial number
- Check System Service Labels to identify the FRU location or attention LEDs

For FRU parts that are not hot-swappable, work through the following procedure:

- Shut down the impaired node
- Erase the data on the caching module (if you are replacing a caching module)

After replacing a component that is not hot-swappable, work through the following procedure:

- Run diagnostic tests on that component
- Transfer the boot image to the boot media (if you are replacing a boot media)
- Restore the controller module to operation after running diagnostics

DM 7100 Series and DM240N FRU parts

Scroll down for more information.

The following FRUs are serviced by Lenovo service personnel. When replacing FRU, servicers should work through the procedures introduced on YouTube or Youku; otherwise, unexpected system damage may occur and impact the customer's data. Always capture the node's console output to a text file for escalation and review, even if you're using the end-user's computer.

DM7100 Series specific: YouTube, Youku

- Power supply
- Fan module
- Controller module
  - NVDIMM
  - NVDIMM battery
  - System DIMM
  - Boot media
  - RTC battery
  - PCIe card
  - Mezzanine card
  - Caching module (DM7100H only)

DM 7100 Series and DM240N FRU parts

Scroll down for more information.

- Fan module
- Controller module
  - NVDIMM
  - NVDIMM battery
  - System DIMM
  - Boot media
  - RTC battery
  - PCIe card
  - Mezzanine card
  - Caching module (DM7100H only)

DM240N specific: YouTube, Youku

- Disk
- Fan module
- Power supply
- NSM
- RDIMM

Identifying the machine type and serial number

Before taking service action, it is suggested that you identify the machine type and serial number of the DM Series storage system that will be serviced. Issue the `run local sysconfig -a` command to identify the Model Name, Machine Type, and Lenovo System SN.

```
DM5000H::> run local sysconfig -a
Lenovo Release 9.4P1: Fri Jul 20 23:30:57 EDT 2018
System ID: 0537143011 (DM5000H-1); partner ID: 0537142993 (DM5000H-2)
System Serial Number: 721822000139 (DM5000H-1)
System Rev: 20
System Storage Configuration: Multi-Path HA
System ACP Connectivity: Inband Active
All-Flash Optimized: true
Backplane Part Number: 111-02495
Backplane Rev: A4
Backplane Serial Number: 021806001720
slot 0: System Board 1.7 GHz (System Board XXII B2)
Model Name:          DM5000H
Machine Type:        7Y57-CTOWW1
Lenovo System SN:    LSSN1234
Part Number:         111-02493
Revision:            20
Serial Number:       021815016159
BIOS version:        11.2.1
```

Note: You can also get this information through the ONTAP 9.7 GUI. Select **Cluster** → **Overview**, and then click the **System ID** to get the machine type and serial number.

System Service Labels

System service labels are located on the PCIe risers and air duct. The labels show diagrams and the locations of the FRU parts, attention LEDs, and the basic steps used to replace FRU parts.
Click the items to see larger images.

System Service Labels

System service labels are located on the PCIe risers and air duct. The label
locations of the FRU parts, attention LEDs, and the basic steps used to repl
Click the items to see larger images.

