Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x

F13771-16
December 2023

Oracle ZFS Storage Appliance Customer Service Manual, Release OS8.8.x,

F13771-16

Primary Author: Heidi Hall

Contents

1 Getting Started with Servicing the Appliance

5 Servicing the Oracle ZFS Storage ZS7-2 Controller

6 Servicing the Oracle ZFS Storage ZS5-4 Controller

9 Servicing the Oracle ZFS Storage ZS3-4 Controller

10 Servicing the Oracle ZFS Storage ZS3-2 Controller

11 Servicing the Sun ZFS Storage 7x20 Controller

19 Working with Software Updates

20 Upgrading the Software

21 Working with Storage Firmware Updates

22 Working with Platform Firmware Updates

23 Working with Health Checks

24 Deferred Updates

1

Getting Started with Servicing the Appliance

To get started with servicing Oracle ZFS Storage Appliance, use the following sections:

- Introduction to the BUI
- Introduction to the CLI
- BUI Hardware View
- Viewing the Remaining SSD Endurance (BUI)
- Setting a Threshold Alert for SSD Endurance (BUI)
- Viewing CPU Details (BUI)
- Checking Disk Shelf Cable Connections (BUI)
- System Overview
- Chassis Details View
- Disk Shelf Details View
- CLI Hardware View
- Checking Disk Shelf Cable Connections (CLI)
- Configuring the Management Port
- Monitoring SSD Endurance
- Viewing the Remaining SSD Endurance (CLI)
- Setting a Threshold Alert for SSD Endurance (CLI)
- Using Oracle ILOM to Diagnose Hardware Faults
- Understanding Hardware Status
- Working with Problems
- Using Logs
- Using Workflows

To service system hardware components, use the following sections:

- Prerequisites for Servicing Hardware
- Servicing the Oracle ZFS Storage ZS9-2 Controller
- Servicing the Oracle ZFS Storage ZS7-2 Controller
- Servicing the Oracle ZFS Storage ZS5-4 Controller
- Servicing the Oracle ZFS Storage ZS5-2 Controller
- Servicing the Oracle ZFS Storage ZS4-4 Controller
- Servicing the Oracle ZFS Storage ZS3-4 Controller
- Servicing the Oracle ZFS Storage ZS3-2 Controller
- Servicing the Sun ZFS Storage 7x20 Controller

- Servicing the Disk Shelves

To maintain the system, use the following sections:

- Viewing System Disks Status
- Working with Support Bundles
- Performing the Initial Setup
- Restarting the Appliance
- Performing a Factory Reset
- Backing Up the Configuration
- Working with Software Updates
- Upgrading the Software
- Working with Storage Firmware Updates
- Working with Platform Firmware Updates
- Working with Health Checks
- Deferred Updates
- Rolling Back System Software
- Removing an Update Package

Introduction to the BUI

The browser user interface (BUI) is the graphical tool for administration of Oracle ZFS Storage Appliance. The BUI provides an intuitive environment for performing administration tasks, visualizing concepts, and analyzing performance data.

The management software is designed to be fully featured and functional on a variety of web browsers, as described in Browser User Interface (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*. Direct your browser to the system using either the IP address or host name that you assigned to the NET-0 port during initial configuration as follows: `https://`*`ipaddress`*`:215` or `https://`*`hostname`*`:215`. The login screen appears.

Click **Help** at the top of any screen to access the context-sensitive online help. Some BUI screens offer additional links to appliance online help within the page itself. These links might be more context-sensitive than the **Help** link at the top of the screen. You can also access the online help at the host name or IP address of the appliance software:

`https://`*`host-name`*`:215/help/`*`language-code`* (English example: `https://hostname:215/help/en`)

`https://`*`IP-address`*`:215/help/`*`language-code`* (English example: `https://192.0.2.1:215/help/en`)

Introduction to the CLI

The Command Line Interface (CLI) provides the same capabilities that the BUI provides, while also providing a powerful scripting environment for performing repetitive tasks.

If you are not sure what your choices are or what to enter in a particular situation, use the following methods to get help:

- Tab completion. If you have typed a partial command, you can type a tab character to see what should come next. For example, if you are setting the value of a property that has a fixed set of possible values, type a tab character after the = symbol to show the list of possible values. You can type a tab character in the middle of a command name and the CLI will complete the command name or type two tab characters to complete the command name and also list the possible arguments of the command.
- The `help` command. Enter `help` to see a list of commands that can be used in the current context, or enter *topic* `help` to get help on a different topic. Enter `help general` for descriptions of how to use the tab character and for information about how to navigate within a long command line and how to move backward and forward through your command history.

The following example demonstrates using both the `help` command and tab completion to show arguments that you can use with the `help` command:

```
hostname:> he tab
           tab

builtins    commands    general    help    properties    script
```

BUI Hardware View

The hardware views in the BUI provide interactive illustrations that enable you to browse through the Oracle ZFS Storage Appliance system, including the attached disk shelf components.

The icons used in the hardware view are described in the following table.

Table 1-1 Icons in the BUI Hardware View

Icon	Description	Icon	Description
	Show more information about this component		Show a more detailed view of this component
	When flashing, the locator LED on the hardware component is lit		Leave this detailed view
	Hardware component is ok (green)		Power off, reboot, or diagnostic reboot
	Hardware component is faulted (amber)		Offline this disk
	Disk drive can be safely removed (blue on amber)		Port is active
	Hardware component is not available (grey)		Port is inactive

Viewing the Remaining SSD Endurance (BUI)

Use the following procedure to view the remaining SSD endurance.

1. From the **Maintenance** menu, select **Hardware**.
2. Click **Show Details** to view controller details.

3. Click **Disk**.
4. Hover over a disk row and click its information icon . The disk details are displayed, including the remaining SSD endurance.

5. Click **OK**.

Setting a Threshold Alert for SSD Endurance (BUI)

Use the following procedure to set a threshold alert for SSD endurance.

1. From the **Configuration** menu, select **Alerts**, then **Threshold Alerts**.
2. Click the add icon .
3. In the **Threshold** field, select **SSD: Endurance on any device** and specify its parameters.

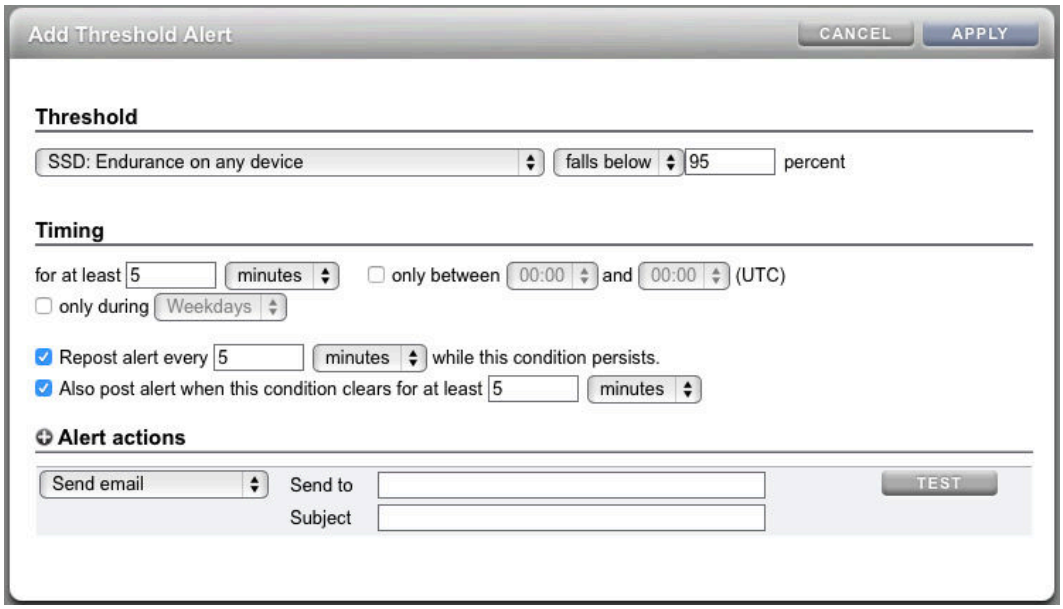

4. Specify the parameters for **Timing** and **Alert actions**.
5. Click **TEST** to verify the settings.
6. Click **APPLY**.

Viewing CPU Details (BUI)

Use the following procedure to view CPU details.

1. From the **Maintenance** menu, select **Hardware**.
2. Click **Show Details** to view controller details.
3. Click **CPU**.
4. Hover over a CPU row and click its information icon. The CPU details are displayed.

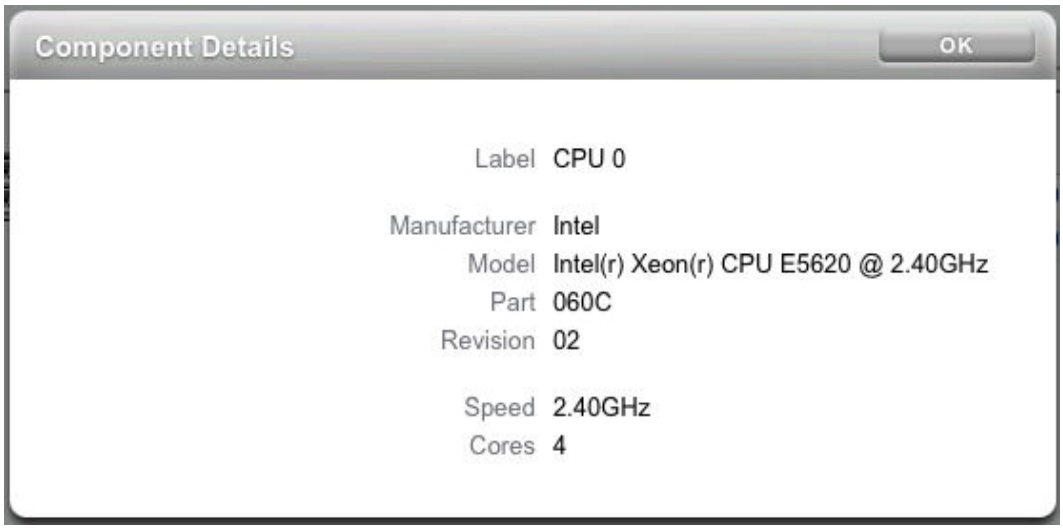

5. Click **OK**.

Checking Disk Shelf Cable Connections (BUI)

The disk shelf cable checking feature detects such SAS cabling errors as a cable connected to the wrong port on the controller or disk shelf, a missing cable in the redundant path, and a faulty cable. Errors and their resolution can be found in the Active Problems area of the software. Use the following procedure to first check the cable connections on the disk shelf, and then on the controller.

1. From the **Maintenance** menu, select **Hardware**.
2. Click the right-arrow icon next to the disk shelf name.
3. Click **Slot**.
4. Click a slot row for an I/O Module (IOM) or SIM to highlight its location in the image to the left.

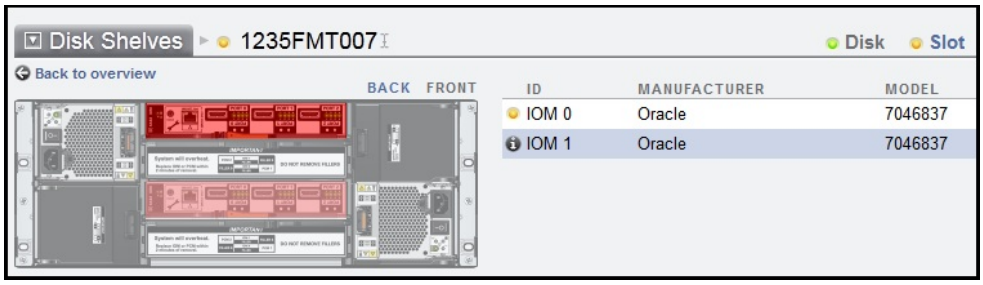

5. Click its information icon to view its **Component Details**. The **Connections** section contains the following information:
   - **Port ID/Location –** Port ID and location within the slot. The boxes to the right of the port ID indicate both the port location and its status:
     - Gray – Connected
     - Clear – Not connected
     - Amber – Faulted
   - **Terminal –** Device ID, label, and corresponding port on the terminating chassis.

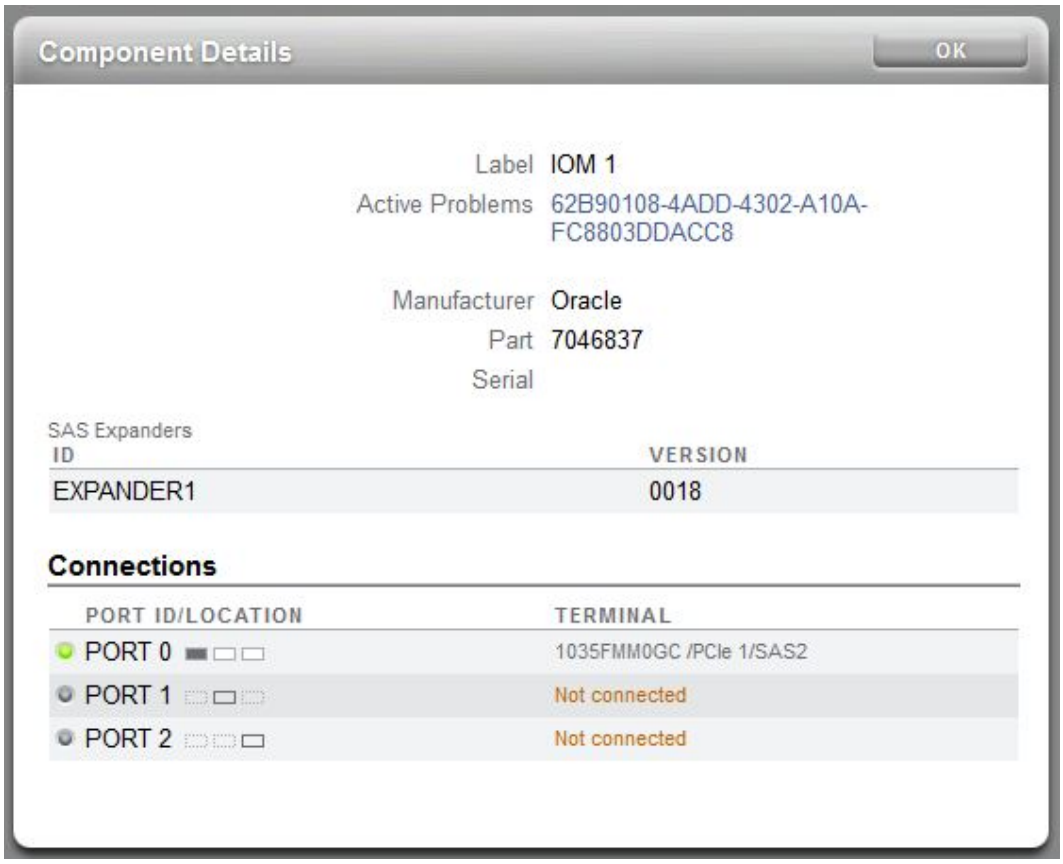

6. If there is a fault, click the **Active Problems** link to view details, including the recommended action. Fixed disk shelf cabling faults are automatically cleared. To return to the **Maintenance: Hardware** screen, click **Hardware**.
7. If there is not a fault, click **OK** in the **Components Details** dialog box. The following figure shows component details for a controller SAS HBA slot.

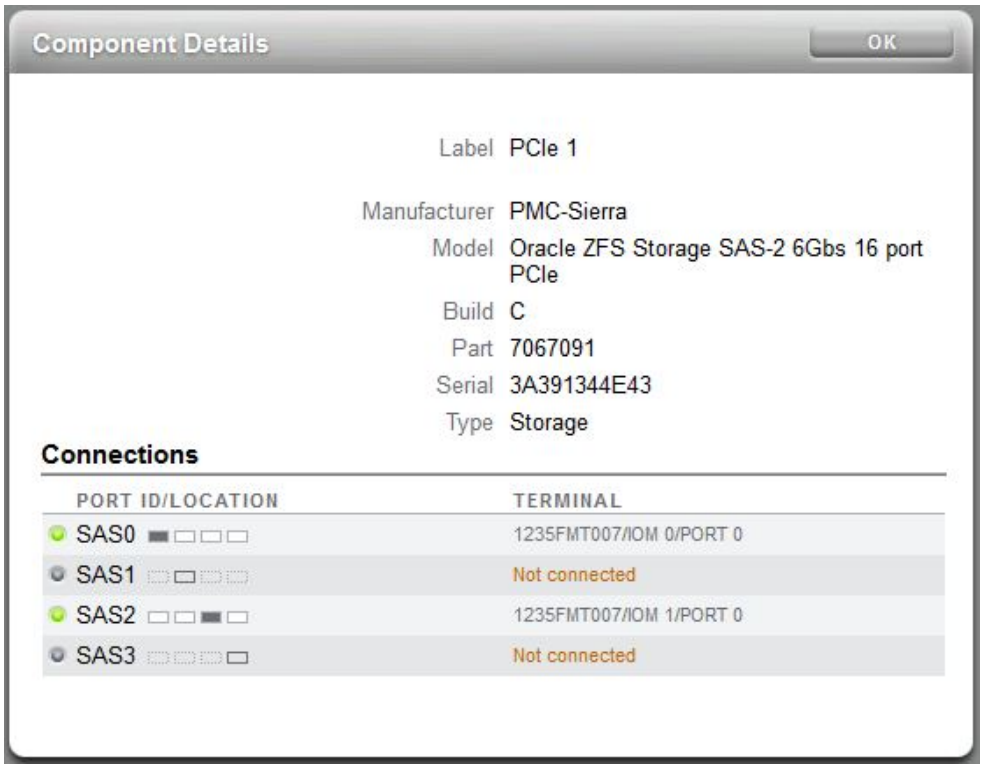

8. Repeat steps 4 through 7 for each applicable slot.
9. Click **Back to overview** to return to the **Maintenance: Hardware** screen.

10. To check the SAS cable connections on the controller, click the right-arrow icon next to the controller name.
11. Repeat steps 3 through 8, except choose a SAS HBA in step 4.

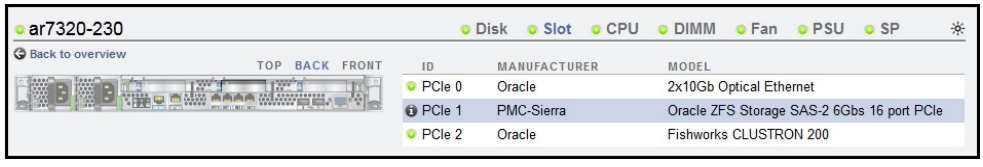

System Overview

Click the Sun/Oracle logo in the upper left corner of the BUI to display a **System Info** pop-up window with information such as last updated date, serial number, and release.

Below the Sun/Oracle logo, the power icon presents a dialog box to either power off, reboot (power cycle), or reboot Oracle ZFS Storage Appliance with diagnostics.

Do not select the diagnostic reboot option unless Oracle Service personnel instruct you to select it. The diagnostic reboot can take a long time to complete and can have adverse results if not performed properly. Do not perform a diagnostic reboot when operations are occurring that affect the entire system, such as upgrading firmware, executing commands, or configuring or unconfiguring storage. The diagnostic reboot option is not available from the system chassis power icon.

The **Maintenance: Hardware** page shows a summary of the system chassis and any attached disk shelves (on supported systems).

- System Chassis
- Disk Shelves

System Chassis

The top half of the **Maintenance: Hardware** page shows the primary system chassis. The chassis name is initially set to the Oracle ZFS Storage Appliance name during installation. To change the chassis name, use the **System Name** field on the **Configuration: Services: System Identity** page. For more information, see System Identity Configuration in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

To the left of the chassis name is an icon that indicates whether any components in the chassis are faulted. See the descriptions of status icons in BUI Hardware View. When you hover over the status icon, the icon becomes an information icon. Click the information icon to display a **Component Details** pop-up window for the chassis.

Do one of the following to display a new page with more detail about the chassis, as described in Chassis Details View:

- Click the right-arrow icon to the left of the status indicator icon.
- Click the chassis illustration to the left of the right-arrow icon.
- Click the **Show Details** link below the illustration.

At the top right of the system chassis section is the locator icon. Click the locator icon to turn the chassis locator LED on or off. The icon flashes when the locator LED is on.

The power icon presents a dialog box to either power off or reboot (power cycle) the appliance.

The following table describes the information that is shown in this summary view.

Table 1-2 System Chassis Properties

Property	Description
Manufacturer	Manufacturer of the system
Model	System model name
Serial	System chassis hardware serial number
Processors	Count and description of processors in the system
Memory	Total memory in the system
System	Size and number of system disks used for the system image
Data	Size and number of data disks in the system chassis. This property is valid only for standalone systems. If no data disks are present, a hyphen character (−) is displayed.
Cache	Size and number of read-optimized cache devices in the system chassis. This property is valid only for expandable systems that support additional disk shelves. If no cache devices are installed in the system chassis, a hyphen character (−) is displayed.
Log	Size and number of log disks in the system chassis. This property is valid only for standalone systems. If no log devices are installed in the system chassis, a hyphen character (−) is displayed.
Total	Total size and count of all disks in the system

Related Topics

- Chassis Details View

Disk Shelves

A list of disk shelves, if supported, is displayed below the system chassis section on the **Maintenance: Hardware** page. The illustration in the **Disk Shelves** section changes to represent the front of the disk shelf that is currently selected in the list.

Do one of the following to display a new page with more detail about that disk shelf, as described in Disk Shelf Details View:

- Double-click an entry in the list.
- Click the right-arrow icon at the far left of each entry in the list.

To the right of the right-arrow icon is an icon that indicates the status of that disk shelf. When you hover over an entry, the status icon becomes an information icon. Click the information icon to display a **Component Details** pop-up window for that disk shelf.

At the far right of each entry in the list is a locator icon. Click the locator icon to turn the locator LED for that disk shelf on or off. The icon flashes when the locator LED is on.

The following table describes the information that is shown in the entry for each disk shelf.

Table 1-3 Disk Shelf Properties

Property	Description
Name	Name of the disk shelf, used in faults and alerts. The name is initially set to the serial number of the disk shelf. To change the name, click on the name in the entry.
MFR/MODEL	The disk shelf manufacturer and the model number.
Type	Device type: HDD or SSD.
RPM	Disk revolutions per minute.
Data	Total size of all data disks within the disk shelf.
Cache	Total size of all read-optimized cache devices within the drive shelf. If there are no cache devices in the shelf, then a hyphen character (−) is displayed. Read cache devices are supported in DE2-24 disk shelves with a minimum software version. For software requirements and slot configuration rules, see Disk Shelf Configurations.
Log	Total size of all write-optimized cache devices within the drive shelf. If there are no log devices in the shelf, then a hyphen character (−) is displayed.
Paths	Total number of I/O paths to the disk shelf. This value normally should be 2 because only configurations with multiple paths to all disks are supported. Click the information icon to show each HBA slot and the state of each path. If the disks within the disk shelf are not currently configured as part of a storage pool, complete path information will not be available, though two paths to the chassis are displayed.

Related Topics

- Disk Shelf Details View

Chassis Details View

To view chassis details, from the **Maintenance** menu, select **Hardware**, and click the right-arrow icon, or the chassis illustration to the left of the right-arrow icon, or the **Show Details** link beneath the illustration.

On the chassis details page, click the **Top**, **Back**, and **Front** buttons above the chassis illustration to change the view of the chassis. Click the left-arrow icon to return to the chassis overview page.

The top of this chassis details view shows the chassis status indicator icon and the chassis name on the left, and the chassis locator icon on the right. Between the chassis name and the locator icon are the following tabs:

- Disk
- Slot
- CPU (controller only)
- DIMM (Memory) (controller only)
- Fan
- PSU (Power supply unit)

- SP (Service processor) (controller only)

Information shown for the service processor is different from the information shown for other components. See Service Processor.

Except for the **SP** view, each component type view presents a list of all physical locations within the chassis where these components can be present. Each entry in the list has a status indicator icon. See descriptions of the status indicator icons in BUI Hardware View. In each of these component views, the currently selected component is highlighted in the illustration. Click a component within the image to select the corresponding entry in the list, or click an entry in the list to highlight the component in the illustration. Double-click an entry in the list, or click the information icon, to show more information about that component, including any **Active Problems**. Click the ID of an active problem to see details about the problem.

The **Disk** tab shows a list of each HDD or SSD. Some entries have an offline icon that you can click to offline this disk. This option is only available for disks that are part of a configured storage pool (including the system pool). Offlining a disk prevents the system from reading or writing to that disk. Faulted devices are already avoided, so this option should only be required if a disk is exhibiting performance problems that do not result in pathological failure. It is not possible to offline a disk that would prevent access to data (for example, offlining both halves of a mirror). If the device is an active hot spare, this action will also give the option of detaching the hot spare completely. Once a hot spare is detached, it cannot be activated except through another fault or hotplug event. When the disk is offlined, the offline icon is replaced by the not-available icon.

> **Note:**
>
> Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check whether an upgrade is in progress, from the **Maintenance** BUI menu, select **System**, or navigate to `maintenance system updates` in the CLI.

Except for the **SP** view, components listed on the tabs report a subset of the properties described in the following table.

Table 1-4 Chassis Component Properties

Property	Description
Build	Manufacturing build identifier. This is used to identify a particular location or batch where the component was manufactured.
Connections (Port ID/ Location, Terminal)	If your appliance and software version support disk shelf cable checking, this section is displayed and contains details for the connection defined by the local port and terminating chassis/port.
Cores	Number of CPU cores.
Endurance	The lifetime percentage remaining on an SSD. Lifetime starts at 100% and decreases with disk usage.
FMRI	Fault managed resource identifier (FMRI) for the component. This is an internal identifier used to identify the component within faults and is intended for service personnel.

Table 1-4 (Cont.) Chassis Component Properties

Property	Description
GUID	Hardware global unique identifier.
HW Crypto	Disk has onboard hardware cryptographic support.
ID	Human-readable identifier for this component within the chassis. This is typically, but not necessarily, equivalent to the label printed on the physical chassis.
Last Update	The date and time of the last endurance reading of the SSD.
MFR or Manufacturer	Component manufacturer.
Model	Component model.
Part	Component part number, or core factory part number. The orderable part number may differ, depending on whether a component is for replacement or expansion, and whether it is part of a larger assembly. Your service provider should be able to refer you to the appropriate orderable part. For components without part numbers, the model number should be used instead.
Revision	Firmware or hardware revision of the component.
RPM	Disk revolutions per minute.
Serial	Component serial number.
Size	Total memory or storage, in bytes.
Speed	Processor speed, in gigahertz.
Type	Disk type. Can be one of System, Data, Log, Cache, Spare, or Peer. When a spare is active, it will be displayed as Spare [A]. When a disk type is shown as Peer, it indicates that the disk is configured as part of a pool that is owned by the cluster peer.

Disk Shelf Details View

To view details for a particular disk shelf, from the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon, or double-click the entry for that disk shelf in the **Disk Shelves** list. Click the **Disk Shelves** label to show the list of disk shelf names and switch to the details view for a different disk shelf. Click the left-arrow icon to return to the **Maintenance: Hardware** page.

The details view of a disk shelf provides information for each device installed in the disk shelf.

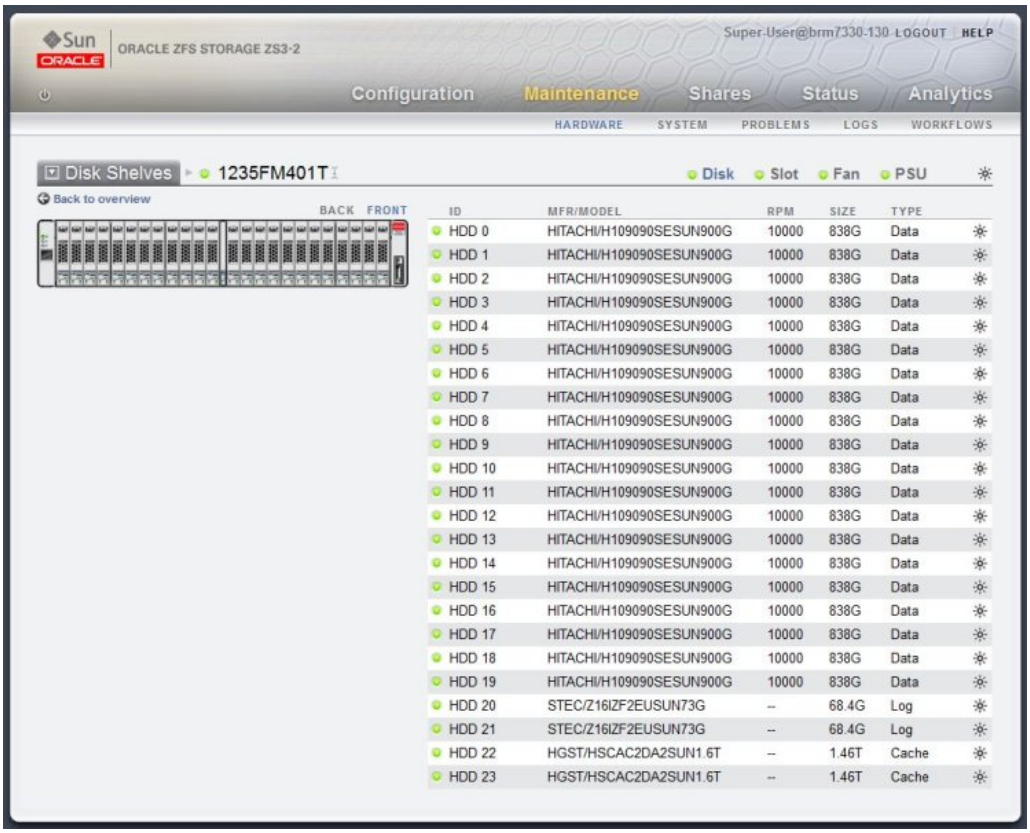

The top of the disk shelf details view shows the disk shelf status indicator icon and the disk shelf name on the left, and the disk shelf locator icon on the right. Between the chassis name and the locator icon are the following tabs:

- Disk
- Slot
- Fan
- PSU (Power supply unit)

The **Disk** tab shows a list of each HDD or SSD. Some entries have an offline icon that you can click to offline this disk. This option is only available for disks that are part of a configured storage pool (including the system pool). Offlining a disk prevents the system from reading or writing to that disk. Faulted devices are already avoided, so this option should only be required if a disk is exhibiting performance problems that do not result in pathological failure. It is not possible to offline a disk that would prevent access to data (for example, offlining both halves of a mirror). If the device is an active hot spare, this action will also give the option of detaching the hot spare completely. Once a hot spare is detached, it cannot be activated except through another fault or hotplug event. When the disk is offlined, the offline icon is replaced by the not-available icon .

> **Note:**
>
> Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check whether an upgrade is in progress, from the **Maintenance** BUI menu, select **System**, or navigate to `maintenance system updates` in the CLI.

Components listed on the tabs report a subset of the properties described in the following table.

Table 1-5 Disk Device Detail Properties

Property	Description
ID	Unique identifier of the device.
MFR or Manufacturer	Component manufacturer.
Model	Component model.
Part	Component part number, or core factory part number. The orderable part number may differ, depending on whether a component is for replacement or expansion, and whether it is part of a larger assembly. Your service provider should be able to refer you to the appropriate orderable part. For components without part numbers, the model number should be used instead.
RPM	Disk revolutions per minute.
Size	Total memory or storage, in bytes.
Type	Disk type. Can be one of System, Data, Log, Cache, Meta, Spare, or Peer. **Read cache devices are supported in DE2 disk shelves with a minimum software version. For software requirements and slot configuration rules, see Disk Shelf Configurations.**

InfiniBand Host Controller Adapters

InfiniBand Host Controller Adapters (HCA) report additional properties for the list of available ports:

Table 1-6 Additional Properties for InfiniBand Host Controller Adapters

Action	Description
State	When "active", the active-port icon is displayed. Other valid port states ("down", "init", and "arm") are denoted by the inactive-port icon. Mousing over the port icon will display the current port state in the tip pop-up.
GUID	The hardware assigned port GUID.
Speed	The current port speed enabled: Single Data Rate (SDR), Dual Data Rate (DDR) or Quad Data Rate (QDR).

Service Processor

The **SP** tab on the chassis details view presents a set of network properties that can be configured from the storage appliance. The following table describes the properties that you

can set to control the behavior of the service processor network management port. Changing multiple values in conflicting ways (such as changing static IP assignments while in DHCP mode) has undefined behavior.

Table 1-7 Properties for the Service Processor Network Management Port

Property	Description
MAC Address	Hardware MAC address. This is read-only.
IP Address Source	Either DHCP or Static. Controls whether DHCP should be used on the interface.
IP Address	IPv4 Address, when using static IP configuration. IPv6 is not supported.
Subnet Mask	Dotted decimal subnet, when using static IP configuration.
Default Gateway	IPv4 default gateway address.

The **SP** view also provides a reset icon that you can click to reset the service processor.

CLI Hardware View

Hardware status details are available in the CLI under the `maintenance hardware` section. Use the `show` command to list the status of all components. The `list` command will list available chassis, which can be selected and then viewed using `show`.

```
hostname:> maintenance hardware show
Chassis:

              NAME              STATE   MANUFACTURER  MODEL

chassis-000   hostname          ok      Oracle        Oracle ZFS Storage ZS9-2

cpu-000       CPU 0             ok      Intel         Intel(r) Xeon(r) Platinum
8358 CPU @ 2.60GHz
cpu-001       CPU 1             ok      Intel         Intel(r) Xeon(r) Platinum
8358 CPU @ 2.60GHz
disk-000      NVME 0            ok      INTEL         SSDPF2KX038T9S
disk-001      NVME 1            ok      INTEL         SSDPF2KX038T9S
disk-002      NVME 2            absent  -             -
disk-003      NVME 3            absent  -             -
disk-004      NVME 4            absent  -             -
disk-005      NVME 5            absent  -             -
disk-006      NVME 6            absent  -             -
disk-007      NVME 7            absent  -             -
disk-008      NVME 8            absent  -             -
disk-009      NVME 9            absent  -             -
disk-010      NVME 10           absent  -             -
disk-011      NVME 11           absent  -             -
fan-000       FM 0              ok      Oracle        8204006
fan-001       FM 1              ok      Oracle        8204006
fan-002       FM 2              ok      Oracle        8204006
fan-003       FM 3              ok      Oracle        8204006
memory-000    DIMM 0/13         ok      Samsung       65536MB DDR4 SDRAM DIMM
memory-001    DIMM 0/12         ok      Samsung       65536MB DDR4 SDRAM DIMM
```

```
memory-002   DIMM 0/15          ok      Samsung       65536MB DDR4 SDRAM DIMM
memory-003   DIMM 0/14          ok      Samsung       65536MB DDR4 SDRAM DIMM
memory-004   DIMM 0/9           ok      Samsung       65536MB DDR4 SDRAM DIMM
memory-005   DIMM 0/8           ok      Samsung       65536MB DDR4 SDRAM DIMM
...
```

The 5th, 6th, and 7th columns for serial number (`SERIAL`), revolutions per minute (`RPM`), and type (`TYPE`) have been truncated in the previous example, as has the length of this list.

Component Properties

If a particular component is selected, detailed information about its properties is reported. The following properties are supported, with the corresponding BUI property name. For a description of a particular property, see the earlier descriptions.

Table 1-8 Component CLI Properties and BUI Equivalent Properties

CLI Property	BUI Property
build	Build
cores	Cores
device	N/A
faulted	(status indicator)
label	Label
locate (writable)	(status indicator)
manufacturer	Manufacturer
model	Model
offline (writable)	(status indicator)
part	Part
present	(status indicator)
revision	Revision
serial	Serial
size	Size
speed	Speed
state	Connections (status)
terminal	Connections - Terminal: Device ID, label, and corresponding port on terminating chassis
type	(combined with use)
use	Type

When viewing a disk that is active as a hot spare, the `detach` command is also available.

Viewing CPU Details (CLI)

Use the following procedure to view CPU details.

- To view CPU details, use the following CLI commands:

```
hostname:maintenance hardware> select chassis-000
hostname:maintenance chassis-000> select cpu
hostname:maintenance chassis-000 cpu> select cpu-000
hostname:maintenance chassis-000 cpu-000> show
Properties:
                        label = CPU 0
                      present = true
                      faulted = false
                 manufacturer = Intel
                        model = Intel(r) Xeon(r) Platinum 8358 CPU @ 2.60GHz
                         part = SRKJ1
                       serial = 54-01933CBE58B1C71C
                     revision = CD8068904572302
                        cores = 32
                        speed = 2.60GHz
```

Checking Disk Shelf Cable Connections (CLI)

The disk shelf cable checking feature detects such SAS cabling errors as a cable connected to the wrong port on the controller or disk shelf, a missing cable in the redundant path, and a faulty cable. Errors and their resolution can be found in the Active Problems area of the software. Use the following procedure to first check the cable connections on the disk shelf, and then on the controller.

1. Go to `maintenance hardware` and select the disk shelf chassis.

   ```
   hostname:maintenance hardware> select chassis-001
   hostname:maintenance chassis-001>
   ```

2. Enter `select slot` and then enter `show`.

   ```
   hostname:maintenance chassis-001> select slot
   hostname:maintenance chassis-001 slot> show
   Slots:
              LABEL     STATE          MANUFACTURER   MODEL                SERIAL
   slot-000   IOM 0     faulted        Oracle         7046837
   unknown
   slot-001   IOM 1     ok             Oracle         7046837
   unknown
   ```

3. Select a slot and enter `show` to view its properties and components.

   ```
   hostname:maintenance chassis-001 slot> select slot-000
   hostname:maintenance chassis-001 slot-000> show
   Properties:
                 label = IOM 0
               present = true
               faulted = true
          manufacturer = Oracle
                  part = 7046837
                serial =
                locate = false

   Components:
                LABEL              STATE          TERMINAL
   port-000     PORT0              connected      1035FMM0GC/PCIe 1/SAS0
   port-001     PORT1              unconnected    -
   port-002     PORT2              unconnected    -
   ```

If the `faulted` property is `true`, go to the `maintenance problems` context and enter `show` to view the active problems. Note that fixed disk shelf cabling faults are automatically cleared.

4. To view the port properties, select a port and enter `show`.

   The same properties are displayed as in the previous step. The `state` property shows the connection status. The `terminal` property displays the device ID, label, and corresponding port on the terminating chassis.

   ```
   hostname:maintenance chassis-001 slot-000> select port
   hostname:maintenance chassis-001 slot-000 port> select port-000
   hostname:maintenance chassis-001 slot-000 port-000> show
   Properties:
               label = PORT0
               state = connected
            terminal = 1035FMM0GC/PCIe 1/SAS0
   ```

5. Repeat steps 2 through 4 for each applicable slot.
6. Enter `done` until you are at the `maintenance hardware` prompt.
7. Select the controller chassis.

   ```
   hostname:maintenance hardware> select chassis-000
   hostname:maintenance chassis-000>
   ```

8. Repeat steps 2 through 5, except choose a SAS HBA slot in step 3.

Performing a Diagnostic Reboot (CLI)

> **Caution:**
>
> Only issue the diagnostic reboot command as instructed by Oracle Service personnel.

A diagnostic reboot gathers diagnostic information before power cycling the appliance. This operation could take a long time to complete and could cause adverse results if not performed properly. Do not reboot when system-affecting operations are occurring, such as upgrading firmware, executing commands, and configuring or unconfiguring storage.

- To perform a diagnostic reboot, use the following CLI command:

  ```
  hostname:maintenance system> reboot -d
  ```

Configuring the Management Port

All standalone controllers should have at least one NIC port configured as a management interface. Select the **Allow Admin** option in the BUI to enable BUI connections on port 215 and CLI connections on ssh port 22.

All cluster installations should have at least one NIC port on each controller configured as a management interface as described above. In addition, the NIC instance number must be unique on each controller. For example, nodeA uses igb0 and nodeB uses igb1, so that neither may be used as a cluster data interface. In addition, these interfaces must be locked to the controller by using the **Configuration** BUI menu and selecting **Cluster**. In some cases,

this may require installation of an additional network interface card on each controller in a cluster configuration.

If access to the appliance data interfaces is impossible for any reason, the management network interface will maintain BUI and CLI access. During a cluster takeover, interfaces are taken down on the failed controller. So, locked interface configuration is required to gather diagnostic information from a failed controller.

> **Note:**
>
> Failure to configure locked management interfaces on a cluster may lead to longer than necessary fault diagnosis and resolution times.

Monitoring SSD Endurance

SSD endurance reports the remaining life expectancy of an SSD. Endurance properties are reported on the BUI **Maintenance: Hardware** page and in the CLI `maintenance chassis disk` context. You can set a threshold alert to occur when one or more SSDs exceeds or falls below a specified percentage. For example, you can set a threshold and generate an alert when one or more SSD devices falls below 5% of endurance (remaining lifetime).

For information about threshold alert properties, see Adding Threshold Alerts in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

You can use this feature to monitor the life expectancy of an SSD. For example, if SSD endurance reports 50% after two years, the SSD is estimated to last for another two years, assuming the workload is consistent.

> **Note:**
>
> Do not use reported SSD endurance percentage as an indication to replace the SSD. SSD warranty replacements are made only when a failure is reported.

Related Topics

- Setting a Threshold Alert for SSD Endurance - BUI, CLI
- Viewing the Remaining SSD Endurance - BUI, CLI

Viewing the Remaining SSD Endurance (CLI)

Use the following procedure to view the remaining SSD endurance.

- To list the properties of a disk, use the following CLI commands:

```
hostname maintenance hardware
hostname:maintenance hardware> select chassis-001
hostname:maintenance hardware chassis-001> select disk
hostname:maintenance hardware chassis-001> select disk-015
hostname:maintenance hardware chassis-001 disk-015> list
```

```
Properties:
                          label = HDD 15
                        present = true
                         faulted = false
                   manufacturer = SANDISK
                          model = LB806M---SUN800G
                         serial = 40042896
                       revision = S30E
                           size = 745G
                           type = data
                            use = data
                            rpm = --
                         device = c0t5001E82002630190d0
                      pathcount = 2
                      interface = SAS
                      endurance = 100%
              endurance_updated = 2022-3-17 22:04:14
                         locate = false
                        offline = false
```

As shown above, `disk-015` is at 100%, which indicates a new SSD. The remaining endurance estimate decreases as the disk is used.

Setting a Threshold Alert for SSD Endurance (CLI)

Use the following procedure to set a threshold alert for SSD endurance.

1. To set a threshold alert for SSD endurance, use the following CLI commands:

```
hostname:configuration alerts> thresholds
hostname:configuration alerts thresholds> create
hostname:configuration alerts threshold (uncommitted)> set
statname=ssd.endurance[ssd]
                        statname = ssd.endurance[ssd] (uncommitted)
```

2. To list the properties of the threshold alert, use the following CLI command:

```
hostname:configuration alerts threshold (uncommitted)> list
Properties:
                            uuid = <generated on commit>
                        statname = ssd.endurance[ssd] (uncommitted)
                            type = normal
                           limit = (unset)
                         minpost = 5 minutes
                            days = all
                    window_start = none
                      window_end = 00:00
                       frequency = 5 minutes
                        minclear = 5 minutes
```

> **Note:**
>
> The remaining fields are set the same way you would set them for any other threshold alert.

Using Oracle ILOM to Diagnose Hardware Faults

In rare cases, faults associated with uncorrectable CPU errors are not diagnosable or displayed in the controller. These faults are preserved by and observable on the Oracle Integrated Lights Out Manager (ILOM), which resides in the Service Processor (SP). Connect to Oracle ILOM to diagnose hardware faults that do not appear in the BUI.

Oracle ILOM provides options for network and serial port connectivity. Network connection is the preferred choice because the Oracle ILOM serial port does not always allow adequate means of platform data collection. Network connections are on port 215, and serial connections are on SSH port 22.

> **Caution:**
>
> Failure to configure Oracle ILOM connectivity may lead to longer than necessary hardware fault diagnosis and resolution times.

For information about using Oracle ILOM, refer to the Oracle Integrated Lights Out Manager (ILOM) Documentation Library at https://www.oracle.com/goto/ilom/docs.

To identify the Oracle ILOM firmware version, and to log in to Oracle ILOM to view and clear CPU faults, use the following tasks:

- Identifying the Oracle ILOM Firmware Version - BUI, CLI
- Logging in to Oracle ILOM Using a Local Serial Connection
- Logging in to Oracle ILOM Remotely Using a Web Interface
- Logging in to Oracle ILOM Remotely Using a Command Line Interface
- Viewing and Clearing CPU Faults from Oracle ILOM

Identifying the Oracle ILOM Firmware Version (BUI)

1. Click the logo in the upper-left side of the masthead.
2. In the **System Info** dialog box, locate the **Service Processor** version.

   The **Service Processor** version is the Oracle ILOM firmware version.

Identifying the Oracle ILOM Firmware Version (CLI)

1. Go to `configuration version`.

   ```
   hostname:> configuration version
   ```

2. Enter `get sp_version`.

   The `sp_version` is the Oracle ILOM firmware version.

   ```
   hostname:configuration version> get sp_version
                   sp_version = 3.1.2.54
   ```

Logging in to Oracle ILOM Using a Local Serial Connection

Use the following procedure to log in to Oracle ILOM using a local serial connection.

1. Prepare an administrative client (terminal or terminal emulator) with the serial connection parameters:
   - 8N1: eight data bits, no parity, one stop bit
   - 9600 baud
   - Disable hardware flow control (CTS/RTS)
   - Disable software flow control (XON/XOFF)
2. Connect the administrative client to the SER MGT port on the controller using a serial null modem cable.
3. To establish a connection between your serial console and Oracle ILOM, press **Enter** on the administrative client.

   A login prompt for Oracle ILOM is displayed.
4. Log in to the Oracle CLI using the administrative account name and password.

   Oracle ILOM displays a default command prompt (->), indicating that you have successfully logged in.

Logging in to Oracle ILOM Remotely Using a Web Interface

This procedure requires that you know the administrative account name and password, and that you know the IP address or host name of the controller Service Processor (SP). To improve response times, disable the web browser proxy server, if used.

Use the following procedure to log in to Oracle ILOM remotely using a web interface.

1. Type the controller SP IP address into your web browser address field.

   Example: **`https://172.16.82.26`**
2. Type the administrative user name and password.
3. Click **Log In**.

   The **System Information Summary** page is displayed.

Logging in to Oracle ILOM Remotely Using a Command Line Interface

This procedure requires that you know the administrative account name and password, and that you know the IP address or host name of the controller Service Processor (SP).

Use the following procedure to log in to Oracle ILOM remotely using a command line interface.

1. Using a Secure Shell (SSH) session, log in to Oracle ILOM by specifying your administrative account user name and the IP address or host name of the controller SP.

   Examples:

   `ssh -l` *username host*

   `ssh` *username@host*

where *host* is either the IP address or host name of the controller SP when using DNS. For example: **ssh root@172.16.82.26**

2. Type the password for the administrative account.

   Oracle ILOM displays a default command prompt (->), indicating that you have successfully logged in.

Viewing and Clearing CPU Faults from Oracle ILOM

Log in to the controller as root using the Oracle ILOM CLI.

Use the following procedure to view and clear CPU faults from Oracle ILOM.

1. To view controller faults, use the following CLI command:

   ```
   -> show /SP/faultmgmt
   ```

2. To clear a CPU fault, use the following CLI command:

   ```
   -> set /SYS/MB/Pn clear_fault_action=true
   Are you sure you want to clear /SYS/MB/Pn (y/n)? y
   ```

Understanding Hardware Status

To understand hardware status, use the following sections.

- Locating the Chassis Serial Number
- Locating a Failed Component (BUI)
- Locating a Failed Component (CLI)

Locating the Chassis Serial Number

To obtain support for your storage controller or to order new parts, you need your chassis serial number. There are three ways to locate the chassis serial number.

1. You can find a chassis serial number label on the storage controller front panel on the left side. Another label is on the top of the storage controller.
2. You can click the Sun/Oracle logo in the BUI masthead to obtain the serial number.
3. You can use the following CLI command:

   ```
   hostname: maintenance hardware show
   ```

Locating a Failed Component (BUI)

Use the following procedure to locate a failed component using the BUI.

1. From the **Maintenance** menu, select **Hardware**.
2. Click the right-arrow icon on the storage system or disk shelf which has the fault icon.
3. Locate the fault icon in the lists of hardware components and click it. The image should be updated to show where that component is physically located.

4. Click the locator icon for that component, if the component has it. The LED on the component will begin to flash.

Locating a Failed Component (CLI)

Use the following procedure to locate a failed component using the CLI.

