

Figure 182. Lifting the system cover

3. Squeeze the orange release tabs on the fan module.
4. Lift the fan module away from the system.

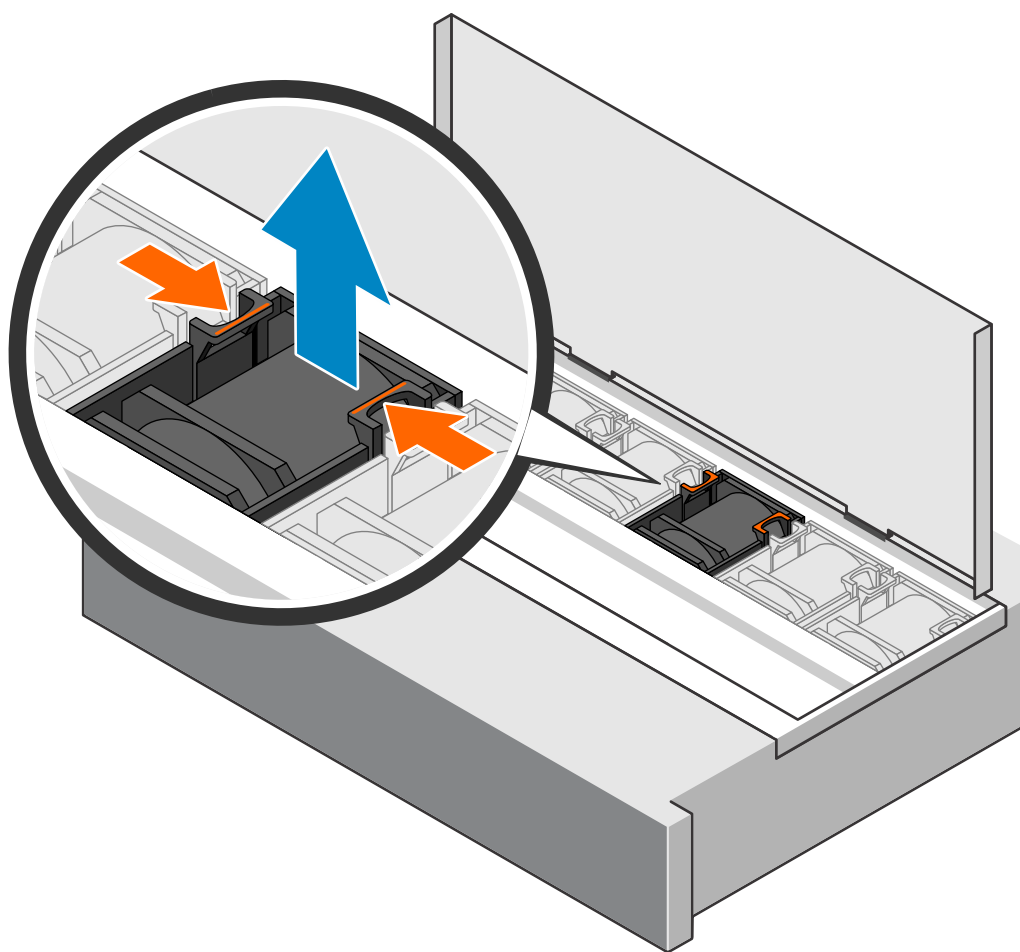


Figure 183. Removing a fan module

Install a fan module

Steps

1. Push the fan module into the empty slot.

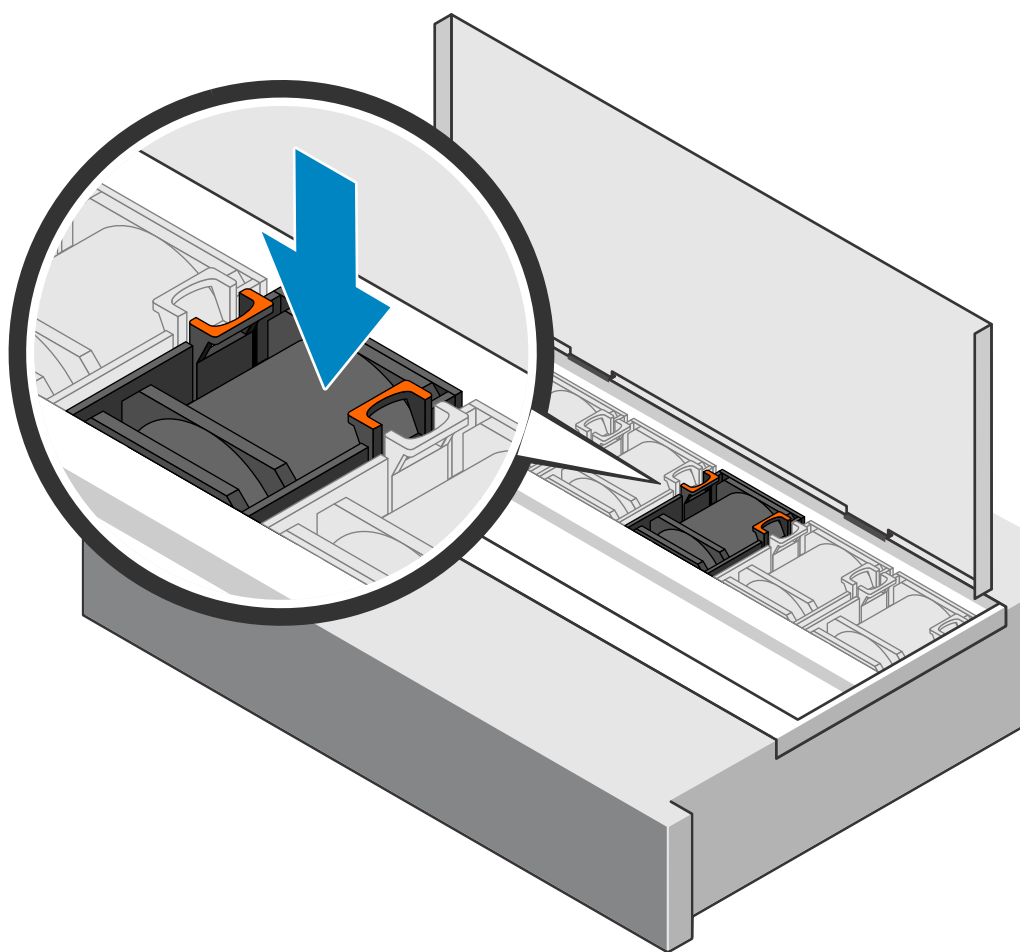


Figure 184. Installing a fan module

2. Close the system cover.

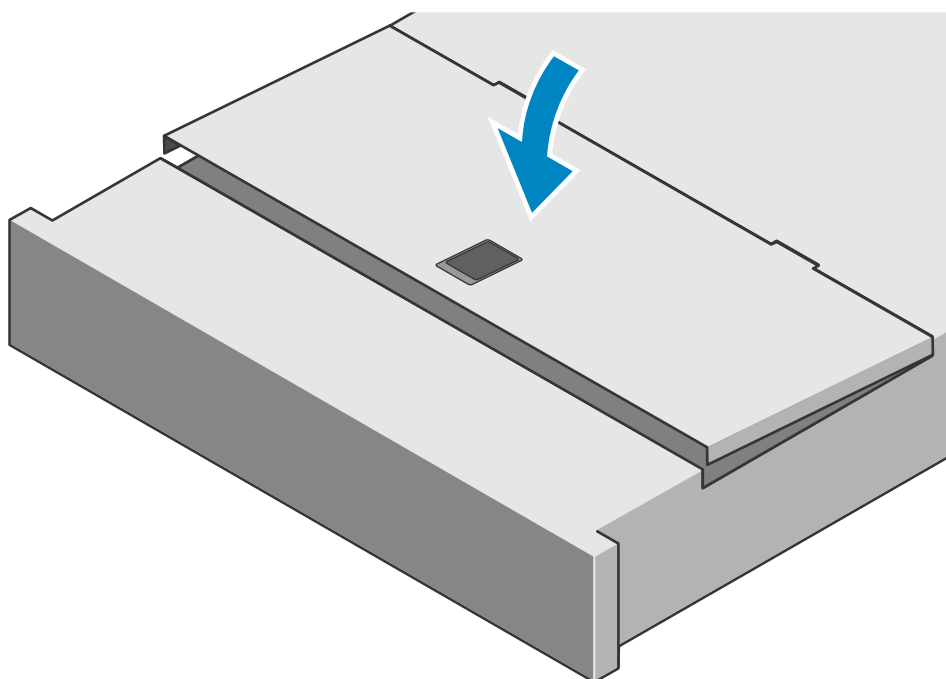


Figure 185. Closing the system cover

3. Push the expansion enclosure into the rack.

Verify the operation of a replacement fan module

Steps

1. From PowerStore Manager, select **Hardware**.
2. Select the appliance where you replaced the fan module.
3. On the **Components** card, under **Internal View**, expand the node that includes the fan module, and then select the relevant **FanModule**.

The status of the replacement fan module should read `Healthy`. If the status is still `Faulted`, wait a few minutes and refresh PowerStore Manager. If the status does not change, ensure that the fan module is correctly seated, or contact your service provider.

Return a faulted part

About this task

For US customers, return defective material within five business days. For International customers, return defective material within 10 business days. The materials required to return your defective part are supplied with the good part shipment.

Steps

1. Package the faulted part in the shipping box that contained the replacement part.
2. Ship the failed part to your service provider as described in the instructions that were included with the replacement part.
3. For more information about returning customer-replaceable parts:
 - a. Open PowerStore Manager.
 - b. Click **Settings** on the upper right of the screen.
 - c. Click **General Support**.
 - d. Under **Drives, Power Supplies, and Other Parts**, click **Return Part**.
 - e. If your screen does not show the Return Part link, contact your service provider for instructions.

Replace a Clock Distribution Board in an NVMe expansion enclosure

Take the following actions to remove the faulted Clock Distribution Board from the NVMe expansion enclosure and install a replacement Clock Distribution Board.

Identify a faulted Clock Distribution Board from PowerStore Manager

Before you replace a Clock Distribution Board, ensure that you have identified its location within the system. Using PowerStore Manager, you can identify and locate a faulted Clock Distribution Board.

Steps

1. From PowerStore Manager, select **Hardware**.
2. Select the appliance that contains the Clock Distribution Board that you need to replace.
3. On the **Components** card, under **Rear View**, expand **ExpansionEnclosure**.
4. Select the relevant **CDB**.
Faulted parts appear in red in the image of the system, and report a status of `Faulted` in the **State** field.

Remove a Clock Distribution Board

Steps

1. Pull the expansion enclosure from the rack until the system cover is accessible.
2. Lift open the system cover.



CAUTION: Do not leave the system cover open for more than two minutes. If you need more time, close the cover and allow the system temperature to stabilize before proceeding.

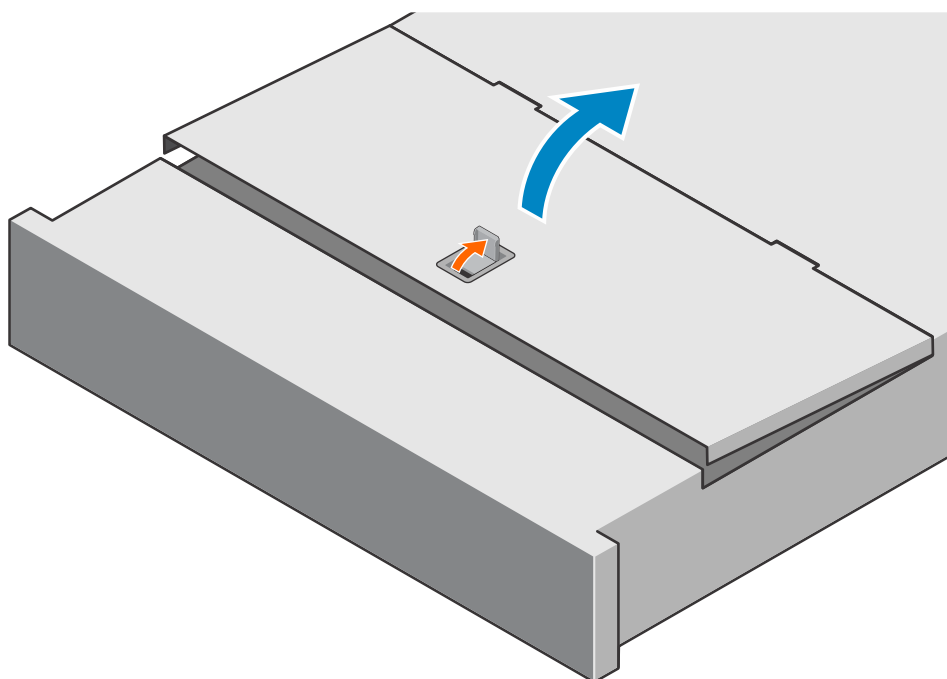


Figure 186. Lifting the system cover

3. Squeeze the orange release tabs on the Clock Distribution Board.
4. Slide the Clock Distribution Board toward the fans and then up out of the system.