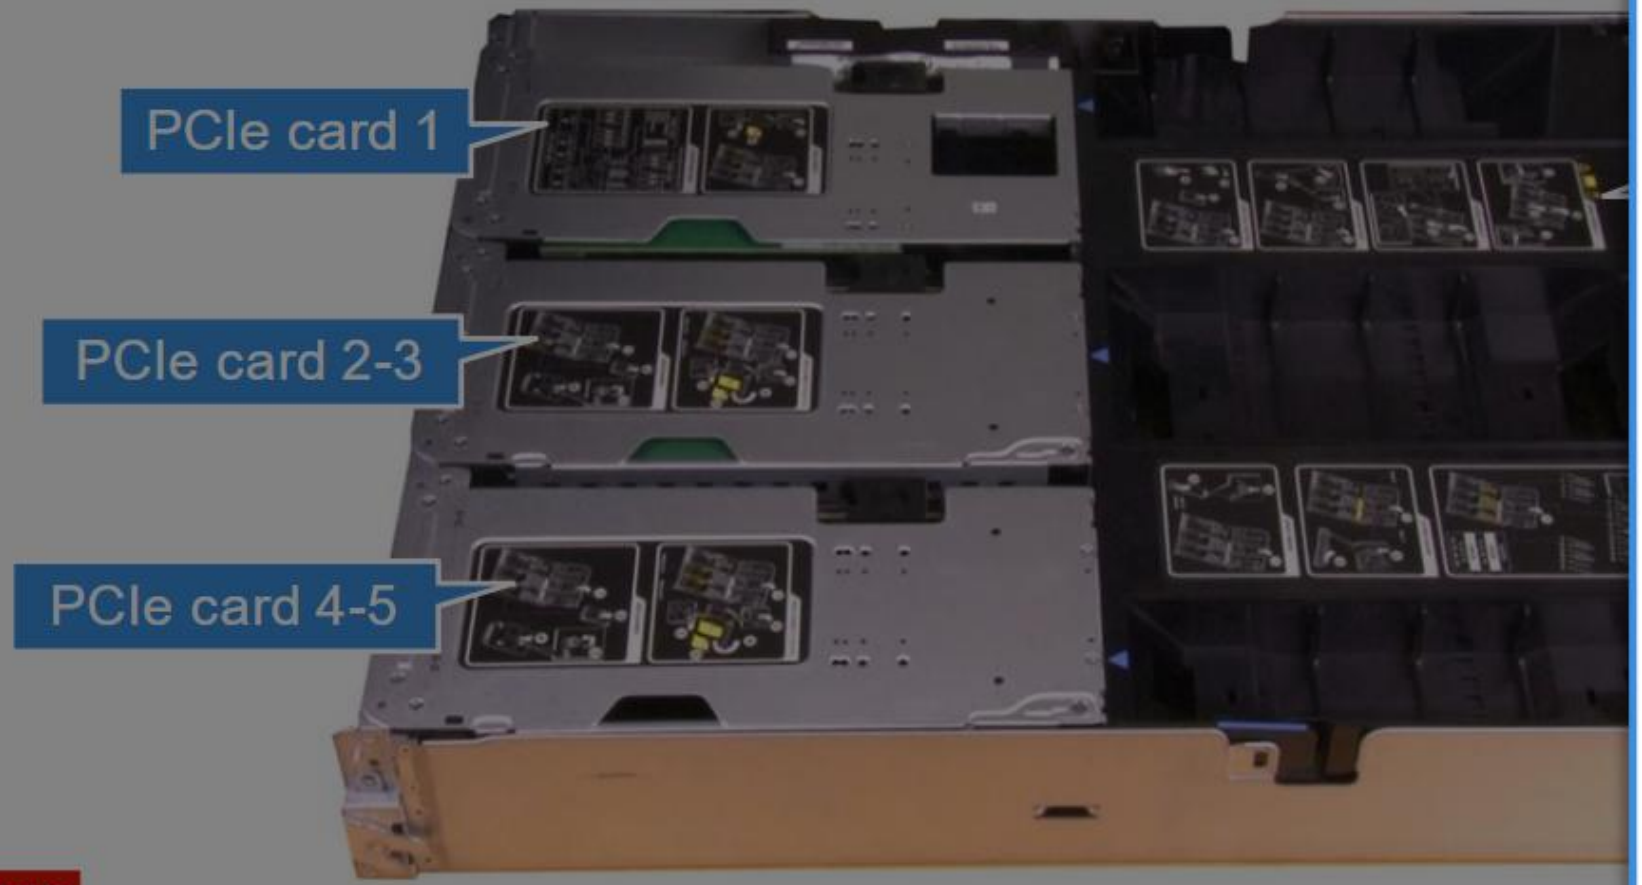

PCIe card 1

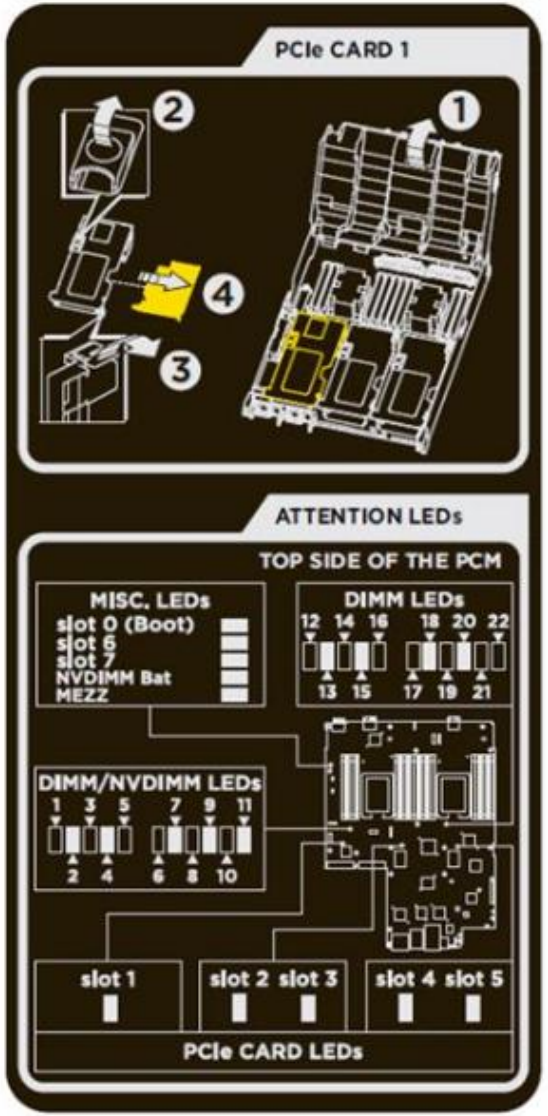

System Service Labels

System service labels are located on the PCIe risers and air duct. The label
locations of the FRU parts, attention LEDs, and the basic steps used to repl
Click the items to see larger images.