1. Go to `maintenance hardware`:

```
hostname:> maintenance hardware
```

2. List the appliance components:

```
hostname:maintenance hardware> list
             NAME        STATE       MODEL                      SERIAL
chassis-000  hostname    ok          Oracle ZFS Storage ZS7-2   unknown
chassis-001  000000000C  faulted     Oracle Storage DE3-24C     000000000C
```

3. Select the chassis and list its components:

```
hostname:maintenance hardware> select chassis-001
hostname:maintenance chassis-001> list
                        disk
                         fan
                         psu
                        slot
```

4. Select the component type and show all available components:

```
hostname:maintenance chassis-001> select disk
hostname:maintenance chassis-001 disk> show
Disks:
          LABEL   STATE       MANUFACTURER  MODEL             SERIAL
disk-000  HDD 0   ok          ST3500630NS   ST3500630NS       9QG1ACNJ
disk-001  HDD 1   faulted     ST3500630NS   ST3500630NS       9QG1A77R
disk-002  HDD 2   ok          ST3500630NS   ST3500630NS       9QG1AC3Z
disk-003  HDD 3   ok          ST3500630NS   ST3500630NS       9QG1ACKW
disk-004  HDD 4   ok          ST3500630NS   ST3500630NS       9QG1ACKF
disk-005  HDD 5   ok          ST3500630NS   ST3500630NS       9QG1ACPM
disk-006  HDD 6   ok          ST3500630NS   ST3500630NS       9QG1ACRR
disk-007  HDD 7   ok          ST3500630NS   ST3500630NS       9QG1ACGD
disk-008  HDD 8   ok          ST3500630NS   ST3500630NS       9QG1ACG4
disk-009  HDD 9   ok          ST3500630NS   ST3500630NS       9QG1ABDZ
disk-010  HDD 10  ok          ST3500630NS   ST3500630NS       9QG1A769
disk-011  HDD 11  ok          ST3500630NS   ST3500630NS       9QG1AC27
disk-012  HDD 12  ok          ST3500630NS   ST3500630NS       9QG1AC41
disk-013  HDD 13  ok          ST3500630NS   ST3500630NS       9QG1ACQ5
disk-014  HDD 14  ok          ST3500630NS   ST3500630NS       9QG1ACKA
disk-015  HDD 15  ok          ST3500630NS   ST3500630NS       9QG1AC5Y
disk-016  HDD 16  ok          ST3500630NS   ST3500630NS       9QG1ACQ2
disk-017  HDD 17  ok          ST3500630NS   ST3500630NS       9QG1A76S
disk-018  HDD 18  ok          ST3500630NS   ST3500630NS       9QG1ACDY
disk-019  HDD 19  ok          ST3500630NS   ST3500630NS       9QG1AC3Y
disk-020  HDD 20  ok          ST3500630NS   ST3500630NS       9QG1ACG6
disk-021  HDD 21  ok          ST3500630NS   ST3500630NS       9QG1AC3X
disk-022  HDD 22  ok          ST3500630NS   ST3500630NS       9QG1ACHL
disk-023  HDD 23  ok          ST3500630NS   ST3500630NS       9QG1ABLW
```

> **Note:**
>
> The RPM (revolutions per minute) of the disk drive is also shown in the output. However, RPM is truncated in the previous example.

5. Select the faulted component and turn on its locator LED, if the component has it:

```
hostname:maintenance chassis-001 disk> select disk-001
hostname:maintenance chassis-001 disk-001> set locate=true
                       locate = true (uncommitted)
hostname:maintenance chassis-001 disk-001> commit
```

Working with Problems

To aid serviceability, Oracle ZFS Storage Appliance detects persistent hardware failures (*faults*) and software failures (*defects*, often included under faults) and reports them as active problems on the **Maintenance: Problems** page in the BUI, and in `maintenance problems` in the CLI.

If the Phone Home service is enabled, active problems are automatically reported to Oracle Support, where a support case might be opened, depending on the service contract and the nature of the fault. Problem notification can be suspended while you are servicing Oracle ZFS Storage Appliance.

The following topics are described in this section:

- Viewing Active Problems
- Repairing Active Problems
- Suspending and Resuming Problem Notification

Viewing Active Problems

The following table shows some example faults as they would be displayed in the **Active Problems** section of the **Maintenance: Problems** page in the BUI. For each problem, Oracle ZFS Storage Appliance reports what happened, when the problem was detected, the severity and type of the problem, and whether the problem has been phoned home. Severity can be **Minor**, **Major**, or **Critical**. Type can be **Alert**, **Defect**, **Error**, or **Fault**. **Phoned Home** is a date and time or **Never**. The table can be sorted by **Date**.

Table 1-9 Example BUI Problem Displays

Date	Description	Type	Phoned Home
2022-09-16 13:56:36	SMART health-monitoring firmware reported that a disk failure is imminent.	Major Fault	Never
2022-09-05 17:42:55	A disk of a different type (cache, log, or data) was inserted into a slot. The newly inserted device must be of the same type.	Minor Fault	Never
2022-08-21 16:40:37	The ZFS pool has experienced currently unrecoverable I/O failures.	Major Error	Never

Table 1-9 (Cont.) Example BUI Problem Displays

Date	Description	Type	Phoned Home
2022-07-16 22:03:22	A memory module is experiencing excessive correctable errors affecting large numbers of pages.	Major Fault	Never

Clicking on a problem shows more information about the problem in the **Problem Details** section of the page, including the impact to the system, affected components, the system's automated response (if any), and the recommended action for the administrator (if any).

To view the affected hardware component for a hardware fault and to optionally turn on its locator LED on Oracle ZFS Storage Appliance, see Locating a Failed Component - BUI, CLI.

The CLI provides similar information, as shown in the following example:

```
hostname:maintenance problems> show
Problems:

COMPONENT    DIAGNOSED            TYPE            DESCRIPTION
problem-000  2022-4-3 20:30:12    Major Fault     A sensor indicates that the
                                                  power supply '1235FM401W/PSU
                                                  01' is not operating
                                                  properly due to some
                                                  external condition.

problem-001  2022-4-3 17:53:58    Major Fault     External sensors indicate
                                                  that the power supply
                                                  'hostname/PSU 1' is no
                                                  longer operating correctly.
```

For more information, select a problem. Only the `uuid`, `diagnosed`, `severity`, `type`, and `description` fields are considered to be stable. Other property values might change in a new release.

```
hostname:maintenance problems> select problem-000
hostname:maintenance problem-000> show
Properties:
                          uuid = uuid
                          code = SENSOR-8000-7L
                     diagnosed = 2022-4-3 20:30:12
                   phoned_home = never
                      severity = Major
                          type = Fault
                           url = https://support.oracle.com/msg/SENSOR-8000-7L
                   description = A sensor indicates that the power supply
                                 '1235FM401W/PSU 01' is not operating properly
                                 due to some external condition.
                        impact = The enclosure may be getting inadequate
                                 power. Subsequent loss of power supplies may
                                 force the enclosure to shutdown.
                      response = None.
                        action = Check to see if the power cord is connected
                                 properly or if there are other conditions
                                 that may be causing inadequate power to be
                                 provided to the indicated power supply.
                                 Please refer to the associated reference
```

```
                                  document at
                                  https://support.oracle.com/msg/SENSOR-8000-7L
                                  for the latest service procedures and
                                  policies regarding this diagnosis.

Components:

component-000  100%  1235FM401W: PSU 01 (degraded)
                     Manufacturer: Oracle
                     Part number: part-number
                     Serial number: serial-number

hostname:maintenance problem-000> select component-000
hostname:maintenance problem-000 component-000> show
Properties:
                     certainty = 100
                        status = degraded
                  chassis_label = 1235FM401W
                component_label = PSU 01
                   manufacturer = Oracle
                           part = part-number
                         serial = serial-number
```

Related Topics

- Persistent logs of all faults, defects, errors, and alerts are available under **Maintenance: Logs** in the BUI, and under `maintenance logs` in the CLI. For more information, see Using Logs.
- Faults and defects are subcategories of alerts. Filter rules can be configured to cause Oracle ZFS Storage Appliance to email administrators or perform other actions when faults are detected. For more information about alerts, see Configuring Alerts in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Repairing Active Problems

Active problems can be a result of a hardware fault or software defect. To repair an active problem, perform the steps described in the suggested action section. For hardware faults, repair typically involves replacing a physical component. For software defects, repair typically involves reconfiguring and restarting the affected service.

After a problem is repaired, the problem no longer appears in the list of active problems.

While the system can detect repairs automatically, in some cases manual intervention is required. If a problem persists after completing the suggested action, contact Oracle support. You might be instructed to mark the problem as repaired. Manually marking a problem as repaired should only be done under the direction of Oracle service personnel or as part of a documented Oracle repair procedure.

Suspending and Resuming Problem Notification

Servicing the appliance can generate false failures. For example, replacing a disk generates FRU remove and Invalid Configuration events, which can generate SRs.

To avoid sending SRs when no problem exists, you can suspend all notifications during the period when you are performing the service.

Suspending Problem Notification

To suspend all notifications, do one of the following:

- BUI – Check the **Suspend Notifications** box at the top of the **Maintenance: Problems** page.
- CLI – Enable the `suspend_notification` property in `maintenance problems`.

```
hostname:maintenance problems> ls
Properties:
          suspend_notification = disabled
                        period =
```

The `period` property is read-only. As in the BUI, it displays the remaining amount of time that notifications will be suspended.

To enable or disable notification suspension, the user must be assigned the maintenance authorization in the Appliance scope.

Notification suspension behaves in the following way:

- All external notifications are suspended, including the following:
  - Phone Home
  - Emails
  - Any user-configured alert actions, as described in Configuring Alerts in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*
- If you suspend notifications for one node of a cluster, notifications are suspended for both cluster nodes.
- While notifications are suspended, events continue to be logged and will be sent when event notification is resumed. See Resuming Problem Notification.
- By default, notifications are suspended for 8 hours, or for a `period` of 480 minutes.
- While notifications are suspended, a persistent minor alert is displayed in the **Active Problems** section of the **Maintenance: Problems** BUI page, or in the `Problems` section of `maintenance problems`: `"The suspending of notifications has started."`

Resuming Problem Notification

While notifications are suspended, events continue to be logged and will be sent when event notification is resumed.

> **Note:**
>
> Before you resume normal problem notification, clear any problem events that should not be sent to Oracle.

Before you resume normal problem notification, the only accumulated events in the **Problems** BUI page or in the `maintenance problems` CLI context should be problems that still need to be corrected and that need to be sent to Oracle for further action.

To end notification suspension and resume normal problem notification prior to the end of the default suspension period, do one of the following:

- BUI – From the **Maintenance** menu, select **Problems**, and clear the **Suspend Notifications** box.
- CLI – Disable the `suspend_notification` property in `maintenance problems`.

To enable or disable notification suspension, the user must be assigned the maintenance authorization in the Appliance scope.

Using Logs

This section describes different types of log entries and how to use logs:

Alert Log Entries

This is the appliance alert log, recording key events of interest during Oracle ZFS Storage Appliance operation.

The following are example alert log entries as they would appear in the BUI:

Table 1-10 Example BUI Alert Log Entry Displays

Time	Event ID	Description	Type
2022-9-16 13:01:56	f18bbad1-8084-4cab-c950-82ef5b8228ea	An I/O path from slot 'PCIe 0' to chassis 'JBOD #1' has been removed.	Major alert
2022-9-16 13:01:51	8fb8688c-08f2-c994-a6a5-ac6e755e53bb	A disk has been inserted into slot 'HDD 4' of chassis 'JBOD #1'.	Minor alert
2022-9-16 13:01:51	446654fc-b898-6da5-e87e-8d23ff12d5d0	A disk has been inserted into slot 'HDD 15' of chassis 'JBOD #1'.	Minor alert

An information icon next to the **Event ID** means that extended information is available. Click the icon to display information below the list of alerts.

Oracle ZFS Storage Appliance can also be configured to send email, raise an SNMP trap, or perform other actions when particular alerts occur. This is configured in the **Alerts** section. All alerts appear in this log, regardless of whether they have actions configured for them. For more information about alerts, see Configuring Alerts in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Fault Log Entries

The fault log records hardware and software faults. This is a useful reference when troubleshooting hardware failure, as timestamps are available for these hardware fault events.

The following are example fault log entries as they would appear in the BUI:

Table 1-11 Example BUI Fault Log Entry Displays

Date	Event ID	Description	Type
2022-9-5 17:42:35	9e46fc0b-b1a4-4e69-f10f-e7dbe80794fe	The device 'HDD 6' has failed or could not be opened.	Major Fault
2022-9-3 19:20:15	d37cb5cd-88a8-6408-e82d-c05576c52279	External sensors indicate that a fan is no longer operating correctly.	Minor Fault
2022-8-21 16:40:48	c91c7b32-83ce-6da8-e51e-a553964bbdbc	The ZFS pool has experienced currently unrecoverable I/O failures.	Major Error

These faults generate alert log entries that use the alert reporting settings (such as sending email), if configured. Faults that require administrator attention appear in **Problems**. For more information, see Working with Problems.

System Log Entries

This is the operating system log, available to read via the appliance interfaces. This may be useful when troubleshooting complex issues, but should only be checked after first examining the alert and fault logs.

The following are example system log entries as they would appear in the BUI:

Table 1-12 Example BUI System Log Entry Displays

Time	Module	Priority	Description
2022-10-11 14:13:38	ntpdate	error	no server suitable for synchronization found
2022-10-11 14:03:52	genunix	notice	SunOS Release 5.11 Version ak/generic@2013.10.10,1-0 64-bit
2022-10-11 14:02:04	genunix	notice	done
2022-10-11 14:02:01	genunix	notice	syncing file systems...
2022-10-11 13:52:16	nxge	warning	WARNING: nxge : ==> nxge_rxdma_databuf_free: DDI

Audit Log Entries

The audit log records user activity events, including login and logout to the BUI and CLI, and administrative actions. If session annotations are used, each audit entry should be noted with a reason (see Configuring Users in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*).

The following are example audit log entries as they would appear in the BUI:

Table 1-13 Example BUI Audit Log Entry Displays

Time	User	Host	Summary	Session Annotation
2022-10-12 05:20:24	root	deimos	Disabled ftp service	
2022-10-12 03:17:05	root	deimos	User logged in	
2022-10-11 22:38:56	root	deimos	Browser session timed out	
2022-10-11 21:13:35	root	<console>	Enabled ftp service	

Phone Home Log Entries

If Phone Home is used, this log shows communication events with Oracle support. For information about Phone Home, see Phone Home Configuration in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

The following is an example phone home entry as it would appear in the BUI:

Table 1-14 Example BUI Phone Home Entry Display

Time	Description	Result
2022-10-12 05:24:09	Uploaded file 'cores/ak.45e5ddd1-ce92-c16e-b5eb-9cb2a8091f1c.tar.gz' to Oracle support	OK

Viewing Logs (BUI)

Use the following procedure to view logs using the BUI. The following log types are available: **Alerts**, **Faults**, **System**, **Audit**, and **Phone Home**.

1. To view logs, from the **Maintenance** menu, select **Logs**.
2. To navigate between the log types, click the **Alerts**, **Faults**, **System**, **Audit**, and **Phone Home** buttons.
3. To scroll through the logs, use the forward and back buttons.

Listing Logs (CLI)

Use the following procedure to list logs using the CLI. Up to 100 recent entries for each log are visible.

- To list available logs and the time stamp of the last log entry for each log type, use the following CLI commands:

```
hostname:> maintenance logs
hostname:maintenance logs> show
Logs:

LOG          ENTRIES    LAST
alert        2          2022-10-16 02:44:04
audit        42         2022-10-16 18:19:53
fltlog       2          2022-10-16 02:44:04
scrk         0          -
system       100        2022-10-16 03:51:01
```

Viewing All Log Entries (CLI)

Use the following procedure to view all log entries.

- To view all log entries, use the following CLI commands. The most recent entries are displayed at the bottom of the list.

```
hostname:maintenance logs> select audit
hostname:maintenance logs audit> list -a

ENTRY      TIME                 SUMMARY
entry-000  2022-1-9 15:13:50    root, 10.154.161.197, User logged in
entry-001  2022-1-9 15:17:42    root, 10.154.161.197, Released resources to
                                cluster peer
entry-002  2022-1-9 15:20:17    root, 10.154.161.197, Cluster takeover
                                requested
entry-003  2022-1-9 15:20:33    root, 10.154.161.197, Transferred ak:/net/igb1
                                to 2917b8aa-0b0a-4b74-f36b-ff0a8d150c3b
...
entry-2077 2022-12-17 05:24:43 osc_agent, 10.80.218.16, User logged in
entry-2078 2022-12-17 05:24:45 osc_agent, 10.80.218.16, User logged out
entry-2079 2022-12-17 05:24:53 osc_agent, 10.80.218.16, User logged in
```

Viewing Groups of Log Entries (CLI)

Use the following procedure to view groups of log entries.

1. To view up to 100 log entries at a time, use the following CLI commands:

```
hostname:maintenance logs> select audit
hostname:maintenance logs audit> show

ENTRY      TIME                 SUMMARY
entry-2874 2022-2-20 18:05:44   root, 10.159.134.211, Beginning system update
                                to ak-nas@2014.02.11,1-0
entry-2875 2022-2-20 18:08:18   root, 10.159.134.211, Rebooting system as part
                                of update to ak-nas@2014.02.11,1-0
entry-2876 2022-2-20 18:08:20   root, 10.159.134.211, User logged out
entry-2877 2022-2-20 18:16:06   root, 10.159.134.211, User logged in
entry-2878 2022-2-20 18:18:20   root, 10.159.134.211, Deleted update media
                                ak-nas@2011.04.24.8.0,1-1.43
entry-2879 2022-2-20 18:34:19   root, 10.159.134.211, Cluster takeover
                                requested
...
entry-2970 2022-2-27 17:40:37   root,
                                dhcp-amer-vpn-rmdc-
anyconnect-10-159-100-157.vpn.oracle.com,
                                User logged out of CLI
```

```
entry-2971 2022-2-27 17:48:23  root, 10.159.100.157, User logged in
entry-2972 2022-2-27 19:05:46  root, 10.159.100.157, Browser session timed
out
entry-2973 2022-2-27 23:53:31  root,
                               dhcp-amer-vpn-rmdc-
anyconnect-10-159-125-20.vpn.oracle.com,
                               User logged in via CLI
```

2. To view the previous or next 100 log entries, use the following CLI commands:

```
hostname:maintenance logs alert> previous

ENTRY      TIME                SUMMARY
entry-000  2022-11-21 15:45:23 67ccd46c-3d4d-eb86-8966-f9e0ec497293, System
                               software update cannot proceed: J4400 and J4500
                               disk shelves are not supported in this
                               release., Minor alert
...
entry-022  2022-1-16 15:06:10  2f840123-221c-49dc-ae26-e5bfe0952599, System
                               software update cannot proceed: see alert log
                               for condition(s) that are preventing upgrade,
                               Minor alert

hostname:maintenance logs alert> list

ENTRY      TIME                SUMMARY
entry-023  2022-1-16 15:06:36  d1dd862d-93ff-6e3d-a51c-fe5f81159a3e, System
                               software update cannot proceed: J4400 and J4500
                               disk shelves are not supported in this
..                            .release., Minor alert

entry-122  2022-2-17 20:34:06  65f79e6b-1a77-6041-9a92-9440dbba4c40, Finished
                               replicating project 'Test1' from appliance
                               'PatSmith'., Minor Alert
```

3. To view a specific group of log entries at one time, for example "5", use the following CLI commands:

```
hostname:maintenance logs> select audit
hostname:maintenance logs audit> list -5

ENTRY      TIME                SUMMARY
entry-721  2022-2-18 22:02:02  <system>, <system>, Request to create
replication package
entry-722  2022-2-18 22:02:08  <system>, <system>, Request to modify
replication package
entry-723  2022-2-18 23:33:32  root, 10.135.69.243, User logged in
entry-724  2022-2-19 00:48:51  root, 10.135.69.243, Browser session timed out
entry-725  2022-2-19 23:30:11  root, User logged in via CLI
```

Viewing Entry Details (CLI)

Use the following procedure to view entry details using the CLI.

- To view details of a log entry, use the following CLI commands:

```
hostname:maintenance logs> select audit
hostname:maintenance logs audit> select entry-000 show
Properties:
                    timestamp = 2022-10-15 00:59:37
                         user = root
```

```
    address = <console>
    summary = Enabled datalink:igb0 service
annotation =
```

> **Note:**
>
> The `annotation` is the session annotation, which can be enabled when configuring users. For information about user configuration, see Configuring Users in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Exporting Logs (BUI)

Use the following procedure to export logs using the BUI. You can export any combination of log types to your local machine by using the BUI **COLLECT** button.

1. To view logs, from the **Maintenance** menu, select **Logs**.
2. Click **COLLECT**.

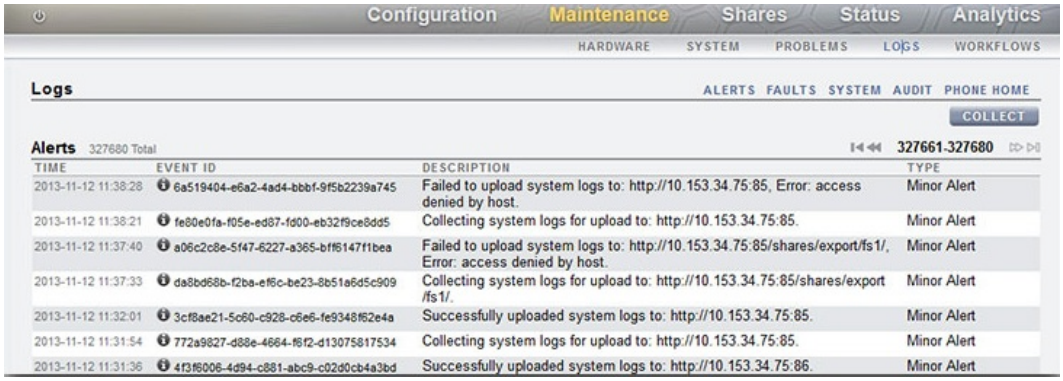

3. In the **Collect Logs** dialog box, choose the log types you want to export.

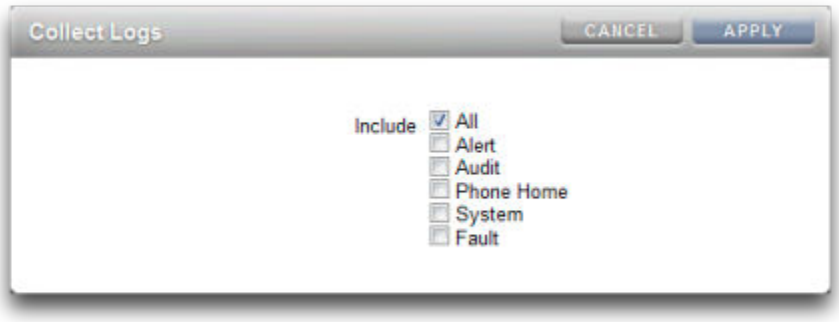

4. To export the logs, click **APPLY**. A `<stdout> gz` file containing all the requested log information is downloaded to your local machine.

Exporting Logs (CLI)

Use the following procedure to export logs using the CLI.

- To export logs, use the following CLI commands. A `<stdout> gz` file containing all the requested log information is downloaded to your local machine.

```
hostname:maintenance logs> show
Logs:

LOG          ENTRIES    LAST
alert        3458       2022-2-18 23:02:29
audit        731        2022-2-20 16:13:04
fltlog       2          2022-2-3 06:29:02
scrk         0          -
system       44         2022-1-14 18:19:59

hostname:maintenance logs> select fltlog
hostname:maintenance logs fltlog> collect

SUNW-MSG-ID: AK-8000-86, TYPE: Defect, VER: 1, SEVERITY: Major
EVENT-TIME: Sun Nov 13 03:50:15 UTC 2022
PLATFORM: Sun-Fire-X4240, CSN: 0000000000,
HOSTNAME: hpc-iwashi-01
SOURCE: ak-diagnosis, REV: 1.0
EVENT-ID: d83655cd-c03d-efff-efde-9c78dd47259e
DESC: The service processor needs to be reset to ensure proper functioning.
AUTO-RESPONSE: None.
IMPACT: Service Processor-controlled functionality, including LEDs,
fault management, and the serial console, may not work correctly.
REC-ACTION: Click the initiate repair button.

SUNW-MSG-ID: FMD-8000-4M, TYPE: Repair, VER: 1, SEVERITY: Major
EVENT-TIME: Sun Nov 13 03:50:15 UTC 2022
PLATFORM: Sun-Fire-X4240, CSN: 0000000000, HOSTNAME: hpc-iwashi-01
SOURCE: ak-diagnosis, REV: 1.0
EVENT-ID: d83655cd-c03d-efff-efde-9c78dd47259e
DESC: All faults associated with an event id have been addressed.
AUTO-RESPONSE: Some system components offlined
because of the original fault may have been brought back online.
IMPACT: Performance degradation of the system
due to the original fault may have been recovered.
REC-ACTION: No action is required.

hostname:maintenance logs fltlog>
```

Using Workflows

A workflow is a script that is uploaded to and managed by Oracle ZFS Storage Appliance by itself. For information on scripting, see Working with CLI Scripting in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*. Workflows can be parameterized and executed in a first-class fashion from either the browser user interface (BUI) or the command line interface (CLI). For information on maintenance workflows, see Maintenance Workflows in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Workflows may also be optionally executed as alert actions or at a designated time. As such, workflows allow for Oracle ZFS Storage Appliance to be extended in ways that capture specific policies and procedures, and can be used to formally encode best practices for a particular organization or application. For information on alert actions, see Configuring Alerts in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

2

Servicing the Oracle ZFS Storage Appliance Hardware Components

The following sections describe maintenance procedures for Oracle ZFS Storage Appliance.

To service the appliance hardware components, use the following sections:

3
Prerequisites for Servicing Hardware

This section describes the prerequisites for servicing hardware. See the following topics and tasks for details:

- Safety Information
- Required Tools and Information
- Powering Off the Controller
- Extending the Controller from the Rack
- Removing the Controller Top Cover
- Powering Off the Disk Shelf

Safety Information

This section contains safety information that you must follow when servicing the storage system. For your protection, observe the following safety precautions when setting up your equipment:

- Do not remove the side panels, or run the storage system with the side panels removed. Hazardous voltage is present that could cause injury. The covers and panels must be in place for proper air flow to prevent equipment damage.
- Follow all cautions, warnings, and instructions marked on the equipment and described in *Important Safety Information for Oracle's Hardware Systems* included with your system.
- Ensure that the voltage and frequency of your power source match the voltage inscribed on the electrical rating label.
- Follow the electrostatic discharge safety practices. Electrostatic discharge (ESD) sensitive devices, such as PCI cards, HDDs, SSDs, and memory cards, require special handling. Circuit boards and HDDs contain electronic components that are extremely sensitive to static electricity. Ordinary amounts of static electricity from clothing or the work environment can destroy the components located on these boards. Do not touch the components without using antistatic precautions, especially along the connector edges.

> **Note:**
>
> The Critical status indicator on the front panel flashes red when powering off and powering on the ZS5-2 or ZS3-2 storage controller. After the controller is powered on, the Power/OK status indicator is steady green, and the red Critical status indicator is off.

Required Tools and Information

The following tools are needed to service the customer-replaceable units (CRUs):

- Antistatic wrist strap - Wear an antistatic wrist strap and use an antistatic mat when handling components such as HDDs or PCI cards. When servicing or removing storage controller components, attach an antistatic strap to your wrist and then to a metal area on the chassis. Following this practice equalizes the electrical potentials between you and the storage controller.
- Antistatic mat - Place static-sensitive components on an antistatic mat.
- No. 2 Phillips screwdriver
- Torx (6 lobe) T15 screwdriver for controllers with a release lever to secure the top cover
- Nonconducting, No. 1 flat-blade screwdriver or equivalent
- Nonconducting stylus or pencil (to power on the storage controller)

Powering Off the Controller

Some replacement tasks require you to power off the controller. Note that there will be a loss of access to the storage unless the system is in a clustered configuration. The recommended methods for powering off a controller depend on whether you have completed initial configuration. Use a power off method from one of the following two situations.

If you have completed initial configuration, power off the controller using one of the following methods:

- Log in to the BUI and click the power icon ⏻ on the left side under the masthead.
- Use Secure Shell (SSH) to log in to Oracle ZFS Storage Appliance, and enter the command `maintenance system poweroff`.
- Use a pen or non-conducting pointed object to press and release the Power button on the front panel.

If you have not completed the initial configuration steps, power off the controller using one of the following methods:

- Use a pen or non-conducting pointed object to press and release the Power button on the front panel.
- Use SSH or a serial connection to log in to the service processor (SP), and then issue the command `stop /SYS`.
- To initiate an emergency shutdown during which all applications and files will be closed abruptly without saving, press and hold the Power button for at least four seconds until the Power/OK status indicator on the front panel flashes, indicating that the storage controller is in standby power mode. To completely remove power, disconnect the AC power cords from the rear panel of the storage controller.

Related Topics

- Powering Off the Disk Shelf

Extending the Controller from the Rack

Use the following procedure to extend the storage controller from the rack. If you cannot extend the controller without removing power, first see Powering Off the Controller.

1. To prevent the rack from tipping forward when the controller is extended, extend all rack anti-tilt mechanisms.
2. Verify that no cables will be damaged or will interfere when the controller is extended.

   Although controller cable management arms (CMAs) are hinged to accommodate extending the controller, you should ensure that all cables and cords are capable of extending. If necessary, label and disconnect cables that are not capable of extending, including cluster cables for clustered controllers.
3. From the front of the storage controller, release the two slide release latches.

   Depending on the controller, either squeeze the latches inward or pull them forward.
4. While holding the slide release latches, slowly pull the storage controller forward until the slide rails latch.

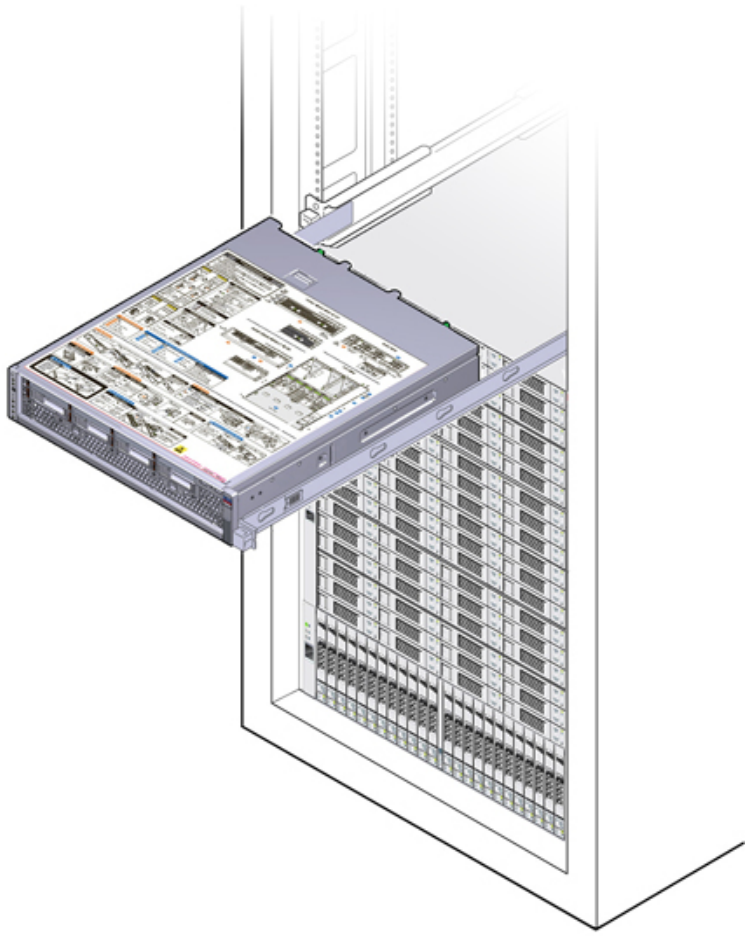

Removing the Controller Top Cover

To access internal controller components for servicing, remove the controller top cover by using the following procedure. Depending on the controller, the top cover is secured with captive screws or a release lever latch.

1. Covers with a release lever latch:

a. If the release lever latch is in the locked position, use a Torx T15 screwdriver to turn the release lever latch clockwise to the unlocked position.

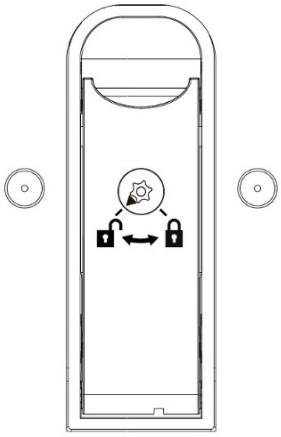

b. Unlatch the top cover [1].

   Lift up on the release lever on top of the controller cover. Lifting the release lever causes the controller cover to slide toward the rear of the chassis for easy removal.

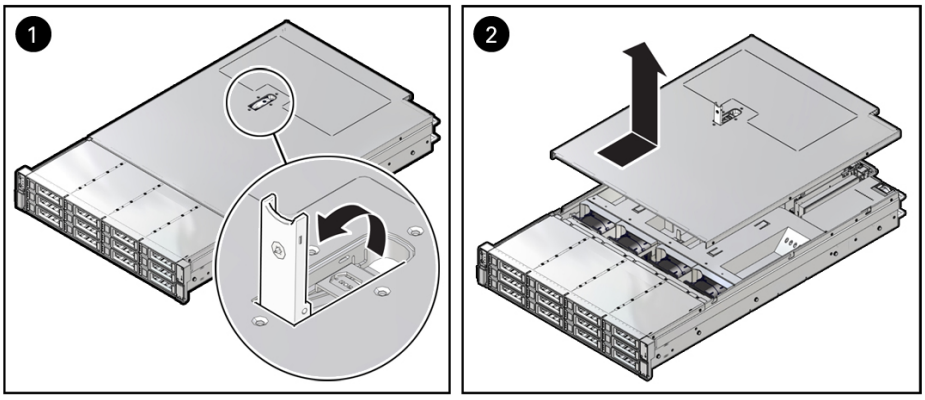

c. Lift up and remove the top cover [2].

2. Covers with captive screws:

   a. Fully loosen the two captive screws at the rear of the top cover [1].

   b. Slide the top cover rearward 0.5 in (13 mm) and lift it straight up and off the chassis [2], and set the top cover aside.

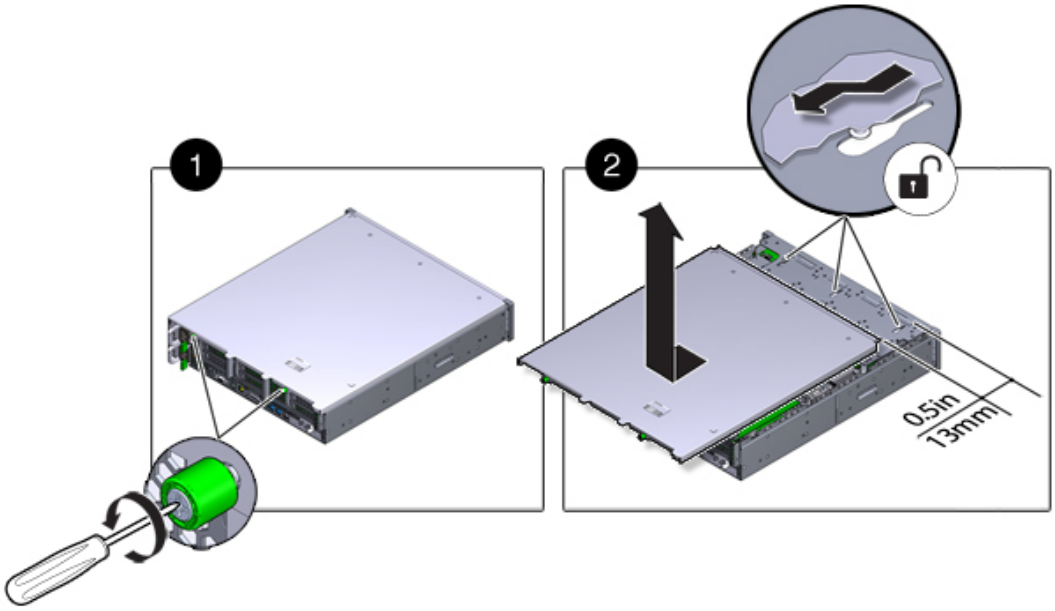

Powering Off the Disk Shelf

Most disk shelf components are hot-swappable; you do not need to remove power when replacing them. Do not remove a component if you do not have an immediate replacement. The disk shelf must not be operated without all components in place.

If a disk shelf is a part of a storage pool, powering off or removing all SAS chains from a disk shelf will trigger a kernel panic on the controller(s) to prevent data loss, unless the shelf is part of an NSPF (no single point of failure) data pool. To avoid this, shut down the controller(s) before decommissioning the shelf, as described in Powering Off the Controller.

1. Stop all input and output to and from the disk shelf.
2. Wait approximately two minutes until all disk activity indicators have stopped flashing.
3. Place the power supply on/off switches to the "0" off position.
4. Disconnect the power cords from the external power source for the cabinet.

> **Note:**
>
> All power cords must be disconnected to completely remove power from the disk shelf.

Related Topics

- Powering Off the Controller
- For details on NSPF profiles, see Storage Pool Concepts in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*

4

Servicing the Oracle ZFS Storage ZS9-2 Controller

Be sure to read sections Getting Started with Servicing the Appliance and Prerequisites for Servicing Hardware.

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see Repairing Active Problems.

To replace Oracle ZFS Storage ZS9-2 hardware components, use the following tasks:

- Replacing an Oracle ZFS Storage ZS9-2 SSD
- Replacing an Oracle ZFS Storage ZS9-2 Fan Module
- Replacing an Oracle ZFS Storage ZS9-2 Power Supply
- Replacing an Oracle ZFS Storage ZS9-2 DIMM
- Replacing or Adding an Oracle ZFS Storage ZS9-2 PCIe Card
- Replacing an Oracle ZFS Storage ZS9-2 System Battery

To understand the Oracle ZFS Storage ZS9-2 controller, use the following topics:

- Oracle ZFS Storage ZS9-2 Configuration Options
- Oracle ZFS Storage ZS9-2 Front and Rear Panel Components
- Oracle ZFS Storage ZS9-2 Specifications
- Oracle ZFS Storage ZS9-2 Internal Components
- Oracle ZFS Storage ZS9-2 PCIe Card Configuration
- Oracle ZFS Storage ZS9-2 Attached Storage

Replacing an Oracle ZFS Storage ZS9-2 SSD

SSDs are hot-swappable and can be removed and installed while the controller is running without affecting other hardware capabilities. Use the following procedure to replace an Oracle ZFS Storage ZS9-2 SSD.

> **Note:**
>
> If you have more than one failed drive, replace only one drive at a time. Removing more than one drive in quick succession will cause a hardware/pool fault. Also, do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** BUI menu, select **System**, or navigate to `maintenance system updates` in the CLI.

1. Identify the faulted SSD.

- If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, and open the detailed view of the controller. The faulted drive has either an amber or a blue-on-amber (ready to remove) status icon. Note the ID of the faulted drive. Click the entry for that drive in the table to highlight the location of the drive in the controller illustration. To view details about the faulted drive, click its information icon and then click **Active Problems**.
- If you are physically at the system, the amber Service Action Required indicator on the SSD should be lit. Additionally, you can click the locator icon in the BUI to light the locator LED on the faulted drive.

2. Verify that the faulted drive is ready to be removed.

> **Caution:**
>
> Removing the drive before it is ready to be removed can result in loss of data.

Use any of the following methods to verify that the faulted drive is ready to be removed:

- On the faulted drive, the Ready to Remove LED is lit as shown in Oracle ZFS Storage ZS9-2 Storage Drive.
- In the BUI:
  - The status icon for the faulted drive in the detail view in **Maintenance: Hardware** has changed from amber to blue-on-amber.
  - The **Component Details** information window (click the information icon) for the faulted drive shows **Yes** as the value for the **Ready to remove** property.
- In the CLI (and RESTful API), the value of the `readytoremove` property for the faulted drive is `true`:

```
hostname:maintenance chassis-000 disk-001> get readytoremove
                readytoremove = true
```

> **Caution:**
>
> If a drive is faulted and there is no indication that the drive is ready to remove, do not remove the drive. Contact Oracle Support.

3. Push the latch release button to open the drive latch [1, 2].

> **Caution:**
>
> The latch is not an ejector. Do not open the latch too far to the right. Doing so can damage the latch.

The drive location in the following figure is for illustration purposes only.

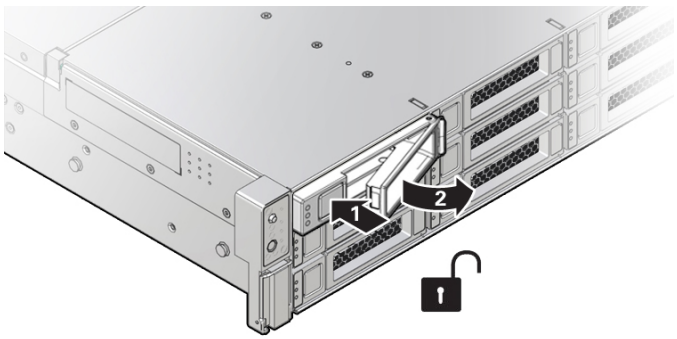

1 Pressing the latch release button.	**2** Opening the latch.

4. Grasp the latch and pull the drive out of the drive slot.

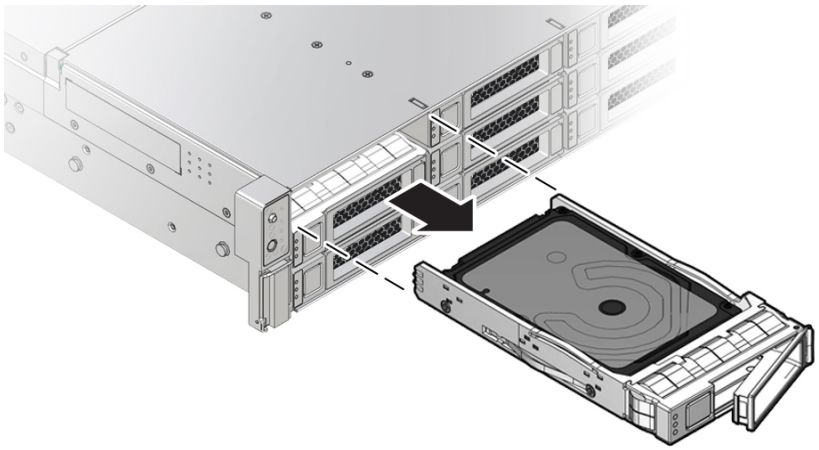

5. After at least 30 seconds, from the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon on the system controller to verify that the software has detected that the drive is not present.
6. On the replacement drive, open the release lever and slide the drive into the same slot until it is fully seated. Close the latch to lock the drive in place.

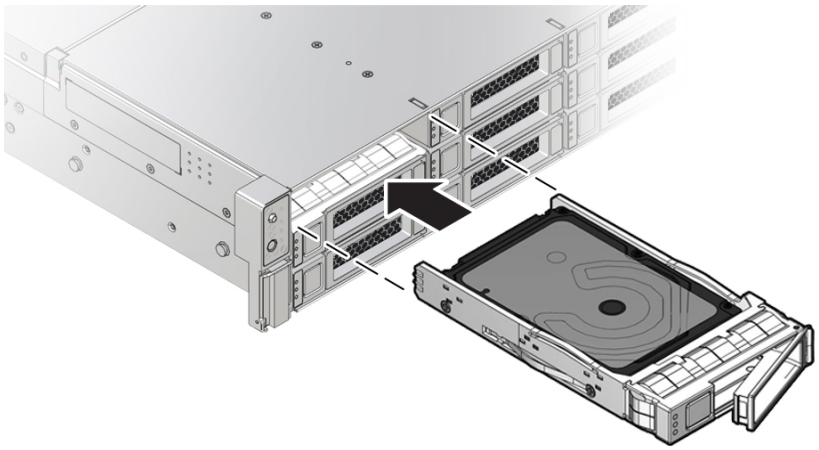

The Oracle ZFS Storage Appliance software automatically detects and configures the new drive.

7. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, and verify that the status icon is green for the newly installed drive. To view details about the new drive, click its information icon.
8. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS9-2 Fan Module

Unpack the replacement fan module and get it ready to insert into the controller chassis as soon as the faulted fan is removed.

You do not have to power off the controller to service fan modules. You should not begin this procedure unless you have a replacement fan module and you are ready to install it immediately.

> **Caution:**
>
> When performing the actual removal and replacement of a fan module in the Oracle ZFS Storage ZS9-2 controller, you must complete the entire procedure within 60 seconds to maintain adequate cooling within the system. Remove and replace only one fan module at a time. Ensure that you have obtained the replacement fan module and that it is ready for installation before starting the replacement procedure. Each fan module contains two fans, with one fan motor per fan. The two fan motors provide separate tachometer signals so that the fan module reports two tach signals to Oracle ILOM. Even if only one fan motor is faulted within the fan module, the Oracle ILOM service processor detects that two fan motors have failed to spin while the fan module is removed. If the fan module is not replaced within 60 seconds of removal, Oracle ILOM will take the protective action to shut down the system to prevent thermal damage.

Fan modules are hot-swappable and can be removed and installed while the controller is running without affecting other hardware capabilities. Use the following procedure to replace an Oracle ZFS Storage ZS9-2 fan module.

1. Identify the faulted fan module.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **Fan**, and note the ID of the faulted fan module, which has an amber icon. To highlight the location of the fan module in the controller picture, click on the fan module ID. To view details about the faulted fan module, click its information icon and then click **Active Problems**.
   - If you are physically at the system, the amber Service Action Required indicator on the fan module should be lit.
2. Extend the controller from the rack as described in Extending the Controller from the Rack.
3. To access the fan modules, remove the top cover as described in Removing the Controller Top Cover.
4. Identify the faulty fan module.

   Each fan module has a fan status indicator (LED) that is located next to the module. If the LED is amber, the fan has failed. The location of the fan status LED is shown in the following figure.

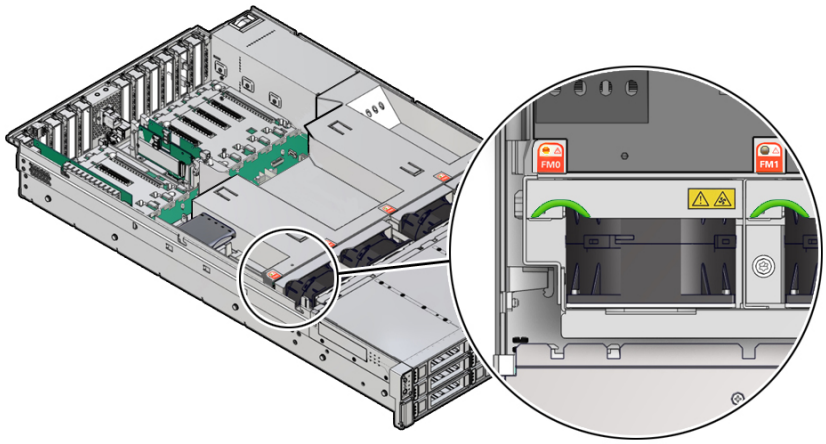

LED Color and State Meaning
Amber – The fan module is faulty. The front Top Fan LED and the front and back panel Fault-Service Required LEDs are also lit if the system detects a fan module fault.
Not Illuminated – The fan module is operational (OK).

5. Using your forefinger and thumb, lift the fan module straight up and out of the chassis and set it aside on an antistatic mat.

> **Caution:**
>
> Excessive movement or rocking of the fan module during removal can cause damage to the internal connector on the fan module board. When removing a fan module, do not rock it back and forth.

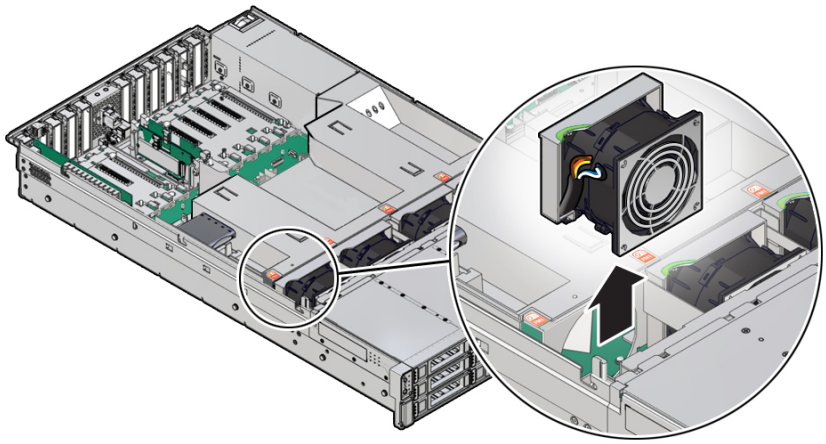

6. Position the replacement fan module so the connector on the bottom of the module is aligned with the connector on the inside of the slot and all keying and labeling are correctly positioned.

   The fan modules are keyed to ensure that they are installed in the correct orientation.