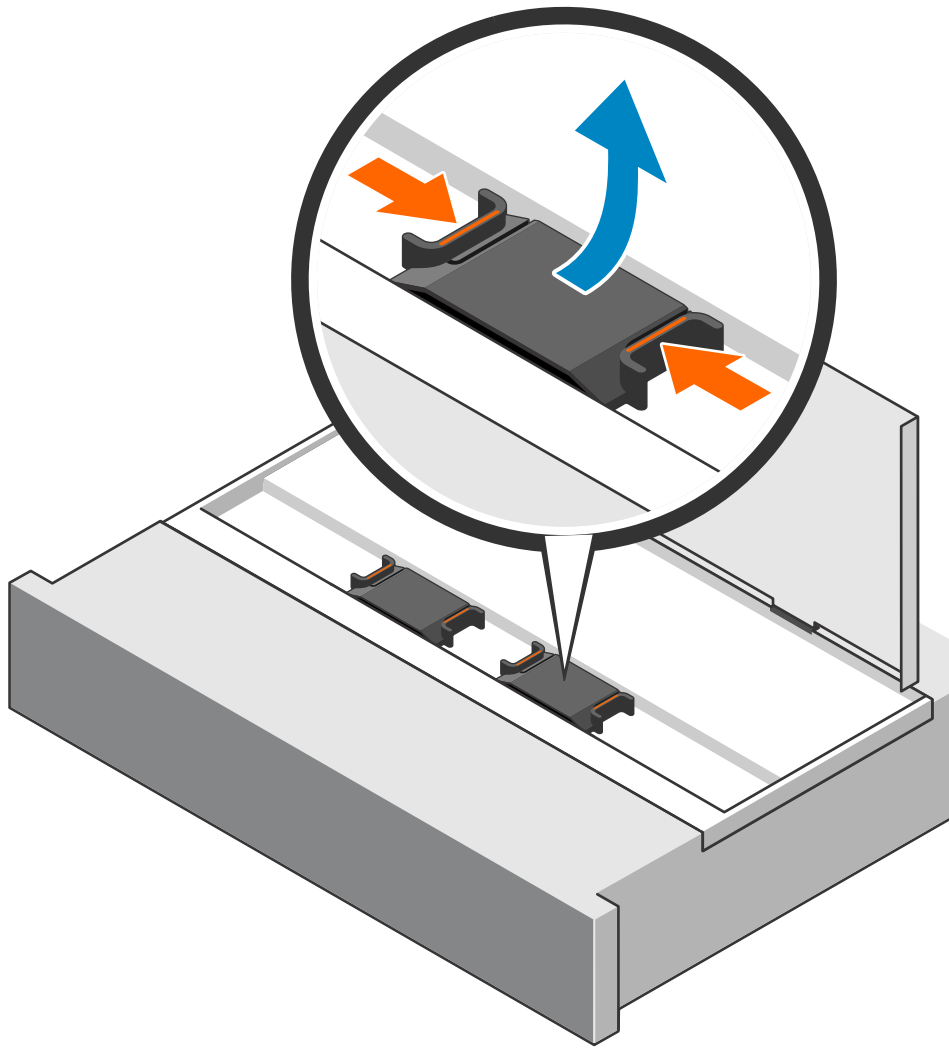


Figure 187. Removing the Clock Distribution Board

Install a Clock Distribution Board

Steps

1. Squeeze the orange tabs and align the Clock Distribution Board with the empty slot.
2. Push the Clock Distribution Board into the empty slot.

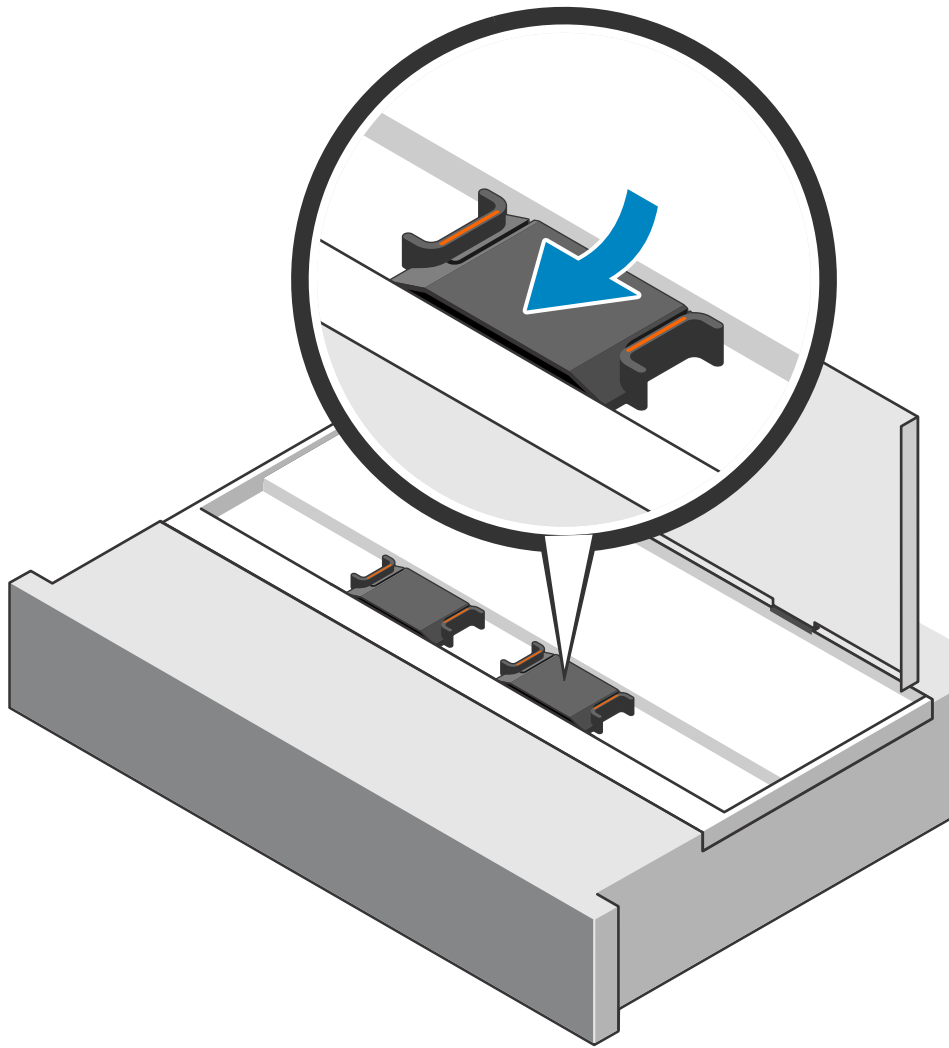


Figure 188. Installing the Clock Distribution Board

3. Close the system cover.

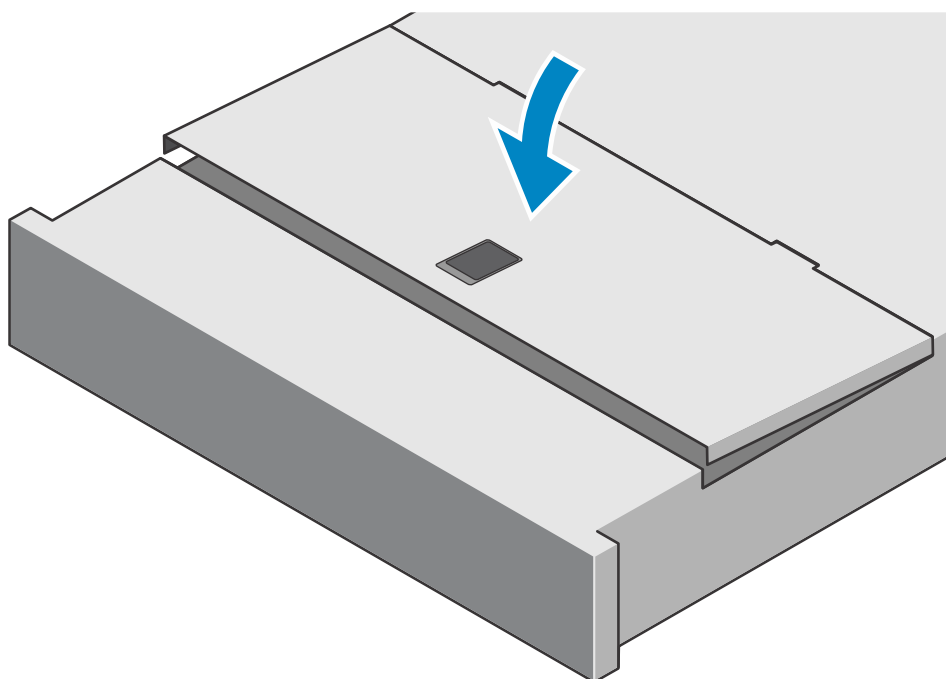


Figure 189. Closing the system cover

4. Push the expansion enclosure into the rack.

Verify the operation of a replacement Clock Distribution Board

Steps

1. From PowerStore Manager, select **Hardware**.
2. Select the appliance where you replaced the Clock Distribution Board.
3. On the **Components** card, under **Rear View**, expand **ExpansionEnclosure**.
4. Select the relevant Clock Distribution Board.

The status of the replacement Clock Distribution Board should read `Healthy`. If the status is still `Faulted`, wait a few minutes and refresh PowerStore Manager. If the status does not change, ensure that the Clock Distribution Board is correctly seated, or contact your service provider.

Return a faulted part

About this task

For US customers, return defective material within five business days. For International customers, return defective material within 10 business days. The materials required to return your defective part are supplied with the good part shipment.

Steps

1. Package the faulted part in the shipping box that contained the replacement part.
2. Ship the failed part to your service provider as described in the instructions that were included with the replacement part.
3. For more information about returning customer-replaceable parts:
 - a. Open PowerStore Manager.
 - b. Click **Settings** on the upper right of the screen.
 - c. Click **General Support**.
 - d. Under **Drives, Power Supplies, and Other Parts**, click **Return Part**.
 - e. If your screen does not show the Return Part link, contact your service provider for instructions.

Replace an Access Module in an NVMe expansion enclosure

Take the following actions to remove the faulted Access Module from the NVMe expansion enclosure and install a replacement Access Module.

Identify a faulted Access Module from PowerStore Manager

Before you replace an Access Module, ensure that you have identified its location within the system. Using PowerStore Manager, you can identify and locate a faulted Access Module.


Steps

1. From PowerStore Manager, select **Hardware**.
2. Select the appliance that contains the Access Module that you need to replace.
3. On the **Components** card, under **Rear View**, expand **ExpansionEnclosure**.
4. Select the relevant **Access Module**.

Faulted parts appear in red in the image of the system, and report a status of `Faulted` in the **State** field.

Remove an Access Module

About this task

 **NOTE:** Access Module 1 is on the top of the NVMe expansion enclosure and Access Module 2 is on the bottom.

Steps

1. Label and remove the cables from the Access Module.
2. Press both orange buttons to release the Access Module latches.

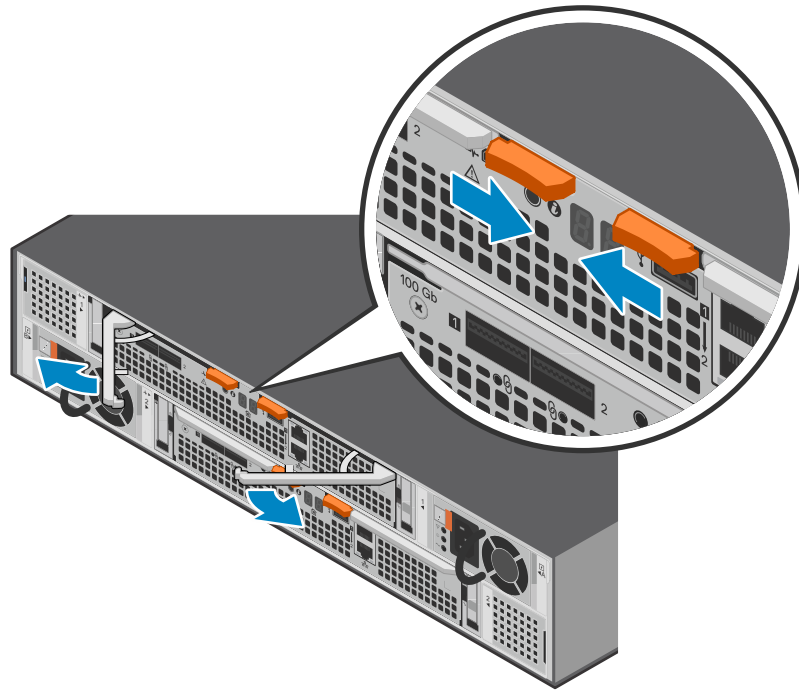


Figure 190. Releasing the Access Module

3. Pull the latches to remove the Access Module from the chassis.

NOTE: The Access Module comes completely out of the chassis. In addition to holding the latches, be prepared to support the Access Module to avoid dropping it.

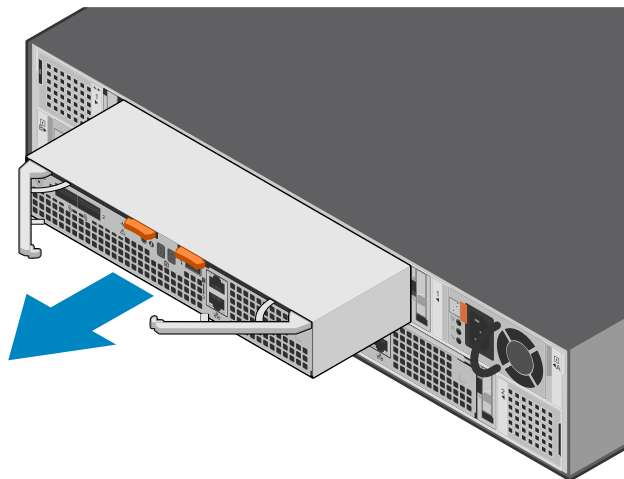


Figure 191. Removing the Access Module

Install an Access Module

Steps

1. Align the Access Module with the empty slot and carefully push it into the slot.

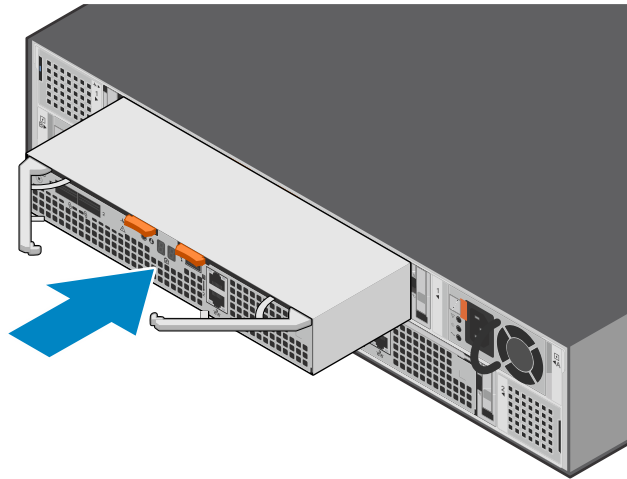


Figure 192. Installing the Access Module

2. Press in on the Access Module latches to lock them into place.

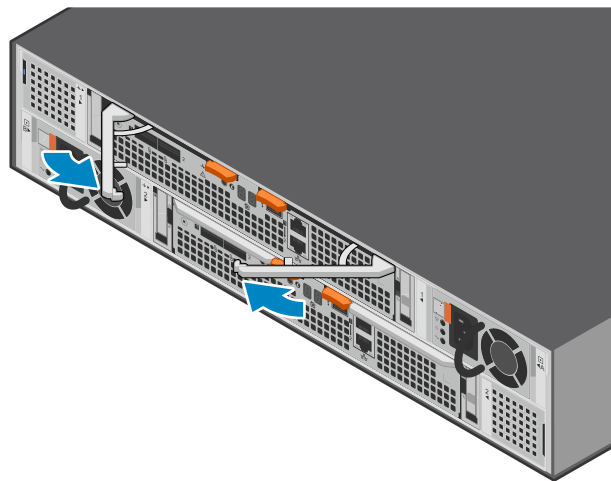


Figure 193. Locking the Access Module into place

3. Connect the cables to the Access Module.

Verify the operation of a replacement Access Module

Steps

1. From PowerStore Manager, select **Hardware**.
2. Select the appliance where you replaced the Access Module.
3. On the **Components** card, under **Rear View**, expand **ExpansionEnclosure**.
4. Select the relevant **Access Module**.