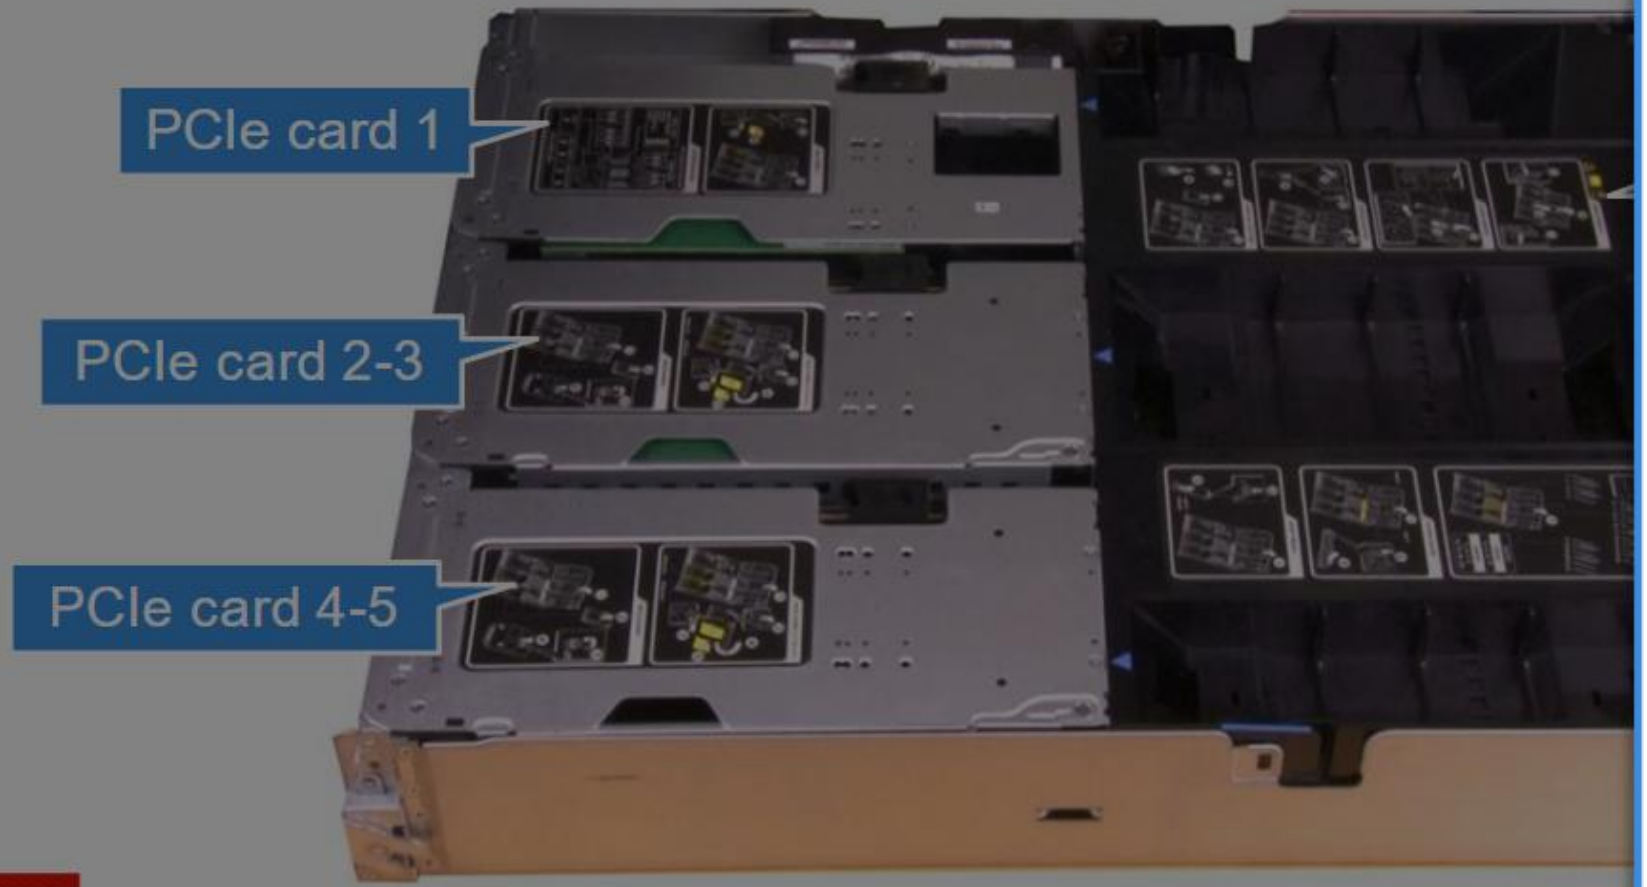

PCIe card 2-3

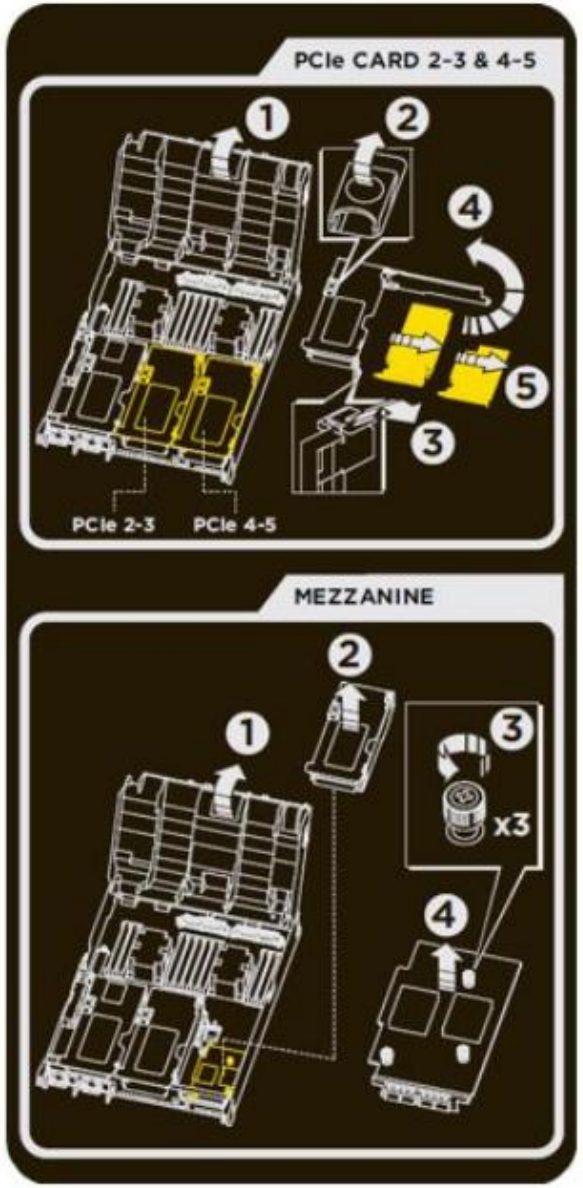

System Service Labels

System service labels are located on the PCIe risers and air duct. The label
locations of the FRU parts, attention LEDs, and the basic steps used to repl
Click the items to see larger images.