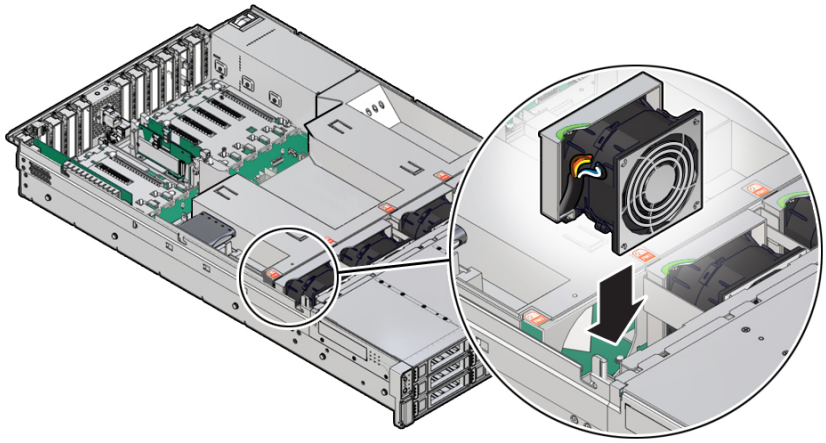

7. Press down on the fan module to fully seat the fan module.
8. Verify that the fan module status indicator (LED) on the replacement fan module is not illuminated.
9. Install the controller top cover.

   See Removing the Controller Top Cover.
10. Verify that the Top Fan Fault LED on the front of the controller, and the Fault-Service Required LEDs on the front and back of the controller are not lit.
11. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **Fan**, and verify that the status icon is green for the newly

installed fan module. To view details about the new fan module, click its information icon .

12. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS9-2 Power Supply

> **Caution:**
>
> If a power supply fails and you do not have a replacement available, do not remove the failed power supply. This maintains the proper air flow and prevents possible equipment damage.

Controllers are equipped with highly redundant, hot-swappable power supplies: Either two A269.2 1200W power supplies or two A271 1400W power supplies. The A269.2 1200W power supply accepts ranges of 100-127 volts AC (VAC) low-line or 200-240 VAC high-line, and the A271 1400W power supply only accepts ranges of 200-240 VAC high-line. High-line power supplies require high-line 200-240 VAC sources.

Use the following procedure to replace an Oracle ZFS Storage ZS9-2 power supply.

> **Caution:**
>
> Mixing A269.2 and A271 power supplies within the same chassis is not supported.

1. Identify the faulted power supply.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **PSU**, and note the ID of the faulted power supply, which has an amber icon. To highlight the location of the power supply in the controller picture, click on the power supply ID. To view details about the faulted power supply, click its information icon and then click **Active Problems**.
   - If you are physically at the system, the amber Service Action Required indicator on the power supply should be lit.

> **Note:**
>
> The fans of a failed power supply might still be spinning when the system is powered on. You can remove a power supply while the fans are still spinning. The fans will stop spinning when you disconnect the power cords.

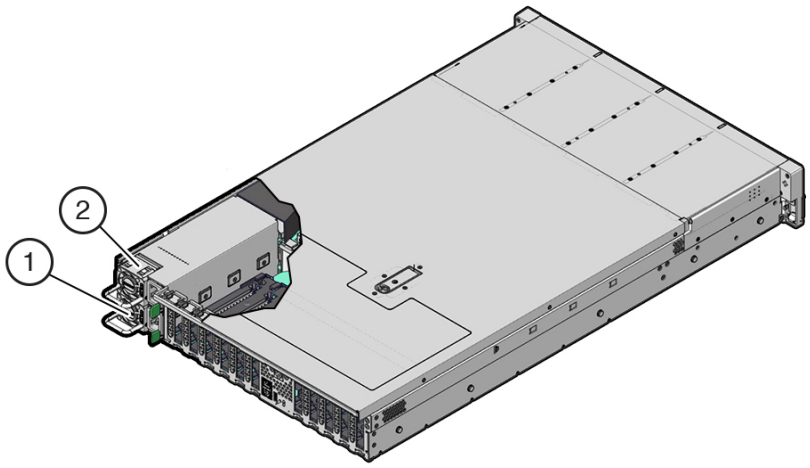

1 Power supply 0	**2** Power supply 1

2. Gain access to the rear of the controller where the failed power supply is located.
3. If the cable management arm (CMA) is installed, disconnect both CMA left-side connectors and move the CMA out of the way.

   For the location of the CMA left-side connectors, see Installing an Oracle ZFS Storage ZS9-2 Cable Management Arm in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*.

> **Caution:**
>
> When disconnecting the CMA left-side connectors, be sure to use your arm to support the CMA so that it does not hang down under its own weight and stress the right-side connectors; otherwise, the CMA could be damaged. You must continue to support the CMA until you have reconnected both of the left-side connectors.

4. Disconnect the power cord from the failed power supply.
5. Grasp the power supply handle and push the power supply latch to the left.

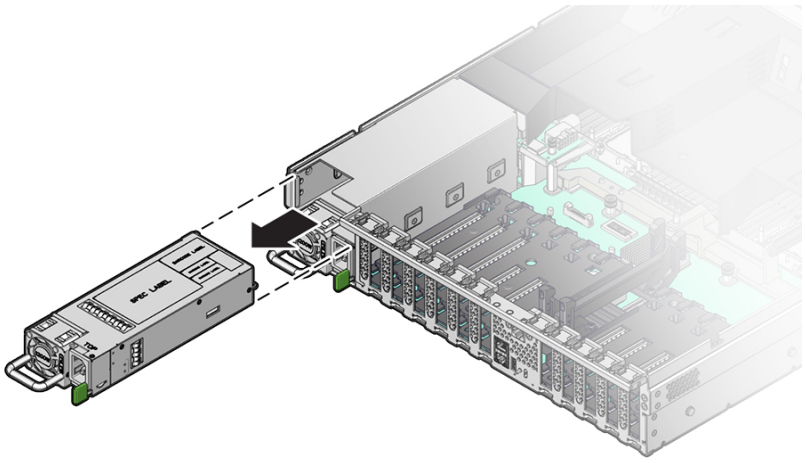

6. Pull the power supply out of the controller.

> **Caution:**
>
> Whenever you remove a power supply, you should replace it with another power supply; otherwise, the controller could overheat due to improper airflow.

7. Remove the replacement power supply from its packaging, and place it on an antistatic mat.

> **Caution:**
>
> Always replace the failed power supply with the same type model of power supply.

8. Align the replacement power supply with the empty power supply slot.
9. Slide the power supply into the slot until it is fully seated.

   You will hear or feel a click when the power supply fully seats.

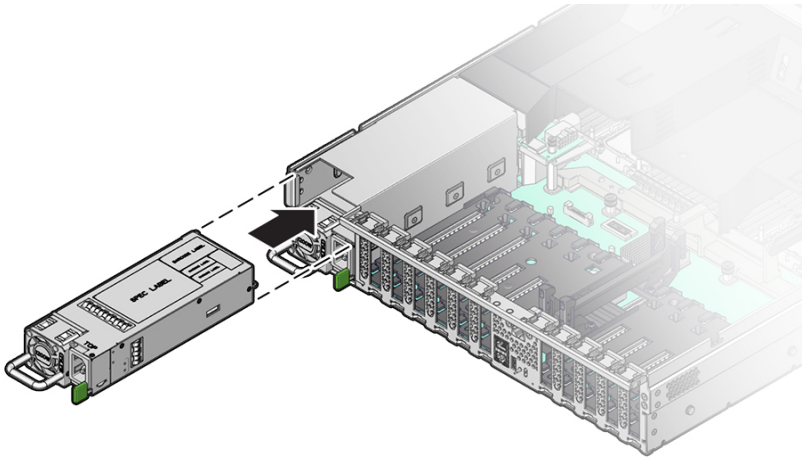

10. Reconnect the power cord to the power supply.
11. Verify that the amber Fault-Service Required LED on the replaced power supply, and the Fault-Service Required LEDs are not lit on the front and back panels.

> **Note:**
>
> After you have replaced Power Supply 0, you might need to reset the Oracle ILOM service processor (SP) to propagate the key identity properties (KIP) data to the new power supply. For instructions on resetting the SP, refer to the *Oracle Integrated Lights Out Manager (ILOM) 5.0 Configuration and Maintenance Guide Firmware Release 5.0.x* at https://www.oracle.com/goto/ilom/docs . Power Supply 1 does not contain KIP data, and therefore does not require an SP reset after replacement.

12. If you disconnected the two CMA left-side connectors, reconnect the connectors.
13. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **PSU**, and verify that the status icon is green for the newly installed power supply. To view details about the new power supply, click its information icon.
14. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS9-2 DIMM

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

Removing a DIMM is a cold-service operation. You *must* power off the appliance before you remove the DIMM. There will be a loss of access to the storage unless the system is in a clustered configuration.

> **Caution:**
>
> Ensure that all power is removed from the controller before removing or installing DIMMs, or damage to the DIMMs might occur. You must disconnect all power cables from the system before performing these procedures.

1. Identify the faulted DIMM.

   From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **DIMM**, and note the ID of the faulted DIMM, which has an amber icon. To highlight the location of the memory module in the controller

picture, click on the module ID. To view details about the faulted module, click its information icon and then click **Active Problems**.

2. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
3. Disconnect the AC power cords from the rear panel of the controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

4. Extend the controller from the rack as described in Extending the Controller from the Rack.
5. Remove the top cover as described in Removing the Controller Top Cover.
6. Remove the air baffle by lifting the baffle up and out of the controller.

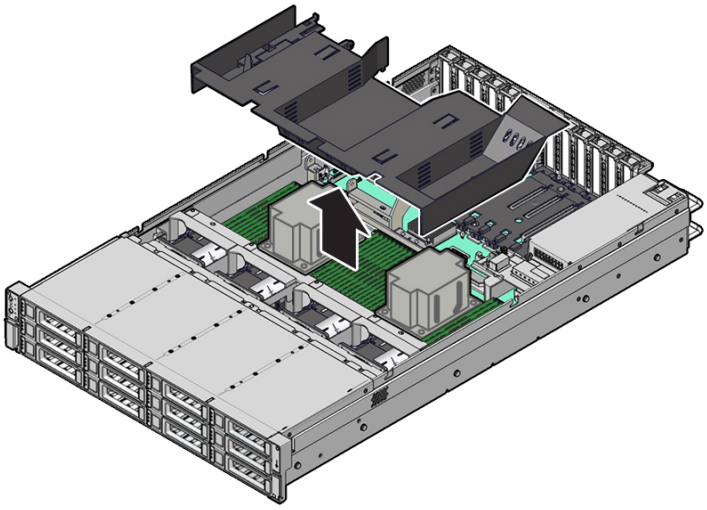

7. Locate the Fault Remind button.

   When the controller Fault Remind button [1] is pressed, an LED located next to the Fault Remind button lights green to indicate that there is sufficient voltage present in the fault remind circuit to light any fault LEDs that were lit due to a component failure. If this LED does not light when you press the Fault Remind button, it is likely that the capacitor powering the fault remind circuit has lost its charge. This can happen if the Fault Remind button is pressed for several minutes with fault LEDs lit or if power has been removed from the controller for more than 15 minutes. The following figure shows the location of the Fault Remind button.

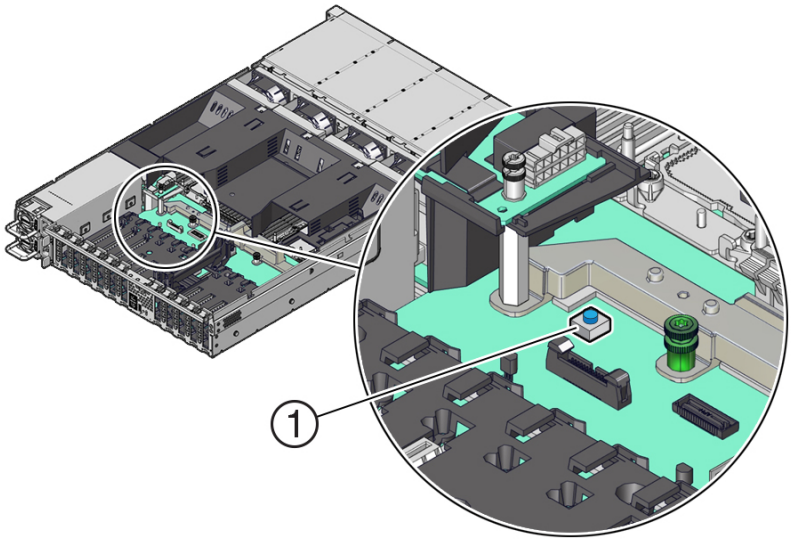

8. Physically identify and note the location of the failed DIMM by pressing the Fault Remind button on the motherboard.

   Failed DIMMs are identified with a corresponding amber LED on the motherboard.

   - If the DIMM Fault LED is off, then the DIMM is operating properly.
   - If the DIMM Fault LED is on (amber), then the DIMM is failed and should be replaced [1].

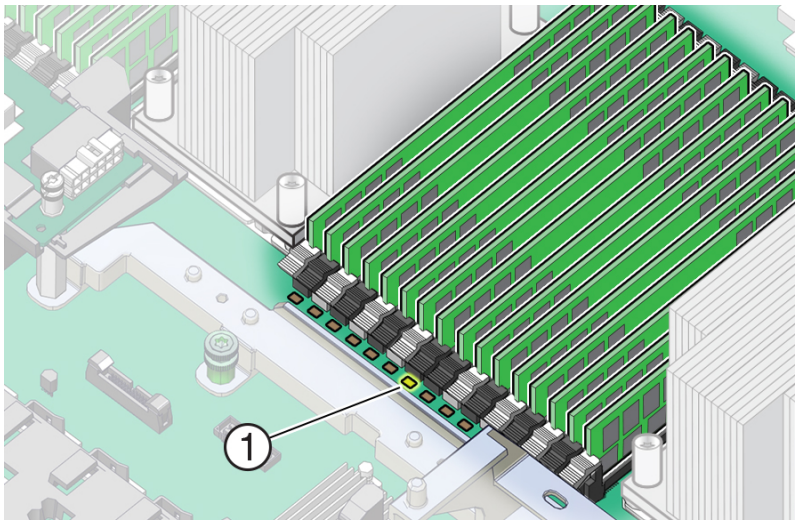

9. Rotate both DIMM slot ejector levers outward as far as they will go.

   The DIMM is partially ejected from the slot.

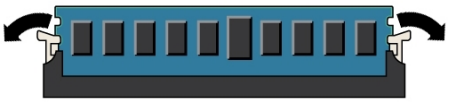

10. Carefully lift the DIMM straight up to remove it from the slot.

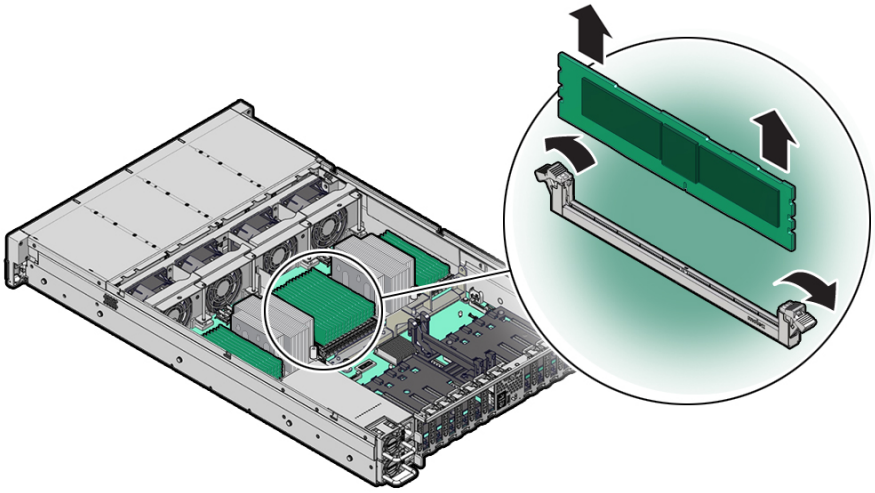

11. Repeat the previous steps until all faulty DIMMs are removed.
12. To install a replacement DIMM, ensure that the DIMM ejector levers at both ends of the DIMM slot are in their fully open position.

    Ensure that the replacement DIMM matches the size, type, and rank of the DIMM it is replacing.

13. Align the notch in the replacement DIMM with the connector key in the connector slot.

    The notch ensures that the DIMM is oriented correctly.
14. Push the DIMM into the connector slot until the ejector tabs lock the DIMM in place.

    If the DIMM does not easily seat into the connector slot, verify that the notch in the DIMM is aligned with the connector key in the connector slot. If the notch is not aligned, damage to the DIMM might occur.

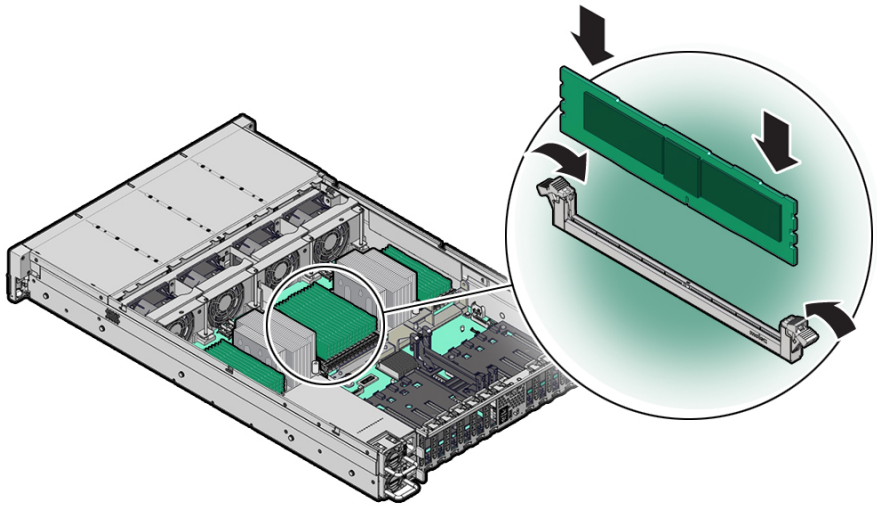

15. Ensure that the levers have risen completely and have locked the DIMM in the slot.

16. Repeat the previous steps until all replacement DIMMs are installed.
17. Install the air baffle by placing it into the controller and lowering it to its inserted position.

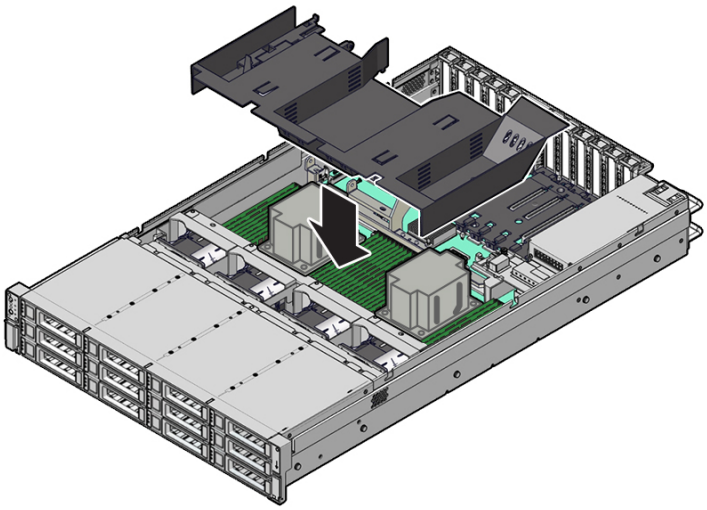

18. Replace the top cover on the chassis.
19. Return the controller to the normal rack position by pulling the green release tabs on the side of each rail while slowly pushing the controller into the rack.

20. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
21. Reconnect the power cords to the power supplies.
22. Verify that standby power is on. Approximately two minutes after plugging in the power cords, the Power/OK status indicator on the front panel starts flashing.
23. After the Power/OK status indicator is flashing, press and release the recessed Power button on the controller front panel.

    The Power/OK status indicator continues flashing during the startup sequence, which takes several minutes. When the system is operational, the indicator remains lit.
24. After the Power/OK status indicator is lit steadily, log in, and from the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **DIMM**, and verify that the status icon is green for the newly installed module. To view details about the new module, click its information icon.
25. If a fault does not self-clear, see Repairing Active Problems.

Replacing or Adding an Oracle ZFS Storage ZS9-2 PCIe Card

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

Replacing or adding a PCIe card is a cold-swap procedure, and you must remove power from the controller. If you are installing a newly released adapter, upgrade your system software before installing the adapter as described in Upgrading the Software. Use the following procedure to replace or add an Oracle ZFS Storage ZS9-2 PCIe card.

1. Identify the PCIe card slot location.
   - If replacing a faulted PCIe card, identify the PCIe card to be replaced.

     From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **Slot**, and note the ID of the faulted PCIe card, which has an amber icon. If replacing a non-faulted PCIe card, note its location. To highlight the location of the card in the controller picture, click on the card ID. To view details about the card, click its information icon and then click **Active Problems**.
   - If adding a new PCIe card, identify the correct slot location as described in Oracle ZFS Storage ZS9-2 PCIe Card Configuration.
2. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
3. Disconnect the AC power cords from the rear panel of the controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

4. Extend the controller from the rack as described in Extending the Controller from the Rack.
5. Remove the top cover as described in Removing the Controller Top Cover.
6. Locate the PCIe card position in the controller.

   There are 10 PCIe slots available for PCIe cards. The following figure shows the PCIe slot numbering (Oracle ZFS Storage ZS9-2 high-end model shown):

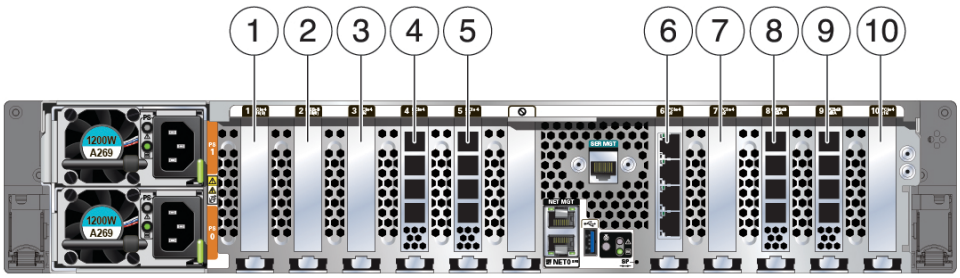

1 PCIe slot 1: Second PCIe option	**6** PCIe slot 6: Cluster interface Ethernet adapter
2 PCIe slot 2: Third PCIe option	**7** PCIe slot 7: First PCIe option
3 PCIe slot 3: Fifth PCIe option	**8** PCIe slot 8: • High-end model: External SAS-3 HBA • Mid-range model: Filler panel; slot not available
4 PCIe slot 4: External SAS-3 HBA	**9** PCIe slot 9: External SAS-3 HBA
5 PCIe slot 5: • High-end model: External SAS-3 HBA • Mid-range model: Filler panel; slot not available	**10** PCIe slot 10: Fourth PCIe option

> **Note:**
>
> All of the PCIe slots comply with the PCI Express 4.0 specification and can accommodate 75 Watt PCIe cards.

7. Disconnect all cables from the PCIe card.

   Note the cable connections to ease proper reconnection of the cables. If the PCIe card has transceivers, remove them before removing the PCIe card from the controller.
8. Rotate the PCIe card locking mechanism to an upright position [1], and then lift up on the PCIe card to disengage it from the motherboard connectors [2].

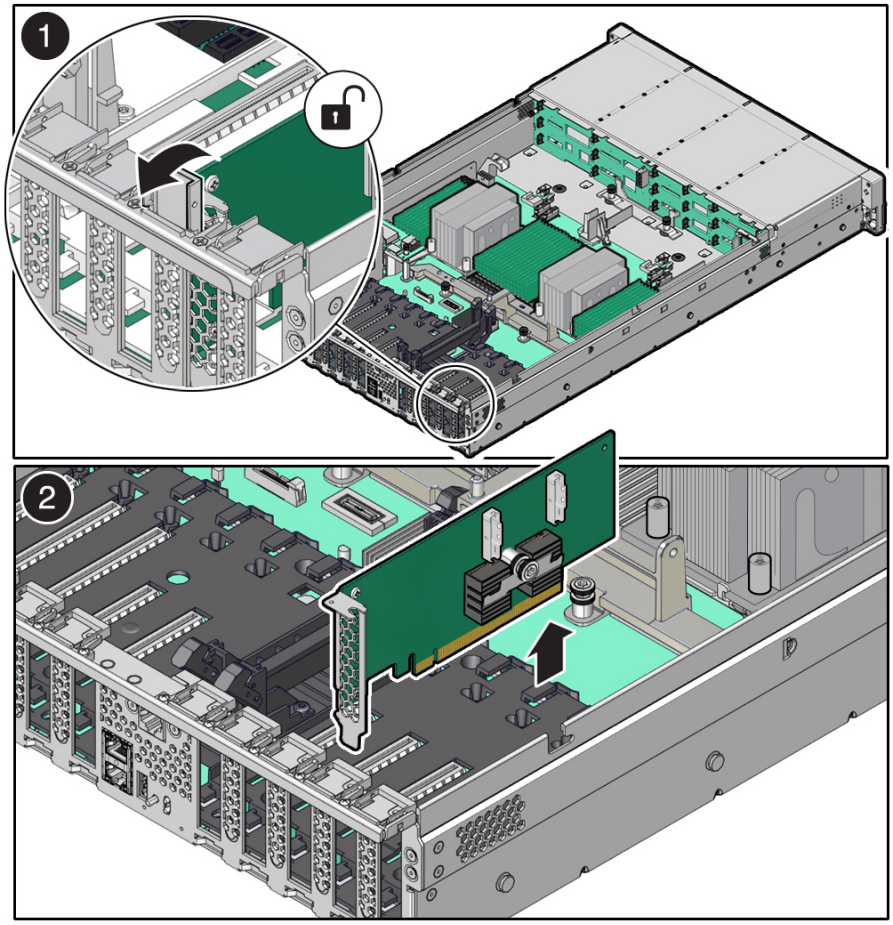

9. Place the PCIe card on an antistatic mat.

> **Caution:**
>
> If you are not immediately inserting a replacement PCIe card into the empty slot, insert a PCIe filler panel in the slot to reduce the possibility of radiated electromagnetic interference (EMI).

10. Insert the PCIe card into the correct slot [1], and rotate the PCIe locking mechanism downward to secure the PCIe card in place [2].

    You will hear or feel a click when the PCIe card is secured into the slot.

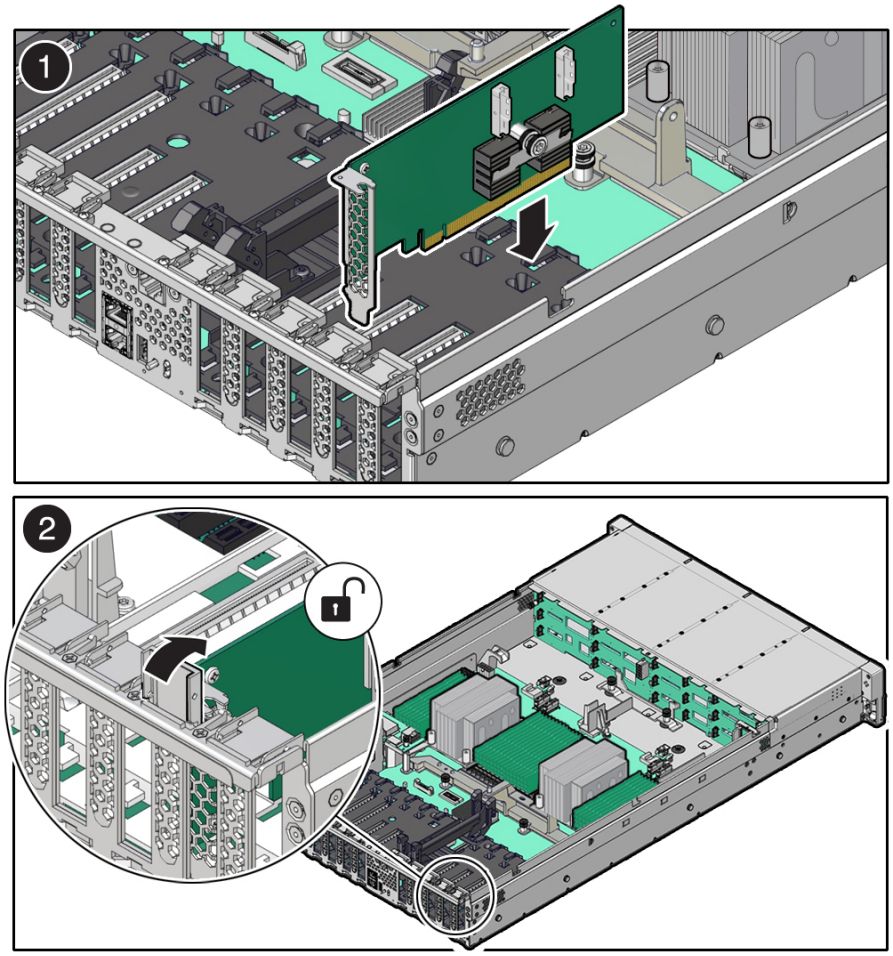

11. Replace the top cover on the chassis.
12. Return the controller to the normal rack position by pulling the green release tabs on the side of each rail while slowly pushing the controller into the rack.

    If the installed PCIe card requires transceivers, install them now.
13. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.

    If you replaced an HBA, connect the disk shelves to its ports as described in the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
14. Reconnect the power cords to the power supplies.
15. Verify that standby power is on. Approximately two minutes after plugging in the power cords, the Power/OK status indicator on the front panel starts flashing.
16. After the Power/OK status indicator is flashing, press and release the recessed Power button on the controller front panel.

    The Power/OK status indicator continues flashing during the startup sequence, which takes several minutes. When the system is operational, the indicator remains lit.
17. After the Power/OK status indicator is lit steadily, log in, and from the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the

controller, click **Slot**, and verify that the status icon is green for the newly installed card. To view details about the new card, click its information icon.

Replacing an Oracle ZFS Storage ZS9-2 System Battery

Replacing a system battery is a cold-swap procedure, and you must remove power from the controller. You might need to replace the battery if the controller fails to maintain the proper time when powered off and disconnected from the network. Use the following procedure to replace an Oracle ZFS Storage ZS9-2 system battery.

1. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. Remove the top cover as described in Removing the Controller Top Cover.
5. To dislodge the battery from its retainer, use your finger to gently push the top edge of the battery away from the retainer.

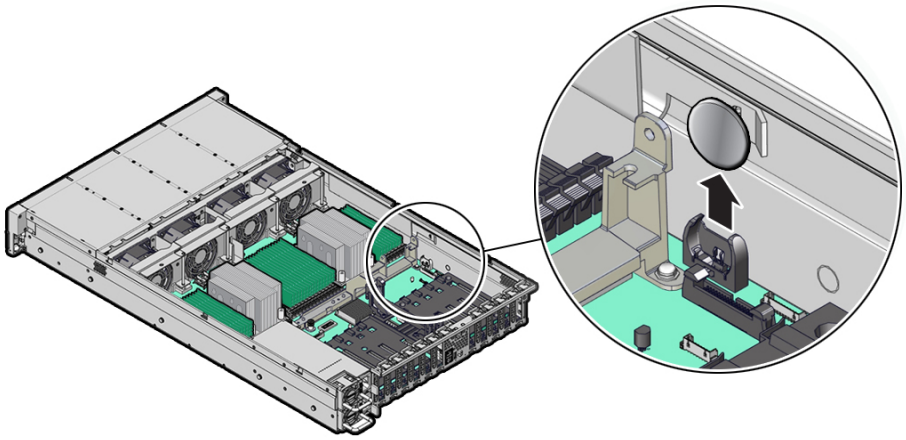

6. Lift the battery up and out of its retainer.
7. Unpack the replacement battery.
8. Press the new battery into the battery retainer with the positive (+) side facing the metal tab that holds it in place.
9. Replace the top cover on the chassis.
10. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the controller into the rack.

11. Connect the power cords to the power supplies.
12. Verify that standby power is on. Approximately two minutes after plugging in the power cords, the Power/OK status indicator on the front panel starts flashing.
13. After the Power/OK status indicator is flashing, press and release the recessed Power button on the controller front panel.

    The Power/OK status indicator continues flashing during the startup sequence, which takes several minutes. When the system is operational, the indicator remains lit.
14. After the Power/OK status indicator is lit steadily, log in and set the time using the steps in the BUI Clock task. For more information, see Setting Clock Synchronization (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Oracle ZFS Storage ZS9-2 Controller Hardware Overview

This section describes the internal and external components of the Oracle ZFS Storage ZS9-2 controller.

- Oracle ZFS Storage ZS9-2 Configuration Options
- Oracle ZFS Storage ZS9-2 Front and Rear Panel Components
- Oracle ZFS Storage ZS9-2 Specifications
- Oracle ZFS Storage ZS9-2 Internal Components
- Oracle ZFS Storage ZS9-2 PCIe Card Configuration
- Oracle ZFS Storage ZS9-2 Attached Storage

Oracle ZFS Storage ZS9-2 Configuration Options

The Oracle ZFS Storage ZS9-2 high-end model incorporates one or two controllers and one to forty-eight disk shelves. Each controller has 10 PCIe Gen4 slots in the rear, and 12 3.5" disk slots in the front, two of which are occupied by system disks. The remaining ten front disk slots have permanent filler panels. The following table describes the configuration options.

CPU	Memory	System Disks	PCIe Slots	HBA
2x32-core, 2.6 GHz	2.0 TB	2x3.5" NVMe	Five dedicated; five available	4-port (4x4) SAS-3 12 Gb/s external

The Oracle ZFS Storage ZS9-2 mid-range model incorporates one or two controllers and one to twenty-four disk shelves. Each controller has 10 PCIe Gen4 slots in the rear, and 12 3.5" disk slots in the front, two of which are occupied by system disks. The remaining ten front disk slots have permanent filler panels. The following table describes the configuration options.

CPU	Memory	System Disks	PCIe Slots	HBA
2x24-core, 2.0 GHz	512 GB or 1.0 TB	2x3.5" NVMe	Three dedicated; five available; two unavailable (permanent filler panels)	4-port (4x4) SAS-3 12 Gb/s external

The most recent component specification is located at https://www.oracle.com/storage/nas/index.html .

Oracle ZFS Storage ZS9-2 Front and Rear Panel Components

Oracle ZFS Storage ZS9-2 Front Panel Components - The Oracle ZFS Storage ZS9-2 controller drive slots and front panel components are shown in the following figure.

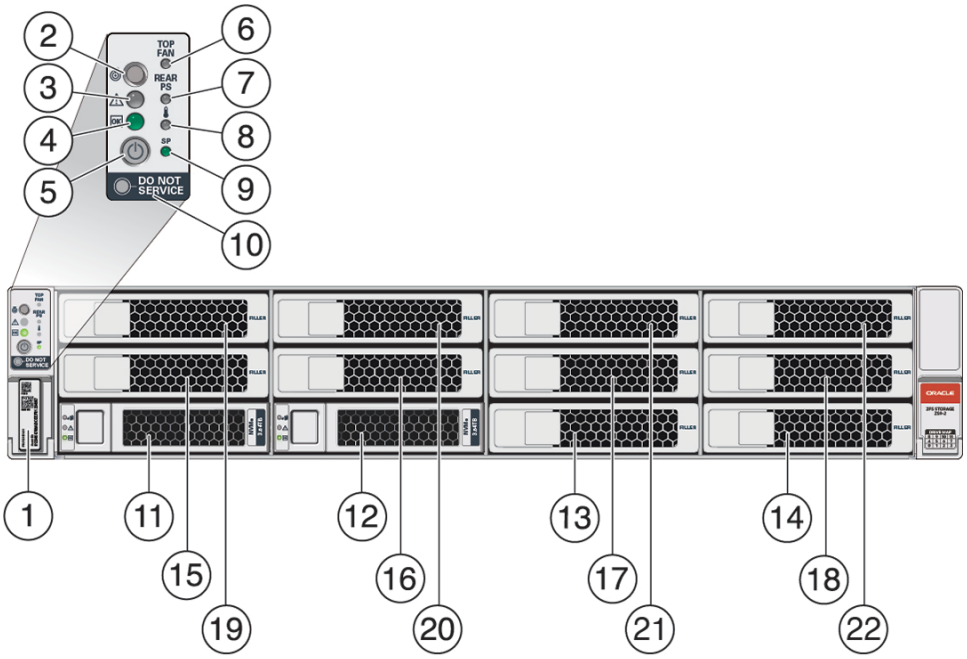

1 RFID tag	**2** System status LED: Locator (white)	**3** System status LED: Service Action Required (amber)
4 System status LED: Power/OK (green)	**5** Power button	**6** Fan service action required (amber)
7 Power Supply (PS) service action required (amber)	**8** Over-temperature warning (amber)	**9** Service Processor (SP) OK (green)
10 Do Not Service indicator	**11** System disk 0	**12** System disk 1
13 Filler panel	**14** Filler panel	**15** Filler panel
16 Filler panel	**17** Filler panel	**18** Filler panel
19 Filler panel	**20** Filler panel	**21** Filler panel

22 Filler panel

Oracle ZFS Storage ZS9-2 Rear Panel Components - The Oracle ZFS Storage ZS9-2 controller PCIe slots and rear panel components are shown in the following figure (Oracle ZFS Storage ZS9-2 high-end model shown).

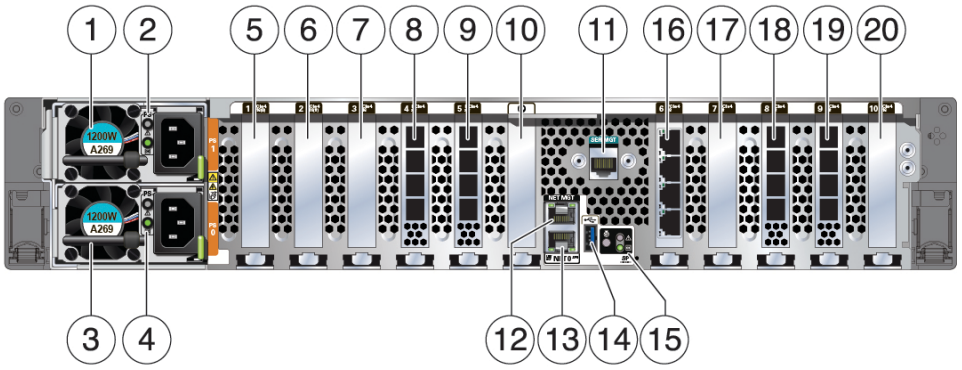

1 Power Supply Unit (PSU) 1	**2** PSU 1 status indicators (top to bottom): • Attention (amber) • Power/OK (green)	**3** PSU 0
4 PSU 0 status indicators (top to bottom): • Attention (amber) • Power/OK (green)	**5** Second PCIe option (slot 1)	**6** Third PCIe option (slot 2)
7 Fifth PCIe option (slot 3)	**8** External SAS-3 HBA (slot 4)	**9** Slot 5: • High-end model: External SAS-3 HBA • Mid-range model: Filler panel; slot not available
10 Not available	**11** Serial management (SER MGT) RJ-45 serial port	**12** Service Processor (SP) NET MGT port
13 Network (NET) 100/1000BASE-T RJ-45 Gigabit Ethernet (GbE) port: NET 0	**14** USB 3.0 port	**15** System status LEDs: • Locate (white) - left • Attention (amber) - right, top • Power/OK (green) - right, bottom
16 Cluster interface Ethernet adapter (slot 6)	**17** First PCIe option (slot 7)	**18** Slot 8: • High-end model: External SAS-3 HBA • Mid-range model: Filler panel; slot not available
19 External SAS-3 HBA (slot 9)	**20** Fourth PCIe option (slot 10)	

Cluster Interface Card Compatibility - Oracle ZFS Storage ZS9-2 controllers employ Ethernet-based clustering via two Ethernet ports in the Oracle Quad Port 10GBASE-T

Ethernet Adapter. All other controller cluster interface cards use serial-based clustering via two serial cluster links, and provide Ethernet connectivity via one link. Therefore, the Oracle ZFS Storage ZS9-2 cluster interface card is not compatible with older controller models.

For information about how to connect cables to form a cluster, see Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

Ethernet Port

The Oracle ZFS Storage ZS9-2 controller has one RJ-45 100/1000BASE-T Gigabit Ethernet (GbE) network connector, labeled NET0, on the rear panel. The port operates at 100 Mbits/sec or 1000 Mbits/sec. Use this port to connect the appliance to the network.

There are two status indicators (LEDs) on this port. These indicators are visible from the rear of the controller.

Table 4-1 Gigabit Ethernet Port Status Indicators

Status Indicator Name	Location	Color	State and Meaning
Activity	Bottom left	Green	ON – Link up OFF – No link FLASHING – Packet activity
Link speed	Bottom right	Green	OFF – 100BASE-T link (if link up) ON – 1000BASE-T link

Network Management Port

The Oracle ZFS Storage ZS9-2 controller has one 10/100/1000BASE-T Ethernet management domain interface, labeled NET MGT. There are two status indicators (LEDs) on this port. These indicators are visible from the rear of the controller. The service processor NET MGT port is the optional connection to the Oracle ILOM service processor. The NET MGT port is configured by default to use Dynamic Host Configuration Protocol (DHCP). The service processor NET MGT port uses an RJ-45 cable for a 10/100/1000BASE-T connection.

Table 4-2 Network Management Port Status Indicators

Status Indicator Name	Location	Color	State and Meaning
Activity	Top left	Green	ON – Link up OFF – No link FLASHING – Packet activity
Link speed	Top right	Green	ON – 1000BASE-T link OFF – 10/100BASE-T link

Serial Management Port

The service processor SER MGT port uses an RJ-45 cable and is the default connection to the Oracle ILOM service processor. This port supports local connections to the controller and only recognizes Oracle ILOM command-line interface (CLI) commands. Typically you connect a terminal or terminal emulator to this port.

> **Note:**
>
> This port does not support network connections.

4x4 SAS HBA Ports

The 4x4 SAS HBAs are installed in PCIe slots 1, 2, 8, and 9 for the Oracle ZFS Storage ZS9-2 high-end model, and slots 1 and 9 for the Oracle ZFS Storage ZS9-2 mid-range model. The HBAs provide connectivity to external disk shelves. The HBA logical ports are numbered 0-3, bottom to top, as shown in the following figure.

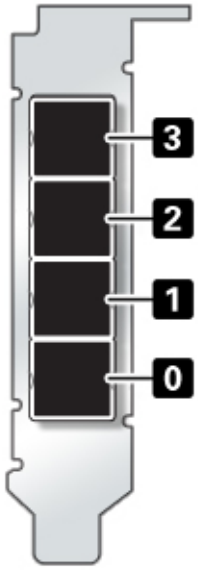

Oracle ZFS Storage ZS9-2 Specifications

Oracle ZFS Storage ZS9-2 Physical Specifications: The Oracle ZFS Storage ZS9-2 controller chassis dimensions are as follows:

Dimension	Measurement	Dimension	Measurement
Height	3.42 in./86.9 mm	Depth	30.04 in./772 mm
Width	18.97 in./482 mm	Weight (Maximum)	48 lb./21.41 kg

Oracle ZFS Storage ZS9-2 Electrical Specifications: The Oracle ZFS Storage ZS9-2 controller has two highly redundant, hot-swappable, rear accessible power supplies with N+N redundancy: Either two A269.2 1200W power supplies or two A271 1400W power supplies. The A269.2 1200W power supply accepts ranges of 100-127 volts AC (VAC) low-line or 200-240 VAC high-line, and the A271 1400W power supply only accepts ranges of 200-240 VAC high-line. High-line power supplies require high-line 200-240 VAC sources. To understand the overall power consumption, see Oracle's online power calculator at http://www.oracle.com/goto/powercalculators.

> **Caution:**
>
> Mixing A269.2 and A271 power supplies within the same chassis is not supported.

Power Supply Parameter	Model A269.2	Model A271
Electrical ratings	100-127/200-240V~, 10/7A, 50/60Hz (x2)	200-240V~, 10A, 50/60Hz (x2)
Voltage (nominal)	100-127 VAC; 200-240 VAC	200-240 VAC
Input current (maximum)	10.0A at 100-127 VAC; 7.0A at 200-240 VAC	10.0A at 200-240 VAC
Frequency (nominal)	50/60 Hz (47-63 Hz range)	50/60 Hz (47-63 Hz range)
Maximum power consumption	800W at AC 100V-127V; 1200W at AC 200V-240V	1400W at AC 200V-240V

Controller Parameter	Corresponding Entry in Power Calculator
Maximum power consumption of controller	Refer to the Power Calculator (https://www.oracle.com/goto/powercalculators) and see Oracle ZFS Storage ZS9-2 for Watt calculations
Maximum heat output of controller	Refer to the Power Calculator (https://www.oracle.com/goto/powercalculators) and see Oracle ZFS Storage ZS9-2 for BTU/hr calculations

Oracle ZFS Storage ZS9-2 Acoustic Noise Emissions: The workplace-dependent noise level of this product is 9.5 Bels A-weighted Operating, and 8.0 Bels A-weighted Idling (measured Sound Power).

Oracle ZFS Storage ZS9-2 Internal Components

The major components of the Oracle ZFS Storage ZS9-2 controller are shown in the following figure and identified in the legend.

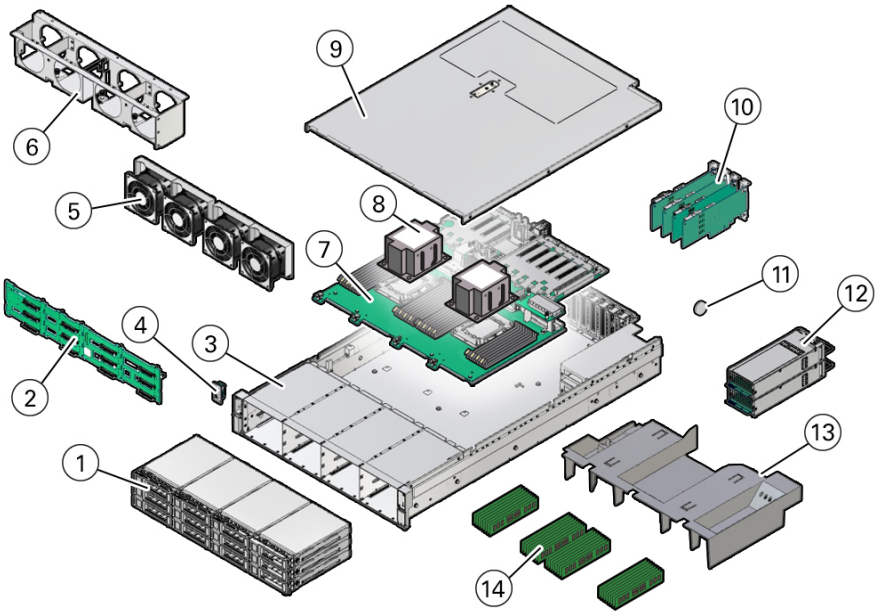

1 Two system disks, ten filler panels	**8** Processors and heatsinks
2 Disk backplane	**9** Top cover
3 System chassis	**10** PCIe cards
4 Front LED indicator module and temperature sensor	**11** System battery
5 Fan modules	**12** Power supplies
6 Fan tray	**13** Air baffle
7 Motherboard assembly	**14** DIMMs

Oracle ZFS Storage ZS9-2 Storage Drive

This section shows the location and explains the functions of the storage drive status indicators.

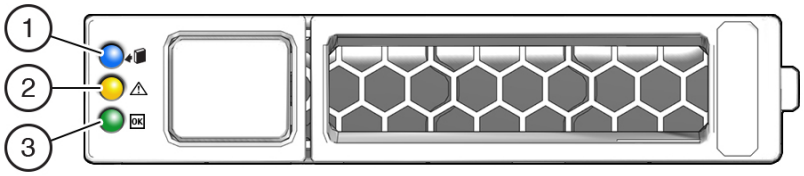

1 Ready to Remove (blue)	**2** Service Action Required (amber)	**3** OK/Activity (green)

Oracle ZFS Storage ZS9-2 Fan Module

The four fan modules are located towards the front of the Oracle ZFS Storage ZS9-2 controller between the CPUs and the drive trays. The fan modules are designated

from left to right. As pictured in the following illustration, the four fan modules are designated as: FMO, FM1, FM2, and FM3.

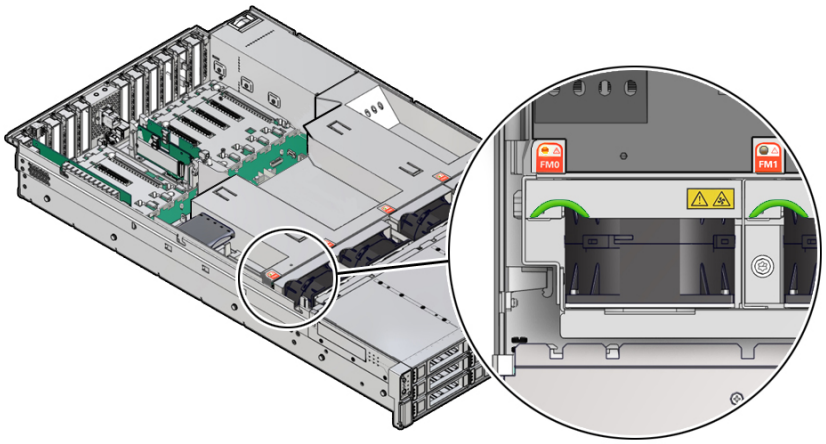

LED Color and State Meaning
Amber – The fan module is faulty. The front Top Fan LED and the front and back panel Fault-Service Required LEDs are also lit if the system detects a fan module fault.
Not Illuminated – The fan module is operational (OK).