The status of the replacement Access Module should read `Healthy`. If the status is still `Faulted`, wait a few minutes and refresh PowerStore Manager. If the status does not change, ensure that the Access Module is correctly seated, or contact your service provider.

Return a faulted part

About this task

For US customers, return defective material within five business days. For International customers, return defective material within 10 business days. The materials required to return your defective part are supplied with the good part shipment.

Steps

1. Package the faulted part in the shipping box that contained the replacement part.
2. Ship the failed part to your service provider as described in the instructions that were included with the replacement part.
3. For more information about returning customer-replaceable parts:
 - a. Open PowerStore Manager.
 - b. Click **Settings** on the upper right of the screen.
 - c. Click **General Support**.
 - d. Under **Drives, Power Supplies, and Other Parts**, click **Return Part**.
 - e. If your screen does not show the Return Part link, contact your service provider for instructions.

Replace a data interface board in an NVMe expansion enclosure

Take the following actions to remove a faulted data interface board (DIB) from an NVMe expansion enclosure and install a replacement DIB.

Identify a faulted DIB from PowerStore Manager

Before you replace a DIB, ensure that you have identified its location within the system. Using PowerStore Manager, you can identify and locate a faulted DIB.

Steps

1. From PowerStore Manager, select **Hardware**.
2. Select the appliance that contains the DIB that you need to replace.
3. On the **Components** card, under **Rear View**, expand **ExpansionEnclosure**.
4. Select the relevant **DIB**.

Faulted parts appear in red in the image of the system, and report a status of `Faulted` in the **State** field.

Removing a DIB

Steps

1. Remove the Access Module as described in [Remove an Access Module](#).
2. Press down on both orange buttons to release the DIB latches.

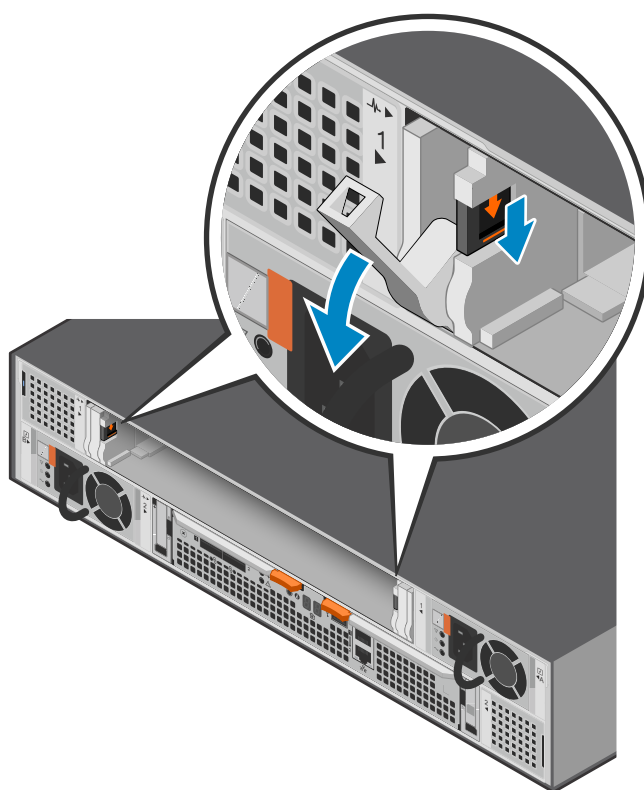


Figure 194. Releasing the DIB

3. Press down firmly on the latches, and then pull the latches to remove the DIB from the chassis.

NOTE: The DIB comes completely out of the chassis. In addition to holding the latches, be prepared to support the DIB to avoid dropping it.

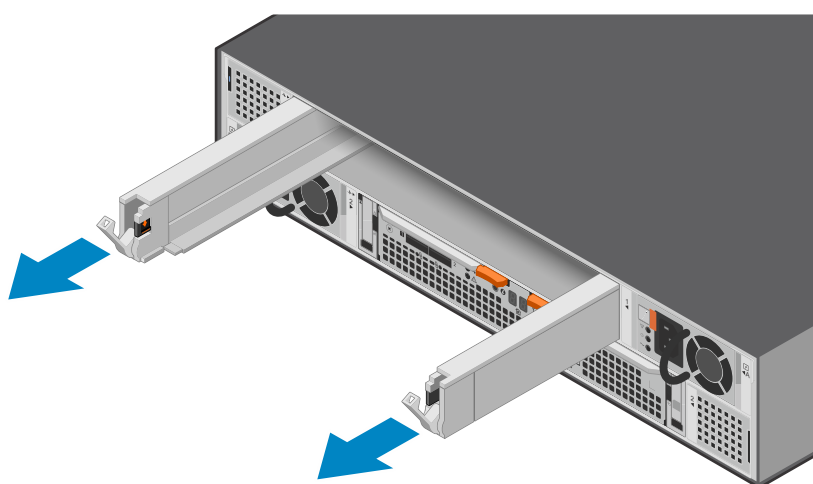


Figure 195. Removing the DIB

Replacing a DIB

Steps

1. Align the DIB with the empty slot and carefully push it into the slot until the latches are engaged and start to rise.

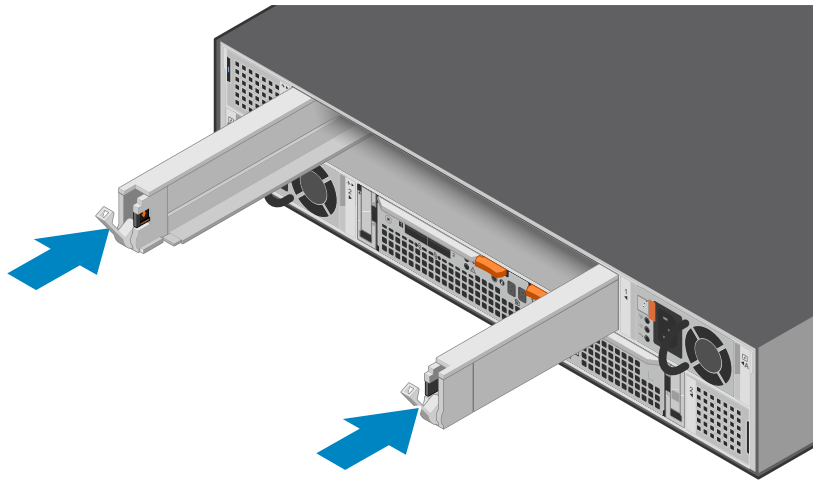


Figure 196. Installing the DIB

2. Press up on the DIB latches to lock them into place.

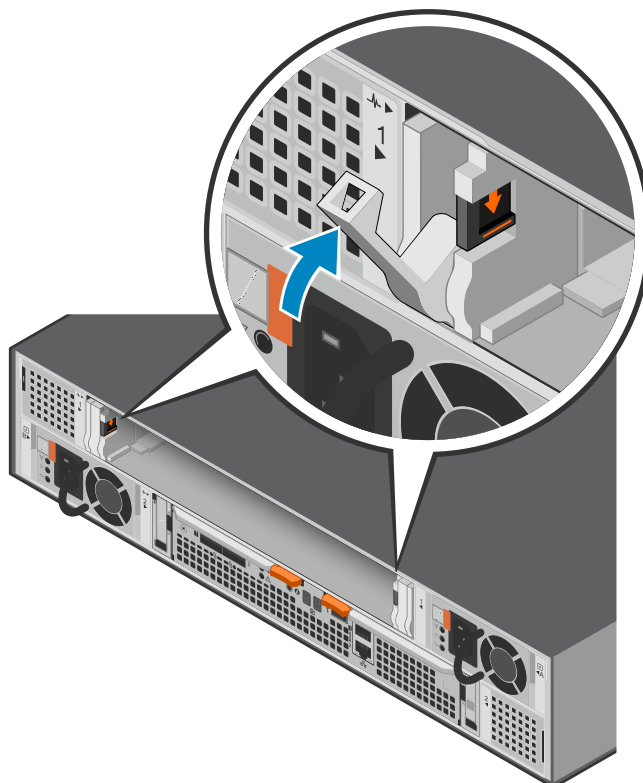


Figure 197. Locking the DIB into place

3. Replace the Access Module as described in [Install an Access Module](#).

Verify the operation of a replacement DIB

Steps

1. From PowerStore Manager, select **Hardware**.
2. Select the appliance where you replaced the DIB.
3. On the **Components** card, under **Rear View**, expand **ExpansionEnclosure**.
4. Select the relevant **DIB**.

The status of the replacement DIB should read `Healthy`. If the status is still `Faulted`, wait a few minutes and refresh PowerStore Manager. If the status does not change, ensure that the DIB is correctly seated, or contact your service provider.

Return a faulted part

About this task


For US customers, return defective material within five business days. For International customers, return defective material within 10 business days. The materials required to return your defective part are supplied with the good part shipment.

Steps


1. Package the faulted part in the shipping box that contained the replacement part.
2. Ship the failed part to your service provider as described in the instructions that were included with the replacement part.
3. For more information about returning customer-replaceable parts:
 - a. Open PowerStore Manager.
 - b. Click **Settings** on the upper right of the screen.
 - c. Click **General Support**.
 - d. Under **Drives, Power Supplies, and Other Parts**, click **Return Part**.
 - e. If your screen does not show the Return Part link, contact your service provider for instructions.

Replace a dual inline memory module (DIMM)

Take the following actions to remove the faulted DIMM and install the replacement DIMM into the system.

 **NOTE:** The DIMMs must stay in their original position. Do not move any DIMMs to a different slot.

Before you begin

 **CAUTION:** Before starting this procedure, use the PowerStore Manager Hardware view and Alerts view to verify that the appliance and peer node are healthy with no outstanding alerts. If multiple nodes need to be removed while performing this procedure, repeat this verification for each affected node before proceeding to remove the next node. If necessary, contact your service provider before starting the replacement procedure.

Identify a faulted DIMM from PowerStore Manager

Before you replace a DIMM, ensure that you have identified its location within the system. Using PowerStore Manager, you can identify and locate a faulted DIMM.

Steps

1. From PowerStore Manager, select **Hardware**.
2. Select the appliance that includes the DIMM that you need to replace.
3. On the **Components** card, under **Internal View**, expand **ExpansionEnclosure**.
4. Select the relevant **DIMM**.

Faulted parts appear in red in the image of the system, and report a status of `Faulted` in the **State** field.

5. You can also identify a faulted DIMM by using the following commands:

To display the DIMM information:

```
svc_diag list --expansion_hardware --sub_option dimm
```

To display the status of the NVMe expansion enclosure:

```
svc_diag list --expansion_hardware --sub_option status
```

To display the verbose output of all of the hardware in the NVMe expansion enclosure including health and status.

```
svc_diag list --expansion_hardware
```

NOTE: The `svc_diag list` command takes a few minutes to run.

Remove an Access Module

About this task

NOTE: Access Module 1 is on the top of the NVMe expansion enclosure and Access Module 2 is on the bottom.

Steps

1. Label and remove the cables from the Access Module.
2. Press both orange buttons to release the Access Module latches.

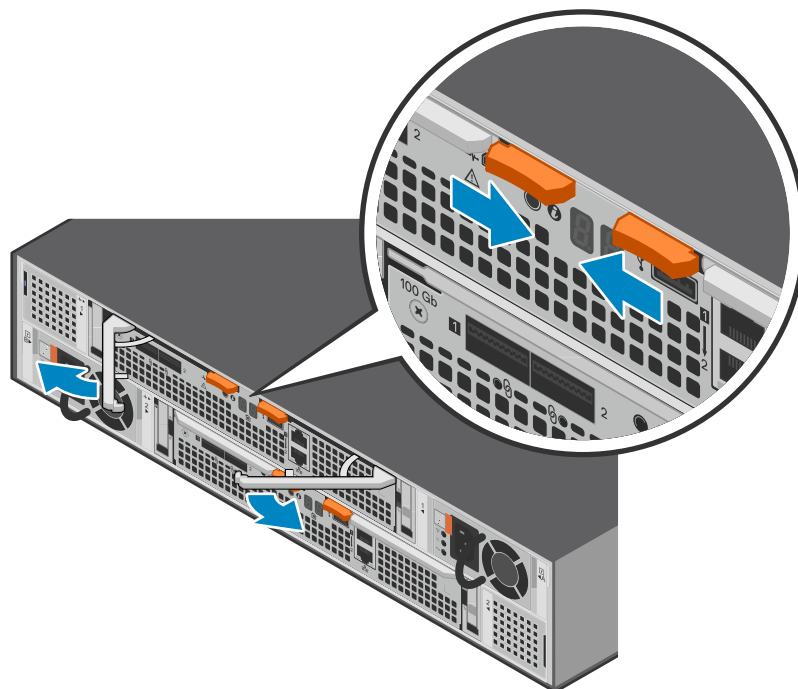


Figure 198. Releasing the Access Module

3. Pull the latches to remove the Access Module from the chassis.

NOTE: The Access Module comes completely out of the chassis. In addition to holding the latches, be prepared to support the Access Module to avoid dropping it.