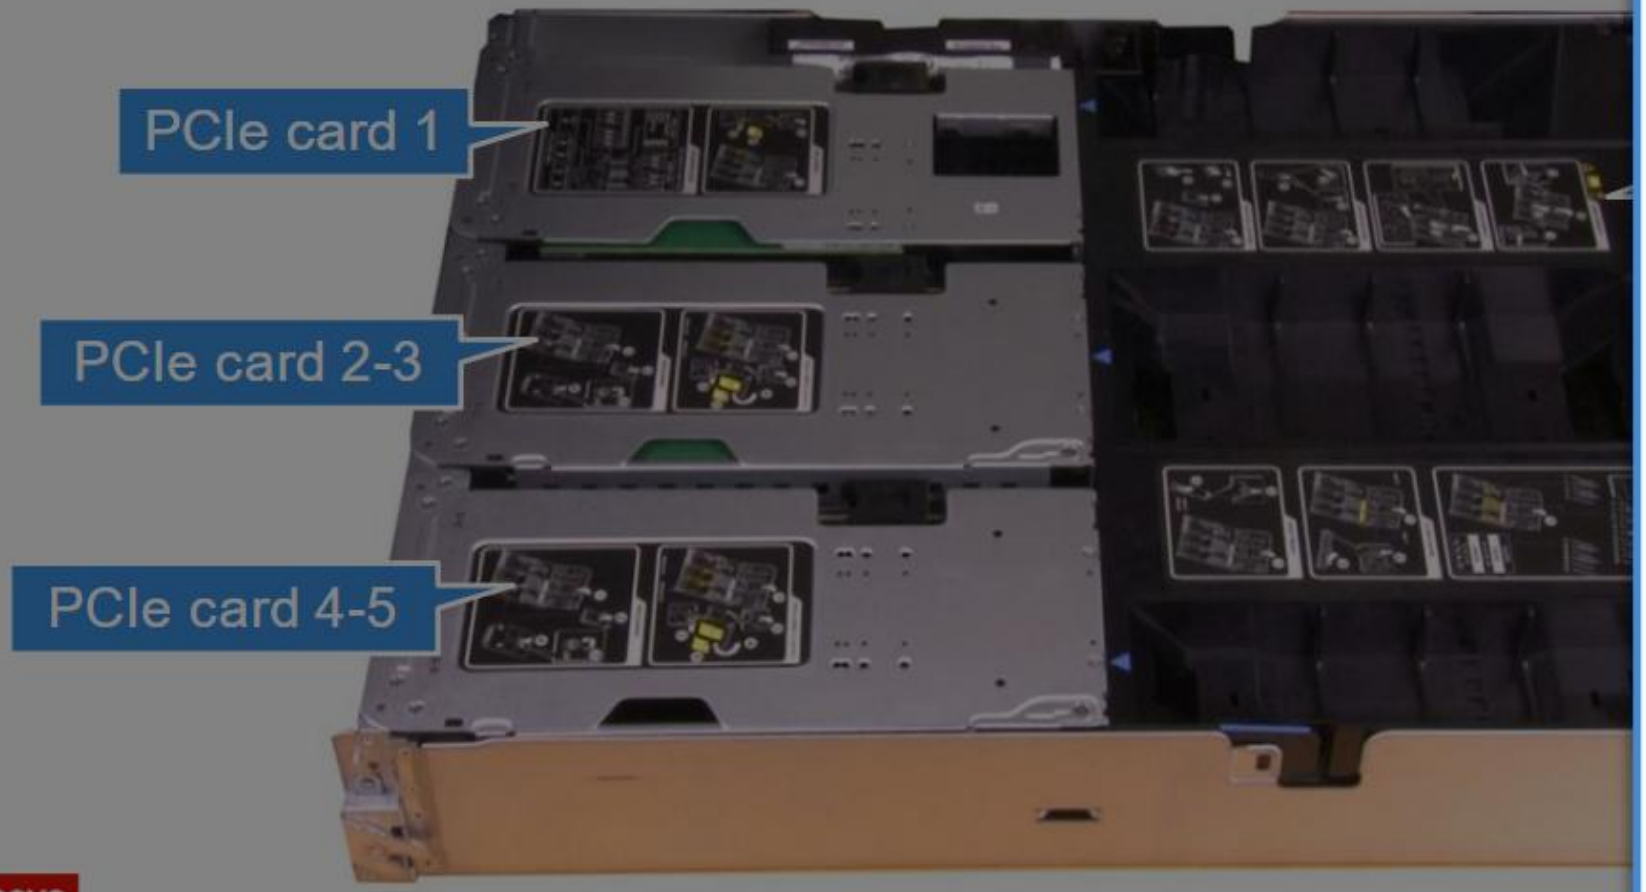

PCIe card 4-5

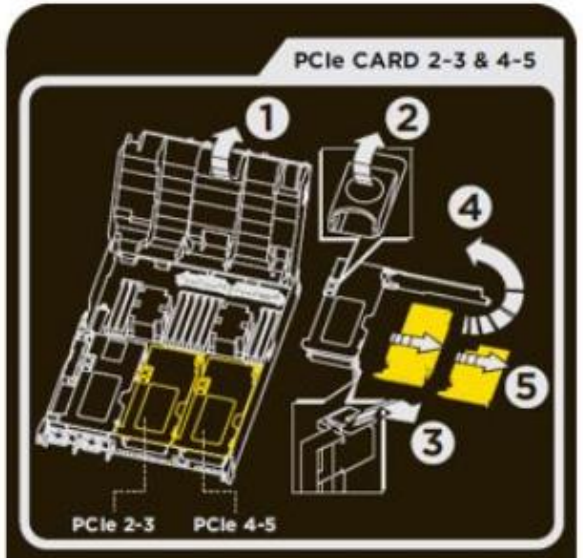

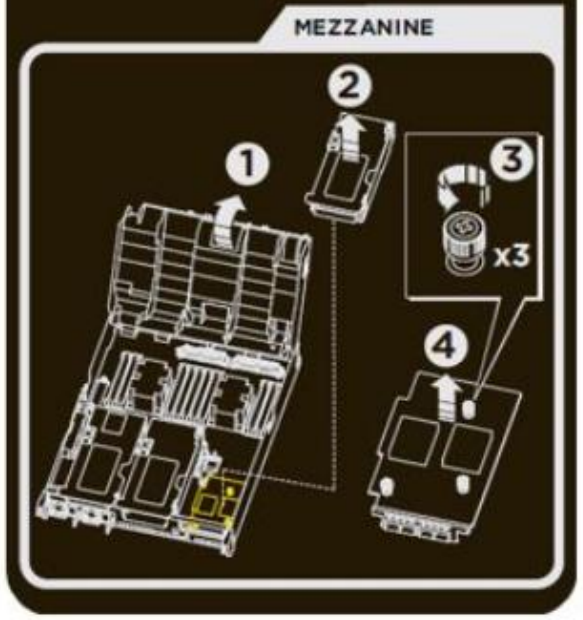

Lenovo

System Service Labels

System service labels are located on the PCIe risers and air duct. The label
locations of the FRU parts, attention LEDs, and the basic steps used to repl
Click the items to see larger images.