Oracle ZFS Storage ZS9-2 Power Supply

The two power supplies are located on the right side of the controller (from the front of the controller) and are designated from top to bottom. The power supplies are accessible from the back of the controller. As indicated in the following drawing, the top designation is PS 1 and the bottom designation is PS 0.

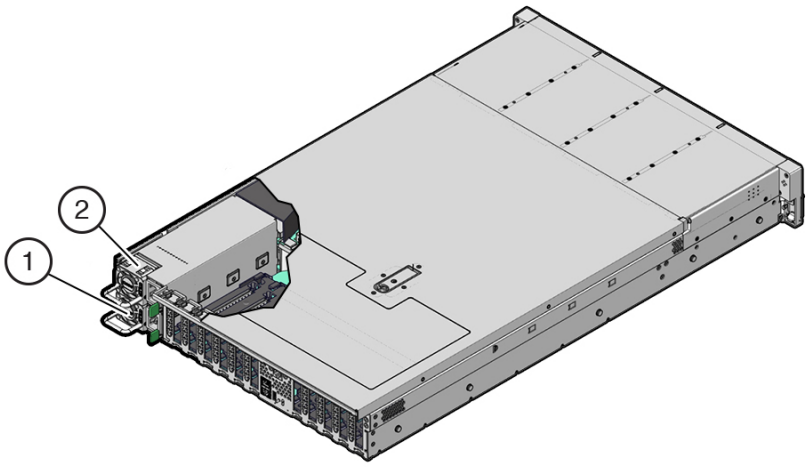

1 Power supply 0	**2** Power supply 1

The two power supplies ensure redundancy by allowing the controller to operate when one of the power supplies fails. However, when the controller is operating on one power supply, redundancy no longer exists, and the risk for an unexpected shut down and a loss of data is high. When a power supply or any component that is part of a redundant configuration fails, replace it immediately.

1 Power supply (PS) 1
2 PS 1 status indicators: Fault-Service Required LED: amber; AC power input and DC power output OK LED: green
3 PS 0
4 PS 0 status indicators: Fault-Service Required LED: amber; AC power input and DC power output OK LED: green

Oracle ZFS Storage ZS9-2 DIMM

Each processor, P0 and P1, has 16 DIMM slots organized into 8 memory channels. Each memory channel contains two DIMM slots: a black DIMM slot (channel slot 0) and a white DIMM slot (channel slot 1).

The physical layout of the DIMMs and processor(s) is shown in the following figure. When viewing the controller from the front, processor 0 (P0) is on the left.

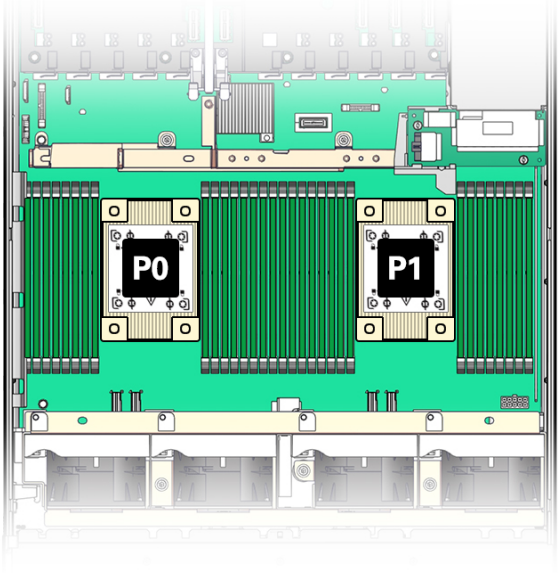

The following table shows memory channels and their associated DIMM slots.

Memory Channels	DIMM Slot 0 (Black)	DIMM Slot 1 (White)
0	D9	D8
1	D11	D10
2	D13	D12
3	D15	D14
4	D6	D7
5	D4	D5
6	D2	D3
7	D0	D1

As shown in the following figure, DIMM slots are labeled D0 through D7 to the left of each processor, and D8 through D15 to the right of each processor. For each processor, the black slots are: D0, D2, D4, D7, D9, D11, D13, and D15. Similarly, for each processor, the white slots are: D1, D3, D5, D6, D8, D10, D12, and D14.

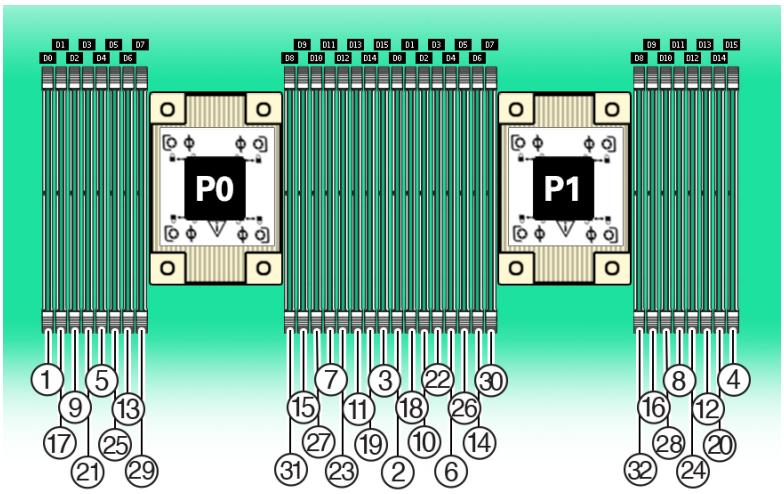

The following table describes the supported memory configurations for Oracle ZFS Storage ZS9-2. The mid-range model supports 8 or 16 DIMMs, and the high-end model supports 32 DIMMs. All DIMMs are 64 GB capacity. The DIMM locations in the table apply to both P0 and P1 processors; for example, P0/D2,P0/D6,P0/D9,P0/D13 and P1/D2,P1/D6,P1/D9,P1/D13.

Configuration	Total DIMMs Installed	DIMM Slots Populated	DRAM Capacity	Description
Oracle Storage ZS9-2 mid-range model, low-memory configuration	8	D2, D6, D9, D13	512 GB	Black sockets are populated
Oracle Storage ZS9-2 mid-range model, high-memory configuration	16	D2, D6, D9, D13 D3, D7, D8, D12	1024 GB	Black sockets are populated White sockets are populated
Oracle Storage ZS9-2 high-end model	32	D0, D2, D4, D6, D9, D11, D13, D15 D1, D3, D5, D7, D8, D10, D12, D14	2048 GB	Black sockets are populated White sockets are populated

Oracle ZFS Storage ZS9-2 Storage Drives

Only two of the storage drive slots are used for system boot purposes, which include redundant mirrored drives. These are used and designated consecutively from left, bottom (System Disk 0) to right (System Disk 1). All other storage drive slots include filler panels as is indicated in the following drawing. Drive slots are numbered from left to right, and from bottom to top.

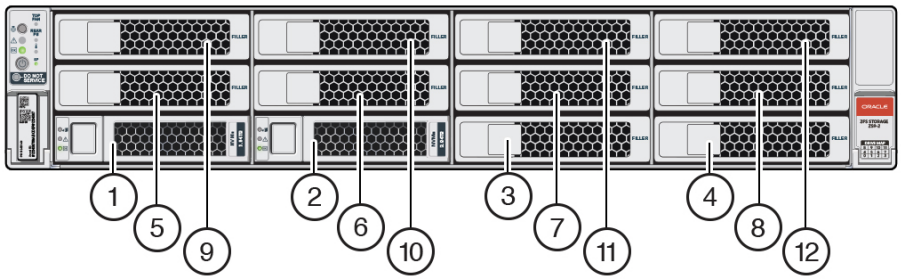

1 System Disk 0	**2** System Disk 1	**3** Filler panel
4 Filler panel	**5** Filler panel	**6** Filler panel
7 Filler panel	**8** Filler panel	**9** Filler panel
10 Filler panel	**11** Filler panel	**12** Filler panel

Oracle ZFS Storage ZS9-2 CPUs

The two CPU sockets are located under the air baffle behind the fan modules and ahead of the PCIe slots. As viewed from the front of the controller, the leftmost socket is designated P0 and the rightmost socket is designated P1.

- If the processor fault LED is off, the processor is operating properly.
- If the processor fault LED is on (amber), the processor has failed and must be replaced.

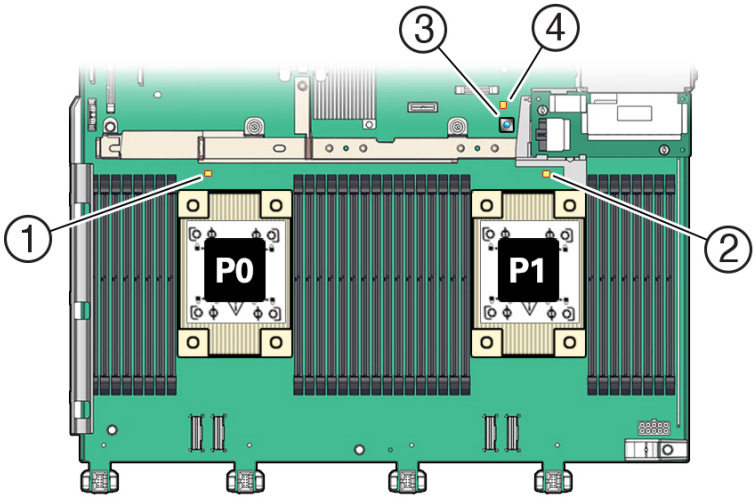

1 Processor 0 fault LED	**3** Fault Remind button
2 Processor 1 fault LED	**4** Fault Remind circuit good LED

Oracle ZFS Storage ZS9-2 PCIe Card Configuration

The 10 PCIe slots are located inside the controller at the back. As viewed from the back of the controller, the slots are divided into two groups, a group of five to the left of the built-in

network ports and rear indicators and a group of five to the right side of the network ports and indicators. The PCIe slots are designated from left to right (back view) with the five slots on the left side designated as PCIe Slot 1 to PCIe Slot 5, and the five slots on the right side designated as PCIe Slot 6 to PCIe Slot 10. The following figure shows the Oracle ZFS Storage ZS9-2 high-end model.

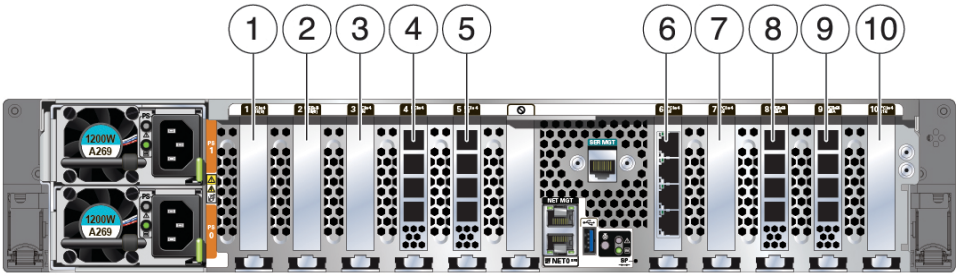

1 Second PCIe option	**2** Third PCIe option	**3** Fifth PCIe option
4 External SAS-3 HBA	**5** HBA or filler panel • High-end model: External SAS-3 HBA • Mid-range model: Filler panel; slot not available	**6** Cluster interface Ethernet adapter
7 First PCIe option	**8** HBA or filler panel • High-end model: External SAS-3 HBA • Mid-range model: Filler panel; slot not available	**9** External SAS-3 HBA
10 Fourth PCIe option		

PCIe Slot Width and Capabilities

The controller contains 10 low-profile PCIe 4.0 card slots, which have the following width and capabilities:

- Slots 3 and 10 have x16 connectors with x16 electrical interface
- Slots 1 and 2 have x16 connectors with x8 electrical interface
- Slots 4, 5, 6, 7, 8, and 9 have x8 connectors (x8 electrical interface)

> **Note:**
>
> While slots 1 and 2 can accommodate a x16-width card, the capability of the card will be limited by the x8 electrical interface.

PCIe Card Slot Population Order

Each base unit Oracle ZFS Storage ZS9-2 is always loaded with the following PCIe adapters or fillers:

- Slot 4: 4x4 SAS-3 HBA

- Slot 5: 4x4 SAS-3 HBA high-end model; filler panel mid-range model (slot not available)
- Slot 6: Cluster interface Ethernet adapter
- Slot 8: 4x4 SAS-3 HBA high-end model; filler panel mid-range model (slot not available)
- Slot 9: 4x4 SAS-3 HBA

All allowable option Oracle ZFS Storage ZS9-2 PCIe cards should be installed in the remaining card slot locations and filled in accordance to the following priority and sequence as follows:

1. Install optional dual port 100GbE NIC cards into slot 3, then slot 10.
2. Install optional dual port 25GbE NIC cards into slot 7, then slot 1, then slot 2, then slot 10, then slot 3.
3. Install optional dual port 40GbE NIC cards into slot 7, then slot 1, then slot 2, then slot 10, then slot 3.
4. Install optional dual port 32Gb FC HBA cards into slot 7, then slot 1, then slot 2, then slot 10, then slot 3.
5. Install optional quad port 10GBASE-T NIC cards into slot 7, then slot 1, then slot 2, then slot 10, then slot 3.

Oracle ZFS Storage ZS9-2 Attached Storage

The Oracle ZFS Storage ZS9-2 controller connects to external storage through 4-port (4x4) SAS HBAs. You can attach from 1 to 48 disk shelves (up to 1,152 drives) to the Oracle ZFS Storage ZS9-2 high-end controller, or 1 to 24 disk shelves (up to 576 drives) to the Oracle ZFS Storage ZS9-2 mid-range controller, including disk shelf configurations with fully populated flash-based SSD data devices. While both the Oracle Storage Drive Enclosure DE3-24 and DE3-12 disk shelves connect to the controller with active optical cables (AOCs), Oracle Storage Drive Enclosure DE2-24 disk shelves connect with copper cables. For disk shelf configuration details, see Disk Shelf Configurations.

The Oracle ZFS Storage ZS9-2 controller supports the following devices in Oracle Storage Drive Enclosure DE2-24 disk shelves:

LFF HDDs

- H7280A520SUN8.0T
- H7280B520SUN8.0T
- H7280B524SUN8.0T

SFF HDDs

- H101860SFSUN600G
- H101812SFSUN1.2T
- H1018124FSUN1.2T

Log Cache Devices

- HSCAC2DA6SUN200G
- HBCAC2DH6SUN200G
- HBSAC2DH6SUN200G

- HPCAC2DH6ORA200G
- HPSAC2DH6ORA200G

Read Cache Devices

- HSCAC2DA2SUN1.6T
- HBCAC2DH2SUN3.2T
- HBSAC2DH2SUN3.2T

5

Servicing the Oracle ZFS Storage ZS7-2 Controller

Be sure to read sections Getting Started with Servicing the Appliance and Prerequisites for Servicing Hardware.

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see Repairing Active Problems.

To replace Oracle ZFS Storage ZS7-2 hardware components, use the following tasks:

- Replacing an Oracle ZFS Storage ZS7-2 HDD
- Replacing an Oracle ZFS Storage ZS7-2 Fan Module
- Replacing an Oracle ZFS Storage ZS7-2 Power Supply
- Replacing an Oracle ZFS Storage ZS7-2 DIMM
- Replacing or Adding an Oracle ZFS Storage ZS7-2 PCIe Card
- Replacing an Oracle ZFS Storage ZS7-2 System Battery

To understand the Oracle ZFS Storage ZS7-2 controller, use the following topics:

- Oracle ZFS Storage ZS7-2 Configuration Options
- Oracle ZFS Storage ZS7-2 Front and Rear Panel Components
- Oracle ZFS Storage ZS7-2 Specifications
- Oracle ZFS Storage ZS7-2 Internal Components
- Oracle ZFS Storage ZS7-2 CPUs
- Oracle ZFS Storage ZS7-2 PCIe Card Configuration
- Oracle ZFS Storage ZS7-2 Attached Storage

Replacing an Oracle ZFS Storage ZS7-2 HDD

HDDs are hot-swappable and can be removed and installed while the controller is running without affecting other hardware capabilities. Use the following procedure to replace an Oracle ZFS Storage ZS7-2 HDD.

> **Note:**
>
> If you have more than one failed drive, replace only one drive at a time. Removing more than one drive in quick succession will cause a hardware/pool fault. Also, do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** BUI menu, select **System**, or navigate to `maintenance system updates` in the CLI.

1. Identify the faulted HDD.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, and open the detailed view of the controller. The faulted drive has either an amber or a blue-on-amber (ready to remove) status icon. Note the ID of the faulted drive. Click the entry for that drive in the table to highlight the location of the drive in the controller illustration. To view details about the faulted drive, click its information icon and then click **Active Problems**.
   - If you are physically at the system, the amber Service Action Required indicator on the HDD should be lit. Additionally, you can click the locator icon in the BUI to light the locator LED on the faulted drive.
2. Verify that the faulted drive is ready to be removed.

   > **Caution:**
   >
   > Removing the drive before it is ready to be removed can result in loss of data.

   Use any of the following methods to verify that the faulted drive is ready to be removed:
   - On the faulted drive, the Ready to Remove LED is lit as shown in Oracle ZFS Storage ZS7-2 Storage Drive.
   - In the BUI:
     - The status icon for the faulted drive in the detail view in **Maintenance: Hardware** has changed from amber to blue-on-amber.
     - The **Component Details** information window (click the information icon) for the faulted drive shows **Yes** as the value for the **Ready to remove** property.
   - In the CLI (and RESTful API), the value of the `readytoremove` property for the faulted drive is `true`:

     ```
     hostname:maintenance chassis-000 disk-001> get readytoremove
                     readytoremove = true
     ```
3. On the drive you want to remove, push the drive release button to open the latch.
4. Push the latch release button to open the drive latch [1, 2].

   The drive location in the following figure is for illustration purposes only.

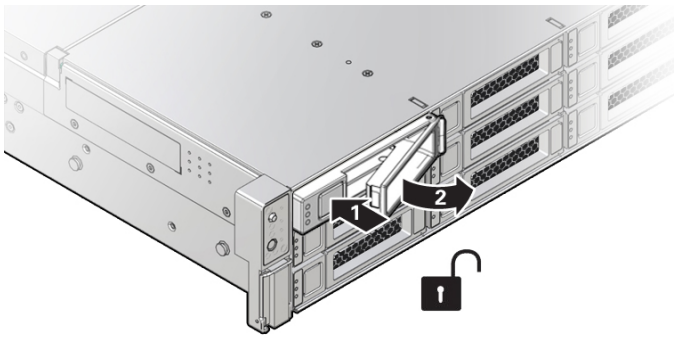

1 Pressing the latch release button	**2** Opening the latch

> **Caution:**
>
> The latch is not an ejector. Do not open the latch too far to the right. Doing so can damage the latch.

5. Grasp the latch and pull the drive out of the drive slot.

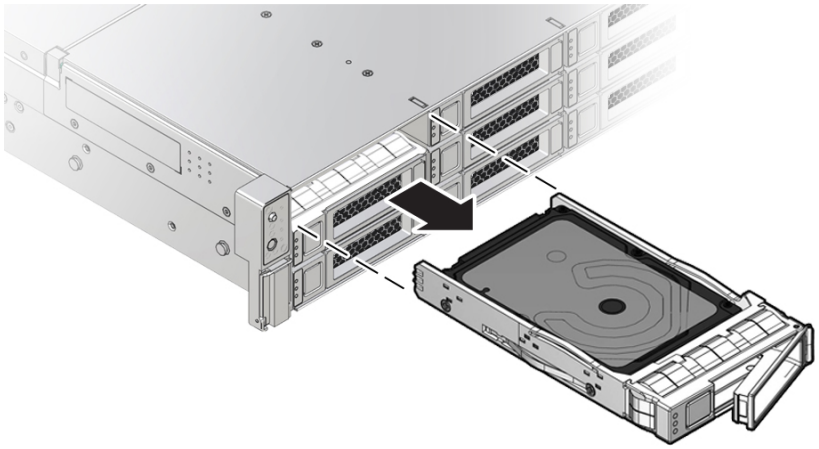

6. After at least 30 seconds, from the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon on the system controller to verify that the software has detected that the drive is not present.
7. On the replacement drive, open the release lever and slide the drive into the same slot until it is fully seated. Close the latch to lock the drive in place.

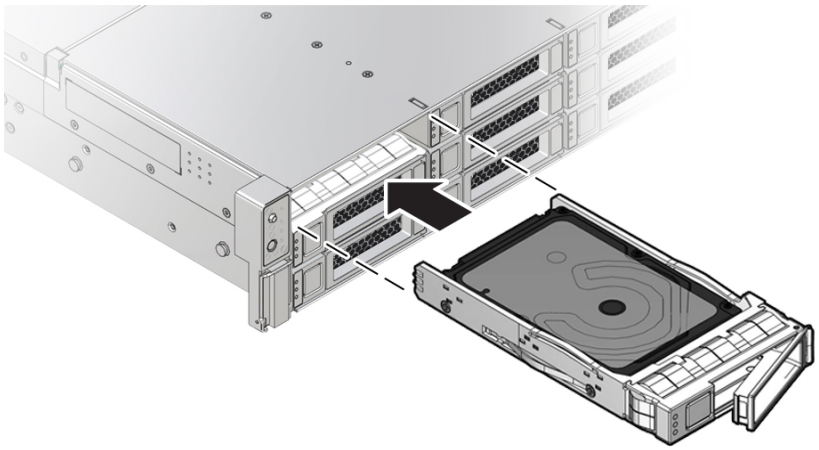

The Oracle ZFS Storage Appliance software automatically detects and configures the new drive.

8. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, and verify that the status icon is green for the newly installed drive. To view details about the new drive, click its information icon .
9. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS7-2 Fan Module

Unpack the replacement fan module and get it ready to insert into the controller chassis as soon as the faulted fan is removed.

You do not have to power off the controller to service fan modules. You should not begin this procedure unless you have a replacement fan module and you are ready to install it immediately.

> **Caution:**
>
> When performing the actual removal and replacement of a fan module in the Oracle ZFS Storage ZS7-2 controller, you must complete the entire procedure within 20 seconds to maintain adequate cooling within the system. Remove and replace only one fan module at a time. Ensure that you have obtained the replacement fan module and that is it is ready for installation before starting the replacement procedure. If the fan module is not replaced within 20 seconds of removal, Oracle ILOM will take the protective action to shut down the system to prevent thermal damage.

Fan modules are hot-swappable and can be removed and installed while the controller is running without affecting other hardware capabilities. Use the following procedure to replace an Oracle ZFS Storage ZS7-2 fan module.

1. Identify the faulted fan module.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **Fan**, and

note the ID of the faulted fan module, which has an amber icon. To highlight the location of the fan module in the controller picture, click on the fan module ID. To view details about the faulted fan module, click its information icon and then click **Active Problems**.

- If you are physically at the system, the amber Service Action Required indicator on the fan module should be lit.

2. Extend the controller from the rack as described in Extending the Controller from the Rack.
3. To access the fan modules, remove the top cover as described in Removing the Controller Top Cover.
4. Identify the faulty fan module.

   Each fan module has a fan status indicator (LED) that is located next to the module. If the LED is amber, the fan has failed. The location of the fan status LED is shown in the following figure.

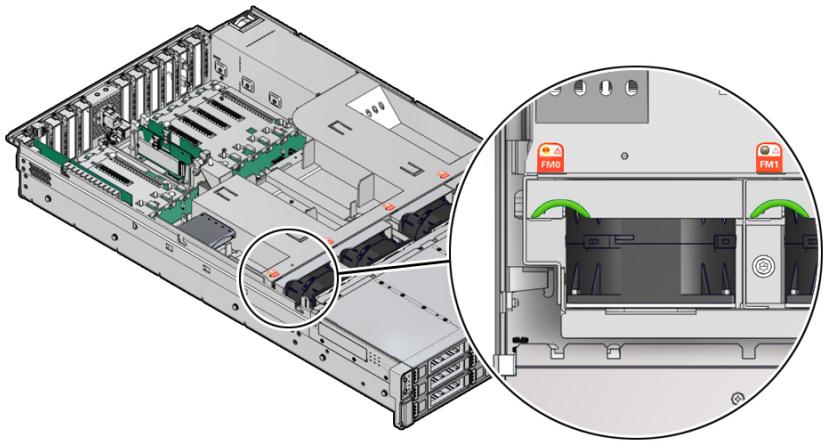

LED Color and State Meaning
Amber – The fan module is faulty. The front Top Fan LED and the front and back panel Fault-Service Required LEDs are also lit if the system detects a fan module fault.
Not Illuminated – The fan module is operational (OK).

5. Using your forefinger and thumb, grasp the green highlighted recess and lift the fan module straight up and out of the chassis. Set it aside on an antistatic mat.

> **Caution:**
>
> Excessive movement or rocking of the fan module during removal can cause damage to the internal connector on the fan module board. When removing a fan module, do not rock it back and forth.

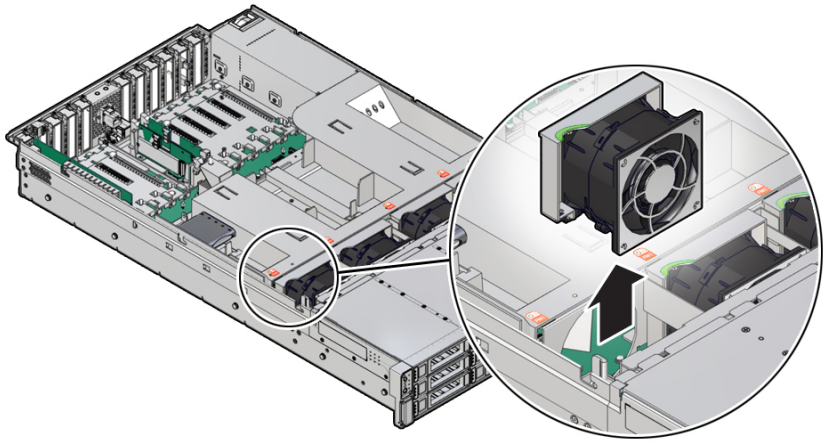

6. Position the replacement fan module so the connector on the bottom of the module is aligned with the connector on the inside of the slot and all keying and labeling are correctly positioned.

   The fan modules are keyed to ensure that they are installed in the correct orientation.

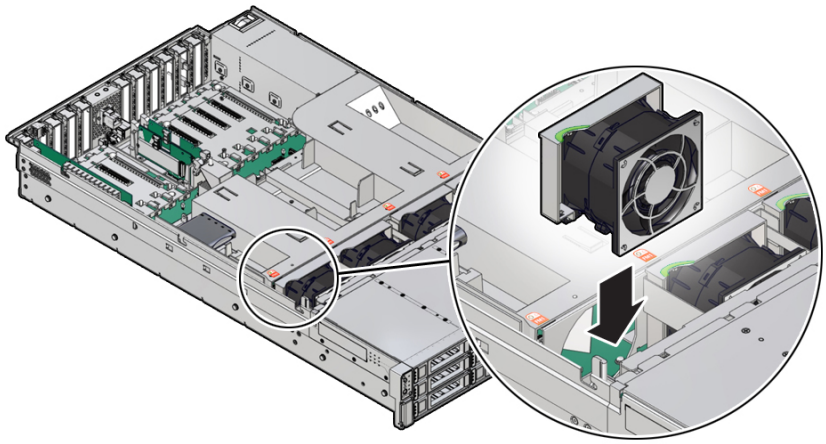

7. Slide the fan module into the slot until it stops.
8. Verify that the fan module status indicator (amber LED) on the replacement fan module is not illuminated.
9. Install the controller top cover.

   See Removing the Controller Top Cover.
10. Verify that the Top Fan Fault LED on the front of the controller and the Fault-Service Required LEDs on the front and back of the controller are not lit.
11. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **Fan**, and verify that the status icon is green for the newly installed fan module. To view details about the new fan module, click its information icon .

12. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS7-2 Power Supply

> **Caution:**
>
> If a power supply fails and you do not have a replacement available, do not remove the failed power supply. This maintains the proper air flow and prevents possible equipment damage.

Controllers are equipped with redundant, hot-swappable power supplies. Use the following procedure to replace an Oracle ZFS Storage ZS7-2 power supply.

1. Identify the faulted power supply.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **PSU**, and note the ID of the faulted power supply, which has an amber icon. To highlight the location of the power supply in the controller picture, click on the power supply ID. To view details about the faulted power supply, click its information icon and then click **Active Problems**.
   - If you are physically at the system, the amber Service Action Required indicator on the power supply should be lit.

> **Note:**
>
> The fans of a failed power supply might still be spinning when the system is powered on. You can remove a power supply while the fans are still spinning.

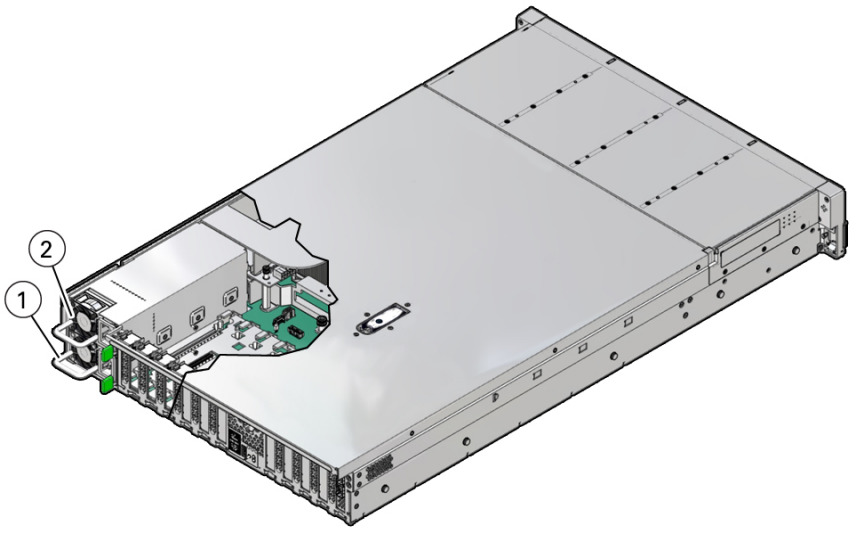

1 Power supply 0	**2** Power supply 1

2. Gain access to the rear of the controller where the failed power supply is located.
3. If the cable management arm (CMA) is installed, disconnect both CMA left-side connectors and move the CMA out of the way.

   For the location of the CMA left-side connectors, see Installing an Oracle ZFS Storage ZS7-2 Cable Management Arm in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*.

> **Caution:**
>
> When disconnecting the CMA left-side connectors, be sure to use your arm to support the CMA so that it does not hang down under its own weight and stress the right-side connectors; otherwise, the CMA could be damaged. You must continue to support the CMA until you have reconnected both of the left-side connectors.

4. Disconnect the power cord from the failed power supply.
5. Grasp the power supply handle and push the green power supply latch to the left.

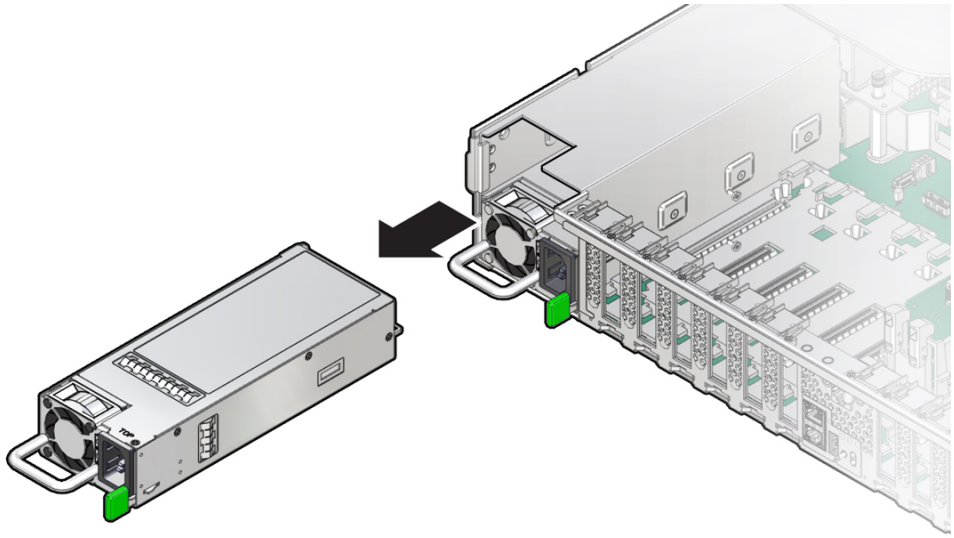

6. Pull the power supply out of the controller.

> **Caution:**
>
> Whenever you remove a power supply, you should replace it with another power supply; otherwise, the controller could overheat due to improper airflow.

7. Remove the replacement power supply from its packaging, and place it on an antistatic mat.

> **Caution:**
>
> Always replace the failed power supply with the same type model of power supply.

8. Align the replacement power supply with the empty power supply slot.
9. Slide the power supply into the slot until it is fully seated.

   You will hear an audible click when the power supply fully seats.

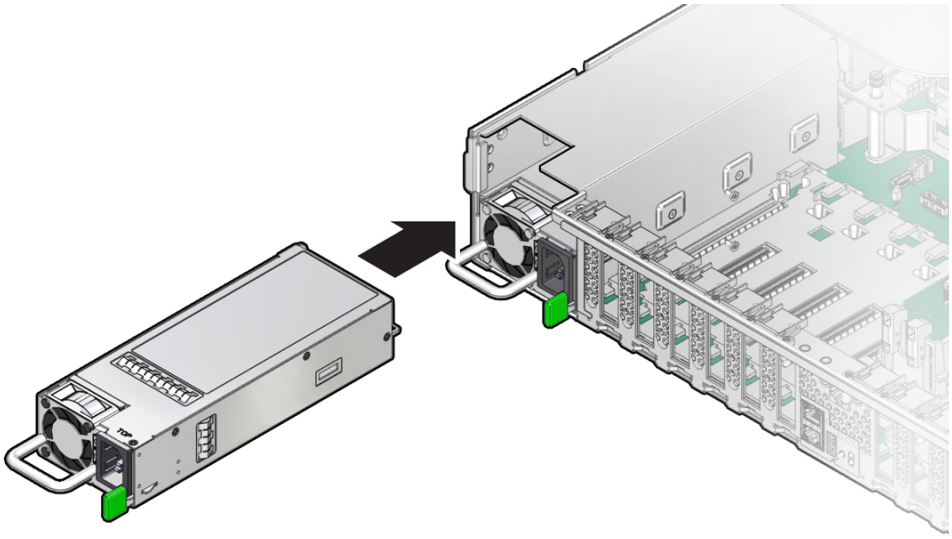

10. Reconnect the power cord to the power supply.
11. Verify that the amber Fault-Service Required LED on the replaced power supply and the Fault-Service Required LEDs are not lit on the front and back panels.

> **Note:**
>
> After you have replaced Power Supply 0, you might need to reset the Oracle ILOM service processor (SP) to propagate the key identity properties (KIP) data to the new power supply. For instructions on resetting the SP, refer to the *Oracle Integrated Lights Out Manager (ILOM) 4.0 Configuration and Maintenance Guide Firmware Release 4.0.x* at https://www.oracle.com/goto/ilom/docs . Power Supply 1 does not contain KIP data, and therefore does not require an SP reset after replacement.

12. If you disconnected the two CMA left-side connectors, reconnect the connectors.
13. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **PSU**, and verify that the status icon is green for the newly installed power supply. To view details about the new power supply, click its information icon .
14. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS7-2 DIMM

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

Removing a DIMM is a cold-service operation. You *must* power off the appliance before you remove the DIMM. There will be a loss of access to the storage unless the system is in a clustered configuration.

1. Identify the faulted DIMM.

   From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **DIMM**, and note the ID of the faulted DIMM, which has an amber icon. To highlight the location of the memory module in the controller picture, click on the module ID. To view details about the faulted module, click its information icon and then click **Active Problems**.
2. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
3. Disconnect the AC power cords from the rear panel of the controller.

   > **Caution:**
   >
   > Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.
4. Extend the controller from the rack as described in Extending the Controller from the Rack.
5. Remove the top cover as described in Removing the Controller Top Cover.
6. Remove the air baffle by lifting the baffle up and out of the controller.

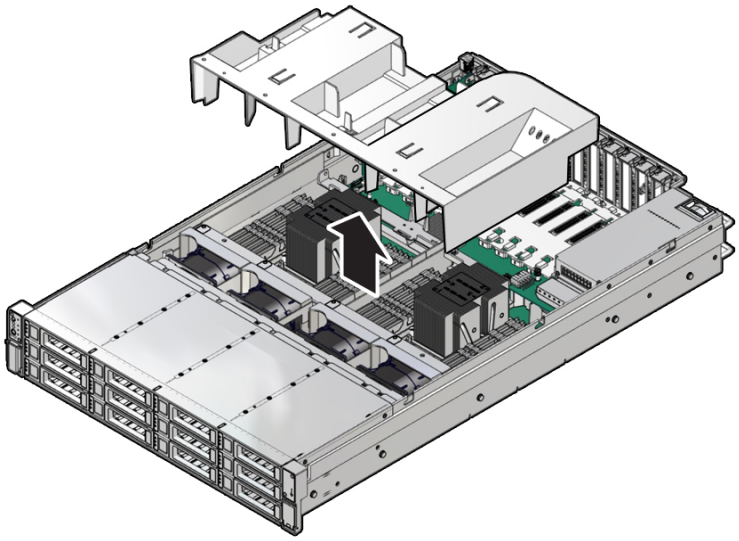

7. Set aside the air baffle.
8. Locate the Fault Remind button.

   When the controller Fault Remind button [1] is pressed, an LED located next to the Fault Remind button lights green to indicate that there is sufficient voltage present in the fault remind circuit to light any fault LEDs that were lit due to a component failure. If this LED does not light when you press the Fault Remind button, it is likely that the capacitor powering the fault remind circuit has lost its charge. This can happen if the Fault Remind button is pressed for several minutes with fault LEDs lit or if power has been removed from the controller for more than 15 minutes.

   The following figure shows the location of the Fault Remind button.

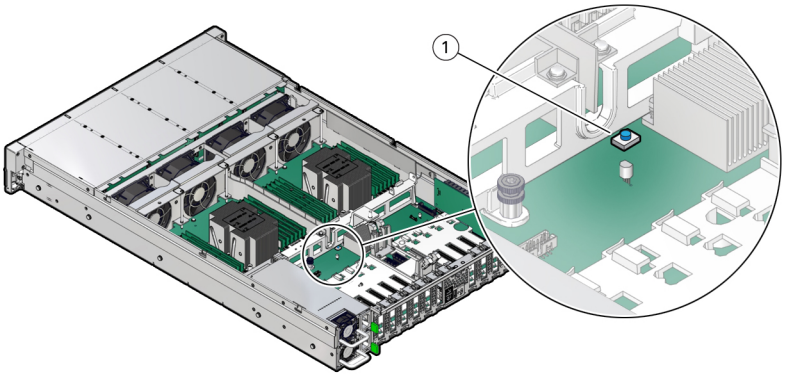

9. Physically identify and note the location of the failed DIMM by pressing the Fault Remind button on the motherboard.

   Failed DIMMs are identified with a corresponding amber LED on the motherboard.

   - If the DIMM Fault LED is off, then the DIMM is operating properly.
   - If the DIMM Fault LED is on (amber), then the DIMM is failed and should be replaced [1].

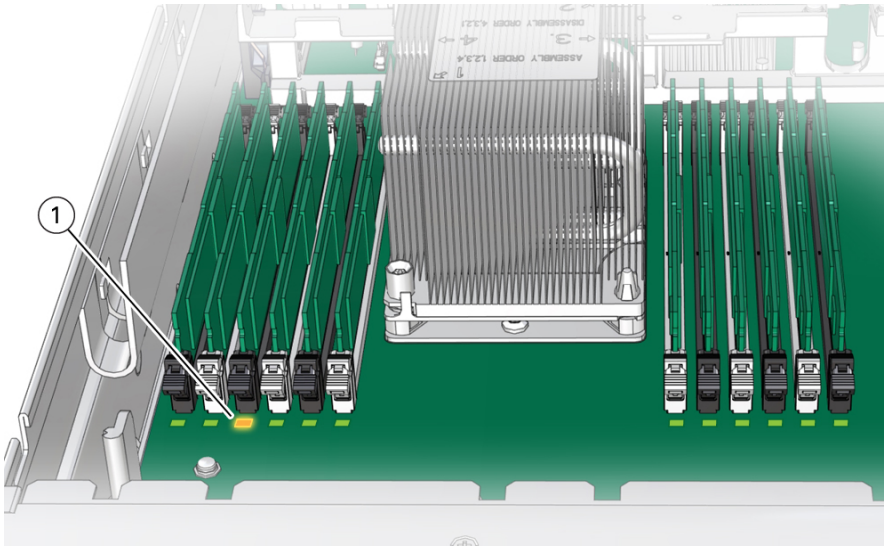

10. Rotate both DIMM slot ejector levers outward as far as they will go.

    The action unlocks the DIMM and expels it from the socket.

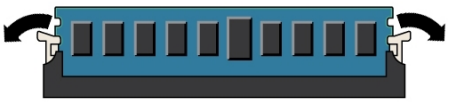

11. Carefully lift the DIMM straight up to remove it from the socket.

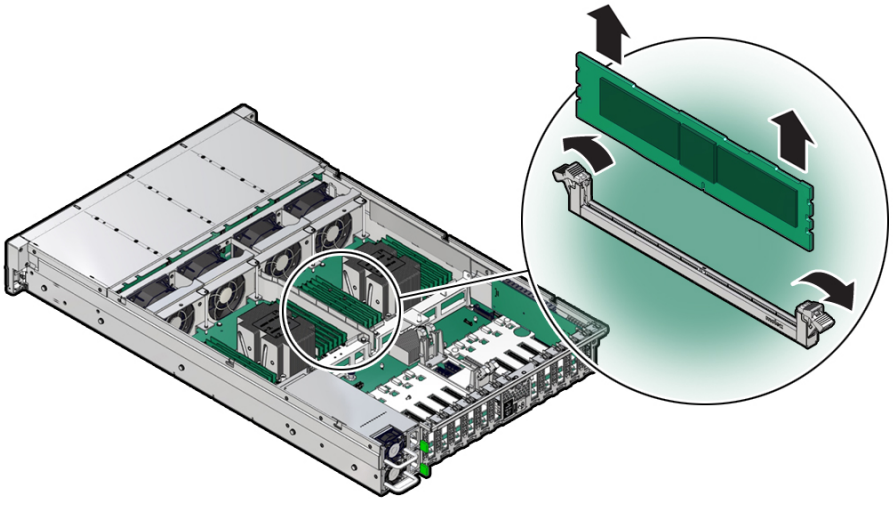

12. Repeat the previous steps until all faulty DIMMs are removed.
13. To install a replacement DIMM, ensure that the DIMM ejector levers at both ends of the DIMM slot are in their fully open position.

14. Align the DIMM with the empty slot.

    Align the notch on the DIMM with the protrusion on the DIMM slot to ensure that the DIMM is installed correctly.

15. Push the DIMM into the connector slot until the ejector tabs lock the DIMM in place.

    If the DIMM does not easily seat into the connector slot, verify that the notch in the DIMM is aligned with the connector key in the connector slot. If the notch is not aligned, damage to the DIMM might occur.

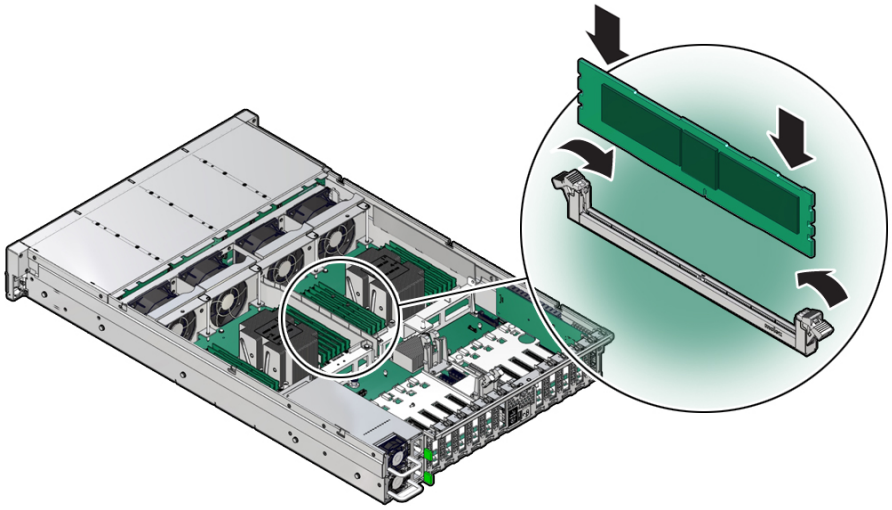

16. Ensure that the levers have risen completely and have locked the DIMM in the slot.

17. Repeat the previous steps until all replacement DIMMs are installed.
18. Install the air baffle by placing it into the controller and lowering it to its inserted position.

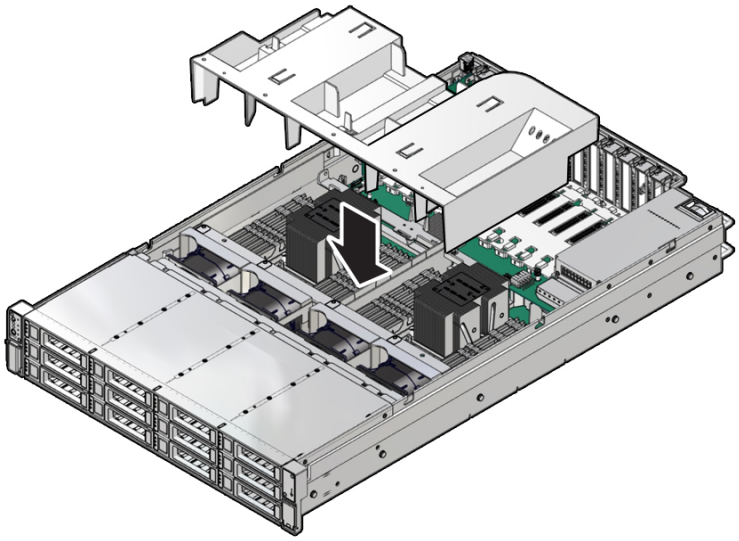

19. Replace the top cover on the chassis.
20. Return the controller to the normal rack position by pulling the green release tabs on the side of each rail while slowly pushing the controller into the rack.
21. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
22. Reconnect the power cords to the power supplies.
23. Verify that standby power is on. Approximately two minutes after plugging in the power cords, the Power/OK status indicator on the front panel starts flashing.
24. After the Power/OK status indicator is flashing, press and release the recessed Power button on the controller front panel.

    The Power/OK status indicator continues flashing during the startup sequence, which takes several minutes. When the system is operational, the indicator remains lit.
25. After the Power/OK status indicator is lit steadily, log in, and from the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **DIMM**, and verify that the status icon is green for the newly installed module. To view details about the new module, click its information icon .
26. If a fault does not self-clear, see Repairing Active Problems.

Replacing or Adding an Oracle ZFS Storage ZS7-2 PCIe Card

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

Replacing or adding a PCIe card is a cold-swap procedure, and you must remove power from the controller. If you are installing a newly released adapter, upgrade your system software before installing the adapter as described in Upgrading the Software. Use the following procedure to replace or add an Oracle ZFS Storage ZS7-2 PCIe card.

1. Identify the PCIe card slot location.
   - If replacing a faulted PCIe card, identify the PCIe card to be replaced.

     From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **Slot**, and note the ID of the faulted PCIe card, which has an amber icon. If replacing a non-faulted PCIe card, note its location. To highlight the location of the card in the controller picture, click on the card ID. To view details about the card, click its information icon and then click **Active Problems**.
   - If adding a new PCIe card, identify the correct slot location as described in Oracle ZFS Storage ZS7-2 PCIe Card Configuration.
2. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
3. Disconnect the AC power cords from the rear panel of the controller.

   > **Caution:**
   >
   > Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.
4. Extend the controller from the rack as described in Extending the Controller from the Rack.
5. Remove the top cover as described in Removing the Controller Top Cover.
6. Locate the PCIe card position in the controller.

   There are 11 PCIe slots available for PCIe cards. The following figure shows the PCIe slot numbering.