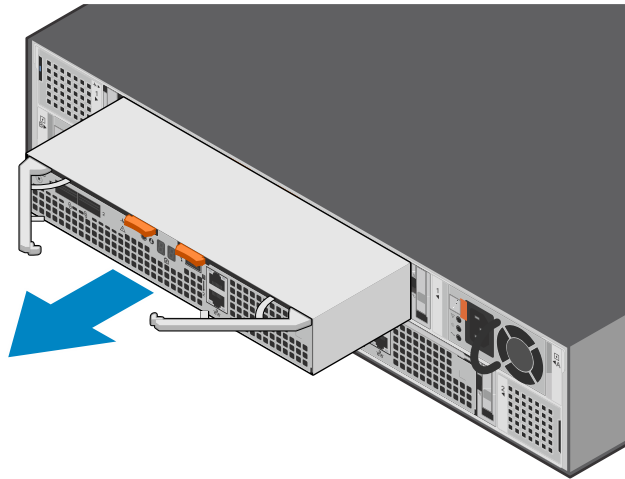


Figure 199. Removing the Access Module

Remove the faulted dual inline memory module

Steps

1. Locate the faulted DIMM in the Access Module by using the picture below as a reference for orientation.

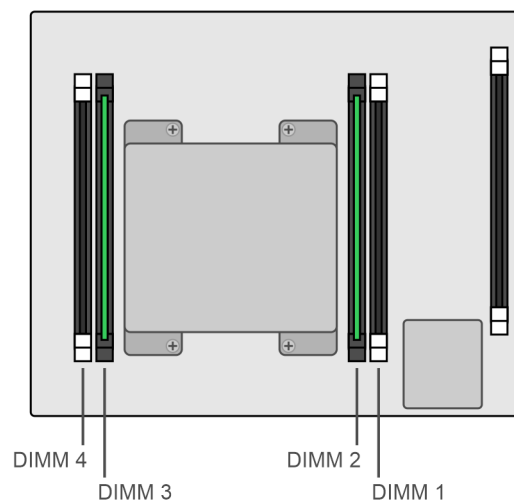


Figure 200. Top view of the Access Module

NOTE: DIMMs are installed in slots 2 and 3.

2. Depress the retaining tabs downward to free the DIMM from its slot.
3. Remove the faulted DIMM.

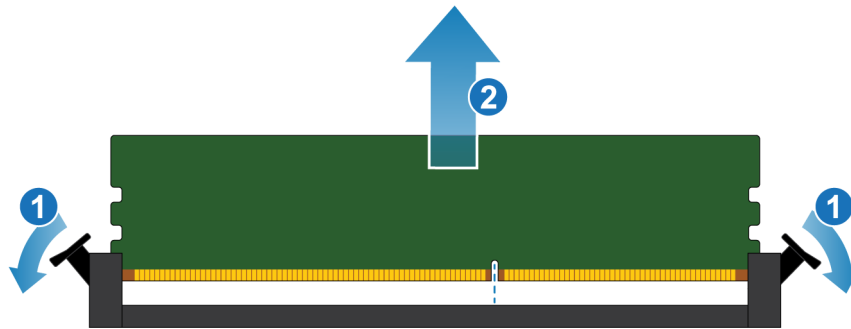


Figure 201. Removing the DIMM

Install the dual inline memory module

Steps

1. Touching only the outside edges of the DIMM, align the DIMM with the connector.
2. Press the DIMM vertically down into the socket using pressure at each end while keeping the leading edge of the DIMM parallel to the connector until it fully seats to the bottom of the socket. When the DIMM engages the contacts in the socket, you will feel resistance, and slightly more force is required to push the module down. During this stage, keep in mind the following precautions:
 - Do not insert the DIMM at an angle.
 - Do not rock the DIMM.
 - Do not insert the DIMM by pushing on one end.
 - Do not seat one end of the DIMM and then the other.

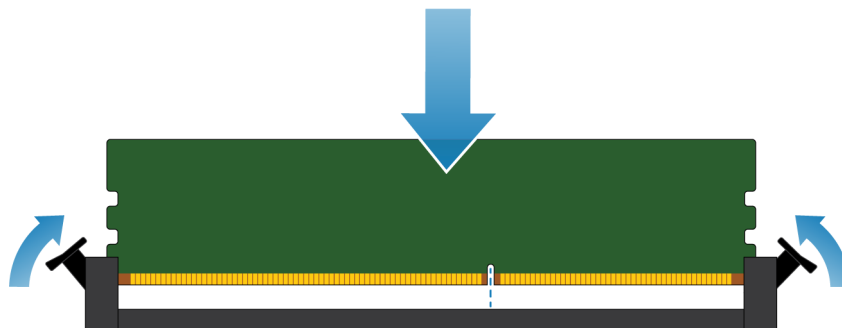


Figure 202. Installing the DIMM

3. Proper DIMM insertion will automatically close the latch ejectors and lock the DIMM into the socket. Verify that the latch ejectors are fully closed and have engaged the notches in the DIMM.

Install an Access Module

Steps

1. Align the Access Module with the empty slot and carefully push it into the slot.

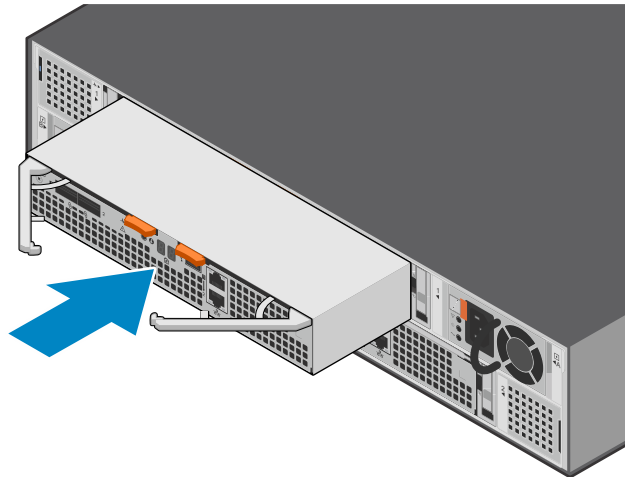


Figure 203. Installing the Access Module

2. Press in on the Access Module latches to lock them into place.

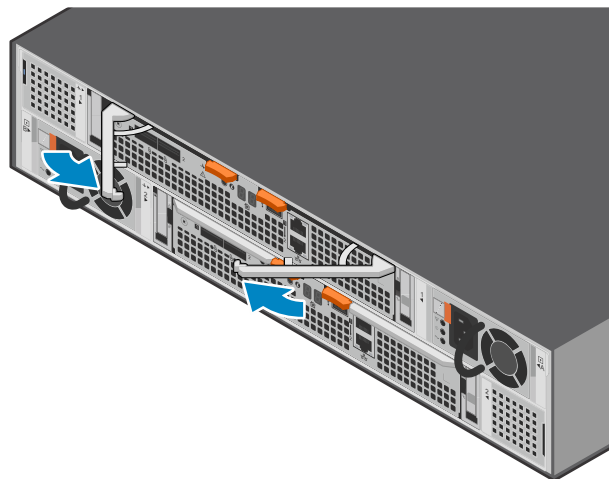


Figure 204. Locking the Access Module into place

3. Connect the cables to the Access Module.

Verify the operation of a replacement DIMM

Steps

1. From PowerStore Manager, select **Hardware**.
2. Select the appliance where you replaced the DIMM.
3. On the **Components** card, under **Internal View**, expand **ExpansionEnclosure**.
4. Select the relevant **DIMM**.

The status of the replacement DIMM should read **Healthy**. If the status is still **Faulted**, wait a few minutes and refresh PowerStore Manager. If the status does not change, ensure that the DIMM is correctly seated, or contact your service provider.

5. You can also verify the operation of a replacement DIMM by using the following commands:

To display the DIMM information:


```
svc_diag list --expansion_hardware --sub_option dimm
```

To display the status of the NVMe expansion enclosure:

```
svc_diag list --expansion_hardware --sub_option status
```

To display the verbose output of all of the hardware in the NVMe expansion enclosure including health and status.

```
svc_diag list --expansion_hardware
```

 **NOTE:** The svc_diag list command takes a few minutes to run.

Return a faulted part

About this task

For US customers, return defective material within five business days. For International customers, return defective material within 10 business days. The materials required to return your defective part are supplied with the good part shipment.

Steps

1. Package the faulted part in the shipping box that contained the replacement part.
2. Ship the failed part to your service provider as described in the instructions that were included with the replacement part.
3. For more information about returning customer-replaceable parts:
 - a. Open PowerStore Manager.
 - b. Click **Settings** on the upper right of the screen.
 - c. Click **General Support**.
 - d. Under **Drives, Power Supplies, and Other Parts**, click **Return Part**.
 - e. If your screen does not show the Return Part link, contact your service provider for instructions.

Safety precautions for handling replaceable units

Review these safety considerations before replacing any parts to avoid damage to your system.

Topics:

- [Handling replaceable units](#)

Handling replaceable units

This section describes the precautions that you must take and the general procedures that you must follow when removing, installing, and storing any replaceable unit.

Avoiding electrostatic discharge (ESD) damage

When replacing or installing hardware units, you can inadvertently damage the sensitive electronic circuits in the equipment by simply touching them.


Electrostatic charge that has accumulated on your body discharges through the circuits. If the air in the work area is very dry, run a humidifier in the work area to help decrease the risk of ESD damage.

Follow these procedures to prevent equipment damage:

- Provide enough room to work on the equipment.
- Clear the work site of any unnecessary materials or materials that naturally build up electrostatic charge, such as foam packaging, foam cups, cellophane wrappers, and similar items.
- Do not remove replacement or upgrade units from their antistatic packaging until you are ready to install them.
- Before you begin service, gather together the ESD kit and all other materials you need.
- Once servicing begins, avoid moving away from the work site; otherwise, you may build up an electrostatic charge.
- Use ESD anti-static gloves or an ESD wristband (with strap). If using an ESD wristband with a strap:
 - Attach the clip of the ESD wristband to the ESD bracket or bare metal on a cabinet/rack or enclosure.
 - Wrap the ESD wristband around your wrist with the metal button against your skin.
 - If a tester is available, test the wristband.
- If an emergency arises and the ESD kit is not available, follow the procedures in Emergency Procedures (without an ESD kit).

Emergency procedures (without an electrostatic discharge kit)

In an emergency when an electrostatic discharge (ESD) kit is not available, use the following precautions to reduce the possibility of an electrostatic discharge by ensuring that your body and the subassembly are at the same electrostatic potential.

 **NOTE:** These precautions are not a substitute for the use of an ESD kit. Follow them only in the event of an emergency.

- Before touching any unit, touch a bare (unpainted) metal surface of the cabinet/rack or enclosure.
- Before removing any unit from its antistatic bag, place one hand firmly on a bare metal surface of the cabinet/rack or enclosure, and at the same time, pick up the unit while it is still sealed in the antistatic bag. At the same time, do not move around the room or touch other furnishings, personnel, or surfaces until you have installed the unit.
- When you remove a unit from the antistatic bag, avoid touching any electronic components and circuits on it.
- If you must move around the room or touch other surfaces before installing a unit, first place the unit back in the antistatic bag. When you are ready again to install the unit, repeat these procedures.

Hardware acclimation times

Units must acclimate to the operating environment before applying power. This requires the unpackaged system or component to reside in the operating environment for up to 16 hours in order to thermally stabilize and prevent condensation.

Table 9. Hardware acclimation times

| Transit/storage environment | | Operating environment temperature | Acclimation time |
|------------------------------|----------------------|--|------------------|
| Temperature | Humidity | - | |
| Nominal 68-72°F (20-22°C) | Nominal 40-55% RH | Nominal 68-72°F (20-22°C) 40-55% RH | 0-1 hour |
| Cold <68°F (20°C) | Dry <30% RH | <86°F (30°C) | 4 hours |
| Cold <68°F (20°C) | Damp ≥30% RH | <86°F (30°C) | 4 hours |
| Hot >72°F (22°C) | Dry <30% RH | <86°F (30°C) | 4 hours |
| Hot >72°F (22°C) | Humid 30-45% RH | <86°F (30°C) | 4 hours |
| | Humid 45-60% RH | <86°F (30°C) | 8 hours |
| | Humid ≥60% RH | <86°F (30°C) | 16 hours |
| Unknown | | <86°F (30°C) | 16 hours |

- If there are signs of condensation after the recommended acclimation time has passed, allow an additional 8 hours to stabilize.
- Systems and components must not experience changes in temperature and humidity that are likely to cause condensation to form on or in that system or component. Do not exceed the shipping and storage temperature gradient of 45°F/hr (25°C/hr).

Removing, installing, or storing replaceable units

Use the following precautions when removing, handling, or storing replaceable units:

CAUTION: Some replaceable units have the majority of their weight in the rear of the component. Ensure that the back end of the replaceable unit is supported while installing or removing it. Dropping a replaceable unit could result in personal injury or damage to the equipment.

NOTE: For a module that must be installed into a slot in an enclosure, examine the rear connectors on the module for any damage before attempting its installation.

CAUTION: A sudden jar, drop, or even a moderate vibration can permanently damage some sensitive replaceable units.

- Do not remove a faulted replaceable unit until you have the replacement available.
- When handling replaceable units, avoid electrostatic discharge (ESD) by wearing ESD anti-static gloves or an ESD wristband with a strap. For additional information, refer to [Avoiding electrostatic discharge \(ESD\) damage](#).
- Avoid touching any exposed electronic components and circuits on the replaceable unit.
- Never use excessive force to remove or install a replaceable unit. Take time to read the instructions carefully.
- Store a replaceable unit in the antistatic bag and the specially designed shipping container in which you received it. Use the antistatic bag and special shipping container when you need to return the replaceable unit.

- Replaceable units must acclimate to the operating environment before applying power. This requires the unpackaged component to reside in the operating environment for up to 16 hours in order to thermally stabilize and prevent condensation. Refer to [Hardware acclimation times](#) to ensure the replaceable unit has thermally stabilized to the operating environment.
- Front bezels should always be attached to ensure EMI compliance. Make sure you reattach the bezel after replacing a component.
- Each I/O module or drive slot should contain a component or filler panel to ensure proper air flow throughout the system.

Unpacking a part

Use these best practices to unpack a part.

Steps

1. Wear ESD gloves or attach an ESD wristband to your wrist and the enclosure in which you are installing the part.
2. Unpack the part and place it on a static-free surface.
3. If the part is a replacement for a faulted part, save the packing material to return the faulted part.

Power control procedures

Learn how to power down and power up the system.