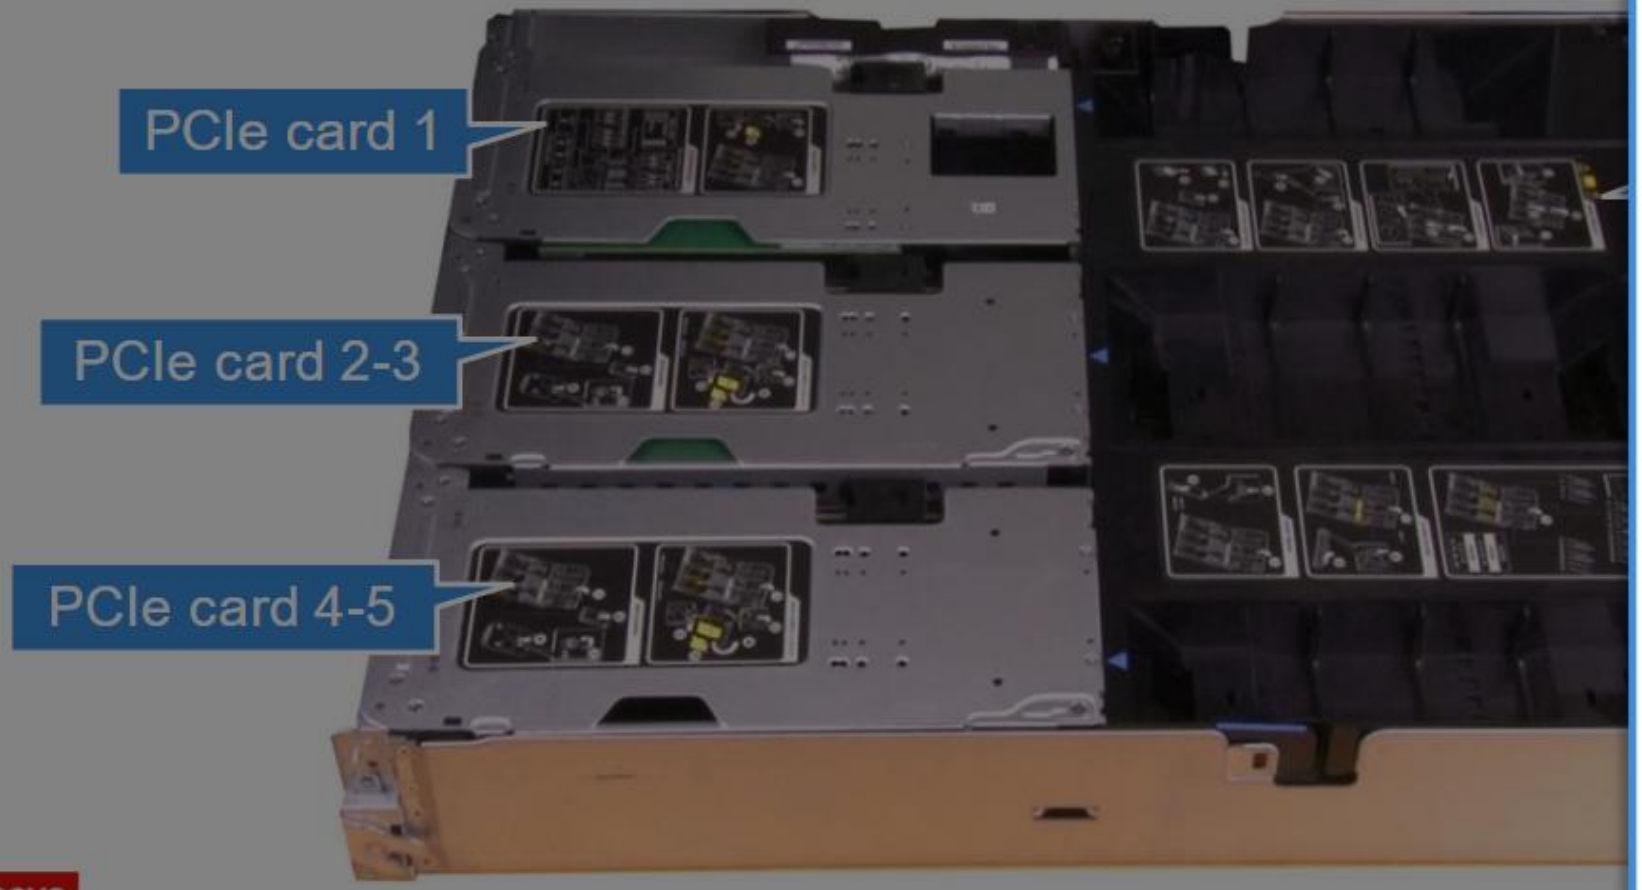

Air duct

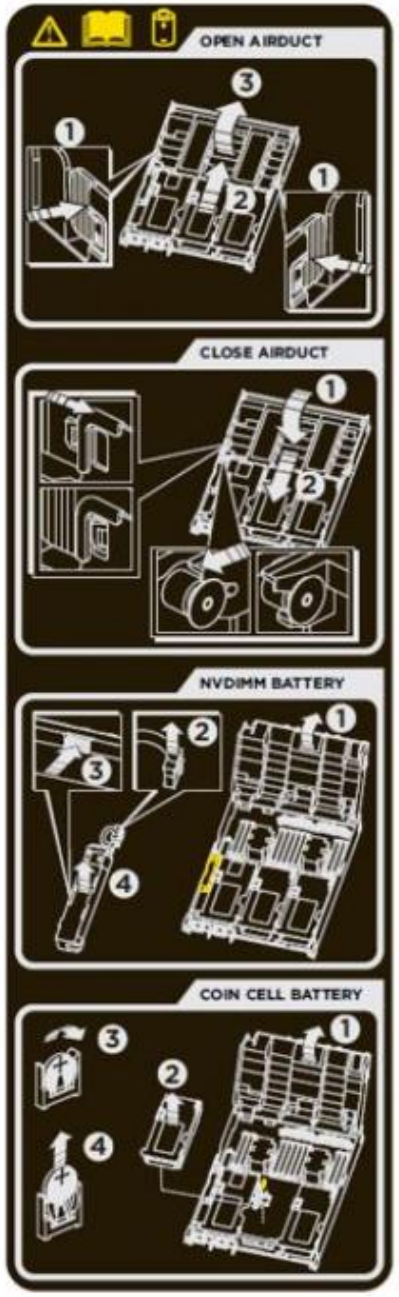

System Service Labels

System service labels are located on the PCIe risers and air duct. The label
locations of the FRU parts, attention LEDs, and the basic steps used to repl
Click the items to see larger images.