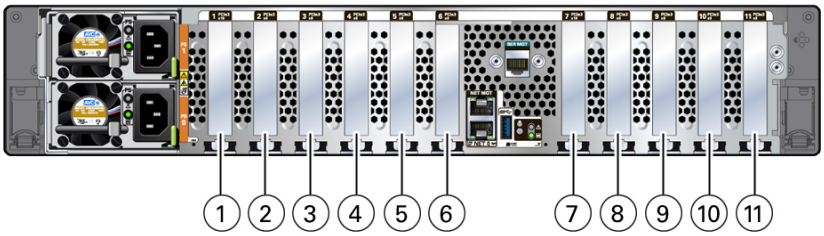

1 PCIe slot 1: Fifth PCIe option	**7** PCIe slot 7: Fourth PCIe option
2 PCIe slot 2: External SAS-3 HBA	**8** PCIe slot 8: Second PCIe option
3 PCIe slot 3: High-end model: External SAS-3 HBA; Mid-range model: Filler panel; slot not available	**9** PCIe slot 9: High-end model: External SAS-3 HBA; Mid-range model: Filler panel; slot not available
4 PCIe slot 4: First PCIe option	**10** PCIe slot 10: External SAS-3 HBA
5 PCIe slot 5: Third PCIe option	**11** PCIe slot 11: Internal SAS-3 HBA
6 PCIe slot 6: Cluster interface card	

> **Note:**
>
> All of the PCIe slots comply with the PCI Express 3.0 specification and can accommodate 25 Watt PCIe3 cards.

7. Unplug all data cables from the PCIe card.

   Note the cable connections in order to ease proper reconnection of the cables. If the PCIe card has transceivers, remove them before removing the PCIe card from the controller.

8. Rotate the PCIe card locking mechanism into an upright position [1], and then lift up on the PCIe card to disengage it from the motherboard connectors [2].

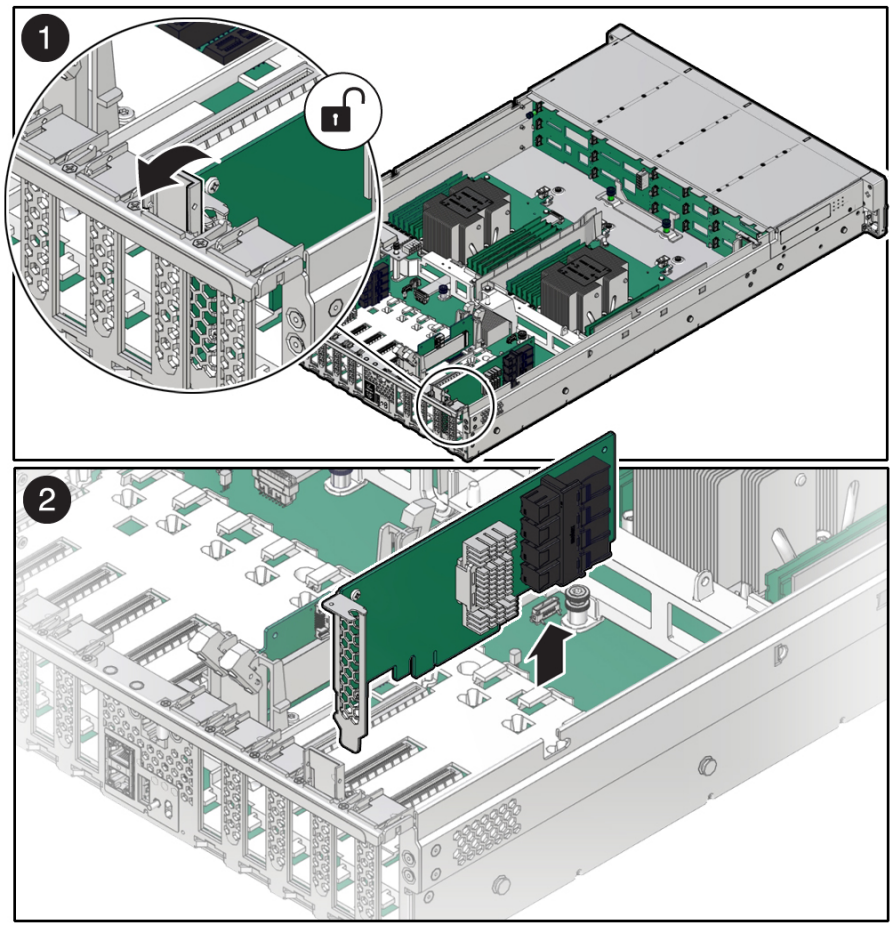

9. Place the PCIe card on an antistatic mat.

> **Caution:**
>
> If you are not immediately inserting a replacement PCIe card into the empty slot, insert a PCIe filler panel in the slot to reduce the possibility of radiated electromagnetic interference (EMI), or the overheating of the controller due to improper airflow.

10. Insert the PCIe card into the correct slot [1], and rotate the PCIe locking mechanism downward to secure the PCIe card in place [2].

    You will hear or feel a click when the PCIe card is secured into the slot.

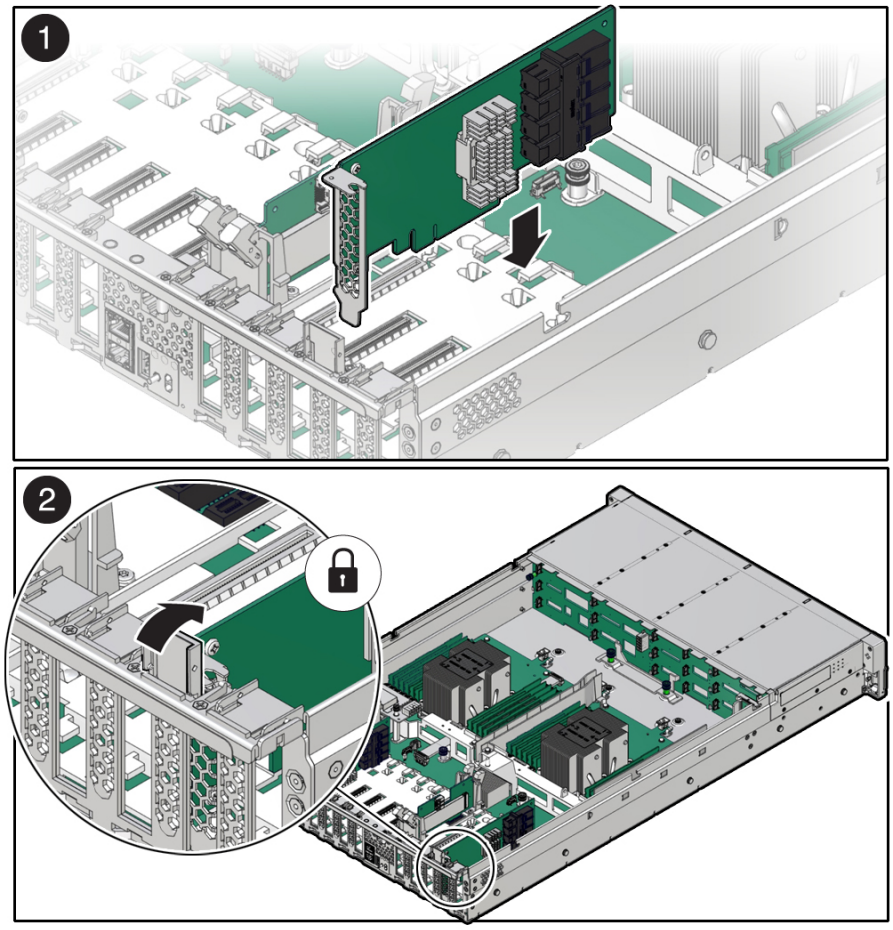

11. Replace the top cover on the chassis.
12. Return the controller to the normal rack position by pulling the green release tabs on the side of each rail while slowly pushing the controller into the rack.

    If the installed PCIe card requires transceivers, install them now.
13. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.

    If you replaced an HBA, connect the disk shelves to its ports as described in the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
14. Reconnect the power cords to the power supplies.
15. Verify that standby power is on. Approximately two minutes after plugging in the power cords, the Power/OK status indicator on the front panel starts flashing.
16. After the Power/OK status indicator is flashing, press and release the recessed Power button on the controller front panel.

    The Power/OK status indicator continues flashing during the startup sequence, which takes several minutes. When the system is operational, the indicator remains lit.
17. After the Power/OK status indicator is lit steadily, log in, and from the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the

controller, click **Slot**, and verify that the status icon is green for the newly installed card. To view details about the new card, click its information icon .

Replacing an Oracle ZFS Storage ZS7-2 System Battery

Replacing a system battery is a cold-swap procedure, and you must remove power from the controller. You might need to replace the battery if the controller fails to maintain the proper time when powered off and disconnected from the network. Use the following procedure to replace an Oracle ZFS Storage ZS7-2 system battery.

1. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. Remove the top cover as described in Removing the Controller Top Cover.
5. To dislodge the battery from its retainer, use your finger to gently push the top edge of the battery away from the retainer.

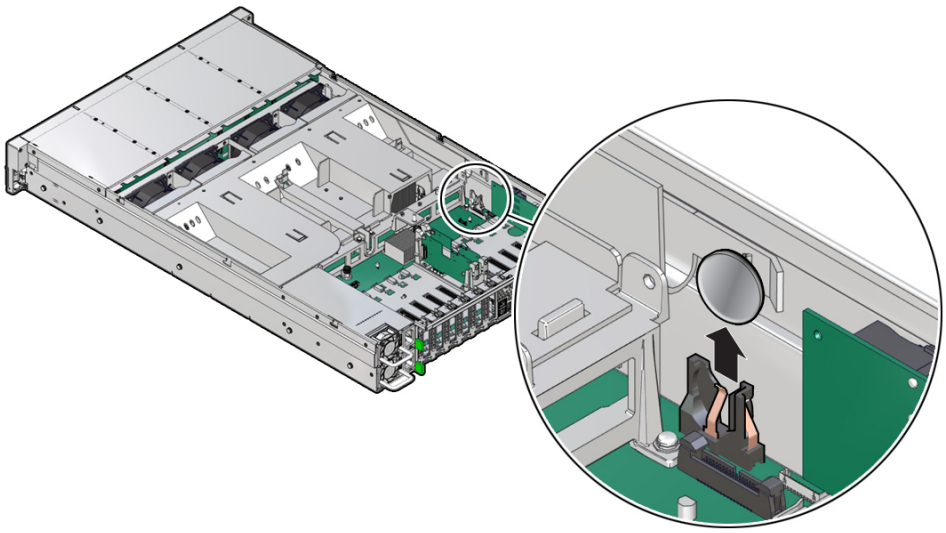

6. Lift the battery up and out of its retainer.
7. Unpack the replacement battery.
8. Press the new battery into the battery retainer with the positive (+) side facing the metal tab that holds it in place.
9. Replace the top cover on the chassis.

10. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the controller into the rack.
11. Connect the power cords to the power supplies.
12. Verify that standby power is on. Approximately two minutes after plugging in the power cords, the Power/OK status indicator on the front panel starts flashing.
13. After the Power/OK status indicator is flashing, press and release the recessed Power button on the controller front panel.

    The Power/OK status indicator continues flashing during the startup sequence, which takes several minutes. When the system is operational, the indicator remains lit.
14. After the Power/OK status indicator is lit steadily, log in and set the time using the steps in the BUI Clock task. For more information, see Setting Clock Synchronization (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Oracle ZFS Storage ZS7-2 Controller Hardware Overview

This section describes the internal and external components of the Oracle ZFS Storage ZS7-2 controller.

- Oracle ZFS Storage ZS7-2 Configuration Options
- Oracle ZFS Storage ZS7-2 Front and Rear Panel Components
- Oracle ZFS Storage ZS7-2 Specifications
- Oracle ZFS Storage ZS7-2 Internal Components
- Oracle ZFS Storage ZS7-2 CPUs
- Oracle ZFS Storage ZS7-2 PCIe Card Configuration
- Oracle ZFS Storage ZS7-2 Attached Storage

Oracle ZFS Storage ZS7-2 Configuration Options

The Oracle ZFS Storage ZS7-2 high-end model incorporates 1 or 2 controllers and 1 to 48 disk shelves (up to 1,152 disk shelf drives). Each controller has 11 PCIe Gen3 slots in the rear, and 12 3.5" disk slots in the front, 2 of which are occupied by system disks. Read cache devices are installed in disk shelves; therefore, the remaining 10 front disk slots have permanent filler panels. The following table describes the configuration options.

CPU	Memory	System Disks	PCIe Slots	HBA
2x24-core, 2.1 GHz	1.5 TB	2x3.5" SAS-3	Six dedicated; five available	4-port (4x4) SAS-3 12 Gb/s external

The Oracle ZFS Storage ZS7-2 mid-range model incorporates 1 or 2 controllers and 1 to 24 disk shelves (up to 576 disk shelf drives). Each controller has 11 PCIe Gen3 slots in the rear, and 12 3.5" disk slots in the front, 2 of which are occupied by system disks. Read cache devices are installed in disk shelves; therefore, the remaining 10 front disk slots have permanent filler panels. The following table describes the configuration options.

CPU	Memory	System Disks	PCIe Slots	HBA
2x18-core, 2.3 GHz	512 GB or 1024 GB	2x3.5" SAS-3	Four dedicated; five available; two unavailable (permanent filler panels)	4-port (4x4) SAS-3 12 Gb/s external

Oracle ZFS Storage ZS7-2 Front and Rear Panel Components

Oracle ZFS Storage ZS7-2 Front Panel Components - The Oracle ZFS Storage ZS7-2 controller drive slots and front panel components are shown in the following figure.

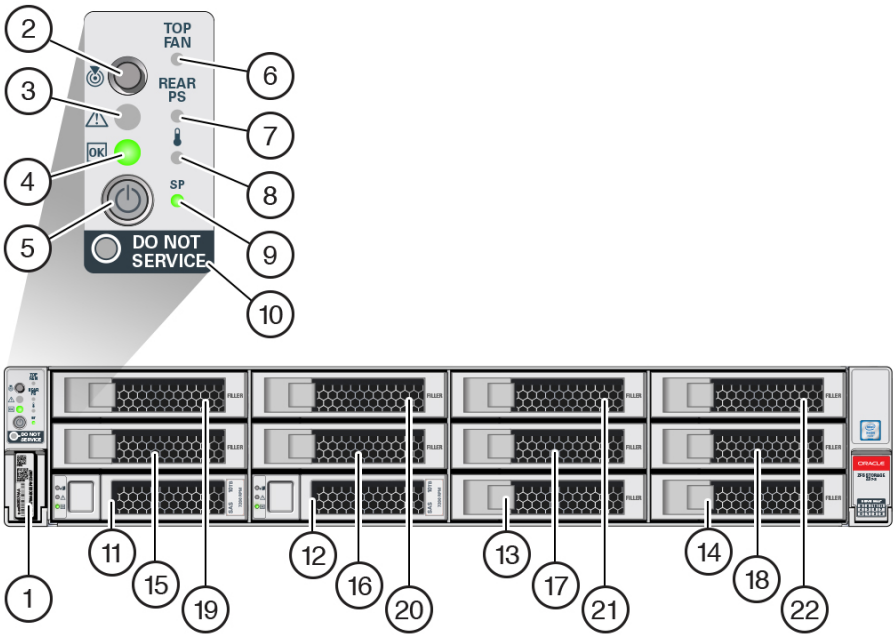

1 RFID tag	**2** System status LED: Locator (white)	**3** System status LED: Service Action Required (amber)
4 System status LED: Power/OK (green)	**5** Power button	**6** Fan service action required (amber)
7 Power Supply (PS) service action required (amber)	**8** Over-temperature warning (amber)	**9** Service Processor (SP) OK (green)
10 Do Not Service indicator	**11** System disk 0	**12** System disk 1
13 Filler panel	**14** Filler panel	**15** Filler panel
16 Filler panel	**17** Filler panel	**18** Filler panel
19 Filler panel	**20** Filler panel	**21** Filler panel
22 Filler panel		

Oracle ZFS Storage ZS7-2 Rear Panel Components - The Oracle ZFS Storage ZS7-2 controller PCIe slots and rear panel components are shown in the following figure.

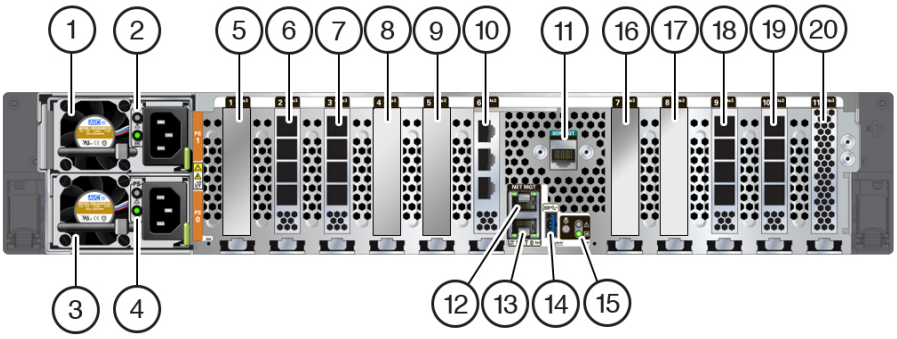

1 Power Supply (PS) Unit 1 (PS 1)	**2** PS 1 status indicators (top to bottom): • Attention (amber) • Power/OK (green)	**3** PS 0
4 PS 0 status indicators (top to bottom): • Attention (amber) • Power/OK (green)	**5** Fifth PCIe option (slot 1)	**6** External SAS-3 HBA (slot 2)
7 Slot 3: • High-end model: External SAS-3 HBA • Mid-range model: Filler panel; slot not available	**8** First PCIe option (slot 4)	**9** Third PCIe option (slot 5)
10 Cluster interface card (slot 6)	**11** Serial management (SER MGT) RJ-45 serial port	**12** Service Processor (SP) NET MGT port
13 Network (NET) 100/1000BASE-T RJ-45 Gigabit Ethernet (GbE) port: NET 0	**14** USB 3.0 port	**15** System status LEDs: • Locate (white) - left • Attention (amber) - right, top • Power/OK (green) - right, bottom
16 Fourth PCIe option (slot 7)	**17** Second PCIe option (slot 8)	**18** Slot 9: • High-end model: External SAS-3 HBA • Mid-range model: Filler panel; slot not available
19 External SAS-3 HBA (slot 10)	**20** Internal SAS-3 HBA (slot 11)	

Cluster Interface Card Compatibility - The Oracle ZFS Storage ZS7-2 controller contains a Version 3 cluster interface card, which is incompatible with Version 2.

For information about how to connect cables to form a cluster, see Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

Ethernet Port

The Oracle ZFS Storage ZS7-2 controller has one RJ-45 100/1000BASE-T Gigabit Ethernet (GbE) network connector, labeled NET0, on the rear panel. The port operates

at 100 Mbits/sec or 1000 Mbits/sec. Use this port to connect the appliance to the network.

There are two status indicators (LEDs) on this port. These indicators are visible from the rear of the controller.

Table 5-1 Gigabit Ethernet Port Status Indicators

Status Indicator Name	Location	Color	State and Meaning
Activity	Bottom left	Green	ON – Link up OFF – No link FLASHING – Packet activity
Link speed	Bottom right	Green	OFF – 100BASE-T link (if link up) ON – 1000BASE-T link

Network Management Port

The Oracle ZFS Storage ZS7-2 controller has one 10/100/1000BASE-T Ethernet management domain interface, labeled NET MGT. There are two status indicators (LEDs) on this port. These indicators are visible from the rear of the controller. The service processor NET MGT port is the optional connection to the Oracle ILOM service processor. The NET MGT port is configured by default to use Dynamic Host Configuration Protocol (DHCP). The service processor NET MGT port uses an RJ-45 cable for a 10/100/1000BASE-T connection.

Table 5-2 Network Management Port Status Indicators

Status Indicator Name	Location	Color	State and Meaning
Activity	Top left	Green	ON – Link up OFF – No link FLASHING – Packet activity
Link speed	Top right	Green	ON – 1000BASE-T link OFF – 10/100BASE-T link

Serial Management Port

The Oracle ZFS Storage ZS7-2 service processor SER MGT port uses an RJ-45 cable and is the default connection to the Oracle ILOM service processor. This port supports local connections to the controller and only recognizes Oracle ILOM command-line interface (CLI) commands. Typically you connect a terminal or terminal emulator to this port.

> **Note:**
>
> This port does not support network connections.

4x4 SAS HBA Ports

The 4x4 SAS HBAs are installed in PCIe slots 2, 3, 9, and 10 for the Oracle ZFS Storage ZS7-2 high-end model, and slots 2 and 10 for the Oracle ZFS Storage ZS7-2 mid-range

model. The HBAs provide connectivity to external disk shelves. The HBA logical ports are numbered 0-3, bottom to top, as shown in the following figure.

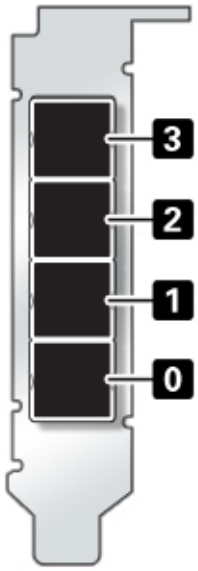

Oracle ZFS Storage ZS7-2 Specifications

Oracle ZFS Storage ZS7-2 Physical Specifications: The Oracle ZFS Storage ZS7-2 controller chassis dimensions are as follows:

Dimension	Measurement	Dimension	Measurement
Height	3.42 in./8.69 cm	Depth	29.9 in./75.9 cm
Width	17.5 in./44.5 cm	Weight (Maximum)	63 lb./28.6 kg

Oracle ZFS Storage ZS7-2 Electrical Specifications: The Oracle ZFS Storage ZS7-2 controller has two hot-swappable rear accessible power supplies with N+N redundancy. To understand the overall power consumption, see Oracle's online power calculator at http://www.oracle.com/goto/powercalculators.

Parameter	AC Power Supplies
Nominal input frequencies	50/60 Hz
Operating input voltage range	100-240 VAC
Rated input current	10 A @ 100 to 127 VAC and 7 A @ 200 to 240 VAC

Parameter	Controllers
Maximum power consumption of controller	Refer to the Power Calculator (https://www.oracle.com/goto/powercalculators) and see Oracle Server X7-2L for Watt calculations
Maximum heat output of controller	Refer to the Power Calculator (https://www.oracle.com/goto/powercalculators) and see Oracle Server X7-2L for BTU/hr calculations

Oracle ZFS Storage ZS7-2 Internal Components

The major components of the Oracle ZFS Storage ZS7-2 controller are shown in the following figure and identified in the legend.

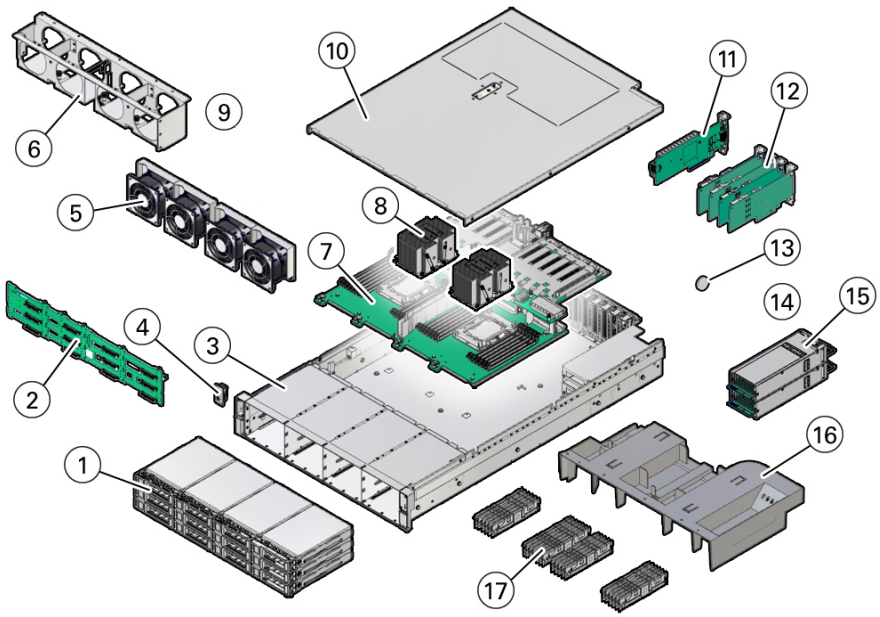

1 Two system disks, ten filler panels	**7** Motherboard assembly	**13** System battery
2 Disk backplane	**8** Processors and heatsinks	**14** N/A
3 System chassis	**9** N/A	**15** Power supplies
4 Front LED indicator module and temperature sensor	**10** Top cover	**16** Air baffle
5 Fan modules	**11** Oracle Storage 12 Gb SAS PCIe HBA, internal	**17** DIMMs
6 Fan tray	**12** PCIe cards	

Oracle ZFS Storage ZS7-2 Storage Drive

This section shows the location and explains the functions of the storage drive status indicators.

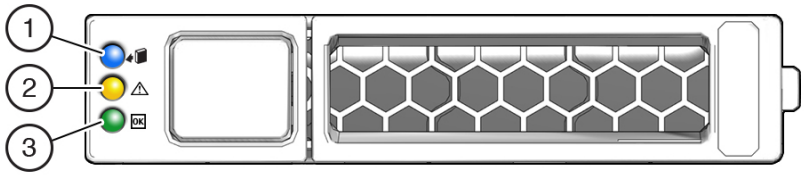

1 Ready to Remove (blue)	**2** Service Action Required (amber)	**3** OK/Activity (green)

Oracle ZFS Storage ZS7-2 Fan Module

The four fan modules are located towards the front of the Oracle ZFS Storage ZS7-2 controller between the CPUs and the drive trays. The fan modules are designated from left to right. As pictured in the following illustration, the four fan modules are designated as: FMO, FM1, FM2, and FM3.

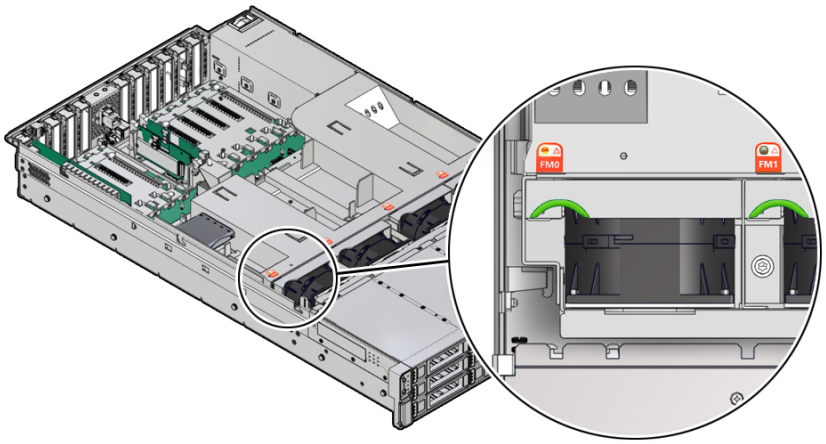

LED Color and State Meaning
Amber – The fan module is faulty. The front Top Fan LED and the front and back panel Fault-Service Required LEDs are also lit if the system detects a fan module fault.
Not Illuminated – The fan module is operational (OK).

Oracle ZFS Storage ZS7-2 Power Supply

The two power supplies are located on the right side of the controller (from the front of the controller) and are designated from top to bottom. The power supplies are accessible from the back of the controller. As indicated in the following drawing, the top designation is PS 1 and the bottom designation is PS 0.

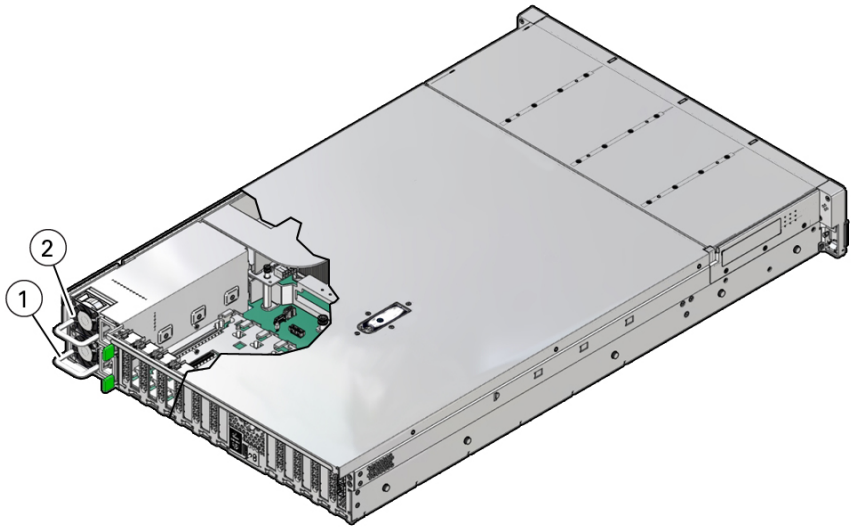

1 Power supply 0	**2** Power supply 1

The two power supplies ensure redundancy by allowing the controller to operate when one of the power supplies fails. However, when the controller is operating on one power supply, redundancy no longer exists, and the risk for an unexpected shut down and a loss of data is high. When a power supply or any component that is part of a redundant configuration fails, replace it immediately.

1 Power supply (PS) 1	**3** PS 0
2 PS 1 status indicators: Fault-Service Required LED: amber; AC power input and DC power output OK LED: green	**4** PS 0 status indicators: Fault-Service Required LED: amber; AC power input and DC power output OK LED: green

Oracle ZFS Storage ZS7-2 DIMM

The physical layout of the DIMMs and processor(s) is shown in the following figure. When viewing the controller from the front, processor 0 (P0) is on the left.

Each processor, P0 and P1, has 12 DIMM slots organized into 6 memory channels. Each memory channel contains two DIMM slots: a black DIMM slot (channel slot 0) and a white DIMM slot (channel slot 1).

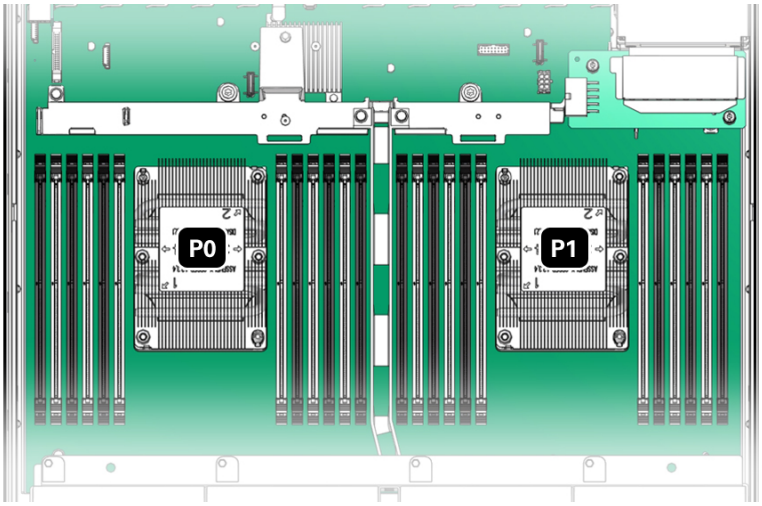

As shown in the following figure, DIMM slots are labeled D0 through D5 to the left of each processor, and D6 through D11 to the right of each processor. For each processor, the black slots are: D0, D2, D4, D7, D9, and D11. Similarly, for each processor, the white slots are: D1, D3, D5, D6, D8, and D10.

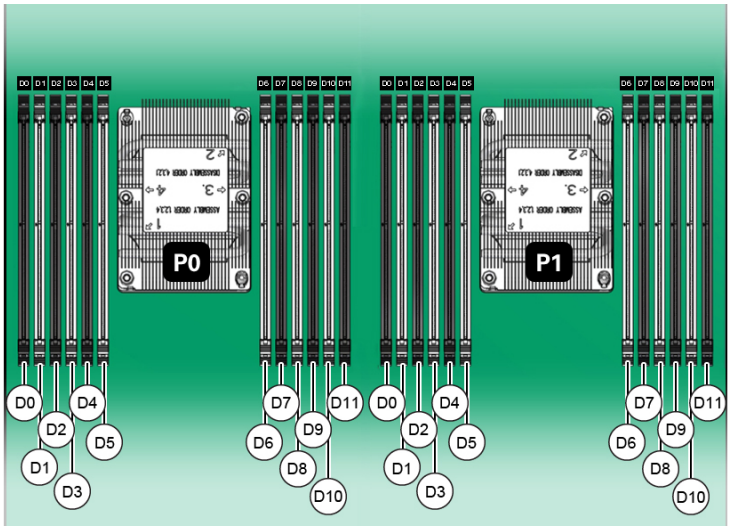

The following table describes the supported memory configurations for Oracle ZFS Storage ZS7-2. The mid-range model supports 8 or 16 DIMMs, and the high-end model supports 24 DIMMs. All DIMMs are 64 GB. The DIMM locations in the table apply to both P0 and P1 processors; for example, P0/D2,P0/D4,P0/D7,P0/D9 and P1/D2,P1/D4,P1/D7,P1/D9.

Configuration	Total DIMMs Installed	DIMM Slots Populated	DRAM Capacity	Description
Oracle Storage ZS7-2 mid-range model, low-memory configuration	8	D2, D4, D7, D9	512 GB	Black sockets are populated
Oracle Storage ZS7-2 mid-range model, high-memory configuration	16	D2, D4, D7, D9 D3, D5, D6, D8	1024 GB	Black sockets are populated White sockets are populated
Oracle Storage ZS7-2 high-end model	24	D0, D2, D4, D7, D9, D11 D1, D3, D5, D6, D8, D10	1536 GB	Black sockets are populated White sockets are populated

Oracle ZFS Storage ZS7-2 Storage Drives

Only two of the storage drive slots are used for system boot purposes, which include redundant mirrored drives. These are used and designated consecutively from left, bottom (System Disk 0) to right (System Disk 1). All other storage drive slots include filler panels as is indicated in the following drawing. Drive slots are numbered from left to right, and from bottom to top.

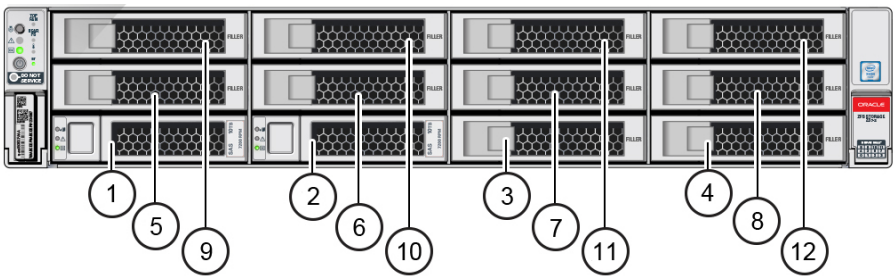

1 System Disk 0	**5** Filler panel	**9** Filler panel
2 System Disk 1	**6** Filler panel	**10** Filler panel
3 Filler panel	**7** Filler panel	**11** Filler panel
4 Filler panel	**8** Filler panel	**12** Filler panel

Oracle ZFS Storage ZS7-2 CPUs

The two CPU sockets are located under the air baffle behind the fan modules and ahead of the PCIe slots. As viewed from the front of the controller, the leftmost socket is designated P0 and the rightmost socket is designated P1.

- If the processor fault LED is off, the processor is operating properly.
- If the processor fault LED is on (amber), the processor has failed and must be replaced.

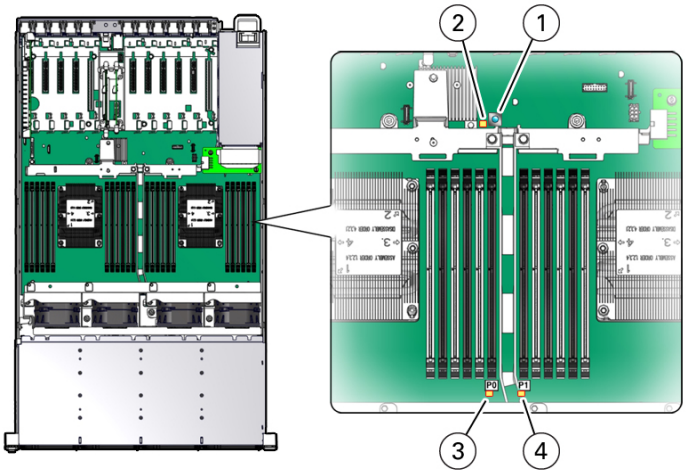

1 Fault Remind button	**3** Processor 0 fault LED
2 Fault Remind LED	**4** Processor 1 fault LED

Oracle ZFS Storage ZS7-2 PCIe Card Configuration

The 11 PCIe slots are located inside the controller at the back. As viewed from the back of the controller, the slots are divided into two groups, a group of six to the left of the built-in network ports and rear indicators and a group of five to the right side of the network ports and indicators. The PCIe slots are designated from left to right (back view) with the six slots on the left side designated as PCIe Slot 1 to PCIe Slot 6, and the five slots on the right side designated as PCIe Slot 7 to PCIe Slot 11. The following figure shows the Oracle ZFS Storage ZS7-2 high-end model.

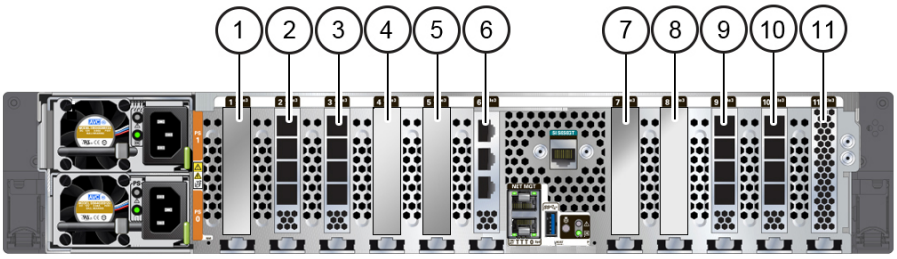

1 PCIe slot 1: Fifth PCIe option	**5** PCIe slot 5: Third PCIe option	**9** PCIe slot 9: • High-end model: External SAS-3 HBA • Mid-range model: Filler panel; slot not available
2 PCIe slot 2: External SAS-3 HBA	**6** PCIe slot 6: Cluster interface card	**10** PCIe slot 10: External SAS-3 HBA

3 PCIe slot 3: • High-end model: External SAS-3 HBA • Mid-range model: Filler panel; slot not available	**7** PCIe slot 7: Fourth PCIe option	**11** PCIe slot 11: Internal SAS-3 HBA
4 PCIe slot 4: First PCIe option	**8** PCIe slot 8: Second PCIe option	

PCIe Slot Width and Capabilities

The Oracle ZFS Storage ZS7-2 controller contains 11 low-profile PCIe 3.0 card slots, which have the following width and capabilities:

- Slot 1 has a x16 connector with x16 electrical interface
- Slot 7 has a x16 connector with x8 electrical interface
- Slots 2, 3, 4, 5, 6, 8, 9, 10 and 11 have x8 connectors (x8 electrical interface)

> **Note:**
>
> While slot 7 can accommodate a x16 width card, the capability of the card will be limited by the x8 electrical interface.

PCIe Card Slot Population Order

Each base unit Oracle ZFS Storage ZS7-2 is always loaded with the following HBAs or fillers:

- Slot 2: 4x4 SAS-3 HBA
- Slot 3: 4x4 SAS-3 HBA high-end model; filler panel mid-range model (slot not available)
- Slot 6: Cluster interface card
- Slot 9: 4x4 SAS-3 HBA high-end model; filler panel mid-range model (slot not available)
- Slot 10: 4x4 SAS-3 HBA
- Slot 11: Internal SAS-3 HBA

All allowable option Oracle ZFS Storage ZS7-2 PCIe cards should be installed in the remaining card slot locations and filled in accordance to the following priority and sequence as follows:

1. Install optional dual port 100GbE NIC cards into slot 1, then slot 7.
2. Install optional dual port QDR InfiniBand (IB) cards into slot 4, then slot 8, then slot 5, then slot 7, then slot 1.
3. Install optional dual port 40GbE NIC cards into slot 4, then slot 8, then slot 5, then slot 7, then slot 1.
4. Install optional dual port 32Gb FC HBA cards into slot 4, then slot 8, then slot 5, then slot 7, then slot 1.
5. Install optional quad port 10GbE NIC cards into slot 4, then slot 8, then slot 5, then slot 7, then slot 1.

6. Install optional quad port 10GBASE-T NIC cards into slot 4, then slot 8, then slot 5, then slot 7, then slot 1.
7. Install optional dual port 25GbE NIC cards into slot 4, then slot 8, then slot 5, then slot 7, then slot 1.

> **Note:**
>
> For Oracle ZFS Storage Appliance Racked System ZS7-2, install optional dual port 25GbE NIC cards into the same slots as described in number (7), except install them first, before any optional dual port QDR IB cards (number (2) in the previous list).

Oracle ZFS Storage ZS7-2 Attached Storage

The Oracle ZFS Storage ZS7-2 controller connects to external storage through 4-port (4x4) SAS HBAs. You can attach from 1 to 48 disk shelves (up to 1,152 drives) to the Oracle ZFS Storage ZS7-2 high-end controller, or 1 to 24 disk shelves (up to 576 drives) to the Oracle ZFS Storage ZS7-2 mid-range controller, including disk shelf configurations with fully populated flash-based SSD data devices. Oracle Storage Drive Enclosure DE3-24 disk shelves connect to the controller with active optical cables (AOCs), and Oracle Storage Drive Enclosure DE2-24 disk shelves (supported with software release OS8.8.4 and later) connect with copper cables. For disk shelf configuration details, see Disk Shelf Configurations.

6

Servicing the Oracle ZFS Storage ZS5-4 Controller

Be sure to read sections Getting Started with Servicing the Appliance and Prerequisites for Servicing Hardware.

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see Repairing Active Problems.

To replace Oracle ZFS Storage ZS5-4 hardware components, use the following tasks:

- Replacing an Oracle ZFS Storage ZS5-4 HDD
- Replacing an Oracle ZFS Storage ZS5-4 Fan Module
- Replacing an Oracle ZFS Storage ZS5-4 Power Supply
- Replacing an Oracle ZFS Storage ZS5-4 DIMM
- Replacing an Oracle ZFS Storage ZS5-4 Memory Riser Card
- Replacing or Adding an Oracle ZFS Storage ZS5-4 PCIe Card
- Replacing an Oracle ZFS Storage ZS5-4 System Battery

To understand the Oracle ZFS Storage ZS5-4 controller, use the following topics:

- Oracle ZFS Storage ZS5-4 Configuration Options
- Oracle ZFS Storage ZS5-4 Front and Rear Panel Components
- Oracle ZFS Storage ZS5-4 Specifications
- Oracle ZFS Storage ZS5-4 Internal Components
- Oracle ZFS Storage ZS5-4 CPUs and Memory Riser Card
- Oracle ZFS Storage ZS5-4 PCIe Card Configuration
- Oracle ZFS Storage ZS5-4 Attached Storage

Replacing an Oracle ZFS Storage ZS5-4 HDD

HDDs are hot-swappable and can be removed and installed while the storage controller is running without affecting other hardware capabilities. Use the following procedure to replace an Oracle ZFS Storage ZS5-4 HDD.

> **Note:**
>
> If you have more than one failed drive, replace only one drive at a time. Removing more than one drive in quick succession will cause a hardware/pool fault. Also, do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** BUI menu, select **System**, or navigate to `maintenance system updates` in the CLI.

1. Identify the faulted HDD.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, and open the detailed view of the controller. The faulted drive has either an amber or a blue-on-amber (ready to remove) status icon. Note the ID of the faulted drive. Click the entry for that drive in the table to highlight the location of the drive in the controller illustration. To view details about the faulted drive, click its information icon and then click **Active Problems**.
   - If you are physically at the system, the amber Service Action Required indicator on the HDD should be lit. Additionally, you can click the locator icon in the BUI to light the locator LED on the faulted drive.
2. Verify that the faulted drive is ready to be removed.

   > **Caution:**
   >
   > Removing the drive before it is ready to be removed can result in loss of data.

   Use any of the following methods to verify that the faulted drive is ready to be removed:
   - On the faulted drive, the Ready to Remove LED is lit as shown in Oracle ZFS Storage ZS5-4 Storage Drive.
   - In the BUI:
     - The status icon for the faulted drive in the detail view in **Maintenance: Hardware** has changed from amber to blue-on-amber.
     - The **Component Details** information window (click the information icon) for the faulted drive shows **Yes** as the value for the **Ready to remove** property.
   - In the CLI (and RESTful API), the value of the `readytoremove` property for the faulted drive is `true`:

     ```
     hostname:maintenance chassis-000 disk-001> get readytoremove
                     readytoremove = true
     ```
3. On the drive you want to remove, push the drive release button [1] to open the latch.
4. Grasp the latch [2] and pull the drive out of the drive slot [3].

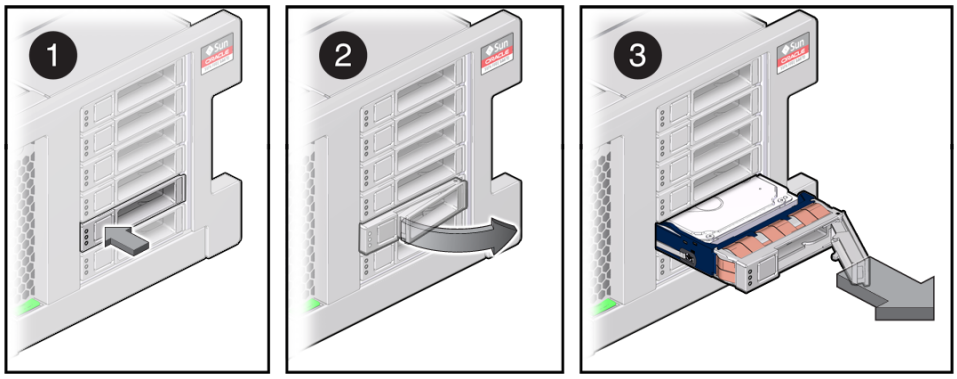

5. After at least 30 seconds, from the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon on the system controller to verify that the software has detected that the drive is not present.
6. On the replacement drive, open the release lever and slide the drive into the same slot [1] until it is fully seated. Close the latch [2] to lock the drive in place.

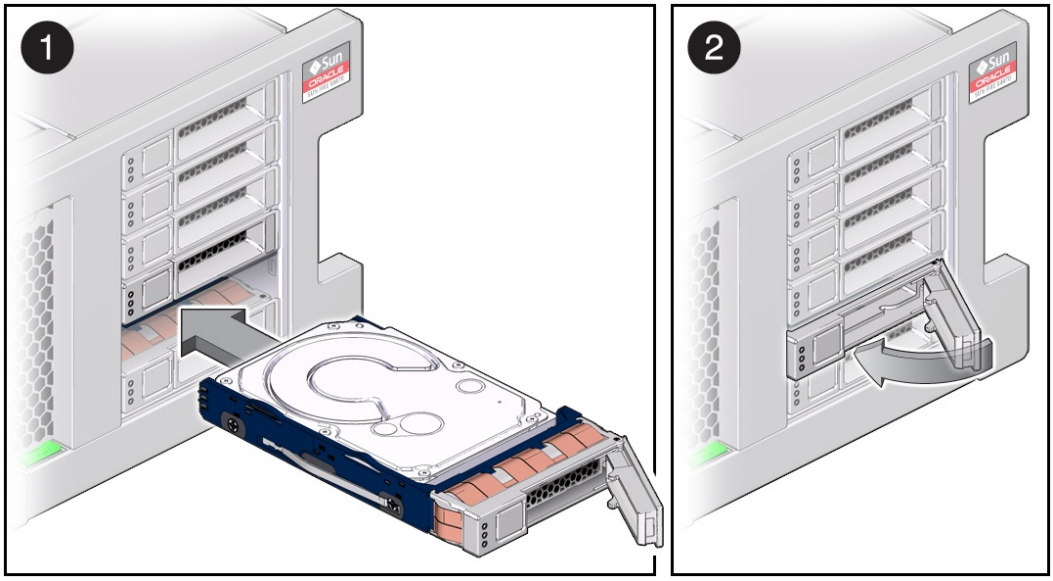

The Oracle ZFS Storage Appliance software automatically detects and configures the new drive.

7. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, and verify that the status icon is green for the newly installed drive. To view details about the new drive, click its information icon .
8. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS5-4 Fan Module

Unpack the replacement fan module and get it ready to insert into the controller chassis as soon as the faulted fan is removed.

> **Caution:**
>
> Operating a controller for an extended period of time with fans removed reduces the effectiveness of the cooling system and could damage the equipment. Therefore, this task needs to be done within a limited period of time.

Fan modules are hot-swappable and can be removed and installed while the storage controller is running without affecting other hardware capabilities. Use the following procedure to replace an Oracle ZFS Storage ZS5-4 fan module.

1. Identify the faulted fan module.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **Fan**, and note the ID of the faulted fan module, which has an amber icon. To highlight the location of the fan module in the controller picture, click on the fan module ID. To view details about the faulted fan module, click its information icon and then click **Active Problems**.
   - If you are physically at the system, the amber Service Action Required indicator on the fan module should be lit.
2. Extend the controller from the rack as described in Extending the Controller from the Rack.
3. Lift the green handle at the top of the fan module to unlock the module, and pull the module straight up and out.