Topics:

- [Power control procedure considerations](#)
- [Power control procedures preview](#)
- [Powering off procedures for PowerStore node](#)
- [Powering on procedures for PowerStore node](#)
- [Rebooting procedures for a PowerStore node](#)
- [Powering off procedures for PowerStore appliances](#)
- [Powering on procedures for PowerStore appliances](#)
- [Powering off procedures for PowerStore cluster](#)
- [Powering on procedures for PowerStore cluster](#)

Power control procedure considerations

Note the following before you get started:

- Powering off a node, appliance, or cluster can take several minutes to complete.
- In a true emergency power off situation, turn the cabinet power switches to the off position to immediately remove power from the all cabinet components.
- Working with hardware may cause electrostatic discharge that could damage your hardware. Before working with any hardware, take precautions around handling replaceable units. See [Safety precautions for handling replaceable units](#).
- If you are relocating or replacing hardware, to help identify associated enclosures when you are ready to cable and power on:
 - Ensure that you make a note of the cabling between enclosures and the appliances. If you used cable labels at the time of initial installation, reconnecting the cables is easier.
 - Ensure that you also record the Dell Service Tag of each enclosure in your cluster.
- A PowerStore X model appliance consists of two ESXi host nodes (node A and node B), running VMware ESXi version 7.0.3, and a controller VM on each that runs the PowerStore operating environment. Powering off a PowerStore X model node or cluster requires operations in the associated VMware vCenter server. See the VMware documentation for more information about the vCenter server commands.
- Nodes in the appliance power on into the same mode they were in before the appliance was powered off. If a node powers on in service mode:
 1. Log in to the appliance from an SSH client.
 2. Run the `svc_rescue_state clear` command to clear the boot mode.
 3. Run the `svc_node reboot` command.
 - For a PowerStore T model node, this command reboots the node itself.
 - For a PowerStore X model node, only the PowerStore controller VM gets rebooted.

Once rebooted, the node returns to normal mode.

NOTE: The `svc_node` service script, with `reboot` and `shutdown` arguments, is typically used in the context of rebooting or powering off a PowerStore T model node.

Use the VMware vCenter server commands to reboot or power off a PowerStore X model node. In PowerStore X model nodes, running the `svc_node` service script with `reboot` or `shutdown` arguments only affects the PowerStore controller VM.

Unless specifically directed by your service provider, for PowerStore X model appliances, use the `svc_node` service script only for the following purposes:

- Go out of service mode (`reboot` argument).
- Power on the node (`power_on` argument).

For more information about the service scripts, see the *PowerStore Service Scripts Guide*.

- If both nodes in an appliance reboot in service mode, always return Node A to normal mode first to avoid management software conflicts. After Node A is operating normally, you can return Node B to normal mode.
- Before powering down an appliance with metro volumes, ensure that the role of the metro volumes on the appliance are all set to non-preferred. Refer to the *Protecting Your Data* guide for details about setting metro volume roles.

Power control procedures preview

CAUTION: Do not power off by pulling cables from the back of the appliance to initiate a shutdown sequence. Use PowerStore Manager or a service script to perform all graceful shutdown operations.

The following table provides a preview of the steps that are required to power off or up the relevant component in your cluster:

Table 10. Power control procedures preview

| Component | Action | PowerStore T model | PowerStore X model |
|-----------|-----------|---|--|
| Node | Power off | Use PowerStore Manager. Or Run a service script. | Use VMware vCenter server. |
| | Power on | <ul style="list-style-type: none"> • If the node was removed from the chassis, reseal the node into the chassis, and reconnect its power cable. • If the node was not removed from the chassis, run a service script. | <ul style="list-style-type: none"> • If the node was removed from the chassis, reseal the node into the chassis, reconnect its power cable, and use VMware vCenter server. • If the node was not removed from the chassis, run a service script and use VMware vCenter server. |
| | Reboot | Use PowerStore Manager. Or Run a service script. | Use VMware vCenter server. |
| Appliance | Power off | Use PowerStore Manager. Or Run a service script. | Use PowerStore Manager. |
| | Power on | If the nodes or expansion enclosures were removed from the chassis, reseal the expansion enclosures and nodes. Reconnect power cables in the right order. | If the nodes or expansion enclosures were removed from the chassis, reseal the expansion enclosures and nodes. Reconnect power cables in the right order, and use VMware vCenter server. |
| | Reboot | Run a service script. | |
| Cluster | Power off | Use PowerStore Manager. Or Run a service script. | Use VMware vCenter server. |
| | Power on | If the nodes or expansion enclosures were removed from the chassis, reseal the expansion enclosures | If the nodes or expansion enclosures were removed from the chassis, reseal the expansion enclosures and nodes. Reconnect power |

Table 10. Power control procedures preview (continued)

| Component | Action | PowerStore T model | PowerStore X model |
|-----------|--------|---|---|
| | | and nodes. Reconnect power cables in the right order. | cables in the right order, and use VMware vCenter server. |

Powering off procedures for PowerStore node

This section includes the following procedures:

- [Power off a PowerStore T model node using PowerStore Manager](#)
- [Power off a PowerStore T model node using a service script](#)
- [Power off a PowerStore X model node](#)

Power off a PowerStore T model node using PowerStore Manager

Prerequisites

Obtain the following information:

- Management IP address of the cluster to log in to PowerStore Manager
- PowerStore Manager user account credentials with administrator privileges

NOTE: Do not power off or reboot a node if the peer node is not operating normally. If the peer node is experiencing any major issues, associated alerts and events appear in PowerStore Manager.

Also, to avoid service interruptions, ensure that there are sufficient and healthy paths from all connected hosts to the peer node.

About this task

Use the following procedure to power off (power down) a PowerStore T model node using PowerStore Manager:

NOTE: If you are unable to access PowerStore Manager, see [Power off a PowerStore T model node using a service script](#).

Steps

1. Under **Hardware**, select the appliance that includes the node you want to power off.
2. On the **Appliance Details** page, select the **Components** card.
3. On the **Components** card, under **Rear View**, expand **Base Enclosure**, and then select the node that you want to power off.
4. Under **More Actions**, select **Power Down**.
5. On the confirmation prompt, enter the service password, and then click **Power Down**.

Next steps

To verify that the node has powered off, check the status of the LEDs in the rear of chassis. Other than the LEDs for the power supply unit, management port, and service port, all LEDs on the node must be OFF. The Unsafe to Remove LED on the active or peer node is ON.

Power off a PowerStore T model node using a service script

Prerequisites

Obtain the following information:

- Management IP address of the appliance that contains the node. In PowerStore Manager, go to **Settings > Networking > Network IPs > Management**. Review the **Management IPs** table to identify the management IP address associated with the appliance.
- Service account credentials

NOTE: Do not power off or reboot a node, if the peer node is not operating normally. If the peer node is experiencing any major issues, associated alerts and events appear in PowerStore Manager.

Also, to avoid service interruptions, ensure that there are sufficient and healthy paths from all connected hosts to the peer node.

About this task

Use the following procedure to power off (power down) a PowerStore T model node using a service script:

Steps

1. Launch an SSH client, and connect to the appliance using the management IP address.

NOTE: External SSH management access must be enabled on the appliance.

2. Enter the username and password that is associated with the service account, and log in.
The login prompt indicates the node that you are logged into. For example, the letter "A" in the prompt [SVC:user@DST5467-A~]\$ indicates that you are logged into node A.
3. Based on the node you are logged into, run one of the following commands:
 - `svc_node shutdown local` to power off the node you are logged into.
 - `svc_node shutdown peer` to power off the peer node.

Next steps

To verify that the node has powered off, check the status of the LEDs in the rear of chassis. Other than the LEDs for the power supply unit, management port, and service port, all LEDs on the node must be OFF. The Unsafe to Remove LED on the active or peer node is ON.

Power off a PowerStore X model node

Prerequisites

Obtain the following information:

- Address of the VMware vCenter server associated with the appliance
- Associated vCenter server account credentials

Ensure that SSH is enabled on the ESXi host and configured to start automatically.

NOTE: Do not power off or reboot a node, if the peer node is not operating normally. If the peer node is experiencing any major issues, associated alerts and events appear in PowerStore Manager. Also, to avoid service interruptions, ensure that there are sufficient and healthy paths from all connected hosts to the peer node.

About this task

Use the following procedure to power off (power down) a PowerStore X model node:

Steps

1. Log in to the associated vCenter server.
2. If there are user VMs that are not configured for automatic migration to the ESXi host of the peer node, migrate them manually.

NOTE: vSphere Remote Office Branch Office Advanced does not support automatic migration.

3. If there are user VMs that cannot migrate to the ESXi host of the peer node, shut them down manually.
4. Place the ESXi host for the PowerStore X model node in Maintenance Mode.
VMware DRS starts to migrate user VMs to the ESXi host of the peer node, and then the PowerStore controller VM is powered off.
5. Monitor and ensure that all relevant user VMs migrate over to the ESXi host of the peer node.

6. Monitor the vCenter server console and confirm that the ESXi host for the PowerStore X model node has entered Maintenance Mode.
7. Using vCenter commands, shut down the ESXi host that is in Maintenance Mode.
8. Monitor the vCenter server console to ensure that the ESXi host has shut down.

Next steps

To verify that the node has powered off, check the status of the LEDs in the rear of chassis. Other than the LEDs for the power supply unit, management port, and service port, all other LEDs on the node must be OFF. The Unsafe to Remove LED on the active or peer node appears ON.

Powering on procedures for PowerStore node

This section includes the following procedures:

- [Power on a PowerStore T model node using a service script](#)
- [Power on a PowerStore T model node by reseating the node](#)
- [Power on a PowerStore X model node using a service script](#)
- [Power on a PowerStore X model node by reseating the node](#)

Power on a PowerStore T model node using a service script

Prerequisites

Obtain the following information:

- Management IP address of the appliance that contains the node. In PowerStore Manager, go to **Settings > Networking > Network IPs > Management**. Review the **Management IPs** table to identify the management IP address associated with the appliance.
- Service account credentials


About this task

Use the following procedure to power on (power up) a PowerStore T model node in scenarios such as:

- You are remote and cannot reseat the node.
- Node was not removed from the chassis.
- Embedded module, I/O module, or 4-port card were replaced.

Steps

1. Launch an SSH client, and connect to the appliance using the management IP address. Since only the peer node is powered on, you are connected directly to the peer node of the appliance.
2. Enter the username and password that is associated with the service account, and log in.
3. Run the following command:
`svc_node power_on`
4. Wait for the node to power on.

 **NOTE:** It may take several minutes for the node to power on.

Power on a PowerStore T model node by reseating the node

About this task

Use the following procedure to power on (power up) a PowerStore T model node when it was removed from the chassis:

Steps

1. Reseat the node into the chassis.
The node will power on automatically.

2. Reconnect the power cable.
3. Wait for the node to complete powering on.

Power on a PowerStore X model node using a service script

Prerequisites

Obtain the following information:

- Address of the VMware vCenter server associated with the appliance
- Associated vCenter server account credentials
- Management IP address of the appliance that contains the node
- Service account credentials

About this task

Use the following procedure to power on (power up) a PowerStore X model node in scenarios such as:

- You are remote and cannot reseal the node.
- Node was not removed from the chassis.
- Embedded module, I/O module, or 4-port card were replaced.

Steps


1. Launch an SSH client, and connect to the appliance using the management IP address. Since only the peer node is powered on, you are connected directly to the peer node of the appliance.
2. Enter the username and password that is associated with the service account, and log in.
3. Run the following command:

```
svc_node power_on
```

This command powers on the ESXi host for the intended node. Wait for the ESXi host to complete powering on.

4. Log in to the associated vCenter server, and exit the ESXi host for the node from Maintenance Mode.

The controller VM powers on automatically.

 **NOTE:** For primary internal M.2 boot module replacements, skip this step. The controller VM powers on automatically once the ESXi host powers up because the ESXi host is not in Maintenance Mode.

5. Monitor the vCenter server console to ensure that the ESXi host and controller VM have returned to normal operation.

Power on a PowerStore X model node by reseating the node


About this task

Use the following procedure to power on (power up) a PowerStore X model node when it was removed from the chassis:

Steps

1. Reseat the node into the chassis.
The node will power on automatically.
2. Reconnect the power cable.
3. Wait for the ESXi host for the node to complete powering on.
4. Log in to the associated vCenter server, and exit the ESXi host for the node from Maintenance Mode.

The controller VM powers on automatically.

 **NOTE:** For primary internal M.2 boot module replacements, skip this step. The controller VM powers on automatically once the ESXi host powers up because the ESXi host is not in the Maintenance Mode.

5. Monitor the vCenter server console to ensure that the ESXi host and controller VM have returned to normal operation.

Rebooting procedures for a PowerStore node

This section includes the following procedures:

- [Reboot a PowerStore T model node using PowerStore Manager](#)
- [Reboot a PowerStore T model node using a service script](#)
- [Reboot a PowerStore X model node](#)

Reboot a PowerStore T model node using PowerStore Manager

Prerequisites

Obtain the following information:

- Management IP address of the cluster to log in to PowerStore Manager
- PowerStore Manager user account credentials with administrator privileges

NOTE: Do not power off or reboot a node, if the peer node is not operating normally. If the peer node is experiencing any major issues, associated alerts and events appear in PowerStore Manager.

Also, to avoid service interruptions, ensure that there are sufficient and healthy paths from all connected hosts to the peer node.

About this task

Use the following procedure to reboot a PowerStore T model node using PowerStore Manager:

Steps

1. Under **Hardware**, select the appliance that includes the node you want to reboot.
2. On the **Appliance Details** page, select the **Components** card.
3. On the **Components** card, under **Rear View**, expand **Base Enclosure**, and then select the node that you want to reboot.
4. Under **More Actions**, select **Reboot**.
5. On the confirmation prompt, select **Confirm you want to reboot the node**, and then click **Reboot**.

Reboot a PowerStore T model node using a service script

Prerequisites

Obtain the following information:

- Management IP address of the appliance that contains the node. In PowerStore Manager, go to **Settings > Networking > Network IPs > Management**. Review the **Management IPs** table to identify the management IP address associated with the appliance.
- Service account credentials

NOTE: Do not power off or reboot a node, if the peer node is not operating normally. If the peer node is experiencing any major issues, associated alerts and events appear in PowerStore Manager.

Also, to avoid service interruptions, ensure that there are sufficient and healthy paths from all connected hosts to the peer node.

About this task

Use the following procedure to reboot a PowerStore T model node using a service script:

Steps

1. Launch an SSH client, and connect to the appliance using the management IP address.
2. Enter the username and password for the service account to log in.
The login prompt indicates the node that you are logged into. For example, the letter "A" in the prompt [SVC:user@FNM12345678910-A~] \$ indicates that you are logged into node A.

3. Based on the node you are logged into, run one of the following commands:

- `svc_node reboot local` to reboot the node you are logged into.
- `svc_node reboot peer` to reboot the peer node.

For more information, see the *PowerStore Service Scripts Guide*.

Reboot a PowerStore X model node

Prerequisites

Obtain the following information:

- Address of the VMware vCenter server associated with the appliance
- Associated vCenter server account credentials

Ensure that SSH is enabled on the ESXi host and configured to start automatically.