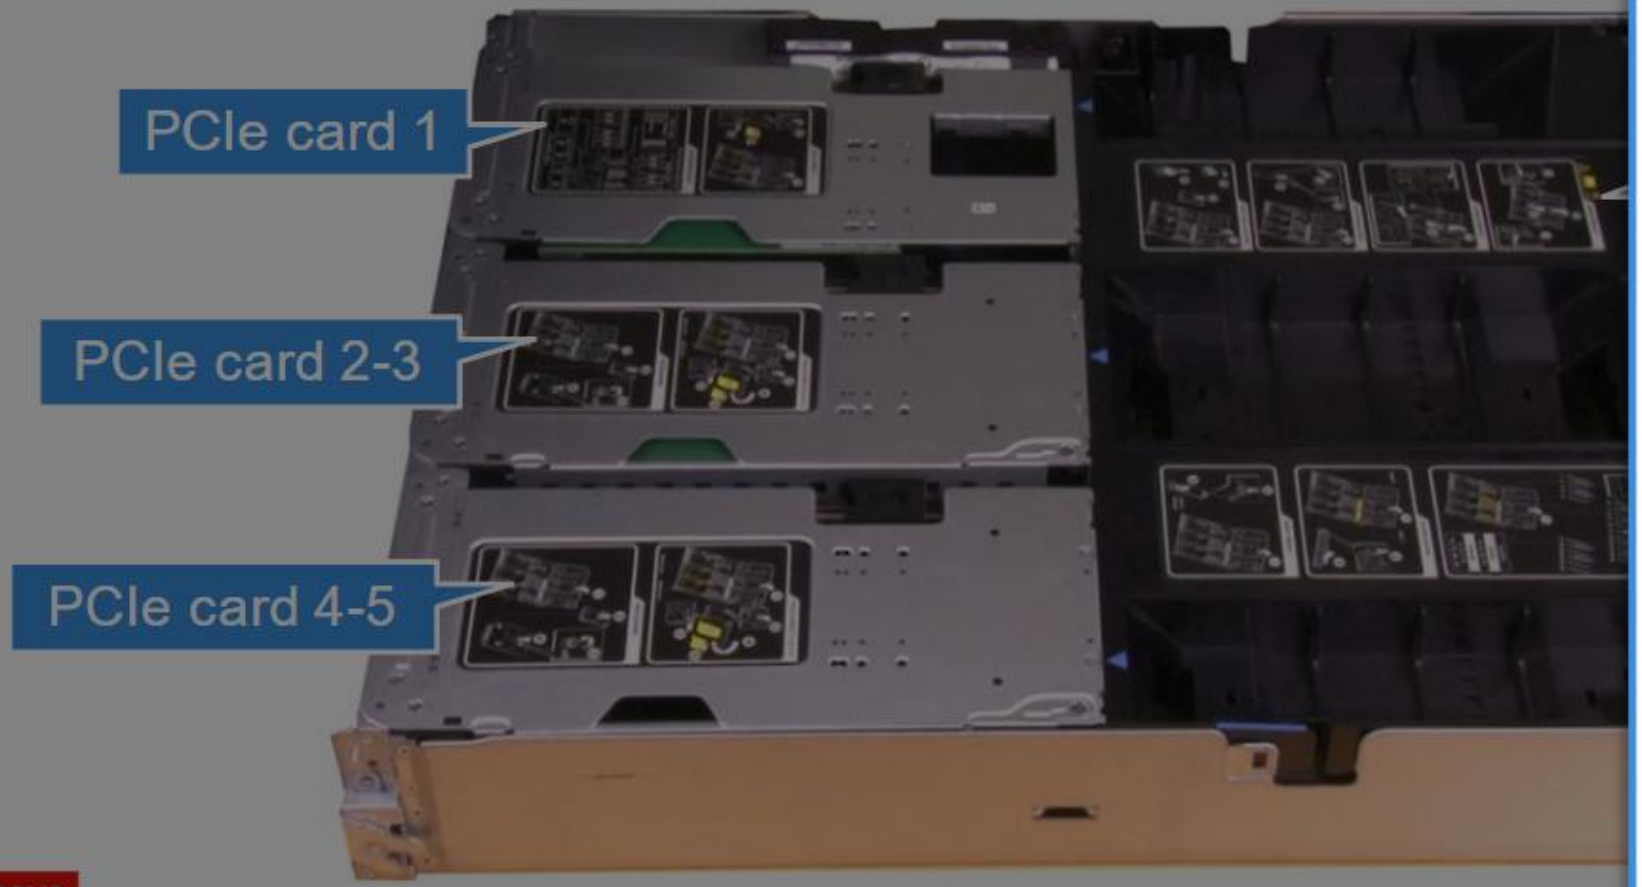

NVDIMM/DIMM

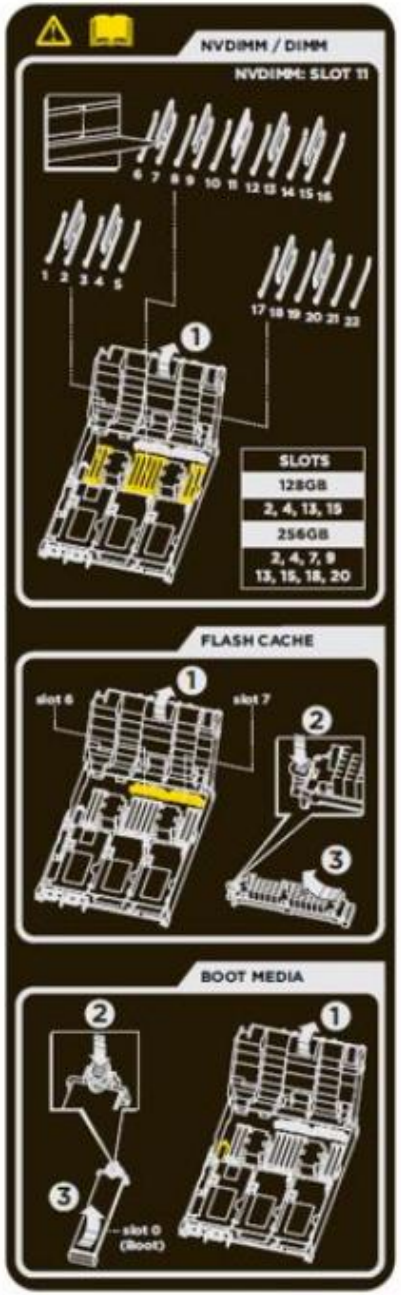

Lenovo

Shutting down an impaired node

Work through the following procedure to shut down an impaired node before replacing a controller or its components.

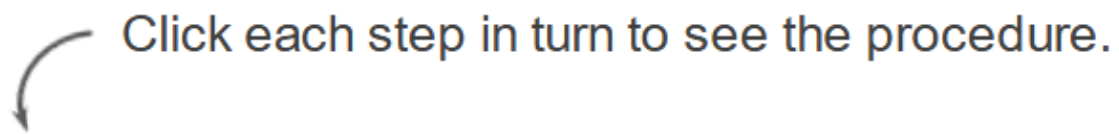

Step 1 2

Lenovo

Shutting down an impaired node

To shut down an impaired node, you must determine its status and, if necessary, take over the node so that the healthy node will continue to serve data from the impaired node's storage.

- If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

Example:

- The following AutoSupport message suppresses automatic case creation for two hours:
  - `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

Shutting down an impaired node

- Disable automatic giveback from the console of the healthy node:
  - `storage failover modify -node local -auto-giveback false`
- Take the impaired node to the LOADER prompt:

If the impaired node is displaying	Then…
The LOADER prompt	Go to the next step.
`Waiting for giveback…`	Press **Ctrl-C**, and then respond with `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired node: • For an HA pair, take over the impaired node from the healthy node ○ `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired node shows `Waiting for giveback...`, press **Ctrl-C**, and then respond with `y`.

Step 1—2

Erasing data on a caching module

Although data on the caching module is encrypted, you might want to erase any data from the impaired caching module and verify that the caching module has no data:

- Use the following command to erase data on the caching module:
  - `system controller flash-cache secure-erase run [-slot] slot#`
- Use the following command to verify that data has been erased from the caching module:
  - `system controller flash-cache secure-erase show -node node_name`
- The output should display the caching module status as `erased`.

Running system-level diagnostics

It is suggested that you run diagnostic tests for specific components and subsystems whenever the controller is being replaced. All commands in the diagnostic procedures are issued from the node where the component is being replaced.

- Enter the `system node halt -node` *`node_name`* command to reboot the node if the node to be serviced is not at the LOADER prompt.
- At the LOADER prompt, enter the `boot_diags` command to access the special drivers specifically designed for system-level diagnostics.
- Select **Scan System** from the displayed menu to enable the diagnostics tests.
- Select **Test system** from the displayed menu to run the diagnostics tests.
  - Select **Test Memory** from the displayed menu if you only want to test the DIMM.
  - Select **NVDIMM Test** from the displayed menu if you only want to test the NVDIMM.
- Based on the results of the preceding step:
  - If the scan show problems, correct the issue, and then rerun the scan.
  - If the scan reported no failures, select **Reboot** from the menu to reboot the system.