> **Caution:**
>
> Excessive movement or rocking of the fan module during removal can cause damage to the internal connector on the fan module board. When removing a fan module, do not rock it back and forth.

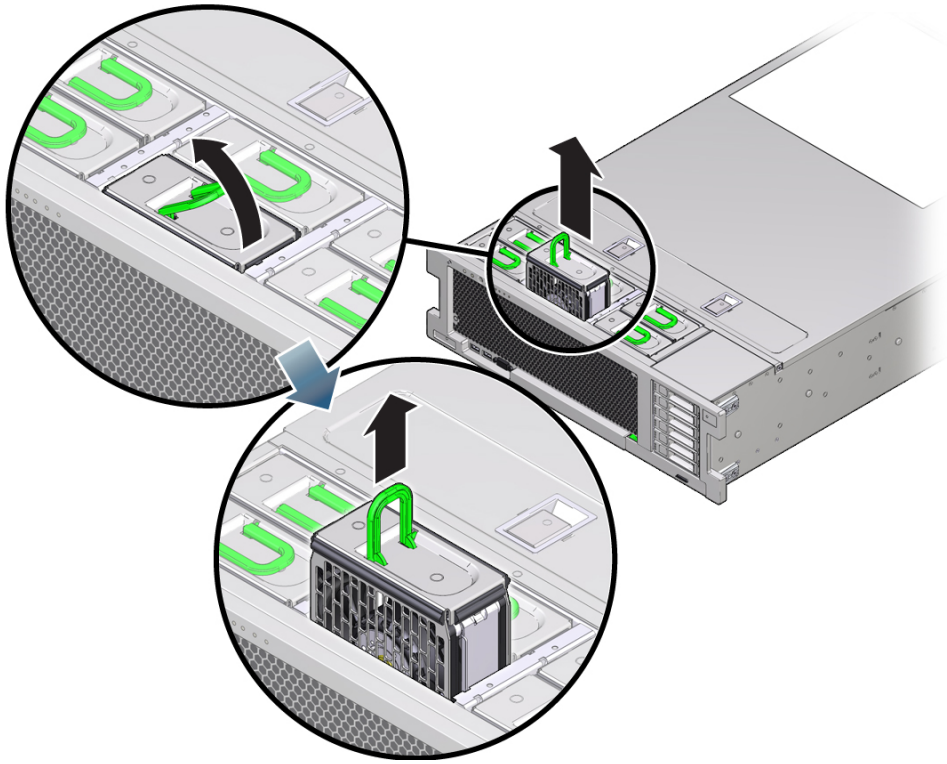

4. Position the replacement fan module so the connector on the bottom of the module is aligned with the connector on the inside of the slot and all keying and labeling are correctly positioned.

   The fan modules are keyed to ensure that they are installed in the correct orientation.

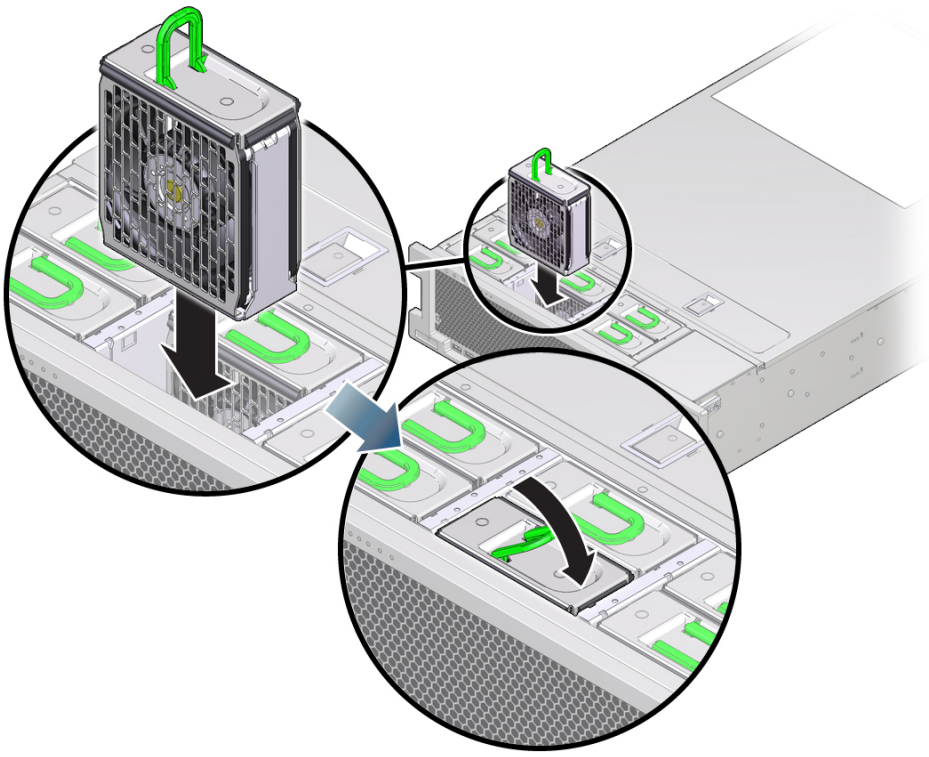

5. Slide the fan module into the slot until it stops.
6. Press downward on the top of the fan module, on the **Press Here to Latch** label, until the fan module is fully seated.
7. Replace the top cover on the chassis.
8. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.
9. Verify that the green OK indicator is lit.
10. Verify that the Top Fan indicator and the Service Action Required indicator on the controller front panel are not lit.
11. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **Fan**, and verify that the status icon is green for the newly installed fan module. To view details about the new fan module, click its information icon .
12. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS5-4 Power Supply

> **Caution:**
>
> If a power supply fails and you do not have a replacement available, do not remove the failed power supply. This maintains the proper air flow and prevents possible equipment damage.

Storage controllers are equipped with redundant, hot-swappable power supplies. Use the following procedure to replace an Oracle ZFS Storage ZS5-4 power supply.

1. Identify the faulted power supply.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **PSU**, and note the ID of the faulted power supply, which has an amber icon. To highlight the location of the power supply in the controller picture, click on the power supply ID. To view details about the faulted power supply, click its information icon and then click **Active Problems**.
   - If you are physically at the system, the amber Service Action Required indicator on the power supply should be lit.

   > **Note:**
   >
   > The fans of a failed power supply might still be spinning when the system is powered on. You can remove a power supply while the fans are still spinning.

2. Lift up on the power cord retaining clip to disengage it from the cord.
3. Disconnect the power cord from the faulted power supply.
4. To unlock the power supply handle, squeeze the release latch.

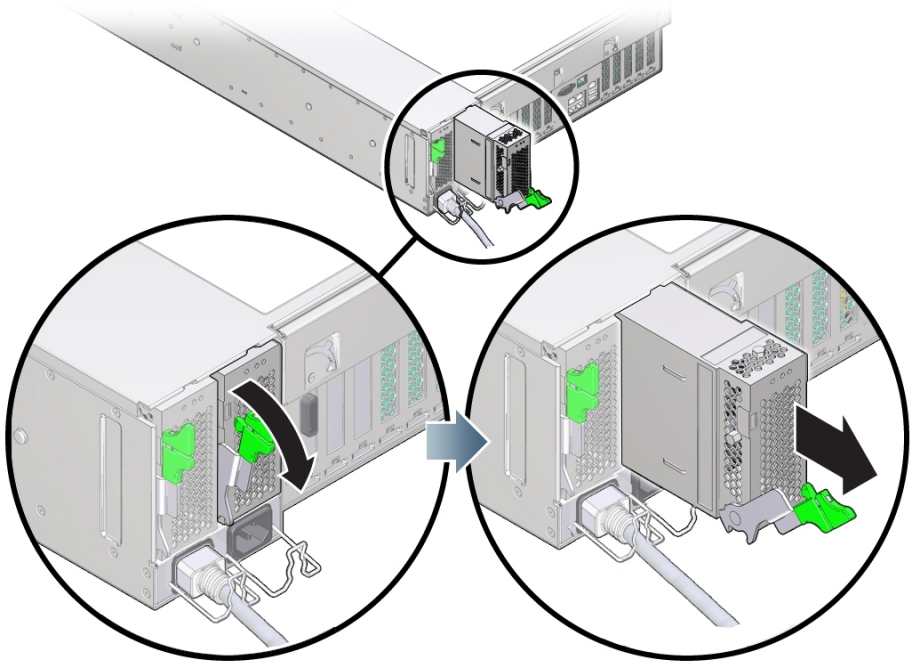

5. To disengage the power supply, rotate the release latch fully downward.

   This action disengages the power supply from the internal power supply backplane connector.
6. Shut down the controller using one of the power-off methods described in Powering Off the Controller.

7. Ensure that the replacement supply handle is in the open position.
8. Align the power supply with the empty power supply bay.
9. Slide the power supply into the bay [1] until it stops.
10. To seat the power supply, rotate the release handle upward [2] until it locks in place.

> **Note:**
>
> The latch at the hinge of the handle must engage the slot at the bottom of the power supply bay.

This action of rotating the handle upward draws the power supply into the controller and engages it with the internal connector.

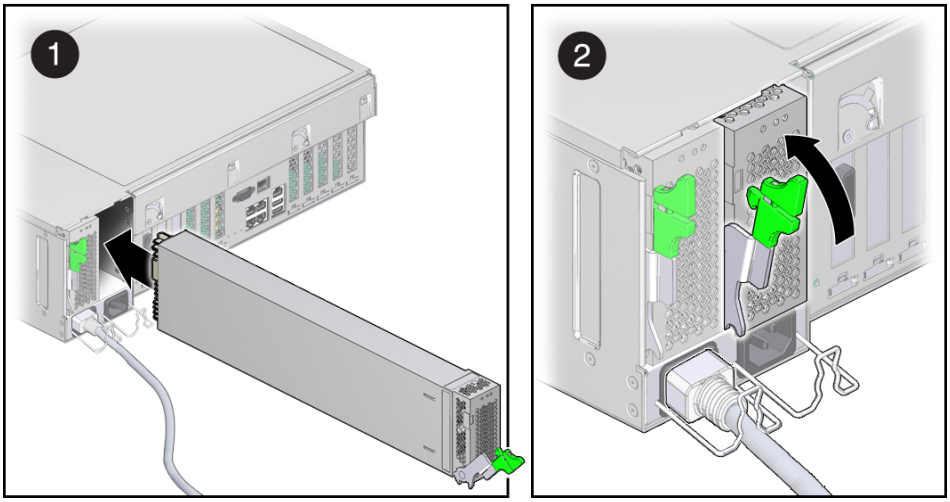

11. Connect the power cord to the power supply.
12. Verify that the green AC Present status indicator is lit.
13. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **PSU**, and verify that the status icon is green for the newly installed power supply. To view details about the new power supply, click its information icon .
14. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS5-4 DIMM

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

Removing a DIMM is a cold-service operation. You *must* power off the appliance before you remove the DIMM. There will be a loss of access to the storage unless the system is in a clustered configuration.

1. Identify the faulted DIMM.

   From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **DIMM**, and note the ID of the faulted DIMM, which has an amber icon. To highlight the location of the memory module in the controller picture, click on the module ID. To view details about the faulted module, click its information icon and then click **Active Problems**.
2. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
3. Disconnect the AC power cords from the rear panel of the storage controller.

   > **Caution:**
   >
   > Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.
4. Extend the controller from the rack as described in Extending the Controller from the Rack.
5. Remove the top cover as described in Removing the Controller Top Cover.
6. Remove the memory riser card in which the faulted DIMM resides. See Replacing an Oracle ZFS Storage ZS5-4 Memory Riser Card.
7. Rotate both DIMM slot ejector levers outward as far as they will go.

   The action unlocks the DIMM and expels it from the socket.

8. Carefully lift the DIMM straight up to remove it from the socket.

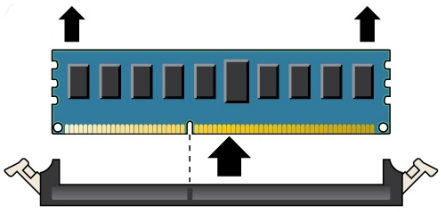

9. Repeat the above steps until all faulty DIMMs are removed.
10. Ensure that the DIMM ejector levers at both ends of the DIMM slot are in their fully open position.

11. Align the DIMM with the empty slot.

    Align the notch on the DIMM with the protrusion on the DIMM slot to ensure that the DIMM is installed correctly.
12. Gently and evenly push the DIMM into the slot until the ejector levers rise and lock the DIMM into place.

    The levers rise as the DIMM is pushed further into the slot.

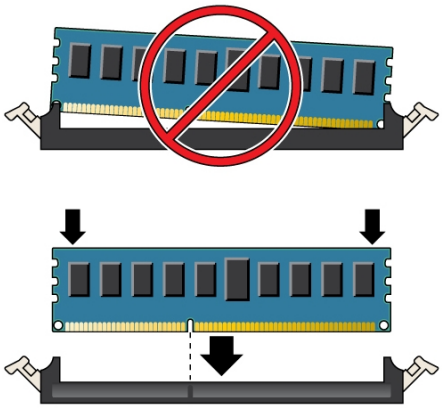

13. Ensure that the levers have risen completely and have locked the DIMM in the slot.

14. Install the replacement memory riser card as described in Replacing an Oracle ZFS Storage ZS5-4 Memory Riser Card.
15. Replace the top cover on the chassis.
16. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.
17. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
18. Reconnect the power cords to the power supplies.
19. Verify that standby power is on. Approximately two minutes after plugging in the power cords, the Power/OK status indicator on the front panel flashes.
20. Press and release the recessed Power button on the storage controller front panel.

    The Power/OK status indicator next to the Power button lights and remains lit.
21. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **DIMM**, and verify that the status icon is green for the newly installed module. To view details about the new module, click its information icon .
22. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS5-4 Memory Riser Card

Use the following procedure to replace an Oracle ZFS Storage ZS5-4 memory riser card.

1. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. Remove the top cover as described in Removing the Controller Top Cover.
5. To remove the memory riser card, pull the handles upward to disengage the connector from the motherboard, and carefully lift the memory riser straight up and out of the controller.

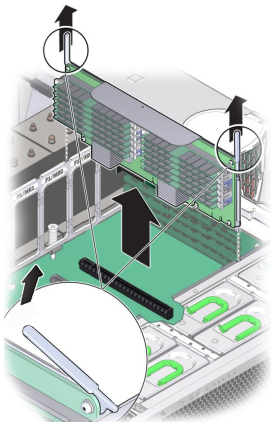

The memory riser card is designed to fit tightly into the memory riser connector on the motherboard. The handles on the memory riser card provide leverage against the chassis to disengage the card's edge connector from the connector on the motherboard.

6. On the replacement memory riser card, ensure that all populated and unpopulated DIMM slot ejector levers are in the closed and locked position.

> **Caution:**
>
> Open DIMM ejector levers can break off during the installation of the memory riser card. All populated and unpopulated DIMM slot ejector levers on the memory riser card must be in the fully closed and locked position before installing the card in the controller.

7. Ensure that all populated and unpopulated DIMM slot ejector levers are in the closed and locked position.

> **Caution:**
>
> The ejector levers are only used to remove a memory riser card and are not used to install a card.

8. Position the memory riser card over the slot on the motherboard.

   The DIMMs on the card must face to the left (when positioned at the front of the controller).

9. Lower the memory riser card into the controller and set it on the slot on the motherboard.

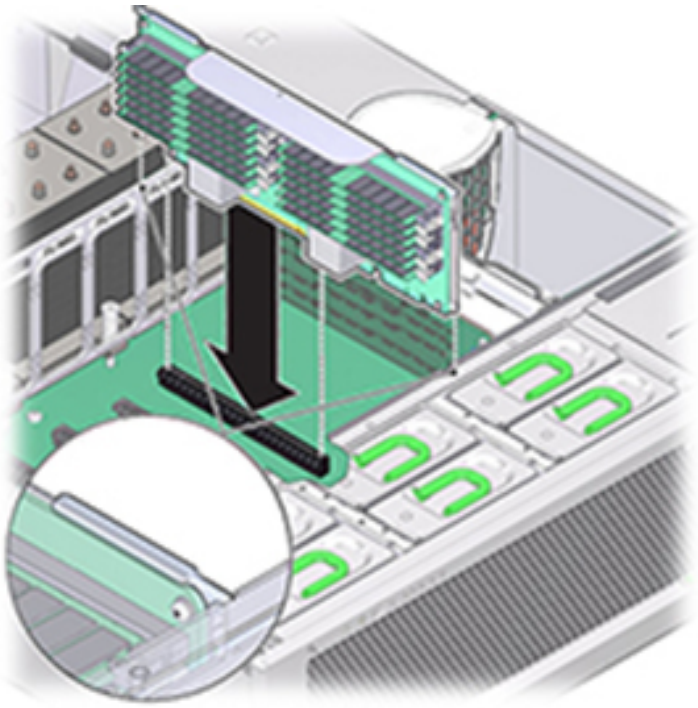

10. Ensure that the memory riser card connector is aligned with the slot.
11. Firmly push down on the metal bracket on top of the card, to seat the card within the connector on the motherboard.
12. Replace the top cover on the chassis.
13. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.
14. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
15. Reconnect the power cords to the power supplies.
16. Verify that standby power is on, indicated by the Power/OK status indicator flashing on the front panel about two minutes after the power cords are plugged in.
17. Press and release the recessed Power button on the storage controller front panel.
18. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller and then click **DIMM** to verify that the status icon is green for the newly installed memory riser card.
19. If a fault does not self-clear, see Repairing Active Problems.

Replacing or Adding an Oracle ZFS Storage ZS5-4 PCIe Card

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

Replacing or adding a PCIe card is a cold-swap procedure, and you must remove power from the storage controller. If you are installing a newly released HBA, upgrade your system software before installing the HBA as described in Upgrading the Software. All HBAs must be of the same type. Use the following procedure to replace or add a ZS5-4 PCIe card.

1. Identify the PCIe card slot location.
   - If replacing a faulted PCIe card, identify the PCIe card to be replaced.

     From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **Slot**, and note the ID of the faulted PCIe card, which has an amber icon. If replacing a non-faulted PCIe card, note its location. To highlight the location of the card in the controller picture, click on the card ID. To view details about the card, click its information icon and then click **Active Problems**.
   - If adding a new PCIe card, identify the correct slot location as described in Oracle ZFS Storage ZS5-4 PCIe Card Configuration.
2. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
3. Disconnect the AC power cords from the rear panel of the storage controller.

   > **Caution:**
   >
   > Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.
4. Extend the controller from the rack as described in Extending the Controller from the Rack.
5. Remove the top cover as described in Removing the Controller Top Cover.
6. Locate the PCIe card position in the storage controller.
7. To disengage the PCIe card retaining bar, push down on the bar, move it away from the controller back wall, and lift it to an upright position.

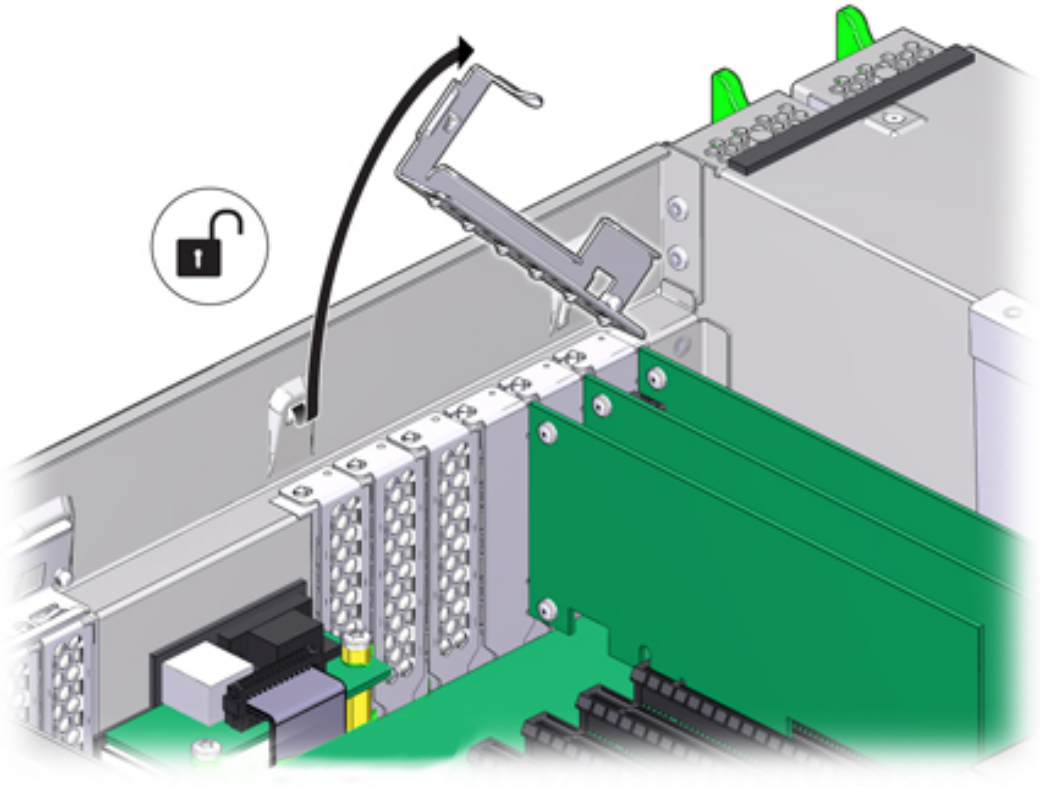

8. Carefully remove the PCIe card from the PCIe card slot.

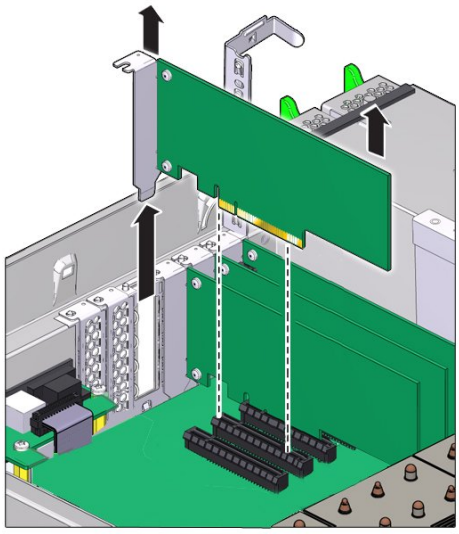

> **Caution:**
>
> Whenever you remove a PCIe card, you should replace it with another PCIe card or a filler panel; otherwise, the controller might overheat due to improper airflow.

9. Ensure the PCIe card retaining bar is open.
10. Install the replacement PCIe card into the PCIe card slot.

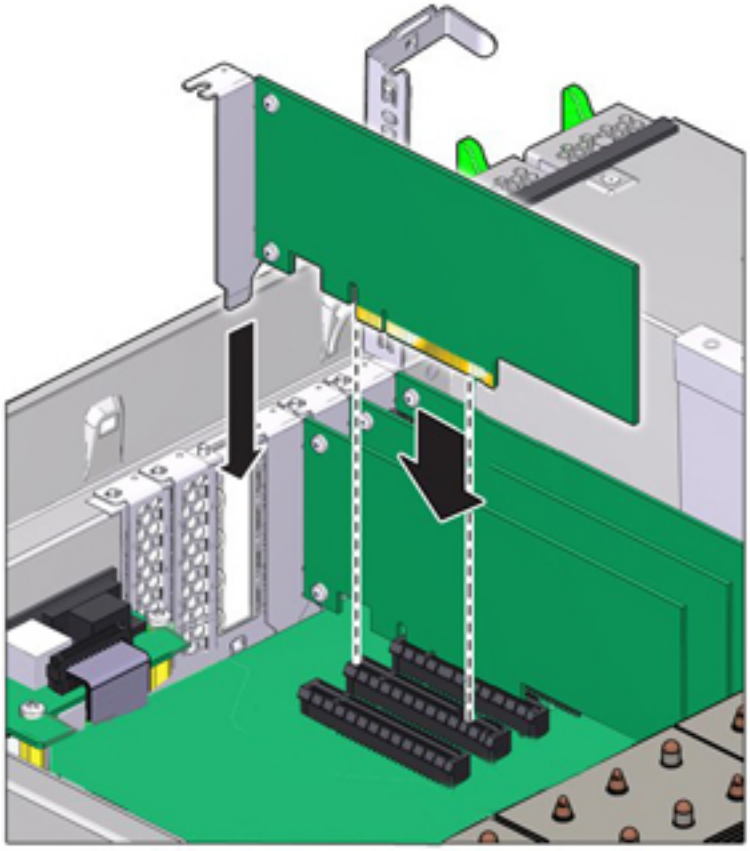

11. Return the PCIe card slot crossbar to its closed and locked position.

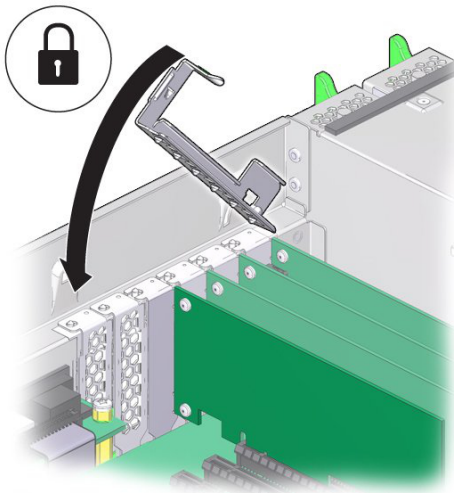

12. Replace the top cover on the chassis.
13. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.
14. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
15. Reconnect the power cords to the power supplies.
16. Verify that standby power is on. Approximately two minutes after plugging in the power cords, the Power/OK status indicator on the front panel flashes.
17. Press and release the recessed Power button on the storage controller front panel.

    The Power/OK status indicator next to the Power button lights and remains lit.
18. Connect data cables to the PCIe card, routing them through the cable management arm, if used.
19. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **Slot**, and verify that the status icon is green for the newly installed card. To view details about the new card, click its information icon .
20. If you replaced an HBA, connect the disk shelves to its ports as described in the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x. If you installed a new HBA, install the disk shelf as described in Installation Prerequisites and Hardware Overviews in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*, and connect the expansion storage as described in the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.

Replacing an Oracle ZFS Storage ZS5-4 System Battery

Replacing a system battery is a cold-swap procedure, and you must remove power from the storage controller. You might need to replace the battery if the storage controller fails to maintain the proper time when powered off and disconnected from the network. You will need a small (No.1 flat-blade) non-metallic screwdriver or equivalent. Use the following procedure to replace an Oracle ZFS Storage ZS5-4 system battery.

1. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. Remove the top cover as described in Removing the Controller Top Cover.

   The battery [1] is located toward the back of the controller between PCIe slot 1 and the power supply side wall.

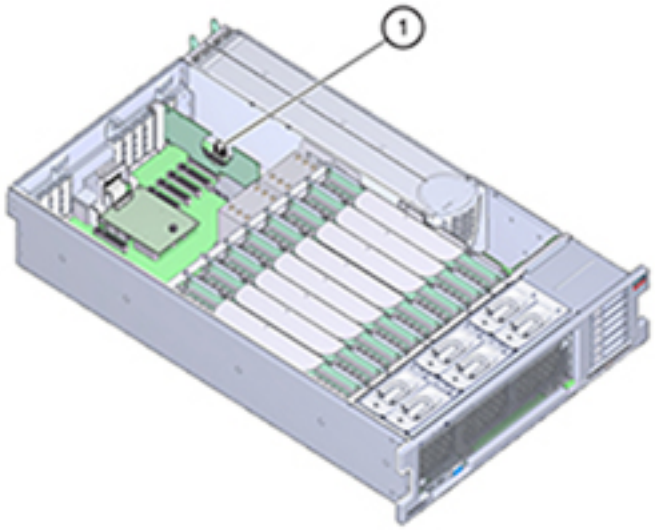

5. If necessary, remove the PCIe card in slot 1 to access the battery. See Replacing or Adding an Oracle ZFS Storage ZS5-4 PCIe Card.
6. Remove the battery from the battery holder by pushing the back of the battery, negative (-) side, toward the metal tab, on the positive (+) side, and lifting the battery up and out of the battery holder.

   Take care not to deform the tab when pushing against it.

> **Caution:**
>
> Do not deform the metal tab on the positive (+) side of the battery. The metal tab maintains the positive battery connection and secures the battery in the holder.

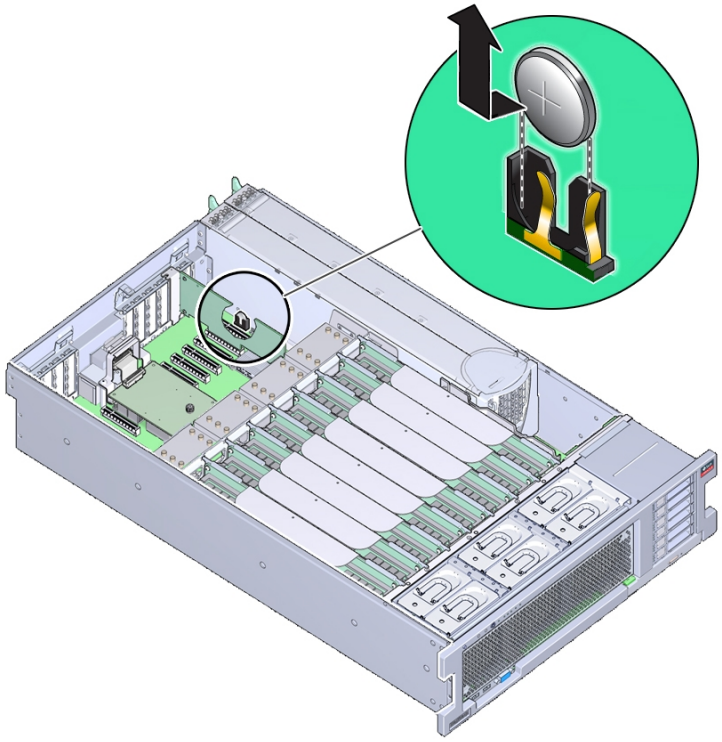

7. Press the new battery into the battery holder with the positive (+) side facing the metal tab that holds it in place.

   The positive side of the battery is marked with a plus sign (cross).

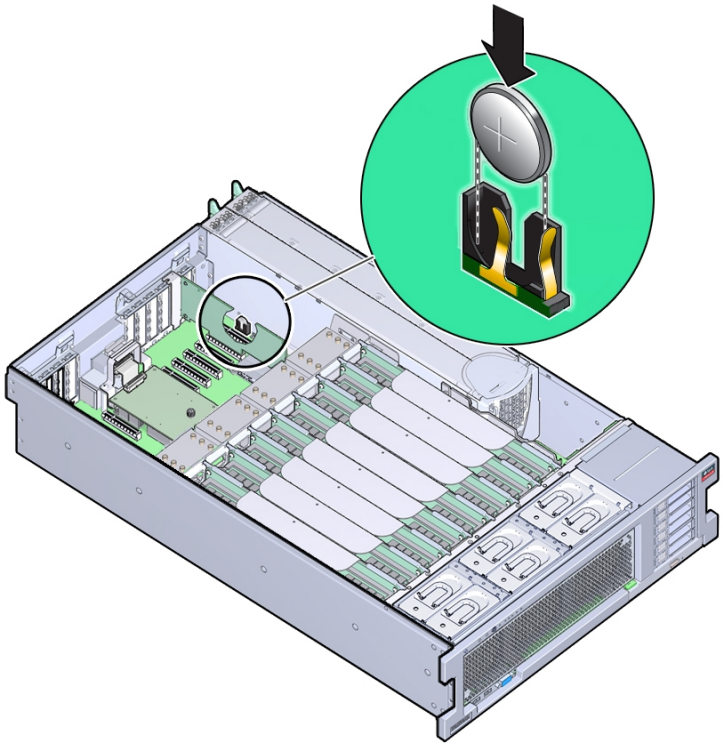

8. If you removed the PCIe card from slot 1, re-install it. See Replacing or Adding an Oracle ZFS Storage ZS5-4 PCIe Card.
9. Replace the top cover on the chassis.
10. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.
11. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
12. Connect the power cords to the power supplies.
13. Verify that standby power is on, indicated by the Power/OK status indicator flashing on the front panel about two minutes after the power cords are plugged in.
14. Press and release the recessed Power button on the storage controller front panel. The Power/OK status indicator next to the Power button lights and remains lit.
15. If you re-installed a PCIe card, connect the data cables to the PCIe card, routing them through the cable management arm, if used.
16. When the system has finished booting, log in and set the time using the steps in the BUI Clock task. For more information, see Setting Clock Synchronization (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Oracle ZFS Storage ZS5-4 Controller Hardware Overview

This section describes the internal and external components of the Oracle ZFS Storage ZS5-4 controller.

- Oracle ZFS Storage ZS5-4 Configuration Options
- Oracle ZFS Storage ZS5-4 Front and Rear Panel Components
- Oracle ZFS Storage ZS5-4 Specifications
- Oracle ZFS Storage ZS5-4 Internal Components
- Oracle ZFS Storage ZS5-4 CPUs and Memory Riser Card
- Oracle ZFS Storage ZS5-4 PCIe Card Configuration
- Oracle ZFS Storage ZS5-4 Attached Storage

Oracle ZFS Storage ZS5-4 Configuration Options

Oracle ZFS Storage ZS5-4 incorporates one or two controllers and one to thirty-six disk shelves. Each controller has eleven PCIe Gen3 slots. Each controller includes six 2.5" disk slots that accommodate two system disks. Read cache devices are installed in disk shelves; therefore, the remaining four disk slots have permanent fillers. The following table describes the configuration options.

CPU	Memory	System Disks	PCIe Slots	HBA
4x18-core, 2.6 GHz	1536 GB	2x2.5" SAS-3	Four dedicated; seven available	4-port (4x4) SAS-3 12 Gb/s external 4-port (4x4) SAS-2 6 Gb/s external

Oracle ZFS Storage ZS5-4 Front and Rear Panel Components

Oracle ZFS Storage ZS5-4 Front Panel Components - The Oracle ZFS Storage ZS5-4 controller drive slots and front panel components are shown in the following figure.

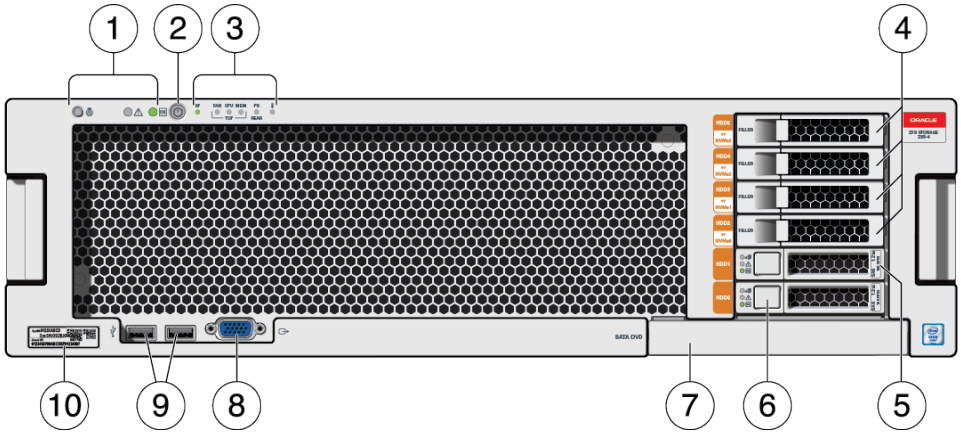

1 LEDs (left to right) • Locator • Service action required • Power/OK	**2** Power button	**3** Alert indicators (left to right) • Service Processor (SP) OK/ Fault • Fan/ CPU/ Memory service action required • Power Supply (PS) service action required • Over-temperature warning
4 Permanent HDD filler panels	**5** System disk 1	**6** System disk 0
7 DVD filler panel	**8** DB-15 VGA port	**9** Two USB 2.0 ports
10 RFID tag		

Oracle ZFS Storage ZS5-4 Rear Panel Components - The Oracle ZFS Storage ZS5-4 controller PCIe slots and rear panel components are shown in the following figure.

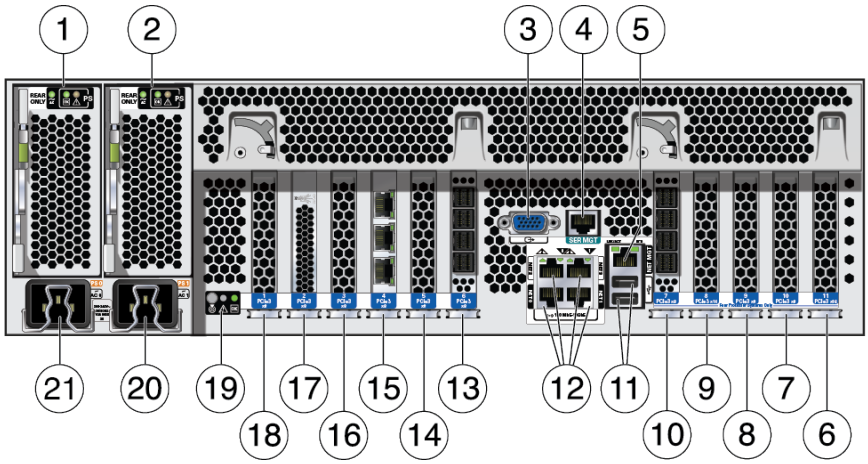

1 Power Supply Unit (PSU) 0	**2** PSU 1	**3** DB-15 VGA port
4 SER MGT port	**5** Service Processor (SP) NET MGT port	**6** First PCIe (slot 11)
7 Fifth PCIe (slot 10)	**8** Third 4x4 SAS HBA or seventh PCIe (slot 9)	**9** Second PCIe (slot 8)
10 4X4 SAS HBA (slot 7)	**11** Two USB 2.0 ports	**12** Network (NET) 10 Gb Ethernet ports, NET 0 - NET 3
13 4X4 SAS HBA (slot 6)	**14** Third PCIe (slot 5)	**15** Cluster interface card (slot 4)
16 Fourth 4X4 SAS HBA or sixth PCIe (slot 3)	**17** Internal SAS HBA (slot 2)	**18** Fourth PCIe (slot 1)
19 System status LEDs (left to right) • Power/OK (green) • Attention (amber) • Locate (white)	**20** PSU 1 AC inlet	**21** PSU 0 AC inlet

Cluster Interface Card Compatibility - The Oracle ZFS Storage ZS5-4 controller contains a Version 3 cluster interface card, which is incompatible with Version 2.

For information about how to connect cables to form a cluster, see Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

Ethernet Ports

The Oracle ZFS Storage ZS5-4 controller has four RJ-45 10-Gigabit Ethernet (10GbE) network connectors, labeled NET0, NET1, NET2, NET3 in a quadrant (where NET0 and NET1 are bottom left to right, NET2 is above NET0 and NET3 is above NET1) on the rear panel. The ports operate at 100 Mbits/sec, 1000 Mbit/sec, or 10-Gbits/sec. Use these ports to connect the appliance to the network.

There are two status indicators (LEDs) on each port. These indicators are visible from the rear of the controller.

Table 6-1 Gigabit Ethernet Port Status Indicators

Status Indicator Name	Location	Color	State and Meaning
Activity	Top left	Green	ON – No link OFF– No activity FLASHING – Packet activity
Link speed	Top right	Bicolored: Amber/Green	OFF – 100BASE-T link Amber ON – 1000BASE-T link Green ON – 10GBASE-T link

Network Management Port

The Oracle ZFS Storage ZS5-4 controller has one 10/100/1000BASE-T Ethernet management domain interface, labeled NET MGT. There are two status indicators (LEDs) on this port. These indicators are visible from the rear of the controller. The

service processor NET MGT port is the optional connection to the Oracle ILOM service processor. The NET MGT port is configured by default to use Dynamic Host Configuration Protocol (DHCP). The service processor NET MGT port uses an RJ-45 cable for a 10/100/1000BASE-T connection.

Table 6-2 Network Management Port Status Indicators

Status Indicator Name	Location	Color	State and Meaning
Activity	Top left	Amber	ON – Link up OFF – No link or down link FLASHING – Packet activity
Link speed	Top right	Green	Green ON – 1000BASE-T link OFF – 10/100BASE-T link

Serial Management Port

The Oracle ZFS Storage ZS5-4 service processor SER MGT port uses an RJ-45 cable and is the default connection to the Oracle ILOM service processor. This port supports local connections to the controller and only recognizes Oracle ILOM command-line interface (CLI) commands. Typically you connect a terminal or terminal emulator to this port.

> **Note:**
>
> This port does not support network connections.

4x4 SAS HBA Ports

The 4x4 SAS HBA, installed in PCIe slot 6 or slot 7 of the Oracle ZFS Storage ZS5-4 controller, provides connectivity to external disk shelves. The HBA logical ports are numbered 3-0, top to bottom, as shown in the following figure.

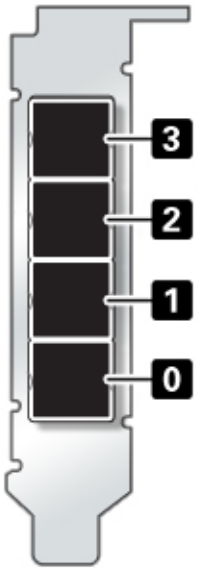

Oracle ZFS Storage ZS5-4 Specifications

Oracle ZFS Storage ZS5-4 Physical Specifications: The Oracle ZFS Storage ZS5-4 controller chassis dimensions are as follows:

Dimension	Measurement	Dimension	Measurement
Height	5.1 in./129.9 mm	Depth	28.8 in./732.0 mm
Width	17.2 in./436.5 mm	Weight (Maximum)	88 lb./40 kg

Oracle ZFS Storage ZS5-4 Electrical Specifications: The Oracle ZFS Storage ZS5-4 controller has two hot-swappable rear accessible power supplies with N+N redundancy. To understand the overall power consumption, see Oracle's online power calculator at http://www.oracle.com/goto/powercalculators .

Parameter	AC Power Supplies
Nominal input frequencies	50/60 Hz
Operating input voltage range	200-240 VAC
Rated input current	10 A @ 200 to 240 VAC
Maximum power consumption	200W
Maximum heat output	6824 BTU/hr

Oracle ZFS Storage ZS5-4 Acoustic Noise Emissions: In compliance with the requirements defined in ISO 9296, the workplace-dependent noise level of this product is less than 75 dBA (bystander position, max. ambient).

Oracle ZFS Storage ZS5-4 Internal Components

The components of the Oracle ZFS Storage ZS5-4 controller are shown in the following figure and identified in the legend.

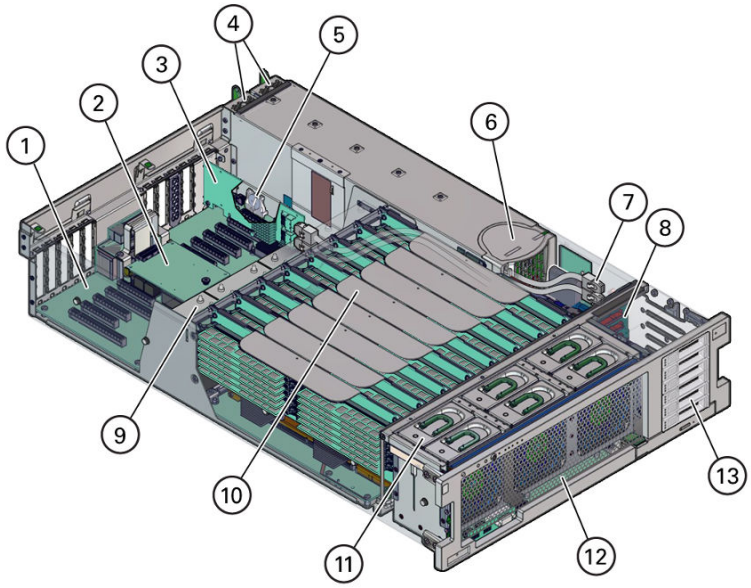

1 Motherboard	**8** Storage drive backplane board
2 SP card	**9** Heatsinks and CPUs (4)
3 HBA card	**10** Memory riser cards (8)
4 Power supplies (2)	**11** Fan modules (6)
5 System battery	**12** Fan board
6 Power supply backplane	**13** Storage drive slots (2)
7 HBA SAS cables (2)	

Oracle ZFS Storage ZS5-4 Internal Components (Exploded View)

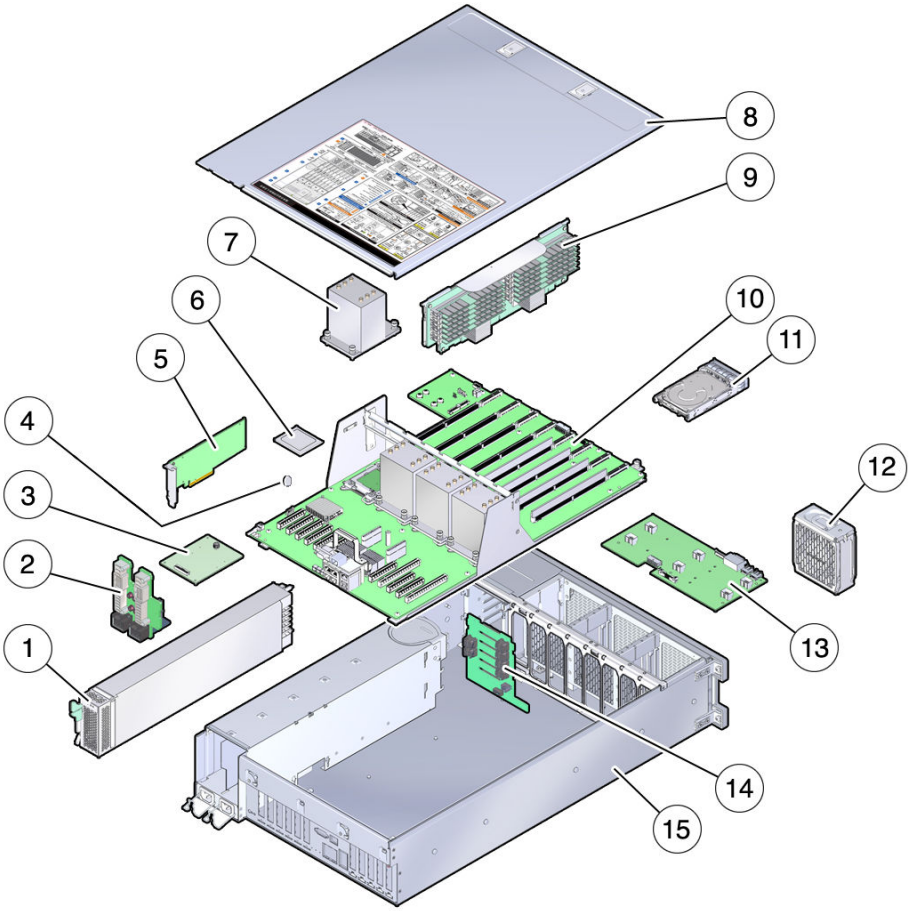

1 Power supplies	**9** Memory riser card
2 Power supply backplane board	**10** Motherboard
3 SP card	**11** Storage drive
4 System battery	**12** Fan module
5 HBA card	**13** Fan board
6 CPU	**14** Storage drive backplane board
7 Heatsink	**15** Controller chassis
8 Cover	

Oracle ZFS Storage ZS5-4 Storage Drive

This section shows the location and explains the functions of the storage drive status indicators.

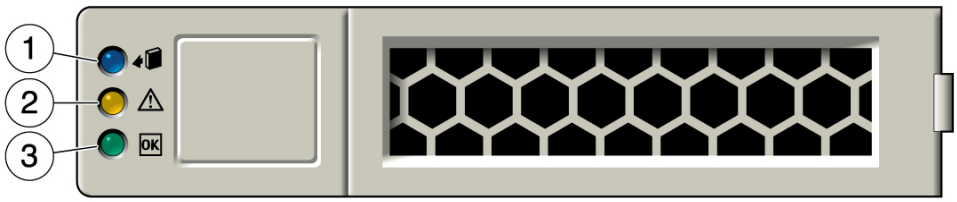

1 Ready to Remove (blue)	**2** Service Action Required (amber)	**3** OK/Activity (green)

Oracle ZFS Storage ZS5-4 Fan Module

The six fan module slots are at the front of the Oracle ZFS Storage ZS5-4 controller and are set in two rows of three slots. The slots are designated from left to right. As pictured in the illustration below, the three front row slots are designated as: FM0, FM1, and FM2. The three back row slots are: FM3, FM4, and FM5.