NOTE: Do not power off or reboot a node, if the peer node is not operating normally. If the peer node is experiencing any major issues, associated alerts and events appear in PowerStore Manager.

Also, to avoid service interruptions, ensure that there are sufficient and healthy paths from all connected hosts to the peer node.

About this task

Use the following procedure to reboot a PowerStore X model node:

Steps

1. Log in to the associated vCenter server.
 2. If there are user VMs that are not configured for automatic migration to the ESXi host of the peer node, migrate them manually.
- NOTE:** vSphere Remote Office Branch Office Advanced does not support automatic migration.
3. If there are user VMs that cannot migrate to the ESXi host of the peer node, shut them down manually.
 4. Place the ESXi host for the node in Maintenance Mode.
VMware DRS starts to migrate user VMs to the ESXi host of the peer node, and then the PowerStore controller VM is powered off.
 5. Monitor and ensure that all relevant user VMs migrate over to the ESXi host of the peer node.
 6. Once all VMs are migrated over and node goes into Maintenance Mode, reboot the ESXi host using vCenter commands.
 7. Monitor the vCenter server console to ensure that the ESXi host has rebooted, and then exit the ESXi host from Maintenance Mode.
The controller VM powers on automatically.
 8. Monitor the vCenter server console to ensure that the ESXi host and controller VM have returned to normal operation.

Powering off procedures for PowerStore appliances

This section includes the following procedures:

- [Power off a PowerStore T model appliance](#)
- [Power off a PowerStore X model appliance](#)

Power off a PowerStore T model appliance

Prerequisites

- Do not power off the appliance if you are replacing a hardware component. Identify the node that includes the faulted hardware component, and power off only that node. For more information, see [Power off a PowerStore T model node using PowerStore Manager](#).
- Powering off an appliance results in the mapped hosts losing access to the data on the appliance. Before you begin, ensure that you temporarily disconnect host access from all storage resources.
- Obtain the following information:
 - Management IP address of the appliance. In PowerStore Manager, go to **Settings > Networking > Network IPs > Management**. Review the **Management IPs** table to identify the management IP address associated with the appliance.
 - Service account credentials
 - Service tags of the appliance
 - If applicable, service tags of the associated expansion enclosures

About this task

Use the following procedure to power off a single PowerStore T model appliance.

To power off all of the appliances in a cluster, see [Power off a PowerStore T model cluster](#):

Steps

1. Log in to PowerStore Manager.
2. Determine the primary appliance by going to **Settings > Cluster > Properties**.
3. If the appliance you are shutting down is the primary appliance:
 - a. Launch an SSH client, and connect to the appliance using the management IP address.
 - b. Enter the username and password that is associated with the service account, and log in.
 - c. Run the following command to specify which appliance you want to become the new primary appliance:

```
svc_cluster_management MoveMasterAppliance <ID number of new primary appliance>
```

4. In PowerStore Manager, under **Hardware**, select the appliance that you want to power off.
5. Under **More Actions**, select **Power Down**.
The **Validation** window opens.
6. Review any errors, warnings, and recommendations. Once the appliance passes all of the validation checks, click **Next**.
The **Active Objects** window opens.
7. Review the list of objects on the appliance that had I/O activity during the last five minutes.
8. Click **Next**.
The **Confirm** window opens.
9. Enter the service password, and click **Power Down**.
10. Check the status of the LEDs in the rear of chassis to verify that the appliance has powered off. Other than the LEDs for the power supply unit, management port, and service port, all other LEDs on the appliance must be OFF.
11. Wait five minutes, and then disconnect the power cables from the base enclosure.
12. Disconnect the power cables from any associated expansion enclosures.

Power off a PowerStore X model appliance

Prerequisites

- Do not power off the appliance if you are replacing a hardware component. Identify the node that includes the faulted hardware component, and power off only that node. For more information, see [Power off a PowerStore X model node](#).
- Ensure that SSH is enabled on the ESXi host and configured to start automatically.
- Powering off an appliance results in the mapped hosts losing access to the data on the appliance. Before you begin, ensure that you temporarily disconnect host access from all storage resources.
- Obtain the following information:
 - Address of the VMware vCenter server associated with the appliance

- Associated vCenter server account credentials
- Service tags of the appliance and, if applicable, the associated expansion enclosures

About this task

Use the following procedure to power off a single PowerStore X model appliance in a multi-appliance cluster.

To power off all of the appliances in a cluster, see [Power off a PowerStore X model cluster](#).

Steps

1. Log in to PowerStore Manager.
2. Use PowerStore Manager to identify VMs that utilize the iSCSI or FC store of the appliance.
 - a. Select **Storage > Storage Containers**.
 - b. Select the first Storage Container.
 - c. Select **Virtual Volumes**.
 - d. Select the **Show Filters** icon.
 - e. Select **Add Filter > Appliance**.
 - f. Identify the VMs that are using storage on the appliance and the VMs that are running on the appliance.
 - g. Repeat these steps for each Storage Container.
3. In vCenter, power off the VMs identified in the previous step as well as any VMs using VMFS or volume on the appliance.

NOTE: If you migrate compute VM storage objects to another appliance, you can migrate the compute VM to that appliance instead of powering it off.

4. In PowerStore Manager, determine the primary appliance by going to **Settings > Cluster > Properties**.
5. If the appliance you are shutting down is the primary appliance:
 - a. Launch an SSH client, and connect to the appliance using the management IP address.
 - b. Enter the username and password that is associated with the service account, and log in.
 - c. Run the following command to specify which appliance you want to become the new primary appliance:

```
svc_cluster_management MoveMasterAppliance <ID number of new primary appliance>
```

6. In PowerStore Manager, under **Hardware**, select the appliance that you want to power off.
7. Under **More Actions**, select **Power Down**.
The **Validation** window opens.
8. Review any errors, warnings, and recommendations. Once the appliance passes all of the validation checks, click **Next**.
The **Active Objects** window opens.
9. Review the list of objects on the appliance that had I/O activity during the last five minutes.
10. Click **Next**.
The **Confirm** window opens.
11. Enter the service password, and then click **Power Down**.
12. Wait until PowerStore Manager confirms that the appliance is offline to ensure that the ESXi hosts have shut down.
13. Monitor the vCenter server console to ensure that the ESXi hosts have shut down.
14. Wait for 5 minutes, and then, if required, disconnect the power cables from the base enclosure and expansion enclosures.

Powering on procedures for PowerStore appliances

This section includes the following procedures:

- [Power on a PowerStore T model appliance](#)
- [Power on a PowerStore X model appliance](#)

Power on a PowerStore T model appliance

About this task

Use the following procedure to power on a PowerStore T model appliance:

Steps

1. If nodes were removed, reseal the nodes into the base enclosure chassis.
2. If applicable, ensure that expansion enclosures are also reseated into the cabinet.
3. If applicable, reconnect the power cables to each associated expansion enclosure in an ascending order, such as the following:
 - Expansion enclosure 0
 - Expansion enclosure 1
 - Expansion enclosure 2

The power status LEDs on each expansion enclosure turns on when the power cable is connected.

4. Reconnect the power cables to node A first, and then node B.
The Node Power LEDs on both nodes turn on when the power cable is connected.

Power on a PowerStore X model appliance

Prerequisites

Obtain the following information:

- Address of the VMware vCenter server associated with the appliance
- Associated vCenter server account credentials
- Service tags of the appliance and, if applicable, the associated expansion enclosures
- Management IP address of the cluster and the service account credentials

About this task


Use the following procedure to power on a PowerStore X model appliance in a multi appliance cluster that already has one appliance powered on:

Steps

1. If nodes were removed, reseal the nodes back into the base enclosure chassis.
2. If applicable, ensure that expansion enclosures are also reseated into the cabinet.
3. If applicable, reconnect the power cables to each associated expansion enclosure in an ascending order, such as the following:
 - Expansion enclosure 0
 - Expansion enclosure 1
 - Expansion enclosure 2

The power status LEDs on each expansion enclosure turns on when the power cable is connected.

4. Reconnect the power cables to node A first, and then node B.
The Node Power LEDs on both nodes turn on when a power cable is connected to either node. ESXi hosts for both the nodes power on in Maintenance Mode.
5. The ESXi hosts for both nodes exit Maintenance Mode automatically, and the controller VMs for both nodes automatically power on. Wait for the task to complete from vCenter.

 **NOTE:** The PowerStore controller VM can take around 10-15 minutes to power on.

6. Power on user VMs on the ESXi hosts for both nodes.
7. Verify that the PowerStore X model cluster is operating normally in both vCenter server and PowerStore Manager.

Powering off procedures for PowerStore cluster

This section includes the following procedures:

- [Power off a PowerStore T model cluster](#)
- [Power off a PowerStore X model cluster](#)

Power off a PowerStore T model cluster

Prerequisites

- Powering off a cluster results in the mapped hosts losing access to the data on the cluster. Before you begin, ensure that you temporarily disconnect host access from all storage resources.
- Check if any VMs are using the storage from the cluster. It is recommended to power off the VMs before powering off the cluster.
- When the cluster is powered off, you have no access to the GUI, API, or CLI interfaces.

Print the power on instructions to ensure that you have the information you require to power on the cluster in a specific order. You can also find these instructions on <https://www.dell.com/powerstoredocs>.

- Obtain the following information:
 - Management IP address of the cluster
 - Service account credentials
 - Site ID
 - Service tags of the appliances
 - If applicable, service tags of the associated expansion enclosures

About this task

Use the following procedure to power off (power down) a PowerStore T model cluster:

Steps

1. Use one of the following ways to issue a power off command:
 - Using PowerStore Manager:
 - a. Select the **Settings** icon, and then select **Power Down** in the **Cluster** section.
 - b. On the confirmation prompt, enter the service password, and then click **Power Down**.
 - Using service script:
 - a. Launch an SSH client, and connect to the appliance using the management IP address.
 - b. Enter the username and password that is associated with the service account, and log in.
 - c. Run the following command to power off the appliance:

```
svc_cluster shutdown
```
2. Check the status of the process by looking at the Node Power LEDs. The power off process is complete when the Node Power LEDs for all nodes in the cluster are off.
3. After confirming that the cluster has shut down, disconnect the power cables from both nodes in one of the base enclosures in the cluster, if required. Wait a few seconds and confirm that all remaining LEDs have turned off.
4. Disconnect the power cables from each of the associated expansion enclosures to power them down, if required.
5. If your cluster has more than one appliance, repeat steps 3 and 4 to disconnect power from the remaining appliances in the cluster.

Power off a PowerStore X model cluster

Use the following procedure to power off a single appliance cluster or a multi appliance cluster. Powering off a PowerStore X model cluster is not supported via PowerStore Manager.

Prerequisites

- Powering off a cluster results in the mapped hosts losing access to the data on the cluster. Before you begin, ensure that you temporarily disconnect host access from all storage resources.
- When the cluster is powered off, you have no access to the PowerStore Manager, API, or CLI interfaces.

Print the power on instructions to ensure that you have the information you require to power on the cluster in a specific order. You can also find these instructions on <https://www.dell.com/powerstoredocs>.

- Obtain the following information:
 - Site ID
 - Service tags of the appliances, and if applicable, the associated expansion enclosures

- Address of the VMware vCenter server associated with the appliance
- Associated vCenter server account credentials

About this task

For multi appliance clusters, the following table describes which controller VMs must be manually powered off. The remaining controller VMs are automatically shut down when you place the ESXi host in Maintenance Mode.

Table 11. Powering off controller VMs

| Cluster size | Appliance 1 | Appliance 2 | Appliance 3 | Appliance 4 |
|--------------|----------------|------------------|------------------|------------------|
| 1 | Auto power off | | | |
| 2 | Auto power off | Manual power off | | |
| 3 | Auto power off | Auto power off | Manual power off | |
| 4 | Auto power off | Auto power off | Manual power off | Manual power off |

Steps

1. Log in to the associated vCenter server.
2. Shut down all user VMs except for the PowerStore X model controller VMs and vCLS VMs on internal ESXi hosts.
3. If your environment has external compute servers that access the PowerStore X model cluster storage using iSCSI or FC, shut down those VMs.
4. From vSphere, place the secondary ESXi host in Maintenance Mode.
5. Wait for the controller VM to power off. If the controller VM does not power off after 5 minutes, manually power off the VM from vSphere.
When the PowerStore X model controller is powered off, ESXi finishes entering Maintenance Mode.
6. Place the primary ESXi host in Maintenance Mode.
7. Wait for the controller VM to power off. If the controller VM does not power off after 5 minutes, manually power off the VM from vSphere.
When the PowerStore X model controller is powered off, ESXi finishes entering Maintenance Mode.
8. After placing the internal ESXi hosts into Maintenance Mode, use vCenter to shut down each host.

Powering on procedures for PowerStore cluster

This section includes the following procedures:

- [Power on a PowerStore T model cluster](#)
- [Power on a PowerStore X model cluster](#)

Power on a PowerStore T model cluster

About this task

Use the following procedure to power on a PowerStore T model cluster:

Steps

1. If nodes were removed, reseal the nodes into the relevant base enclosure chassis.
2. If applicable, for each appliance in the cluster, ensure that expansion enclosures are also reseated into the cabinet.
3. If applicable, for each appliance in the cluster, reconnect the power cables to each expansion enclosure in the following order:
 - Expansion enclosure 0
 - Expansion enclosure 1
 - Expansion enclosure 2

The power status LEDs on each expansion enclosure turns on when the power cable is connected.

4. For each appliance, reconnect the power cables to node A first, and then node B.
The Node Power LED on each node turns on when the power cable is connected.

Power on a PowerStore X model cluster

Use the following procedure to power on a single appliance cluster or a multi appliance cluster.

Prerequisites

Obtain the following information:

- Address of the VMware vCenter server associated with the appliance.
- Associated vCenter server account credentials
- Service tags of the appliance and, if applicable, the associated expansion enclosures.
- Management IP address of the cluster and the service account credentials

About this task

Use the following procedure to power on a PowerStore X model cluster.

For multi appliance clusters, the following table describes which controller VMs must be manually powered on. The remaining controller VMs are automatically powered on when you remove the ESXi host from Maintenance Mode.