Restoring a controller module to operation after running diagnostics

After completing diagnostics, you must re-cable the system, give back the controller module, and then re-enable automatic giveback.

- Re-cable the system, as needed.
  - If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.
- Return the node to normal operation by giving back its storage:
  - `storage failover giveback -ofnode impaired_node_name`
- If automatic giveback was disabled, re-enable it:
  - `storage failover modify -node local -auto-giveback true`

Transferring a boot image to a boot media

The replacement boot media that you installed will not have a boot image, so you need to transfer a boot image using a USB flash drive. The USB drive should be formatted to FAT32 and have the following items:

- A copy of the same image version of ONTAP that the impaired controller was running. You can download the appropriate image from the Lenovo Support Site.
- If the system is an HA pair, you must have a network connection.

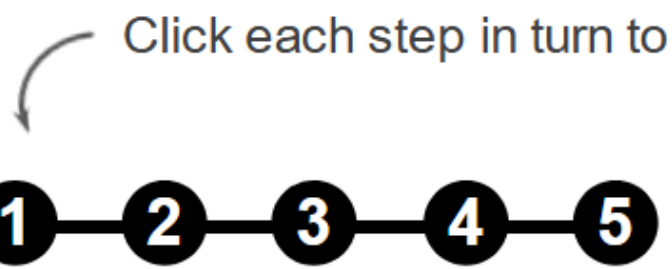

Transferring a boot image to a boot media

- Unzip the service image.
- There will be two folders in the unzipped service image file:
  - boot
  - efi
- Copy the efi folder to the top directory on the USB flash drive.
- The USB flash drive should have the efi folder and the same image version of ONTAP that the impaired controller was running.
- Remove the USB flash drive from your laptop.

Note: If you are extracting the contents using Windows, do not use winzip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

Transferring a boot image to a boot media

- Gently push the controller module halfway into the system.
- Reinstall the cables and cable management device, including the media converters (SFPs or QSFPs) if they were removed.
- Insert the USB flash drive into the USB slot on the controller module (not the micro-USB console port).
- Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.
- Push the controller module all the way into the system, firmly push the cam handle to finish seating the controller module, and then push the cam handle into the closed position.
- The node will begin to boot as soon as it is completely installed in the chassis.
- Interrupt the boot process to stop at the LOADER prompt by pressing **Ctrl-C** when you see `Starting AUTOBOOT press Ctrl-C to abort....`
- If you miss this message, press **Ctrl-C** to boot to Maintenance mode, and then `halt` the node to boot to LOADER.
- The system will begin to boot and then stop at the LOADER prompt.

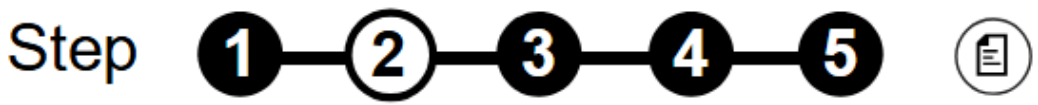

Transferring a boot image to a boot media

- From the LOADER prompt, boot the recovery image from the USB flash drive:
  - `boot_recovery`
  - The image will be downloaded from the USB flash drive.
- When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
- After the image has been installed, start the restoration process:
  - Record the IP address of the impaired node that is displayed on the screen.
  - Press `y` when prompted to restore the backup configuration.
  - Press `y` when prompted to overwrite `/etc/ssh/ssh_host_dsa_key`.
- From the partner node in advanced privilege level, start the configuration synchronization using the IP address recorded in the previous step:
  - `system node restore-backup -node local -target-address` *`impaired_node_IP_address`*
- If the restoration is successful, press `y` on the impaired node when prompted to `use the restored copy?`.
- Press `y` when you see `confirm backup procedure was successful` message, and then press `y` when prompted to reboot the node.

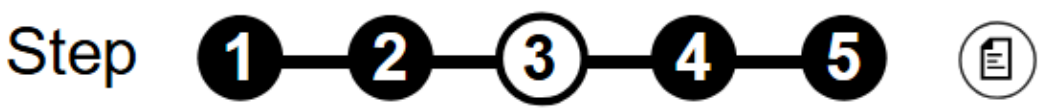

Transferring a boot image to a boot media

Verify that the environmental variables are set as expected.

- Take the node to the LOADER prompt.
  - From the ONTAP prompt, you can issue the following command: `system node halt -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true`.
- Check the environment variable settings with the `printenv` command.
- If an environment variable is not set as expected, modify it with the `setenv` *`environment_variable_name changed_value`* command.
- Save your changes using the `saveenv` command.
- Reboot the node.

Transferring a boot image to a boot media

With the rebooted impaired node displaying the `Waiting for Giveback...` message, perform a giveback from the healthy node:

If the system is in…	Then…
an HA pair	after the impaired node displays the `Waiting for Giveback...` message, perform a giveback from the healthy node. • Perform a giveback from the healthy node by entering the following command: • `storage failover giveback -ofnode` *`partner_node_name`* • The impaired node will take back its storage, finish booting, reboot, and will then be taken over again by the healthy node. • Monitor the progress of the giveback operation by using the `storage failover show-giveback` command. • After the giveback operation is complete, confirm that the HA pair is healthy and that takeover is possible by using the `storage failover show` command. • If you disabled automatic giveback, restore it by using the `storage failover modify` command. • Exit advanced privilege level on the healthy node.

Step 1 2 3 4 5

Lenovo