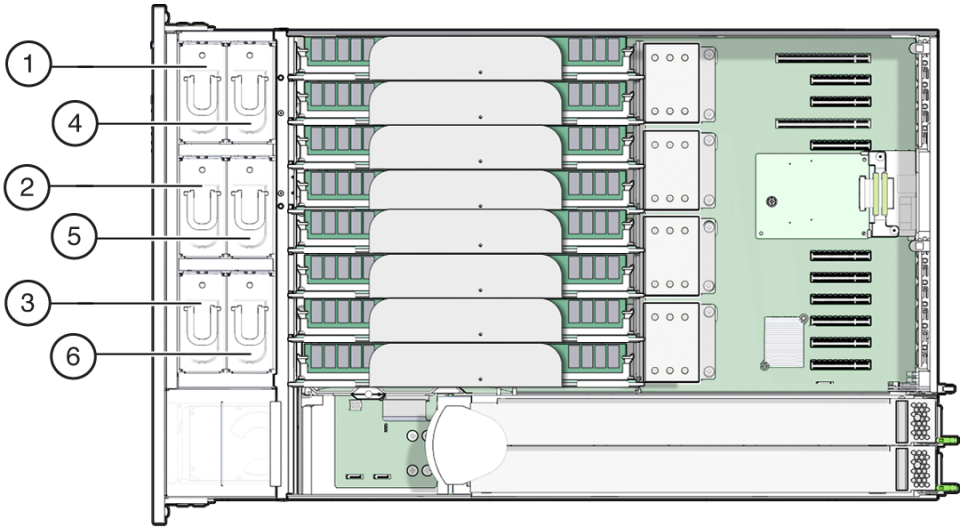

1 Fan Module, FM 0	**3** Fan Module, FM 2	**5** Fan Module, FM 4
2 Fan Module, FM 1	**4** Fan Module, FM 3	**6** Fan Module, FM 5

Oracle ZFS Storage ZS5-4 Power Supply

The two power supply slots are located on the right side of the Oracle ZFS Storage ZS5-4 controller (from the front of the controller) and are designated from right to left. The slots are accessible from the back of the controller. From the back of the controller, the left slot is designated as PS-0, and the right slot is PS-1.

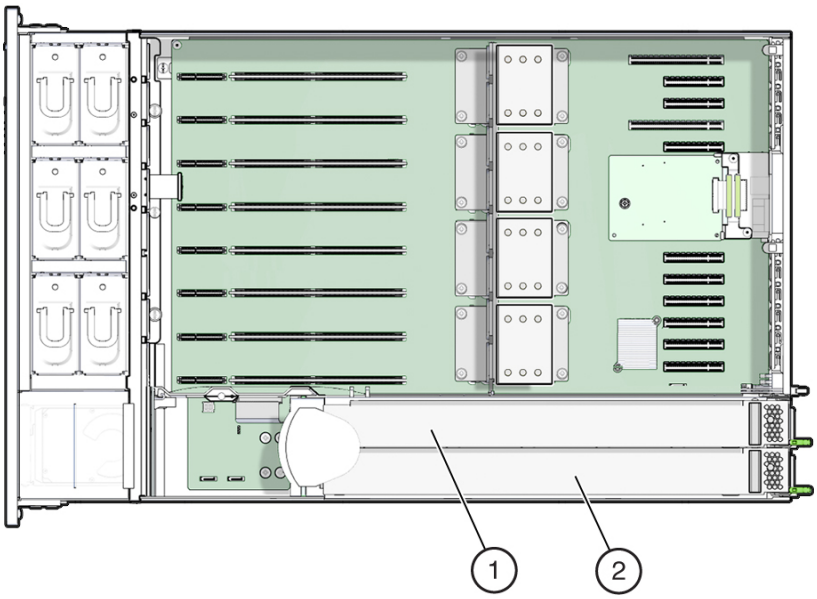

1 PS 1	**2** PS 0

The two power supplies ensure redundancy by allowing the controller to operate when one of the power supplies fails. However, when the controller is operating on one power supply, redundancy no longer exists, and the risk for an unexpected shut down and a loss of data is high. When a power supply or any component that is part of a redundant configuration fails, replace it immediately.

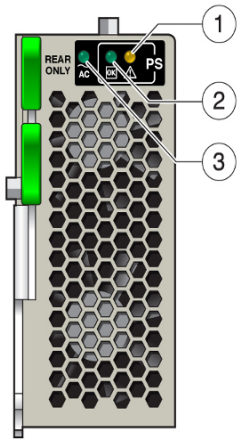

1 Service Action Required (amber)	**2** DC OK (green)	**3** AC OK (green)

Oracle ZFS Storage ZS5-4 DIMM

The DIMM slots are located on the memory riser cards. The DIMMs are arranged in two banks of six slots for a total of 12 slots. The slots are designated numerically from

top to bottom. The left bank of slots are designated as D0–D6. The right bank of slots are designated as D7–D11.

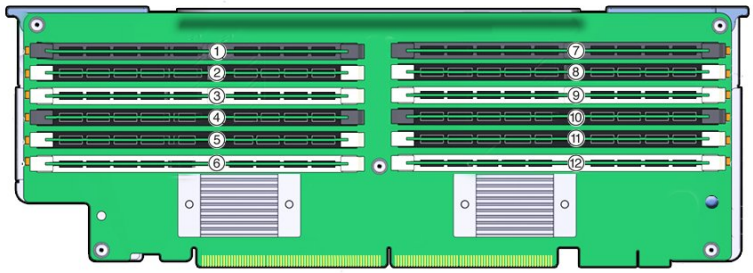

1 Slot D0	**4** Slot D3	**7** Slot D6	**10** Slot D9
2 Slot D1	**5** Slot D4	**8** Slot D7	**11** Slot D10
3 Slot D2	**6** Slot D5	**9** Slot D8	**12** Slot D11

Memory Riser layout:

The following illustration shows the DIMM slot and channel designations on the memory riser card:

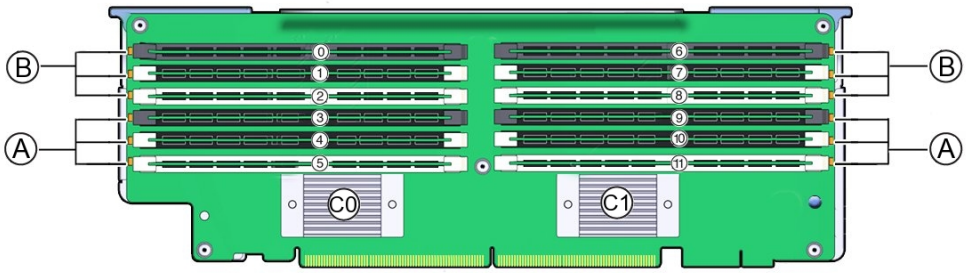

0 - **11** DIMM slot numbers. Board labeling prepends a "D" to the slot number (for example: D0–D11).
B Channel B slots: D0–D2 for memory buffer 0, D6–D8 for memory buffer 1.
A Channel A slots: D3–D5 for memory buffer 0, D9–D11 for memory buffer 1.
C0 Memory buffer 0
C1 Memory buffer 1

Oracle ZFS Storage ZS5-4 Storage Drive and USB

The storage drive slots are on the right side of the Oracle ZFS Storage ZS5-4 controller and are designated consecutively from bottom to top.

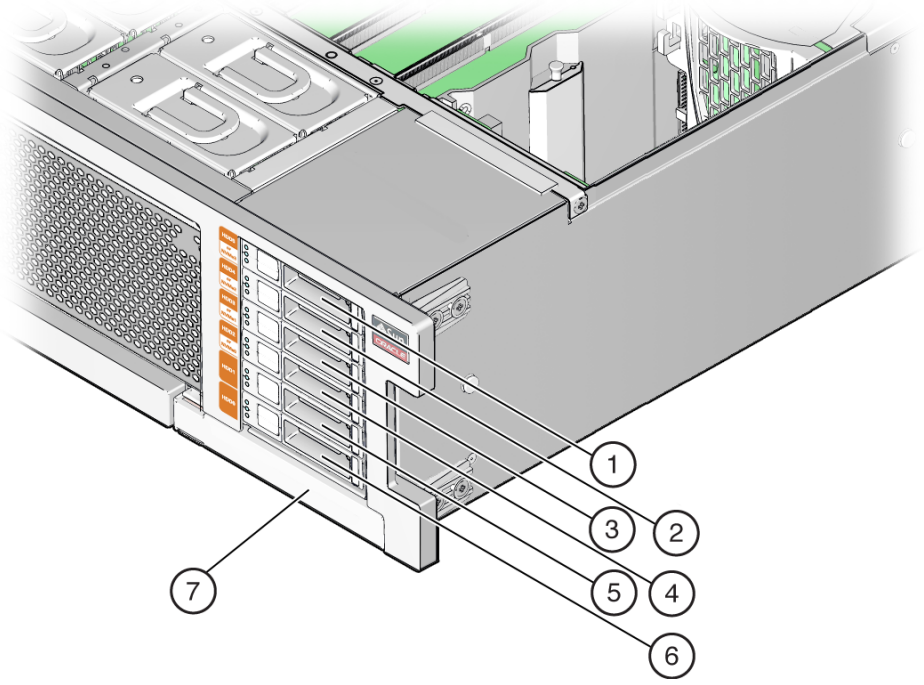

1 Drive Filler Panel	**5** Secondary Boot Drive
2 Drive Filler Panel	**6** Primary Boot Drive
3 Drive Filler Panel	**7** DVD Filler Panel
4 Drive Filler Panel	

Oracle ZFS Storage ZS5-4 CPUs and Memory Riser Card

The four CPU sockets are located in the middle of the Oracle ZFS Storage ZS5-4 controller and are designated consecutively from right to left (from the front of the controller). The rightmost socket is CPU-0 and is designated as P0, and the leftmost socket is CPU-3, designated as P3.

The eight memory riser (MR) card slots are located between the fan module slots and the CPU sockets. Consecutively from right to left, the rightmost slot is slot 0, and the leftmost slot is slot 7.

The slots are also designated by their association with the four CPU sockets (P0-P3). Two slots are assigned to each CPU socket. For example, slots 0 and 1 are paired with CPU socket, P0, and are designated as P0/MR0 and P0/MR1. Slots 2 and 3 are paired with CPU socket, P1 and are designated as P1/MR0 and P1/MR1. This numbering pattern continues for the remaining slots.

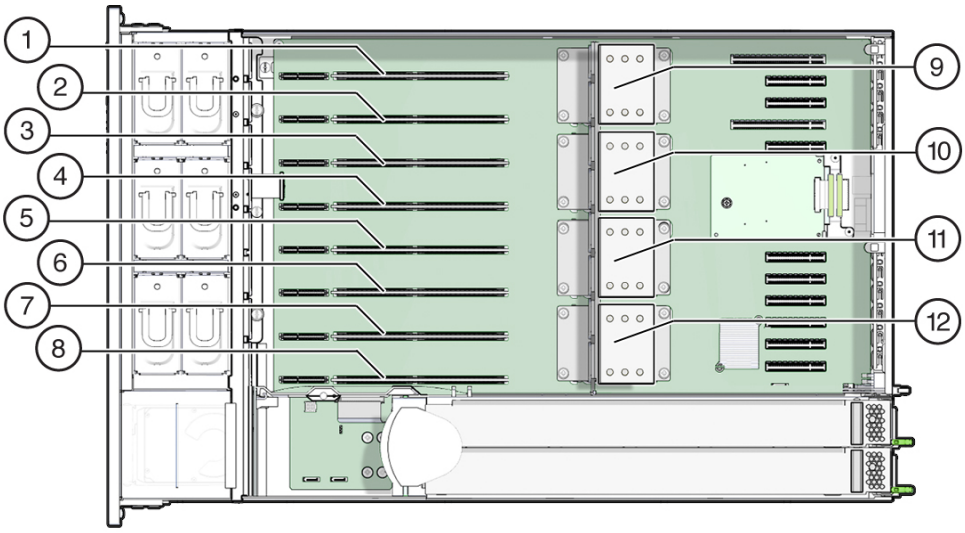

1 MR card slot P3/MR1	**5** MR card slot P1/MR1	**9** CPU-3 (P3)
2 MR card slot P3/MR0	**6** MR card slot P1/MR0	**10** CPU-2 (P2)
3 MR card slot P2/MR1	**7** MR card slot P0/MR1	**11** CPU-1 (P1)
4 MR card slot P2/MR0	**8** MR card slot P0/MR0	**12** CPU-0 (P0)

Memory Riser Card Components

The following illustration shows the components of the Oracle ZFS Storage ZS5-4 memory riser card.

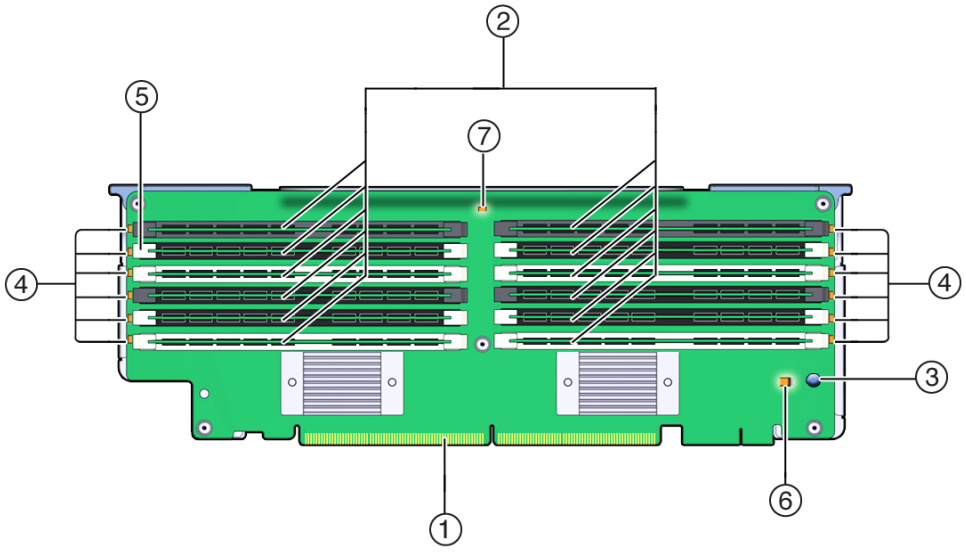

1 Connector	**5** DIMM ejector/locking levers (two on each slot)
2 DIMM slots (12)	**6** Charge status indicator (green)

3 Fault Remind button	**7** Memory riser card fault indicator
4 DIMM fault indicators	

Memory Riser Cards Physical Layout

The Oracle ZFS Storage ZS5-4 memory riser cards are located behind the fan module bay. The designations each memory riser (MR) card and its associated CPU are located inside the Oracle ZFS Storage ZS5-4 controller on labels affixed to the rear MR card bracket. The cards are designated as MR, and the CPUs are designated as P. The slots and CPUs are labeled from right to left when viewed from the front of the controller as follows:

> **Note:**
>
> Each CPU has two assigned card slots (MR0 and MR1).

- P0/MR0 (rightmost slot)
- P0/MR1
- P1/MR0
- P1/MR1
- P2/MR0
- P2/MR1
- P3/MR0
- P3/MR1 (leftmost slot)

The following illustration shows the numbering for the memory riser slots and their associated CPUs as described earlier:

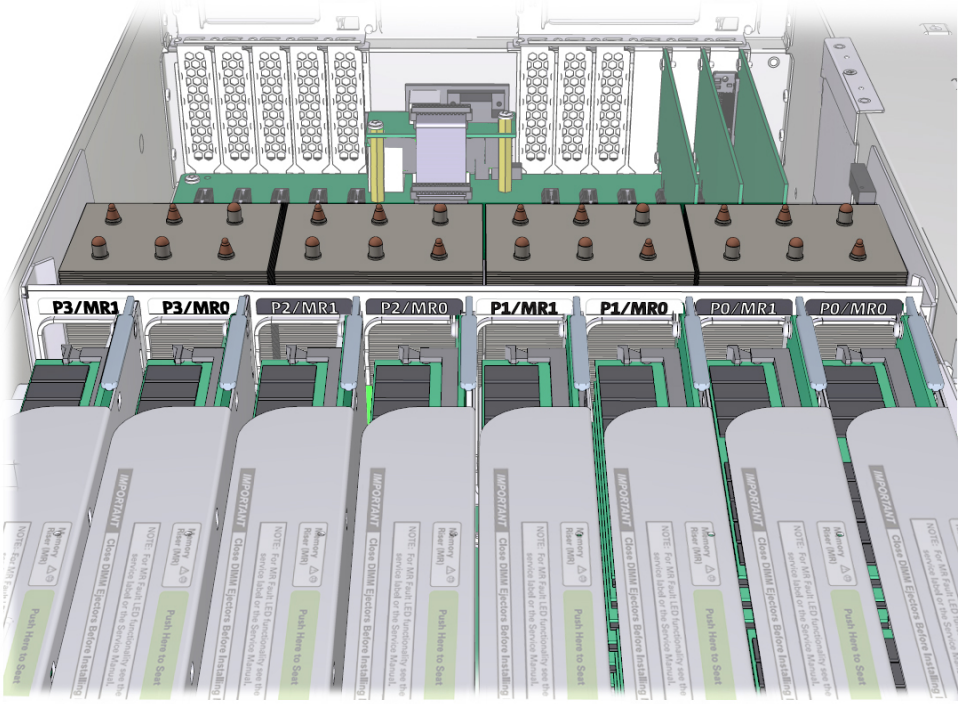

Memory Riser Card Population Rules

The memory riser population rules for the Oracle ZFS Storage ZS5-4 controller are as follows:

1. Each CPU has two dedicated memory riser card slots (MR0 and MR1).
2. The dedicated memory riser slots for each installed CPU must contain a memory riser card. A four CPU system must contain eight memory riser cards.
3. When populating the controller with memory riser cards:
   - First populate riser slot MR0 for each CPU, starting with the lowest numbered CPU (P0).
   - Then populate riser slot MR1 for each CPU, starting with the lowest numbered CPU (P0).

Oracle ZFS Storage ZS5-4 PCIe Card Configuration

The eleven PCIe slots are located inside the Oracle ZFS Storage ZS5-4 controller at the back. As viewed from the front of the controller, the slots are divided into two groups, a group of six on the right of the SP card and a group of five on the left of the SP card. The slots are designated from right to left. The six slots on the right side are designated as PCI-1 to PCI-6. The five slots on the left are designated as PCI-7 to PCI-11.

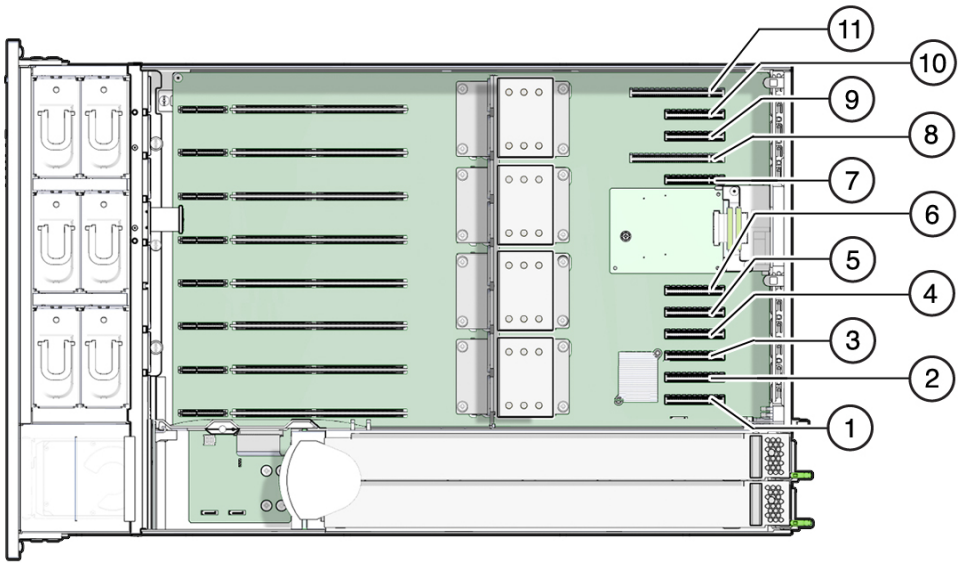

1 PCIe 1	**4** PCIe 4	**7** PCIe 7	**10** PCIe 10
2 PCIe 2	**5** PCIe 5	**8** PCIe 8	**11** PCIe 11
3 PCIe 3	**6** PCIe 6	**9** PCIe 9	

PCIe Slot Length and Capabilities

The Oracle ZFS Storage ZS5-4 controller contains eleven low-profile PCIe 2.0 card slots, which have the following length and capabilities:

- Slots 1, 2, 3, 4, 5, 6, 7, 9 and 10 have x8 connectors (x8 electrical interface)
- Slots 8 and 11 have x16 connectors (x8 or x16 electrical interface)

PCIe Card Slot Population Order

Each base unit Oracle ZFS Storage ZS5-4 is always loaded with the following HBAs in the following PCIe slots:

- Slot 2: Internal SAS-3 HBA
- Slot 4: Cluster interface card
- Slot 6: 4x4 SAS HBA
- Slot 7: 4x4 SAS HBA

All allowable option Oracle ZFS Storage ZS5-4 PCIe cards should be installed in the remaining card slot locations and filled in accordance to the following priority and sequence as follows:

1. Install additional 4x4 SAS HBA cards into slot 9, then slot 3.
2. Install optional dual port QDR InfiniBand (IB) cards into slot 11, then slot 8, then slot 5, then slot 1, then slot 10, then slot 3, then slot 9.
3. Install optional dual port 40GbE NIC cards into slot 11, then slot 8, then slot 5, then slot 1, then slot 10, then slot 3, then slot 9.

4. Install optional dual port 16Gb FC HBA cards into slot 11, then slot 8, then slot 5, then slot 1, then slot 10, then slot 3, then slot 9.
5. Install optional dual port 10GbE SFP+NIC cards into slot 11, then slot 8, then slot 5, then slot 1, then slot 10, then slot 3, then slot 9.
6. Install optional dual port 10GBase-T NIC cards into slot 11, then slot 8, then slot 5, then slot 1, then slot 10, then slot 3, then slot 9.

Oracle ZFS Storage ZS5-4 Attached Storage

The Oracle ZFS Storage ZS5-4 controller connects to external storage through a 4-port (4x4) SAS HBA. You can attach from one to forty-eight disk shelves to the controller, including disk shelf configurations with fully populated flash-based SSD data devices. To improve read performance, the Oracle ZFS Storage ZS5-4 controller does not contain read cache devices, and such devices can be included externally in disk shelves. For disk shelf configuration details, see Disk Shelf Configurations. The Oracle ZFS Storage ZS5-4 controller does not support the Sun Disk Shelf.

7

Servicing the Oracle ZFS Storage ZS5-2 Controller

Be sure to read sections Getting Started with Servicing the Appliance and Prerequisites for Servicing Hardware.

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see Repairing Active Problems.

To replace Oracle ZFS Storage ZS5-2 hardware components, use the following tasks:

- Replacing an Oracle ZFS Storage ZS5-2 HDD
- Replacing an Oracle ZFS Storage ZS5-2 Fan Module
- Replacing an Oracle ZFS Storage ZS5-2 Power Supply
- Replacing an Oracle ZFS Storage ZS5-2 DIMM
- Replacing or Adding an Oracle ZFS Storage ZS5-2 PCIe Card
- Replacing an Oracle ZFS Storage ZS5-2 System Battery

To understand the Oracle ZFS Storage ZS5-2 controller, use the following topics:

- Oracle ZFS Storage ZS5-2 Configuration Options
- Oracle ZFS Storage ZS5-2 Front and Rear Panel Components
- Oracle ZFS Storage ZS5-2 Specifications
- Oracle ZFS Storage ZS5-2 Internal Components
- Oracle ZFS Storage ZS5-2 PCIe Card Configuration
- Oracle ZFS Storage ZS5-2 Cable Management Arm
- Oracle ZFS Storage ZS5-2 Attached Storage

Replacing an Oracle ZFS Storage ZS5-2 HDD

HDDs are hot-swappable and can be removed and installed while the storage controller is running without affecting other hardware capabilities. Use the following procedure to replace an Oracle ZFS Storage ZS5-2 HDD.

> **Note:**
>
> If you have more than one failed drive, replace only one drive at a time. Removing more than one drive in quick succession will cause a hardware/pool fault. Also, do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** BUI menu, select **System**, or navigate to `maintenance system updates` in the CLI.

1. Identify the faulted HDD.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, and open the detailed view of the controller. The faulted drive has either an amber or a blue-on-amber (ready to remove) status icon. Note the ID of the faulted drive. Click the entry for that drive in the table to highlight the location of the drive in the controller illustration. To view details about the faulted drive, click its information icon and then click **Active Problems**.
   - If you are physically at the system, the amber Service Action Required indicator on the HDD should be lit. Additionally, you can click the locator icon in the BUI to light the locator LED on the faulted drive.
2. Verify that the faulted drive is ready to be removed.

   > **Caution:**
   >
   > Removing the drive before it is ready to be removed can result in loss of data.

   Use any of the following methods to verify that the faulted drive is ready to be removed:
   - In the BUI:
     - The status icon for the faulted drive in the detail view in **Maintenance: Hardware** has changed from amber to blue-on-amber.
     - The **Component Details** information window (click the information icon) for the faulted drive shows **Yes** as the value for the **Ready to remove** property.
   - In the CLI (and RESTful API), the value of the `readytoremove` property for the faulted drive is `true`:

     ```
     hostname:maintenance chassis-000 disk-001> get readytoremove
                     readytoremove = true
     ```
3. To remove a faulted drive, press the release button and pivot the release lever open to the right.

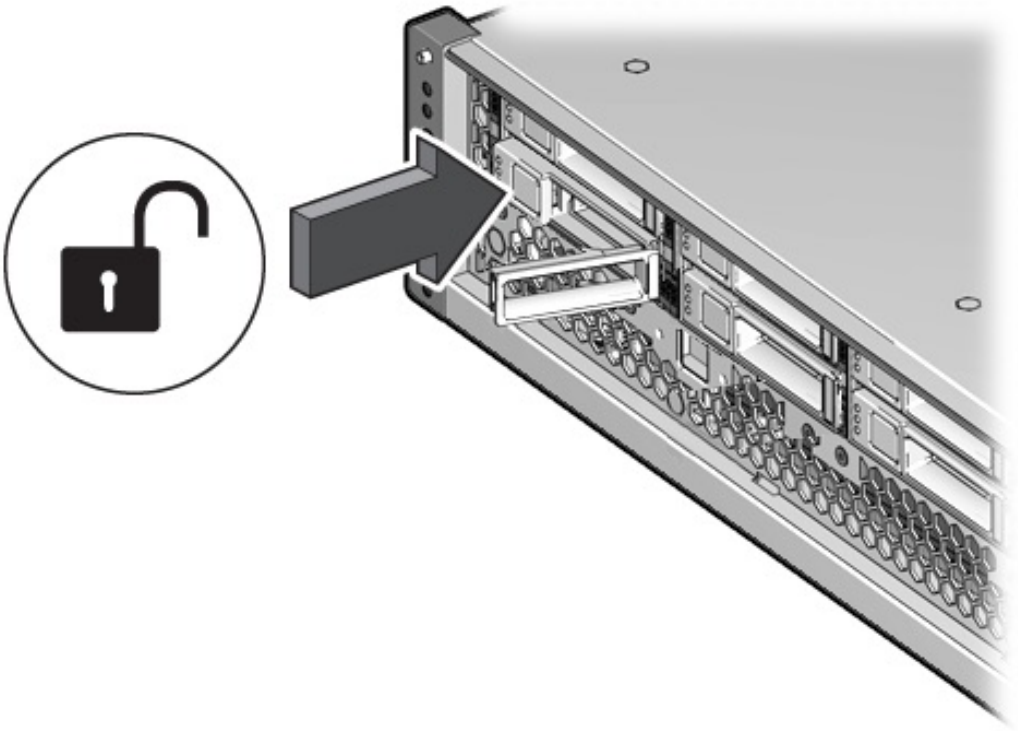

4. Grasp the release lever, pull the drive out of the chassis, and set the drive aside.

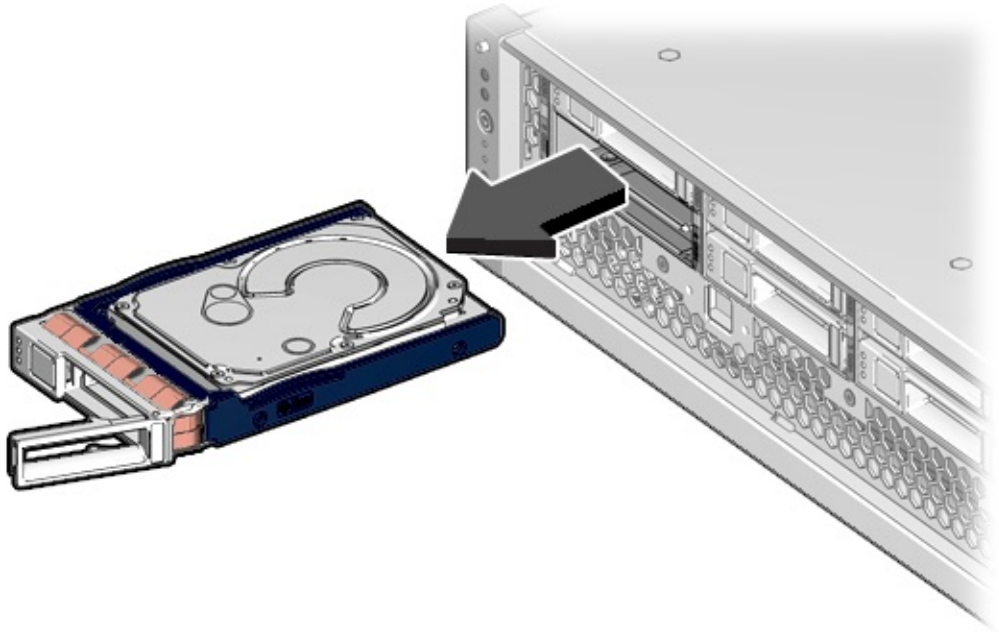

5. After at least 30 seconds, from the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon on the system controller to verify that the software has detected that the drive is not present.
6. To insert the replacement drive, push the release button of the drive.

   The release lever opens.
7. Align the replacement drive with the chassis slot.

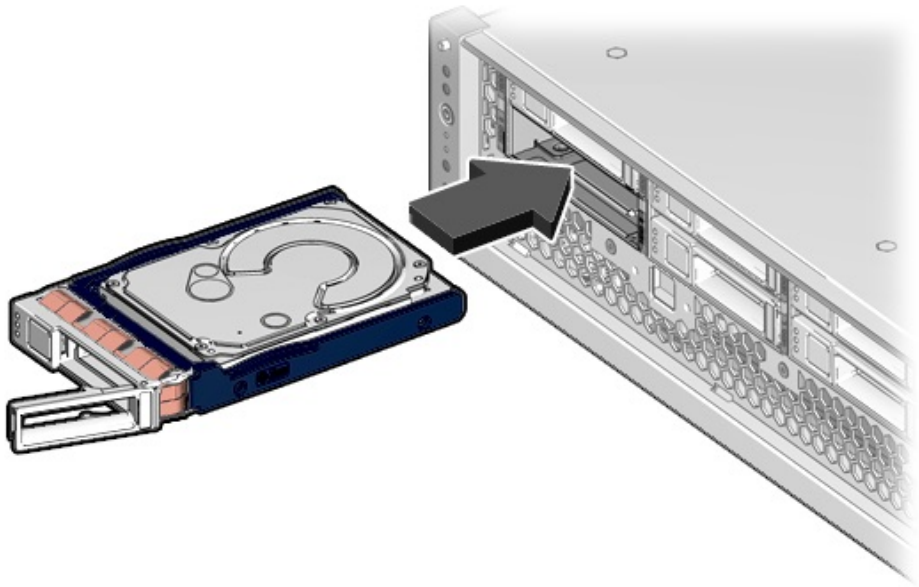

   The connector is at the rear of the drive and the release button is on the left.
8. Slide the drive into the chassis until the release lever moves slightly inward.
9. Press the release lever fully closed until it clicks, securing the drive into the chassis.

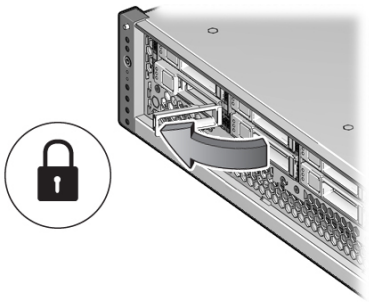

   The Oracle ZFS Storage Appliance software automatically detects and configures the new drive.
10. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, and verify that the status icon is green for the newly installed drive. To view details about the new drive, click its information icon .

11. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS5-2 Fan Module

Removing a fan module is a cold-service operation. You *must* power off the appliance before you remove the fan module. There will be a loss of access to the storage unless the system is in a clustered configuration. Use the following procedure to replace an Oracle ZFS Storage ZS5-2 fan module.

1. Identify the faulted fan module.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **Fan**, and note the ID of the faulted fan module, which has an amber icon. To highlight the location of the fan module in the controller picture, click on the fan module ID. To view details about the faulted fan module, click its information icon and then click **Active Problems**.
   - If you are physically at the system, note that each fan module has a fan fault LED that is located near the module on the motherboard. The amber LED is illuminated for failed fan modules.
2. Power off the storage controller as described in Powering Off the Controller.
3. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

4. Extend the controller from the rack as described in Extending the Controller from the Rack.
5. Remove the top cover as described in Removing the Controller Top Cover.
6. Press the Fault Remind button on the motherboard to illuminate the fan module fault LED.

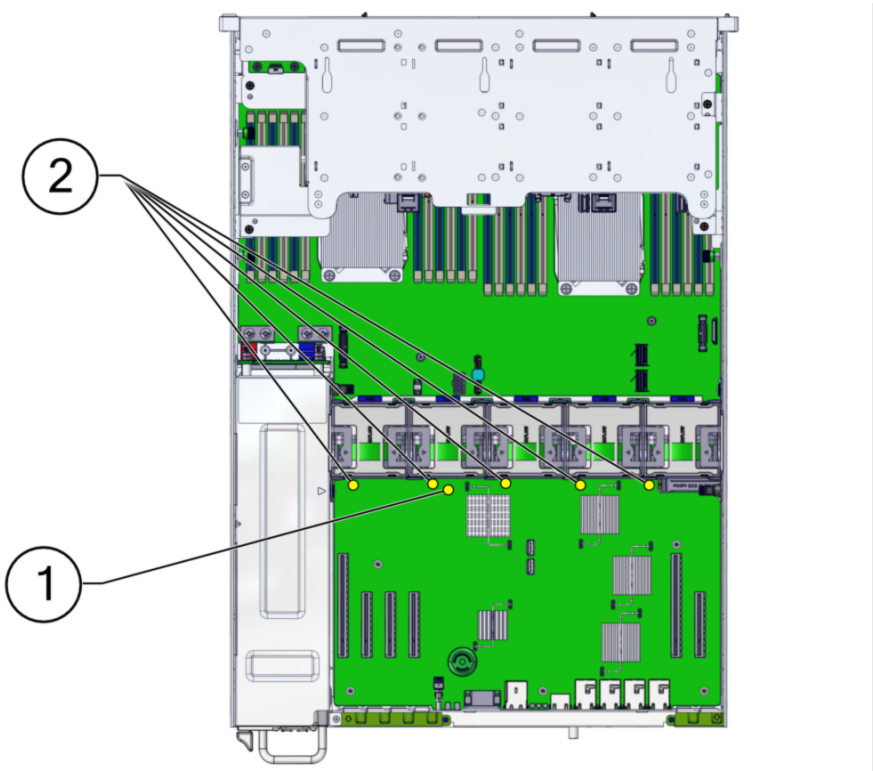

1 Fault Remind button
2 Fan Fault LEDs

7. On the fan module to be replaced, pull the green band straight up to lift the fan module out of the chassis, and set the module aside.

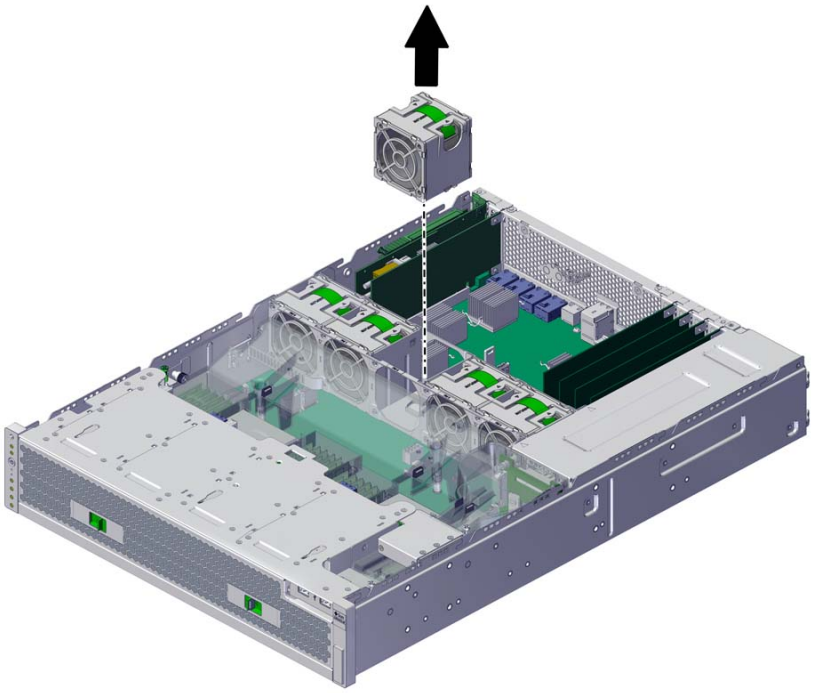

8. Align the replacement fan module with the fan tray slot.

   The green band is on the top of the fan module and the arrow points to the rear of the chassis.

9. Lower the fan module into the fan tray.

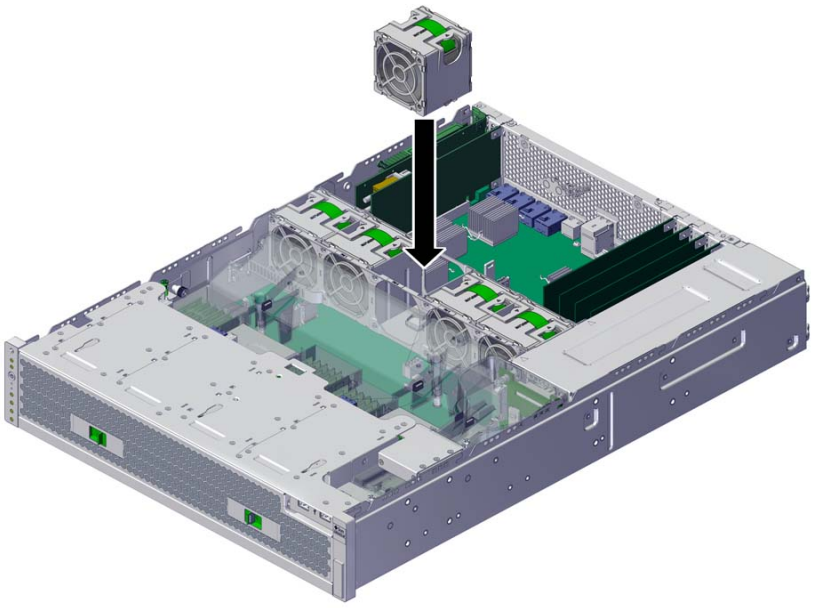

10. Press the fan module down until it clicks securely in the fan tray.

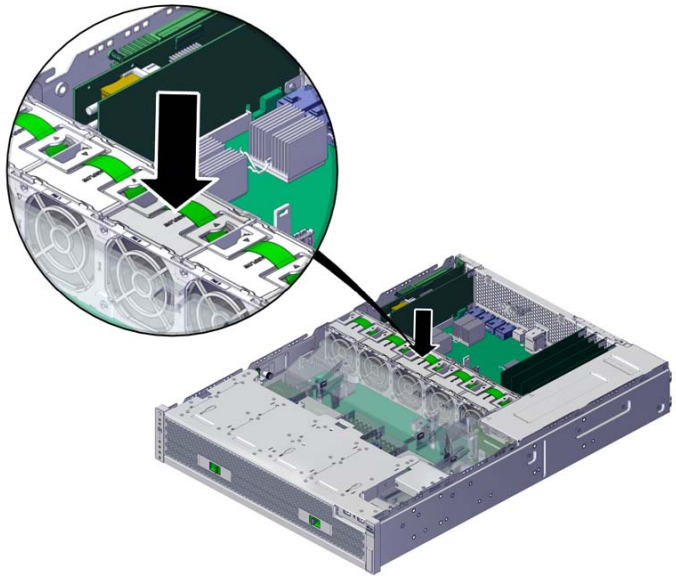

11. Replace the top cover on the chassis, and tighten the two captive screws at the back of the top cover.
12. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.
13. Connect the power cords to the power supplies. Verify that standby power is on, indicated by the Power/OK and SP status indicators flashing green on the front panel.
14. Verify that the SP status and Power/OK indicators go solid green after the boot.
15. Verify that the green OK indicator is lit on the fan module.
16. Verify that the Fan indicator and the Service Action Required indicator on the front panel are not lit.
17. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **Fan**, and verify that the status icon is green for the newly installed fan module. To view details about the new fan module, click its information icon .
18. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS5-2 Power Supply

> **Caution:**
>
> If a power supply fails and you do not have a replacement available, do not remove the failed power supply. This maintains the proper air flow and prevents possible equipment damage.

Storage controllers are equipped with redundant, hot-swappable power supplies. Use the following procedure to replace an Oracle ZFS Storage ZS5-2 power supply.

1. Identify the faulted power supply.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **PSU**, and note the ID of the faulted power supply, which has an amber icon. To highlight the location of the power supply in the controller picture, click on the power supply ID. To view details about the faulted power supply, click its information icon and then click **Active Problems**.
   - If you are physically at the system, the amber Service Action Required indicator on the power supply should be lit.

   > **Note:**
   >
   > The fans of a failed power supply might still be spinning when the system is powered on. You can remove a power supply while the fans are still spinning.

2. Gain access to the rear of the controller where the faulted power supply is located.
3. Press and hold the CMA release tab [a], and rotate the arm out of the way [b].

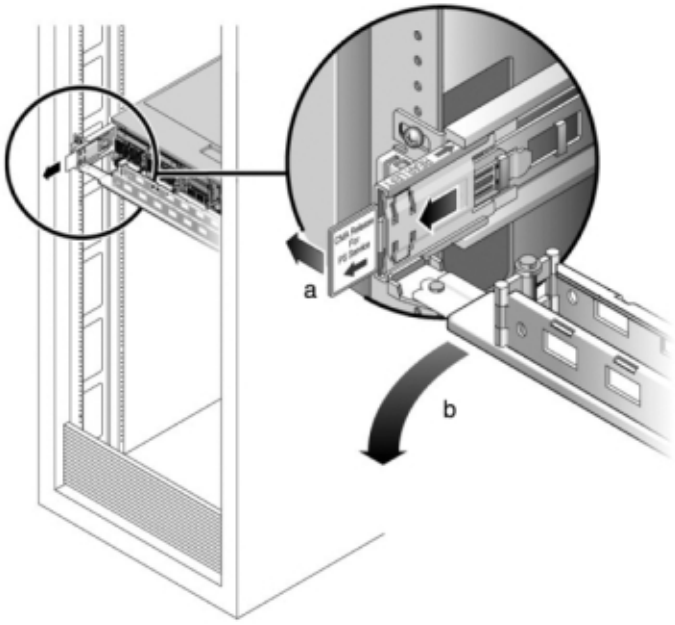

4. Disconnect the power cord from the faulted power supply.
5. Press the release tab toward the right while pulling on the handle.

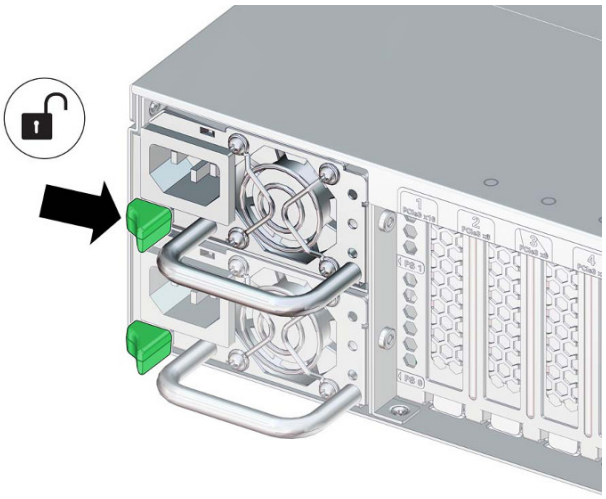

6. Continue to pull on the handle to slide the power supply out of the chassis. Set the power supply aside.

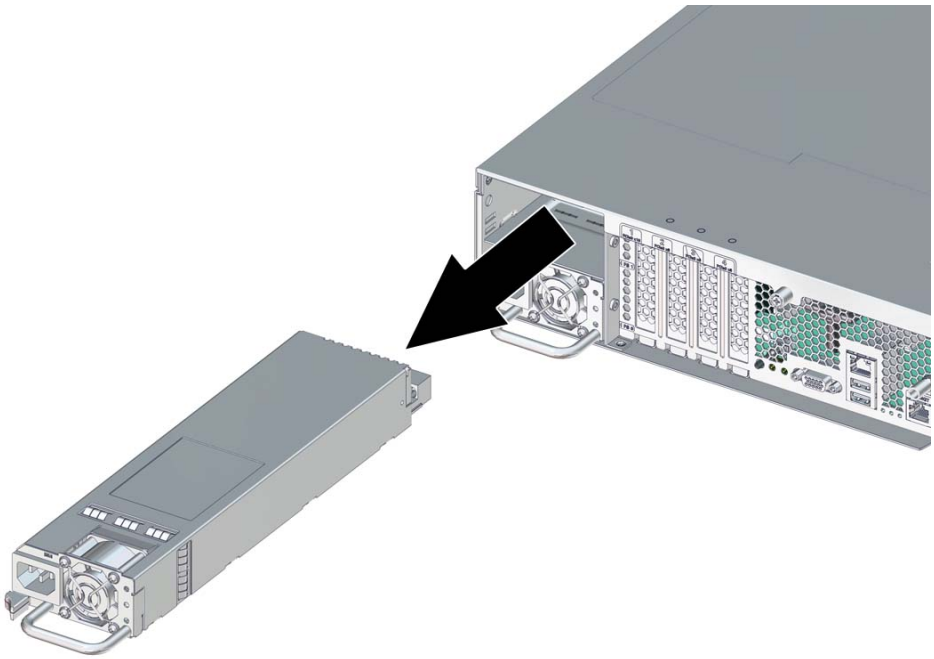

7. Align the replacement power supply with the chassis slot.

   The PDB connector is toward the chassis, and the release tab is on the left.

8. Slide the replacement power supply into the chassis until the release tab clicks.

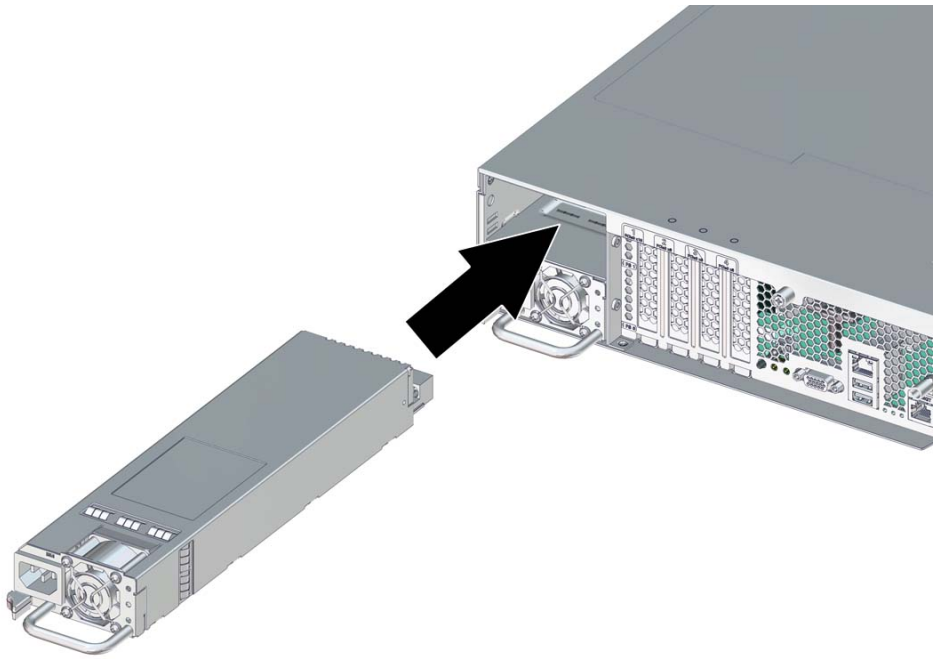

9. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
10. Attach the power cord to the power supply.
11. Verify that the green AC Present status indicator is lit.
12. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **PSU**, and verify that the status icon is green for the newly installed power supply. To view details about the new power supply, click its information icon .
13. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS5-2 DIMM

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

Removing a DIMM is a cold-service operation. You *must* power off the appliance before you remove the DIMM. There will be a loss of access to the storage unless the system is in a clustered configuration.