Table 12. Powering on controller VMs

| Cluster size | Appliance 1 | | Appliance 2 | | Appliance 3 | | Appliance 4 | |
|--------------|-------------------|-------------------|-------------------|-------------------|-------------------|-----------------|-----------------|-----------------|
| | Controller VM A | Controller VM B | Controller VM A | Controller VM B | Controller VM A | Controller VM B | Controller VM A | Controller VM B |
| 1 | 1 Manual power on | 2 Auto power on | | | | | | |
| 2 | 1 Manual power on | 2 Manual power on | 3 Manual power on | 4 Auto power on | | | | |
| 3 | 1 Manual power on | 2 Manual power on | 3 Manual power on | 4 Auto power on | 5 Auto power on | 6 Auto power on | | |
| 4 | 1 Manual power on | 2 Manual power on | 3 Manual power on | 4 Manual power on | 5 Manual power on | 6 Auto power on | 7 Auto power on | 8 Auto power on |

Steps

1. If nodes were removed, reseal the nodes back into the base enclosure chassis.
2. If applicable, ensure that expansion enclosures are also reseated into the cabinet.
3. If applicable, reconnect the power cables to each expansion enclosure in the following order:
 - Expansion enclosure 0
 - Expansion enclosure 1
 - Expansion enclosure 2

The power status LEDs on each expansion enclosure turns on when the power cable is connected.

4. Reconnect the power cables to node A first, and then node B.
The Node Power LEDs on both nodes turn on when a power cable is connected to either node. ESXi hosts for both the nodes power on in Maintenance Mode.
5. From vSphere, exit the ESXi host for node A from Maintenance Mode.
6. If the controller VM does not power on in five minutes, manually power on the VM from vSphere.
7. Exit the ESXi host for node B from Maintenance Mode.
8. If the controller VM does not power on in five minutes, manually power on the VM from vSphere.
9. Repeat steps 4 through 8 for each appliance.
10. Power on user VMs on the ESXi hosts.
11. Verify that the PowerStore X model cluster is operating normally in both vCenter server and PowerStore Manager.

Transferring the internal battery backup module

Learn how to transfer the internal battery backup module from a faulted node to a replacement node.

NOTE: Review the information in [Safety precautions for handling replaceable units](#) before handling replaceable parts.

Topics:

- [Remove the internal battery backup module](#)
- [Install the internal battery backup module](#)

Remove the internal battery backup module

Steps

1. Lift the gray padding away from the internal battery backup module.
2. Remove the clip around the internal battery backup module by pulling the tab on the left side.
3. Disconnect the internal battery backup module cable from the motherboard.

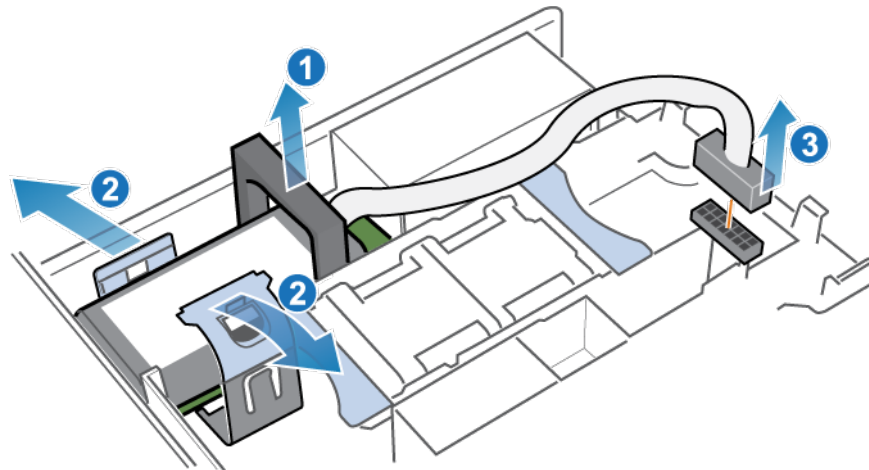


Figure 205. Releasing the BBU

4. Lift the internal battery backup module away from the motherboard.

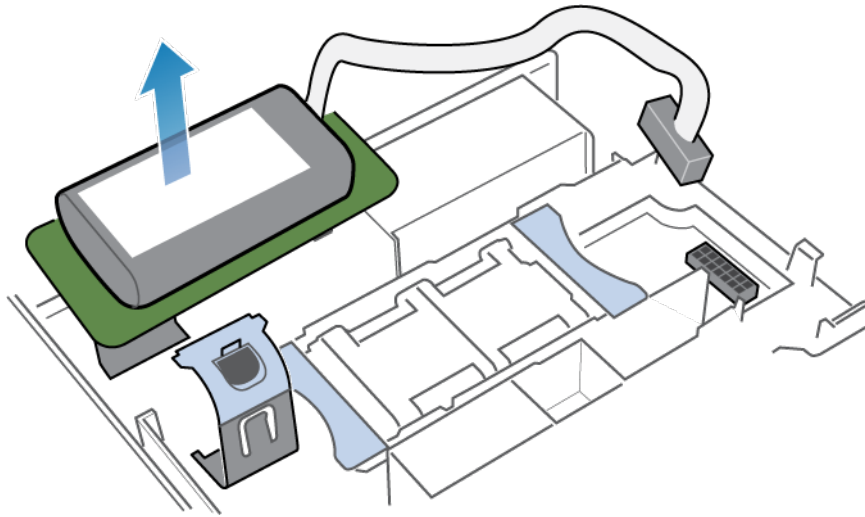


Figure 206. Removing the BBU from the motherboard

Install the internal battery backup module

Install the internal battery backup module into the replacement node.

Steps

1. Place the internal battery backup module into position on the motherboard.

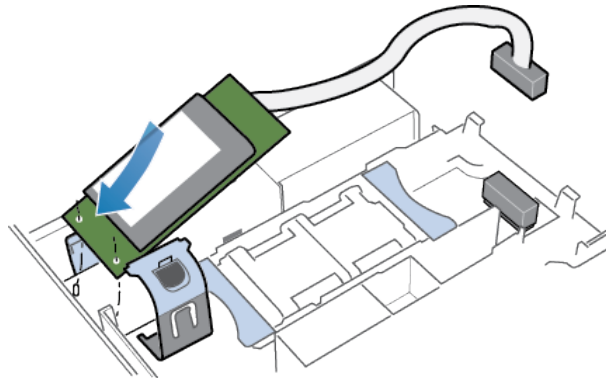


Figure 207. Placing the internal battery backup module into position

2. Close the clip around the center of the internal battery backup module until it clicks into place.
3. Replace the gray padding around the top of the internal battery backup module.
4. Connect the internal battery backup module cable to the motherboard.

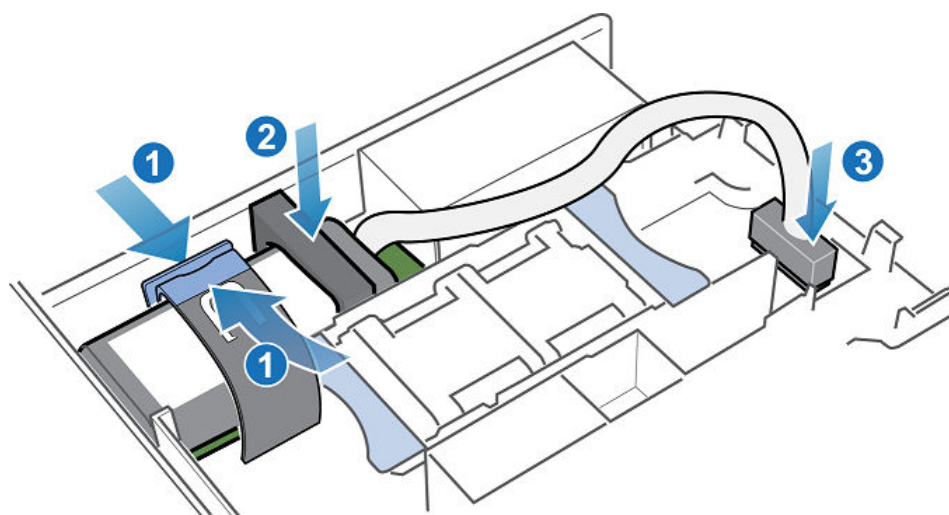


Figure 208. Connecting the internal battery backup module cable to the motherboard

Data collection

Learn how to collect support materials to help troubleshoot the appliances in your system.

Topics:

- [Support materials collection](#)
- [Collect support materials](#)

Support materials collection


You can collect support materials to help troubleshoot the appliances in your system.

Depending on the option you choose, support materials can include system logs, configuration details, and other diagnostic information. Use this information to analyze performance issues, or send it to your service provider so they can diagnose and help you resolve the issues. This process does not collect user data.

You can collect support materials for one or more appliances. When you start a collection, data is always collected at the appliance level. For example, if you request a collection for a volume, the system collects support materials for the appliance that contains the volume. If you request a collection for multiple volumes, the system collects support materials for all appliances that contain the volumes.


You can set a timeframe for collecting support materials. Setting a timeframe can result in smaller and more relevant data collection which is easier to analyze. You can either set a predefined timeframe or set a custom timeframe that suits your needs.

You can also include additional information in the support materials collection from **Advanced collection options**. Collecting additional information can take longer than the default support materials collection, and the size of the resulting data collection is larger. Select this option if your service provider requests it. By default the support materials collection uses the *essentials* profile. Use the `svc_dc` service script to collect support materials for other profiles. See the PowerStore Service Scripts Guide for more information about the `svc_dc` service script and the available profiles.

 **NOTE:** The system can run only one collection job at a time.

You can perform the following actions on a collection of support materials:

- View information about existing collections.
- Upload a collection to support, if remote support through Secure Remote Services is enabled.
- Download a collection to a local client.
- Delete a collection.


 **NOTE:** Some of these operations might not be available if the cluster is operating in a degraded state.


Collect support materials

Steps

1. Select the **Settings** icon, and then select **Gather Support Materials** in the **Support** section.
2. Click **Gather Support Materials**.
3. Type a description of the collection in the **Description** field.
4. Select the timeframe for the data collection.

You can select one of the available options from the **Collection Timeframe** drop-down menu, or select **Custom** and set a timeframe.

 **NOTE:** If you select **Custom** as the timeframe for the data collection, the estimated finish time for the data collection is displayed in the **Collection Timeframe Finish** column of the **Support Materials Library** table.

5. Select the type of support data to collect from the **Object type** drop-down menu.
6. In the **Objects to collect data for:** area, select the check boxes of the appliances from which to collect support data.
7. To send the data collection to support when the job completes, select the **Send materials to Support when finished** check box.
 **NOTE:** This option is available only when Support Connectivity is enabled on the system. You can also send the data collection to support from the **Gather Support Materials** page after the job is completed.
8. Click **Start**.
The data collection is initiated, and the new job appears in the **Support Materials Library** table. You can click the job entry to view its details and progress.

Results

When the job is completed, the job information is updated in the **Support Materials Library** table.

Next steps

After the job is finished, you can download the data collection, send the data collection to support, or delete the data collection.

Support Notifications

Learn how to disable and enable support notifications.

Topics:


- [Disable support notifications](#)
- [Enable support notifications](#)

Disable support notifications

Disable support notifications before performing procedures, such as a software upgrade or support procedure, which can power off or reboot a node in an appliance.

Steps

1. Select the **Settings** icon, and then select **Disable Support Notifications** in the **Support** section.
2. Select the appliance for which you want to disable support notifications and click **Modify**.
3. Select the **Enable Maintenance Mode** check box.
4. In the **Maintenance Window Duration (in hours)** field, type the number of hours to disable support notifications.

 **NOTE:** Specify a time period that is longer than the time it takes to complete the procedure.

5. Click **Apply**.

Results

When support notifications are disabled for an appliance, the **Maintenance Mode** column shows **Enabled**, and **End Time (Cluster Time)** shows the date and time when support notifications are reenabled for the appliance.

Enable support notifications

Enable support notifications after performing a procedure, such as a software upgrade or support procedure.

Steps

1. Select the **Settings** icon, and then select **Disable Support Notifications** in the **Support** section.
2. Select the appliance for which you want to enable support notifications and click **Modify**.
3. Clear the **Enable Maintenance Mode** checkbox.
4. Click **Apply**.

Results

When support notifications are enabled for an appliance, the **Maintenance Mode** column shows **Disabled**.

Add appliances to the cluster

Learn how to add appliances to the cluster.

Topics:

- [Add appliances to the cluster](#)

Add appliances to the cluster

Prerequisites

- You can only have up to four appliances in a cluster.
- Ensure that the cluster is functional and in a healthy state. If any other appliance in the cluster is not operational, you may not be able to add an appliance.
- Ensure that the appliances you are adding are in an unconfigured, original factory settings state.
- Obtain the Service Tag of the appliances you want to add.
- Ensure that you have sufficient number of unused IP addresses for each appliance. For each appliance you want to add, have at least four IP addresses for the management network and three IP addresses for the storage network. Work with your network administrator to provision and obtain more IP addresses, if necessary. To review, or add more IP addresses, select the **Settings** icon, and then select **Network IPs** in the **Networking** section.
- You cannot have PowerStore T model and PowerStore X model appliances in the same cluster.

NOTE: The cluster and appliance you are adding must be running the same PowerStoreOS version before the appliance can be added to the cluster. Clusters on PowerStoreOS 3.0.x and later automatically detect if there is a mismatch between OS versions and provide the option to synchronize during the **Add Appliance** wizard.

For clusters on PowerStoreOS 2.x and earlier where the appliance is running a later version of the PowerStoreOS than the cluster, upgrade the cluster prior to adding the new appliance.

For clusters on 2.x and earlier where the cluster is running a later version of the PowerStoreOS than the appliance, the appliance should be installed in the same rack and use the same switches as the existing cluster, and configured into its own separate cluster. Then, upgrade the OS of the separate cluster. Log in to Support and refer to KB article 000133192 (PowerStore Manager prevents "Add an appliance to an existing cluster" operation from succeeding if code version of appliance to be added does not match...) for more information.

About this task

To add appliances to the cluster:

Steps

1. Under **Hardware**, click **Add** in the **Appliances** tab.
2. Follow the prompts in the **Add Appliance** wizard to select and add the appliances to your cluster.
 - NOTE:** When this process is running, do not run commands, such as adding external hosts or changing CHAP configurations, which can change the state of the cluster.
 - NOTE:** Any operations started while the Add Appliance process is running will not run until the Add Appliance operation is complete.
3. If the PowerStoreOS version of the cluster does not match the PowerStoreOS of the appliance being added, you will be prompted to synchronize software versions before adding the appliance. Click **Synchronize** to initiate the version synchronization.
When the synchronization is complete, you will be returned to the **Add Appliance** wizard. Proceed with the rest of the steps of the wizard.

Remove appliances from the cluster

Learn how to remove appliances from the cluster.

Topics:

- [Removing an appliance from a PowerStore cluster](#)
- [Remove an appliance from a PowerStore T model cluster](#)
- [Remove an appliance from a PowerStore X model cluster](#)
- [Migrate storage objects from an appliance](#)

Removing an appliance from a PowerStore cluster

The steps for removing an appliance from a cluster are different depending on the type of PowerStore system.

- [Remove an appliance from a PowerStore T model cluster](#)
- [Remove an appliance from a PowerStore X model cluster](#)


Remove an appliance from a PowerStore T model cluster

Prerequisites

- Identify the Service Tag of the appliance that you want to remove. For details, see the Hardware Information Guide for PowerStore 1000, 1200, 3000, 3200, 5000, 5200, 7000, 9000, and 9200 or Hardware Information Guide for PowerStore 500T Model .
- If the appliance that you want to remove is not functional, contact your service provider for assistance with removing the appliance.
- Ensure the appliance that you want to is not running SDNAS.
- Ensure that the appliance is not the only appliance in the cluster.
- Stop and remove all running or scheduled import, migration, or replication jobs on the appliance to remove, and remove all remote systems that are associated with the appliance.
- Before starting the appliance removal process, migrate storage objects from the appliance that you want to remove to another appliance in the cluster. To migrate storage objects, see [Migrate storage objects from an appliance](#).

About this task

A factory reset is performed on an appliance when it is removed from a cluster.

 **NOTE:** When the appliance removal process is running, do not run commands that can change the state of the cluster, such as adding external hosts or changing CHAP configurations.


Steps

1. If the appliance you want to remove is the primary appliance in the cluster, run the following command to move the primary appliance to another appliance:

```
svc_cluster_management MovePrimaryAppliance appliance_id
```

2. Disable support notifications on the PowerStore cluster as described in [Disable support notifications](#).
3. Under **Hardware**, select **Appliances**, and find the appliance with the Service Tag identified in Prerequisites.
4. Run the `svc_appliance_provisioning` script to disable resource balancing and prevent objects from being automatically created or placed on the appliance that you want to remove.

- a. Open an SSH client, and connect to the management IP address of the appliance to remove.

 **NOTE:** External SSH management access must be enabled on the appliance.

- b. Type the username and password of the service account to log in to the appliance.
- c. Run the following command to display the names of the appliances in the cluster:

```
svc_appliance_provisioning list
```

- d. Run the following command, where *appliance_name* is the name of the appliance that you want to remove:

```
svc_appliance_provisioning disable appliance_name
```

- e. Run the following command to ensure that the provision status of the appliance is disabled:


```
svc_appliance_provisioning list
```

5. Ensure that Node A is the primary node on the appliance to be removed using PowerStore Manager.
The factory reset of the appliance cannot occur unless Node A is the primary node of the appliance .
 - a. Under **Hardware**, select **Appliances**, and select the appliance that you want to remove.
 - b. Select the **Components** card, select the **Rear View** tab.
 - c. Ensure that Node A is the primary node on the appliance.
If the Node B is the primary node of the appliance, select Node B in the **Rear View** tab, and select **More Actions > Reboot** to reboot Node B and make Node A the primary node.
6. If there still are storage objects on the appliance to be removed, use PowerStore Manager to migrate the storage objects to another appliance in the cluster or remove the storage objects from the appliance.
7. Remove the appliance from the PowerStore T model cluster in PowerStore Manager.

 **NOTE:** The system resets the appliance to original factory settings and powers it off during the removal process.

- a. Under **Hardware**, select **Appliances**.
- b. Select the check box of the appliance to remove.
- c. Click **Remove**.
The **Removing the Appliance** dialog box is displayed.
- d. Click **Remove**.

The appliance is removed from the cluster and reset to its factory default settings. The appliance can be rediscovered and deployed to a new or existing cluster after it is reset to factory default settings.

 **NOTE:** The factory reset of the appliance can take up to 1.5 hours to complete.

Remove an appliance from a PowerStore X model cluster

Prerequisites

- Identify the Service Tag of the appliance that you want to remove. For details, see the Hardware Information Guide for PowerStore 1000, 1200, 3000, 3200, 5000, 5200, 7000, 9000, and 9200.
- If the appliance that you want to remove is not functional, contact your service provider for assistance with removing the appliance.
- Ensure the appliance that you want to remove is not the primary appliance in the cluster or the only appliance in the cluster.
- Stop and remove all running or scheduled import, migration, or replication jobs on the appliance to remove, and remove all remote systems that are associated with the appliance.
- Before starting the appliance removal process, migrate storage objects from the appliance that you want to remove to another appliance in the cluster. To migrate storage objects, see [Migrate storage objects from an appliance](#).

About this task

A factory reset is performed on an appliance when it is removed from a cluster.

NOTE: When the appliance removal process is running, do not run commands that can change the state of the cluster, such as adding external hosts or changing CHAP configurations.

Steps

1. Under **Hardware**, select **Appliances**, and identify the IP address of the appliance with the Service Tag identified in Prerequisites.
2. Under **Settings**, select **Network IPs**, identify the Host Node IPs of the appliance.
3. If the appliance to be removed has VMFS volumes with internal host mappings, migrate the VMFS volumes to another appliance:
 - a. Identify the related compute VMs and VMFS volumes on the appliance.
 - b. Ensure that internal host mappings are defined between the VMFS volumes and the hosts on the destination appliance. If these internal host mappings do not exist, create the host mappings before continuing.
 - c. Manually migrate the VMFS volumes from to the destination appliance using PowerStore Manager.
 - d. vMotion the related compute VMs to the destination appliance.

NOTE: Do not vMotion the vCLS VMs.

 - e. Unmap the VMFS volumes from the appliance to be removed using the PowerCLI.
4. Migrate the virtual volumes for vCLS VMs on the appliance to another appliance in the cluster using PowerStore Manager.
 - a. Under **Storage**, select **Storage Containers**, select the storage container that contains the virtual volumes for the vCLS VMs, and select **Virtual Volumes**.
 - b. Select **Show/Hide Columns** and select the **Appliance** check box to display the appliances on which the virtual volumes are located.
 - c. Identify the virtual volumes for the vCLS VMs that must be migrated.
 - d. Select a virtual volume to migrate and select **Migrate**.

NOTE: You can only migrate one virtual volume at a time.

 - e. Select the appliance to which to migrate the virtual volume and click **Next**.
 - f. Repeat the previous steps for each virtual volume in the storage container.
5. Disable support notifications on the PowerStore cluster as described in [Disable support notifications](#).
6. Run the `svc_appliance_provisioning` script to disable resource balancing and prevent objects from being automatically created or placed on the appliance that you want to remove.
 - a. Open an SSH client, and connect to the management IP address of the appliance to remove.

NOTE: External SSH management access must be enabled on the appliance.

 - b. Type the username and password of the service account to log in to the appliance.
 - c. Run the following command to display the names of the appliances in the cluster:

```
svc_appliance_provisioning list
```

 - d. Run the following command, where *appliance_name* is the name of the appliance that you want to remove:

```
svc_appliance_provisioning disable appliance_name
```

 - e. Run the following command to ensure that the provision status of the appliance is disabled:

```
svc_appliance_provisioning list
```
7. Identify the primary and secondary nodes of the appliance in PowerStore Manager.
 - a. Under **Hardware**, select **Appliances**, and select the appliance to remove.
 - b. On the Appliance Details page, select the **Components** card and select the **Internal View** tab.

The node with the text **(primary)** is the primary node of the appliance, and the node without the text **(primary)** is the secondary node. The names of the nodes are Node A and Node B.
 - c. Under **Settings**, select **Network IPs** and select the **Management** tab.
 - d. Record the **Host Node** IP addresses of Node A and Node B on the appliance.


- If the primary node on the appliance is Node A, the IP address of Node A is for the primary node and the IP address of Node B is for the secondary node.
- If the primary node on the appliance is Node B, the IP address of Node B is for the primary node and the IP address of Node A is for the secondary node.

8. Perform the following steps on the vCenter Server to exclude the ESXi host on the secondary node of the appliance from the ESXi cluster:

- a. If there are user VMs on the secondary node of the appliance, vMotion the user VMs to another appliance in the cluster.

 **CAUTION:** Do not vMotion the vCLS VMs.

- b. Place the ESXi host on the secondary node of the appliance into maintenance mode.
The controller VM is automatically powered off, and the vCLS VMs on the ESXi host are automatically migrated to another host.

 **NOTE:** If the ESXi host is unable to enter maintenance mode because of the vCLS VMs, manually power off the vCLS VMs to force the VMs to migrate to another host.

- c. Create a folder of the type **New Host and Cluster Folder** on the Datacenter.
d. Move the ESXi host on the secondary node of the appliance to the new folder on the Datacenter.
e. Take the ESXi host out of maintenance mode and power on the controller VM.

9. Wait 15 minutes, and then ensure that the secondary node is up and running in PowerStore Manager.

Ensure that there are no active alerts on the secondary node:

- a. Under **Hardware**, select **Appliances**, and select the appliance that you want to remove.
b. Select the **Components** card, select the **Rear View** tab.
c. Expand **BaseEnclosure**, select the secondary node, and ensure that there are no active alerts.

In addition, go to **Monitoring > Events** in PowerStore Manager and ensure that the following event is displayed for the secondary node:


```
Event: NODE_IO_SERVICE_LEVEL.  
Event text: IO service level of the node has changed to ready
```

10. Perform the following steps on the vCenter Server to exclude the ESXi host on the primary node of the appliance from the ESXi cluster:

- a. If there are user VMs on the primary node of the appliance, vMotion the user VMs to another appliance in the cluster.

 **CAUTION:** Do not vMotion the vCLS VMs.

- b. Place the ESXi host on the primary node of the appliance into maintenance mode.
The controller VM is automatically powered off, and the vCLS VMs on the ESXi host are automatically migrated to another host.

 **NOTE:** If the ESXi host is unable to enter maintenance mode because of the vCLS VMs, manually power off the vCLS VMs to force the VMs to migrate to another host.

- c. Move the ESXi host on the secondary node of the appliance to the new folder on the Datacenter, which was created in step 9.
d. Take the ESXi host out of maintenance mode and power on the controller VM.

11. Wait 15 minutes, and then ensure that the primary node is up and running in PowerStore Manager.

Ensure that there are no active alerts on the primary node in PowerStore Manager:

- a. Under **Hardware**, select **Appliances**, and select the appliance that you want to remove.
b. Select the **Components** card, select the **Rear View** tab.
c. Expand **BaseEnclosure**, select the primary node, and ensure that there are no active alerts.

In addition, go to **Monitoring > Events** in PowerStore Manager and ensure that the following event is displayed for the primary node:

```
Event: NODE_IO_SERVICE_LEVEL.  
Event text: IO service level of the node has changed to ready
```

12. Ensure that Node A is the primary node on the appliance to be removed.

The factory reset of the appliance cannot occur unless Node A is the primary node of the appliance.

- a. Under **Hardware**, select **Appliances**, and select the appliance that you want to remove.
b. Select the **Components** card, select the **Internal View** tab.

- c. Ensure that Node A is the primary node on the appliance.

If the Node B is the primary node of the appliance, log in to the appliance using an SSH client and run the `svc_node reboot local` command to reboot Node B and make Node A the primary node.

13. If there still are storage objects on the appliance to be removed, use PowerStore Manager to migrate the storage objects to another appliance in the cluster or remove the storage objects from the appliance.
14. Remove the appliance from the PowerStore X model cluster using PowerStore Manager.

 **NOTE:** The system resets the appliance to original factory settings and powers it off during the removal process.

- a. Under **Hardware**, select **Appliances**.
- b. Select the check box of the appliance to remove.
- c. Click **Remove**.

The **Removing the Appliance** dialog box is displayed.

- d. Click **Remove**.

The appliance is removed from the cluster and reset to its factory default settings. The appliance can be rediscovered and deployed to a new or existing cluster after it is reset to factory default settings.

 **NOTE:** The factory reset of the appliance can take up to four hours to complete.

Migrate storage objects from an appliance

Use appliance storage object migration to move storage resources to another appliance or multiple appliances in the cluster. This feature is applicable if you want to evacuate space from an appliance, power off an appliance, or remove an appliance from a cluster.

About this task

Volumes, volume groups, and vVols are eligible for migration. When you migrate a storage object, all associated snapshots and thin clones are also migrated.

The following storage objects are not eligible for migration:

Table 13. Storage objects ineligible for migration


| Ineligible object | How to make the object eligible |
|--|--|
| File objects | File objects cannot be migrated. |
| Volumes or volume groups in an active import session | Wait for the import session to end. |
| Volumes, volume groups, or vVols in an active internal migration | If the system is migrating objects from the appliance, wait for the migration to end. If the system is migrating objects to the appliance, consider canceling the migration. |
| Offline volumes | The volume is offline due to metadata inconsistencies. Contact your service provider to bring it online. |
| vVols with bound snapshots | Interrupt the VMware operation that created the bound snapshot or wait for the process to complete. |
| vVol bound fast clones | Power off the linked clones of the VM that the vVol belongs to. |

You cannot migrate storage objects from an appliance that is out of space and has entered read-only mode. If an appliance is out of space, you must add more storage capacity or delete storage objects until the appliance has at least 16 GB of free space.

To migrate storage objects to another appliance in the cluster:

Steps

1. Under **Hardware**, select the appliance from which you want to migrate storage objects.
2. Under **More Actions**, select **Migrate**.
3. Follow the prompts in the **Migration** wizard to migrate storage objects to another appliance.

 **NOTE:** The maximum number of storage objects that you can select for a single migration action is 4000.

Results

The LUN ID of a volume changes automatically when a volume is migrated from one PowerStore appliance to another appliance in the same cluster.

Follow these guidelines when migrating a boot volume:

- Power off the connected host before migrating the boot volume. Then, change the LUN ID for the boot volume after the migration is performed and power on the host.
- The recommend Host LUN ID for a boot LUN is 0.
- After migrating a boot from SAN volume, the LUN ID can be changed back to 0.

 **NOTE:** For instructions on changing the LUN, see the PowerStore Host Configuration Guide.

Reinitialize the system


Learn how to reset the entire system to the original, default factory settings.

Topics:

- [Reinitialize the system](#)

Reinitialize the system

Reinitializing the system resets the entire system to the original, default factory settings. You can reinitialize the system using service scripts. To use service scripts to reinitialize the system, refer to the *PowerStore Series Service Scripts Guide*.

 **CAUTION:** Reinitializing the system could result in data loss.