1. Identify the faulted DIMM.

From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **DIMM**, and note the ID of the faulted DIMM, which has an amber icon. To highlight the location of the memory module in the controller picture, click on the module ID. To view details about the faulted module, click its information icon and then click **Active Problems**.

2. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
3. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

4. Extend the controller from the rack as described in Extending the Controller from the Rack.
5. Remove the top cover as described in Removing the Controller Top Cover.
6. Press the Fault Remind button on the motherboard to illuminate the DIMM fault LED.

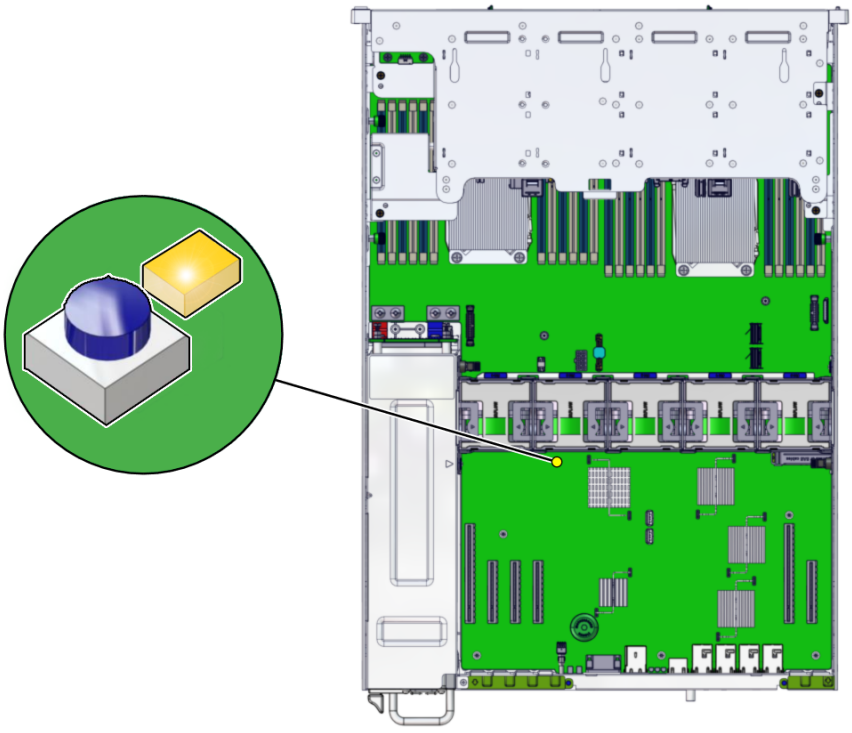

7. To remove the faulted DIMM, press down and out on the release levers at both ends of the DIMM slot. See Oracle ZFS Storage ZS5-2 Internal Components for reference.

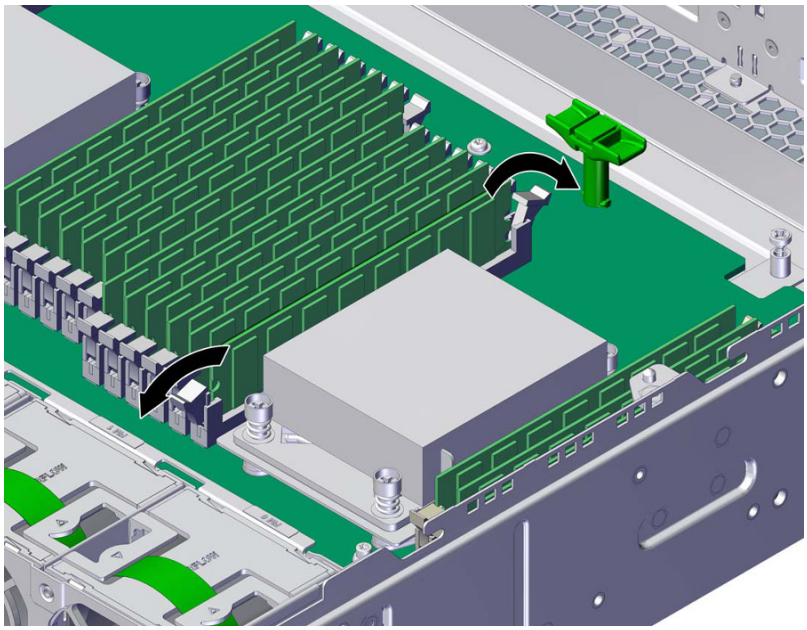

8. Lift the DIMM up and out the slot, and set the DIMM aside.

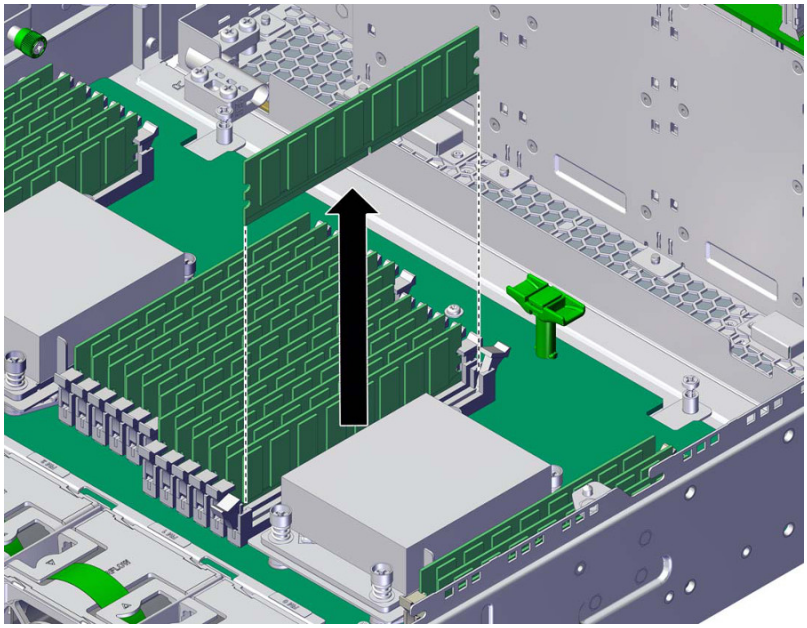

9. Align the replacement DIMM with the chassis slot.

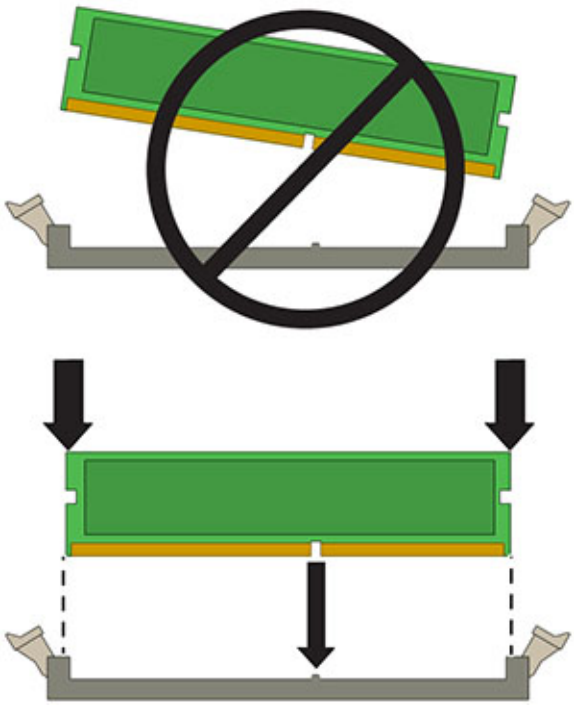

Ensure that the notch in the DIMM lines up with the key in the slot.

10. Insert the DIMM into the slot, pressing firmly so that both release levers click closed.

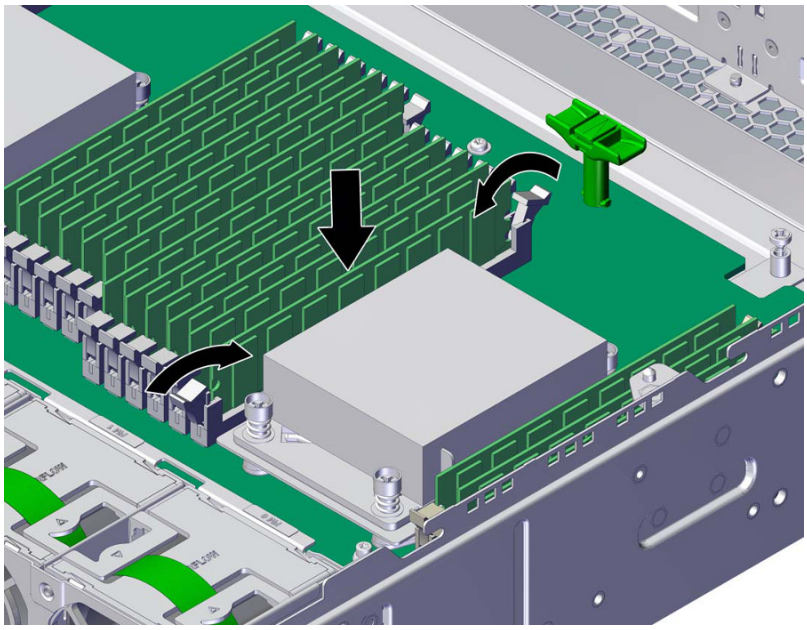

11. Replace the top cover on the chassis, and tighten the two captive screws at the back of the top cover.
12. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.

13. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
14. Reconnect the power cords to the power supplies.
15. Verify that standby power is on. Approximately two minutes after plugging in the power cords, the Power/OK status indicator on the front panel flashes.
16. Press and release the recessed Power button on the storage controller front panel.

    The Power/OK status indicator next to the Power button lights and remains lit.
17. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **DIMM**, and verify that the status icon is green for the newly installed module. To view details about the new module, click its information icon .
18. If a fault does not self-clear, see Repairing Active Problems.

Replacing or Adding an Oracle ZFS Storage ZS5-2 PCIe Card

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

Replacing or adding a PCIe card is a cold-swap procedure, and you must remove power from the Oracle ZFS Storage ZS5-2 controller. If you are installing a newly released HBA, upgrade your system software before installing the HBA as described in Upgrading the Software. All HBAs must be of the same type. Use the following procedure to replace or add an Oracle ZFS Storage ZS5-2 PCIe card.

1. Identify the PCIe card slot location.
   - If replacing a faulted PCIe card, identify the PCIe card to be replaced.

     From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **Slot**, and note the ID of the faulted PCIe card, which has an amber icon. If replacing a non-faulted PCIe card, note its location. To highlight the location of the card in the controller picture, click on the card ID. To view details about the card, click its information icon and then click **Active Problems**.
   - If adding a new PCIe card, identify the correct slot location as described in Oracle ZFS Storage ZS5-2 PCIe Card Configuration.
2. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
3. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

4. If replacing a PCIe card, disconnect the cables attached to it.
5. Extend the controller from the rack as described in Extending the Controller from the Rack.
6. Remove the top cover as described in Removing the Controller Top Cover.
7. To remove a PCIe card, swing the retainer to its fully open position, and lift the PCIe card out of the socket. Place the card aside.

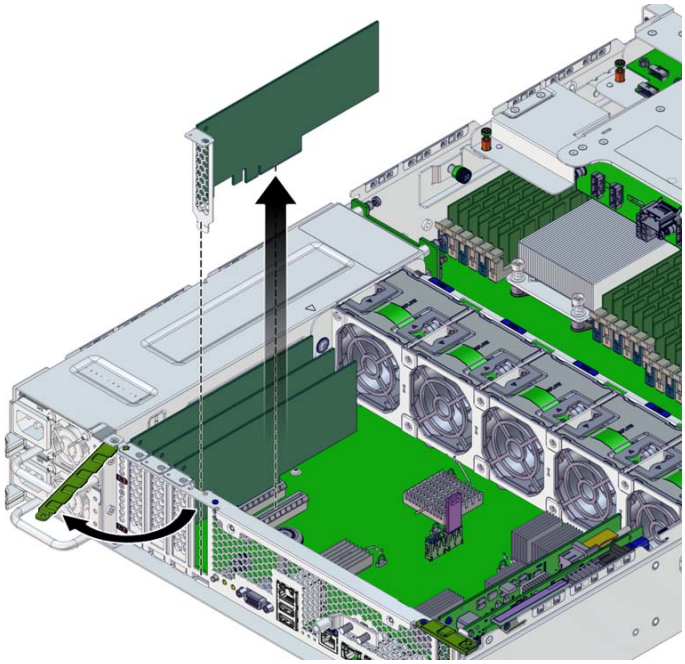

8. Align the replacement or new PCIe card with the chassis slot.
9. Press the PCIe card into the socket.
10. Swing the retainer to its fully closed position.

    If there is resistance, check the alignment of the PCIe card bracket and try again.
11. Replace the top cover on the chassis, and tighten the two captive screws at the back of the top cover.
12. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.
13. Connect cables to the PCIe card, routing them through the cable management arm.
14. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
15. Reconnect the power cords to the power supplies.
16. Verify that standby power is on. Approximately two minutes after plugging in the power cords, the Power/OK status indicator on the front panel flashes.
17. Press and release the recessed Power button on the storage controller front panel.

    The Power/OK status indicator next to the Power button lights and remains lit.

18. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **Slot**, and verify that the status icon is green for the newly installed card. To view details about the new card, click its information icon .
19. If you replaced or installed a new HBA, connect the disk shelves to it as described in the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
20. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS5-2 System Battery

Removing the battery is a cold-service operation. You *must* power off the appliance before you remove the battery. There will be a loss of access to the storage unless the system is in a clustered configuration.

You might need to replace the battery if the storage controller fails to maintain the proper time when powered off and disconnected from the network. You will need a small, non-metallic screwdriver (No. 1 flat-blade) or equivalent. Use the following procedure to replace an Oracle ZFS Storage ZS5-2 battery.

1. Power off the storage controller as described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. Remove the top cover as described in Removing the Controller Top Cover.
5. Remove the air duct.

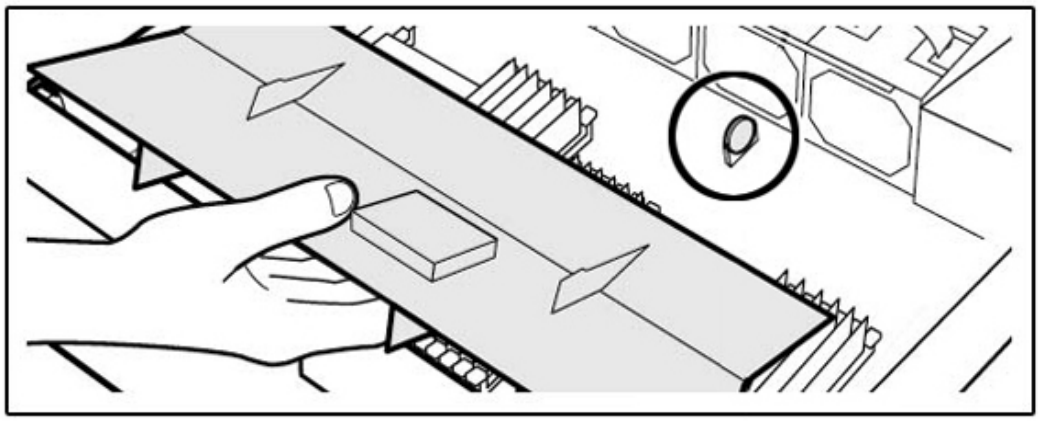

6. To remove the battery, tilt the battery slightly to left [a] and lift it out of the receptacle [b].

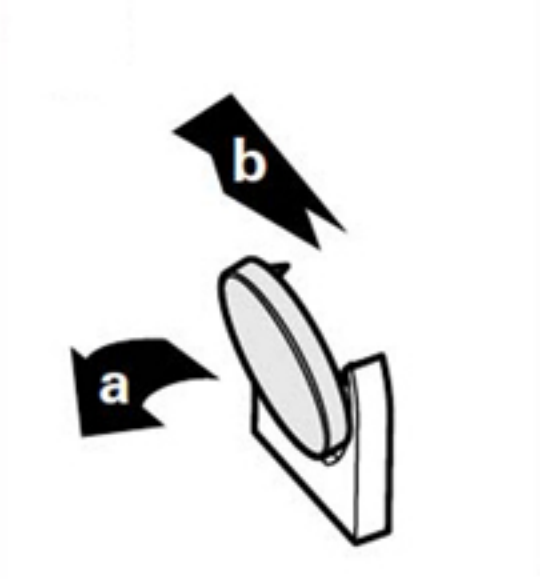

7. Set the battery aside.
8. Align the replacement battery with the receptacle.

The positive (+) side of the battery faces left.

9. Press the battery into the receptacle until it clicks.
10. Install the air duct.
11. Replace the top cover on the chassis, and tighten the two captive screws at the back of the top cover.
12. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.

13. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
14. Connect the power cords to the power supplies.
15. Verify that standby power is on. Approximately two minutes after plugging in the power cords, the Power/OK status indicator on the front panel flashes.
16. Press and release the recessed Power button on the storage controller front panel. The Power/OK status indicator next to the Power button lights and remains lit.
17. When the system has finished booting, log in and set the time using the steps in the BUI Clock task. For more information, see Setting Clock Synchronization (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Oracle ZFS Storage ZS5-2 Controller Hardware Overview

This section describes the internal and external components of the Oracle ZFS Storage ZS5-2 controller.

- Oracle ZFS Storage ZS5-2 Configuration Options
- Oracle ZFS Storage ZS5-2 Front and Rear Panel Components
- Oracle ZFS Storage ZS5-2 Specifications
- Oracle ZFS Storage ZS5-2 Internal Components
- Oracle ZFS Storage ZS5-2 PCIe Card Configuration
- Oracle ZFS Storage ZS5-2 Cable Management Arm
- Oracle ZFS Storage ZS5-2 Attached Storage

Oracle ZFS Storage ZS5-2 Configuration Options

Oracle ZFS Storage ZS5-2 incorporates one or two controllers and one to sixteen disk enclosures. Each controller has six PCIe Gen3 slots. Each controller includes eight 2.5" disk slots that accommodate two system disks. Read cache devices are installed in disk enclosures; therefore, the remaining six disk slots have permanent fillers. The following table describes the configuration options.

CPU	Memory	System Disks	PCIe Slots	HBA
2x18-core, 2.3 GHz	384 GB 768 GB	2x2.5" SAS-3	Two dedicated; four available	4-port (4x4) SAS-3 12 Gb/s external 4-port (4x4) SAS-2 6 Gb/s external

Oracle ZFS Storage ZS5-2 Front and Rear Panel Components

Oracle ZFS Storage ZS5-2 Front Panel Components - The Oracle ZFS Storage ZS5-2 controller drive slots and front panel components are shown in the following figure.

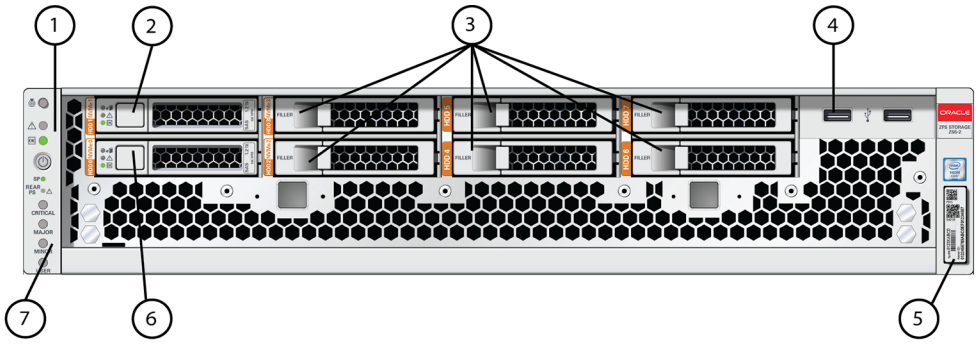

1 LEDs: Locator, Service, Power	**2** System disk 1	**3** Permanent filler panels
4 Two USB 2.0 ports	**5** RFID tag	**6** System disk 0
7 Alarm indicators (not used)		

Oracle ZFS Storage ZS5-2 Rear Panel Components - The Oracle ZFS Storage ZS5-2 controller PCIe slots and rear panel components are shown in the following figure.

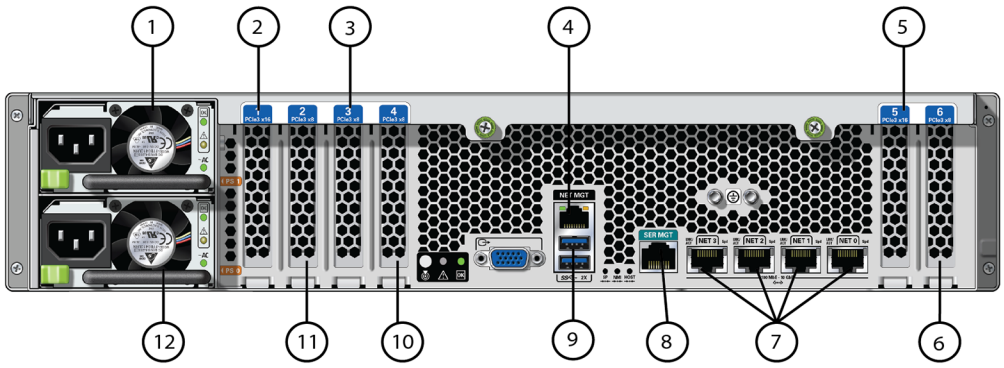

1 Power Supply Unit (PSU) 1	**2** Second PCIe option (slot 1)	**3** Third PCIe option (slot 3)
4 Network Management (NET MGT) 10/100/1000 BASE-T Ethernet port	**5** First PCIe option (slot 5)	**6** Second 4x4 SAS HBA, or fourth PCIe option (slot 6)
7 Network (NET) 10 GbE ports: NET 0, NET 1, NET 2, NET 3	**8** SER MGT port	**9** Two USB 3.0 ports
10 Cluster interface card (slot 4)	**11** First 4x4 SAS HBA (slot 2)	**12** Power Supply Unit (PSU) 0

Cluster Interface Card Compatibility - The Oracle ZFS Storage ZS5-2 controller contains a Version 3 cluster interface card, which is incompatible with Version 2.

For information about how to connect cables to form a cluster, see Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

Ethernet Ports

The Oracle ZFS Storage ZS5-2 controller has four RJ-45 10-Gigabit Ethernet (10GbE) network connectors, labeled NET3, NET 2, NET 1, and NET 0 (left to right) on the rear panel. The ports operate at 100 Mbits/sec, 1000 Mbit/sec, or 10-Gbits/sec. Use these ports to connect the appliance to the network.

There are two status indicators (LEDs) on each port. These indicators are visible from the rear of the controller.

Table 7-1 Gigabit Ethernet Port Status Indicators

Status Indicator Name	Location	Color	State and Meaning
Activity	Top left	Green	ON – No link OFF– No activity FLASHING – Packet activity
Link speed	Top right	Bicolored: Amber/ Green	OFF – 100BASE-T link Amber ON – 1000BASE-T link Green ON – 10GBASE-T link

Network Management Port

The Oracle ZFS Storage ZS5-2 controller has one 10/100/1000BASE-T Ethernet management domain interface, labeled NET MGT. There are two status indicators (LEDs) on this port. These indicators are visible from the rear of the controller. The service processor NET MGT port is the optional connection to the Oracle ILOM service processor. The NET MGT port is configured by default to use Dynamic Host Configuration Protocol (DHCP). The service processor NET MGT port uses an RJ-45 cable for a 10/100/1000BASE-T connection.

Table 7-2 Network Management Port Status Indicators

Status Indicator Name	Location	Color	State and Meaning
Activity	Top left	Amber	ON – Link up OFF – No link or down link FLASHING – Packet activity
Link speed	Top right	Green	Green ON – 1000BASE-T link OFF – 10/100BASE-T link

Serial Management Port

The Oracle ZFS Storage ZS5-2 service processor SER MGT port uses an RJ-45 cable and is the default connection to the Oracle ILOM service processor. This port supports local connections to the controller and only recognizes Oracle ILOM command-line interface (CLI) commands. Typically you connect a terminal or terminal emulator to this port.

> **Note:**
>
> This port does not support network connections.

4x4 SAS HBA Ports

The 4x4 SAS HBA, installed in PCIe slot 2 of the Oracle ZFS Storage ZS5-2 controller, provides connectivity to external disk shelves. The HBA logical ports are numbered 3-0, top to bottom, as shown in the following figure.

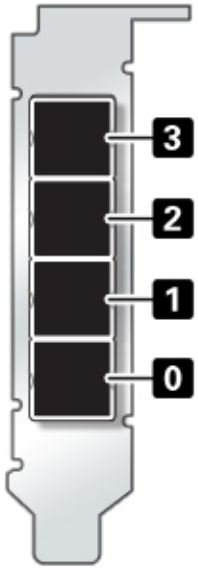

Oracle ZFS Storage ZS5-2 Specifications

Oracle ZFS Storage ZS5-2 Physical Specifications: The Oracle ZFS Storage ZS5-2 controller chassis dimensions are as follows:

Dimension	Measurement	Dimension	Measurement
Height	3.43 in./87.1 mm	Depth	25.2 in./640.0 mm
Width	17.52 in./445 mm	Weight (Maximum)	38.5 lb./17.5 kg

Oracle ZFS Storage ZS5-2 Electrical Specifications: The Oracle ZFS Storage ZS5-2 controller has two N+1 AC redundant hot-swappable AC power supplies, rated 91% efficiency. To understand the overall power consumption, see Oracle's online power calculator at http://www.oracle.com/goto/powercalculators .

Parameter	AC Power Supplies
Nominal input frequencies	50/60 Hz

Parameter	AC Power Supplies
Operating input voltage range	100 VAC (Japan only) 110-127 VAC (90-132 range) 200-240 VAC (180-264 range)
Rated input current	10 A @ 100 VAC 9 A @110-127 VAC 5.5 A @ 200-240 VAC

Oracle ZFS Storage ZS5-2 Acoustic Noise Emissions: In compliance with the requirements defined in ISO 9296, the workplace-dependent noise level of this product is less than 75 dBA.

Oracle ZFS Storage ZS5-2 Internal Components

Oracle ZFS Storage ZS5-2 Storage, Power, and Fan Modules - The Oracle ZFS Storage ZS5-2 internal storage, power, and cooling components are described in the following figure and legend. The Oracle ZFS Storage ZS5-2 chassis contains both customer-replaceable units (CRUs) and field-replaceable units (FRUs) as shown in the following figure. FRUs must be replaced by trained Oracle service technicians.

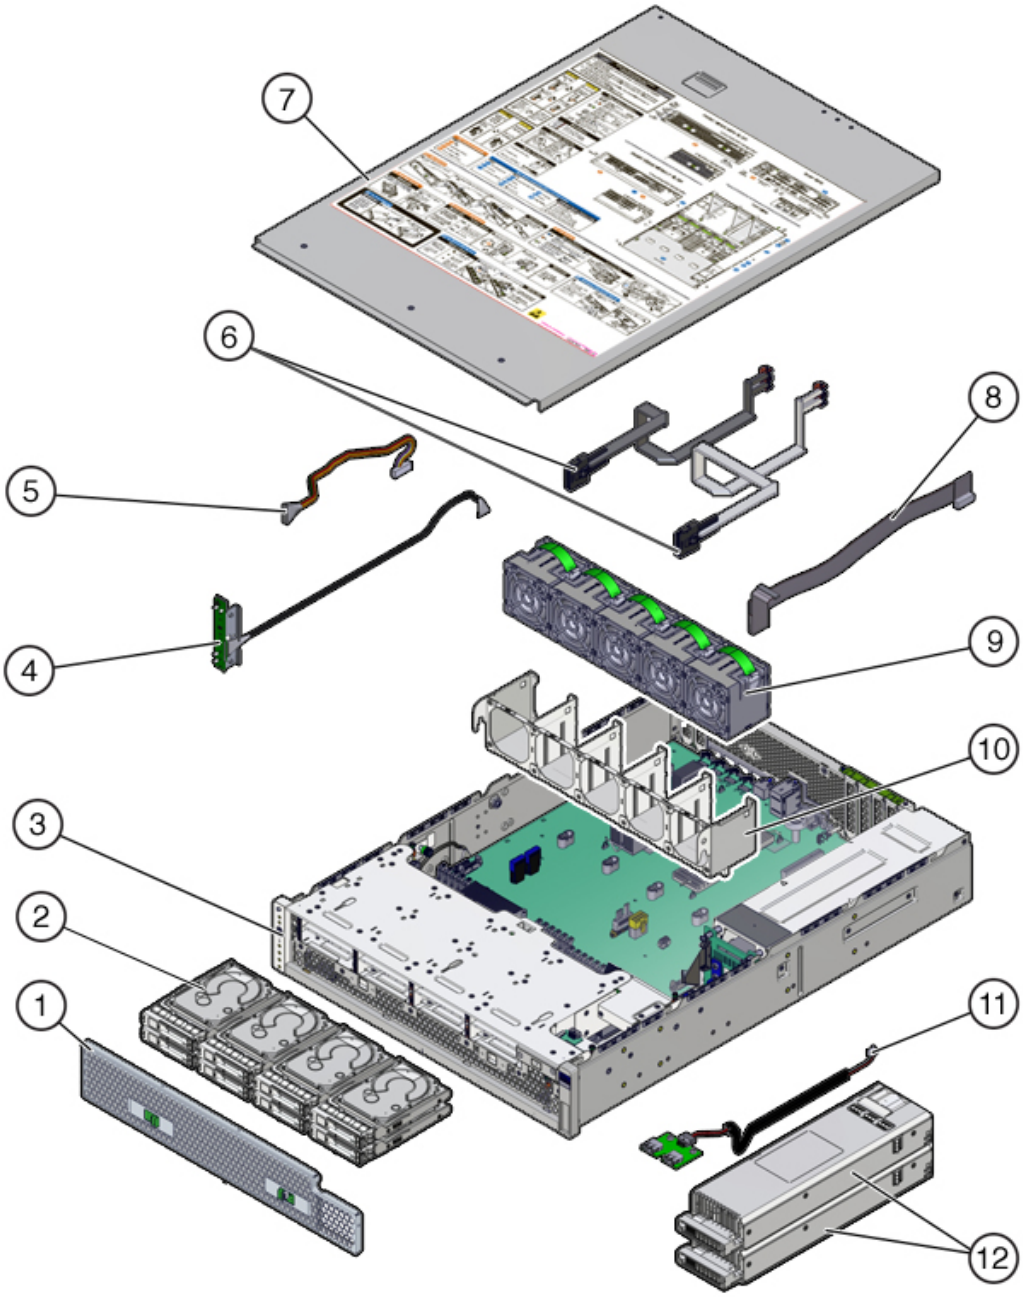

1 Filter	**2** Drives (CRU)	**3** Chassis (FRU)
4 LED board (FRU)	**5** Drive backplane data cable (FRU)	**6** SAS cables (FRU)
7 Top cover	**8** Drive backplane power cable (FRU)	**9** Fan modules (CRU)
10 Fan tray	**11** USB board (FRU)	**12** Power supplies (CRU)

Oracle ZFS Storage ZS5-2 Memory, PCIe Cards, and System Battery - The Oracle ZFS Storage ZS5-2 memory, PCIe cards, and system battery are described in the

following figure and legend. These components are customer-replaceable units (CRUs) as shown in the following figure.

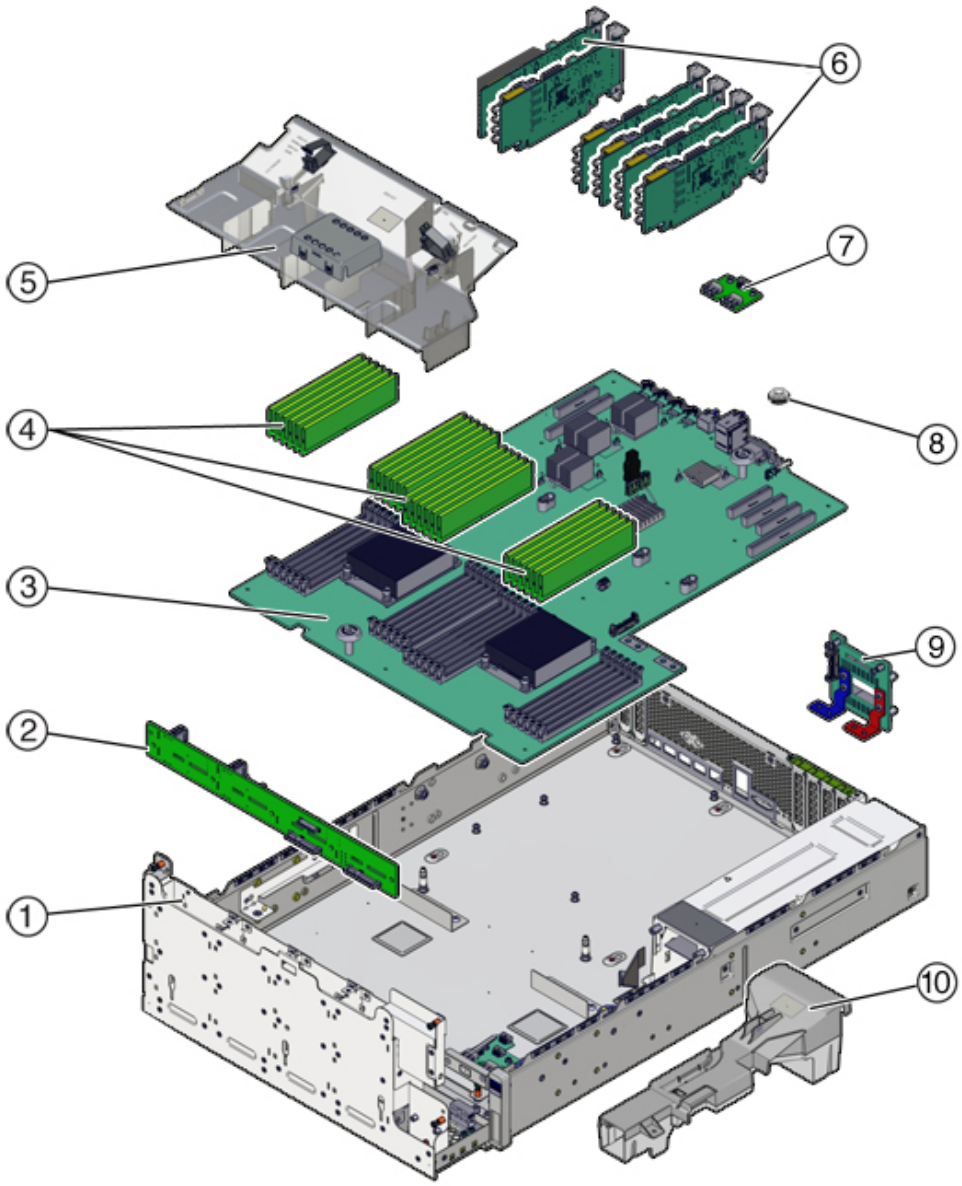

1 Drive cage (FRU)	**2** Drive backplane (FRU)	**3** Motherboard (FRU)
4 DIMMs (CRU)	**5** Air duct	**6** PCIe cards (CRU)
7 USB board (FRU)	**8** System battery (CRU)	**9** Power distribution board (FRU)
10 PSU duct		

Oracle ZFS Storage ZS5-2 DIMM Configuration - The Oracle ZFS Storage ZS5-2 controller supports only fully populated DIMMs, 16 GB for 384 GB or 32 GB for 768 GB. Do not mix DIMM capacities; use either all 16 GB or 32 GB DIMMs.

In this illustration of the DIMM slots, the Oracle ILOM legacy targets for the DIMMs are provided across the top, and the Oracle ILOM `/System` targets for the DIMMs are provided across the bottom.

For example, the slot on the far left has the Oracle ILOM legacy target of `/SYS/MB/P1/D0`. This same slot has the Oracle ILOM `/System` target of `/System/Memory/DIMMs/DIMM_0`.

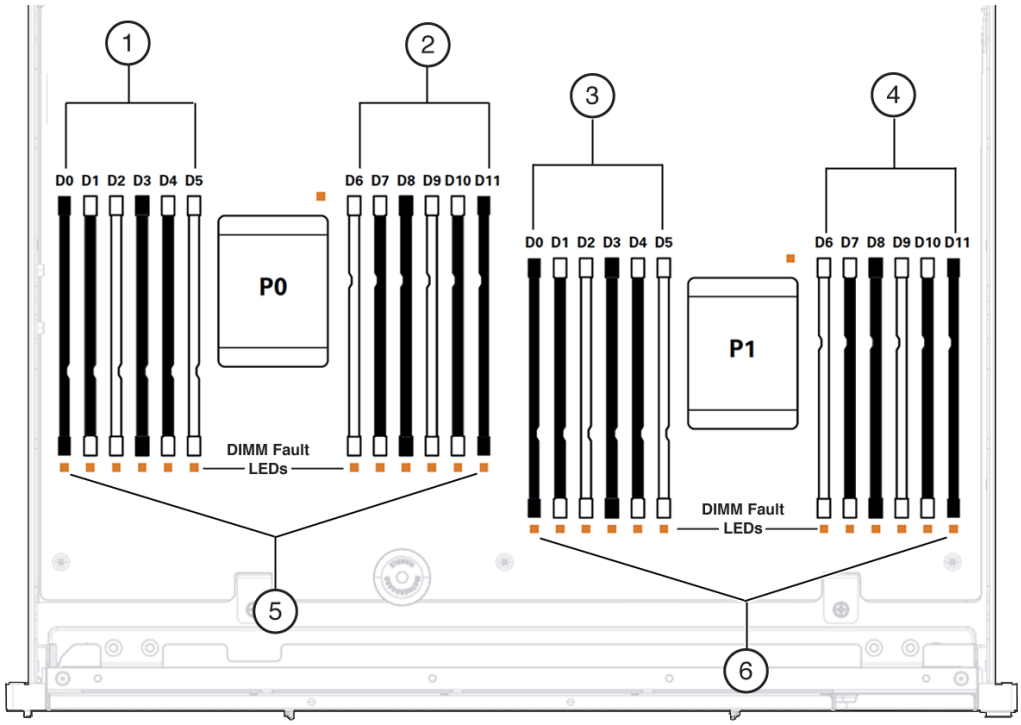

1 /SYS/MB/P0/D(0-5)	**4** /SYS/MB/P1/D(6-11)
2 /SYS/MB/P0/D(6-11)	**5** /System/Memory/DIMMs/DIMM_(0-11)
3 /SYS/MB/P1/D(0-5)	**6** /System/Memory/DIMMs/DIMM_(12-23)

Oracle ZFS Storage ZS5-2 PCIe Card Configuration

The Oracle ZFS Storage ZS5-2 controller is equipped with six PCIe Gen3 expansion slots. One of these slots is always dedicated to an external HBA and one is reserved for a cluster interface card. The remaining slots may be used for an additional external HBA or one of several client-facing adapters.

The Oracle ZFS Storage ZS5-2 base configuration contains the following PCIe cards:

- One 4-port (4x4) SAS external HBA (slot 2)
- One cluster interface card (slot 4)

The following figure shows the PCIe I/O slot numbers.

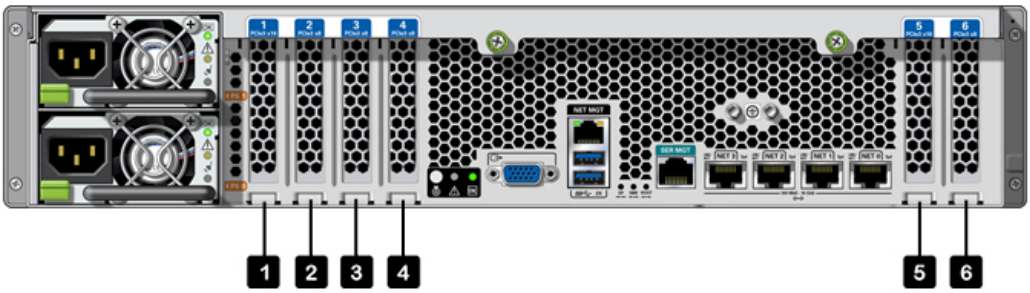

PCIe Slot Order - Install optional PCIe cards in the following order:

1. Install an additional 4x4 SAS HBA card into slot 6.
2. Install optional dual port QDR 40Gb InfiniBand HCA cards into slot 5, then slot 1, then slot 3, then slot 6.
3. Install optional dual port 16Gb FC HBA cards into slot 5, then slot 1, then slot 3, then slot 6.
4. Install optional dual port 40GbE NIC cards into slot 5, then slot 1, then slot 3, then slot 6.
5. Install optional dual port 10Gb Ethernet cards into slot 5, then slot 1, then slot 3, then slot 6.
6. Install optional dual port 10GBASE-T NIC cards into slot 5, then slot 1, then slot 3, then slot 6.

Oracle ZFS Storage ZS5-2 Cable Management Arm

The following figure identifies the components of the Oracle ZFS Storage ZS5-2 cable management arm (CMA). See the CMA installation instructions as described in Installing an Oracle ZFS Storage ZS5-2 Cable Management Arm in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*.

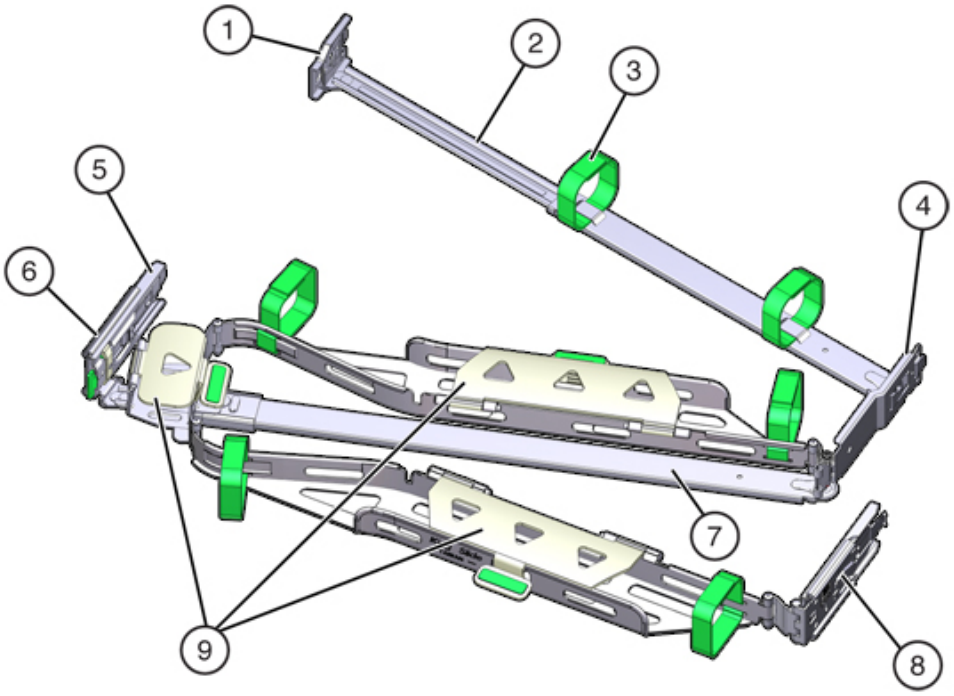

1 Connector A	**6** Connector D
2 Front slide bar	**7** Slide-rail latching bracket (used with connector D)
3 Velcro straps (6)	**8** Rear slide bar
4 Connector B	**9** Cable covers
5 Connector C	

Oracle ZFS Storage ZS5-2 Attached Storage

The Oracle ZFS Storage ZS5-2 controller connects to external storage through a 4-port (4x4) SAS HBA. You can attach from one to sixteen disk shelves to the controller, including disk shelf configurations with fully populated flash-based SSD data devices. To improve read performance, the Oracle ZFS Storage ZS5-2 controller does not contain read cache devices, and such devices can be included externally in disk shelves. For disk shelf configuration details, see Disk Shelf Configurations. The Oracle ZFS Storage ZS5-2 controller does not support the Sun Disk Shelf.

8

Servicing the Oracle ZFS Storage ZS4-4 Controller

Be sure to read sections Getting Started with Servicing the Appliance and Prerequisites for Servicing Hardware.

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see Repairing Active Problems.

To replace Oracle ZFS Storage ZS4-4 hardware components, use the following tasks:

- Replacing an Oracle ZFS Storage ZS4-4 HDD or SSD
- Replacing an Oracle ZFS Storage ZS4-4 Power Supply
- Replacing an Oracle ZFS Storage ZS4-4 Fan Module
- Identifying a Faulted Oracle ZFS Storage ZS4-4 Memory Module
- Replacing an Oracle ZFS Storage ZS4-4 DIMM
- Replacing an Oracle ZFS Storage ZS4-4 Memory Riser Card
- Replacing an Oracle ZFS Storage ZS4-4 PCIe Card
- Replacing an Oracle ZFS Storage ZS4-4 Battery

To understand the Oracle ZFS Storage ZS4-4 controller, use the following topics:

- Oracle ZFS Storage ZS4-4 Base Configuration
- Oracle ZFS Storage ZS4-4 Front and Rear Panel Components
- Oracle ZFS Storage ZS4-4 PCIe I/O Cards
- Oracle ZFS Storage ZS4-4 Internal Components
- Oracle ZFS Storage ZS4-4 Attached Storage

Replacing an Oracle ZFS Storage ZS4-4 HDD or SSD

HDDs and SSDs are hot-swappable and can be removed and installed while the storage controller is running without affecting other hardware capabilities. Use the following procedure to replace an Oracle ZFS Storage ZS4-4 HDD or SSD.

> **Note:**
>
> If you have more than one failed drive, replace only one drive at a time. Removing more than one drive in quick succession will cause a hardware/pool fault. Also, do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** BUI menu, select **System**, or navigate to `maintenance system updates` in the CLI.

1. Identify the faulted HDD or SSD.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, and open the detailed view of the controller. The faulted drive has either an amber or a blue-on-amber (ready to remove) status icon. Note the ID of the faulted drive. Click the entry for that drive in the table to highlight the location of the drive in the controller illustration. To view details about the faulted drive, click its information icon and then click **Active Problems**.
   - If you are physically at the system, the amber Service Action Required indicator on the HDD or SSD should be lit. Additionally, you can click the locator icon in the BUI to light the locator LED on the faulted drive.
2. Verify that the faulted drive is ready to be removed.

   > **Caution:**
   >
   > Removing the drive before it is ready to be removed can result in loss of data.

   Use any of the following methods to verify that the faulted drive is ready to be removed:
   - On the faulted drive, the Ready to Remove LED is lit as shown in System Drive LEDs.
   - In the BUI:
     - The status icon for the faulted drive in the detail view in **Maintenance: Hardware** has changed from amber to blue-on-amber.
     - The **Component Details** information window (click the information icon) for the faulted drive shows **Yes** as the value for the **Ready to remove** property.
   - In the CLI (and RESTful API), the value of the `readytoremove` property for the faulted drive is `true`:

     ```
     hostname:maintenance chassis-000 disk-001> get readytoremove
                     readytoremove = true
     ```
3. On the drive you want to remove, push the drive release button [1] to open the latch.
4. Grasp the latch [2] and pull the drive out of the drive slot [3].

5. After at least 30 seconds, from the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon on the system controller to verify that the software has detected that the drive is not present.
6. On the replacement drive, open the release lever and slide the drive into the same slot [1] until it is fully seated. Close the latch [2] to lock the drive in place.

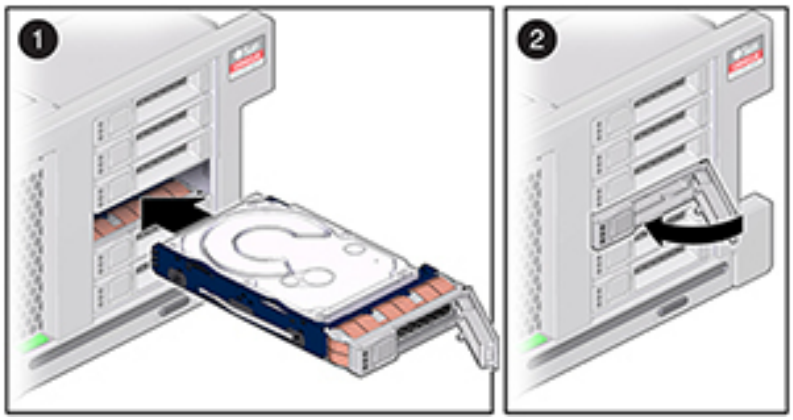

The Oracle ZFS Storage Appliance software automatically detects and configures the new drive.

7. From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, and verify that the status icon is green for the newly installed drive. To view details about the new drive, click its information icon .
8. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS4-4 Power Supply

Oracle ZFS Storage ZS4-4 controllers are equipped with redundant, hot-swappable power supplies. If a power supply fails and you do not have a replacement available, leave the failed power supply installed to ensure proper air flow. A faulted power supply is indicated by an amber-colored status LED. Use the following procedure to replace an Oracle ZFS Storage ZS4-4 power supply.

1. Gain access to the rear of the storage controller where the power supplies are located.

2. Disconnect the power cord from the faulted power supply.
3. To unlock the power supply handle, squeeze the release latch.

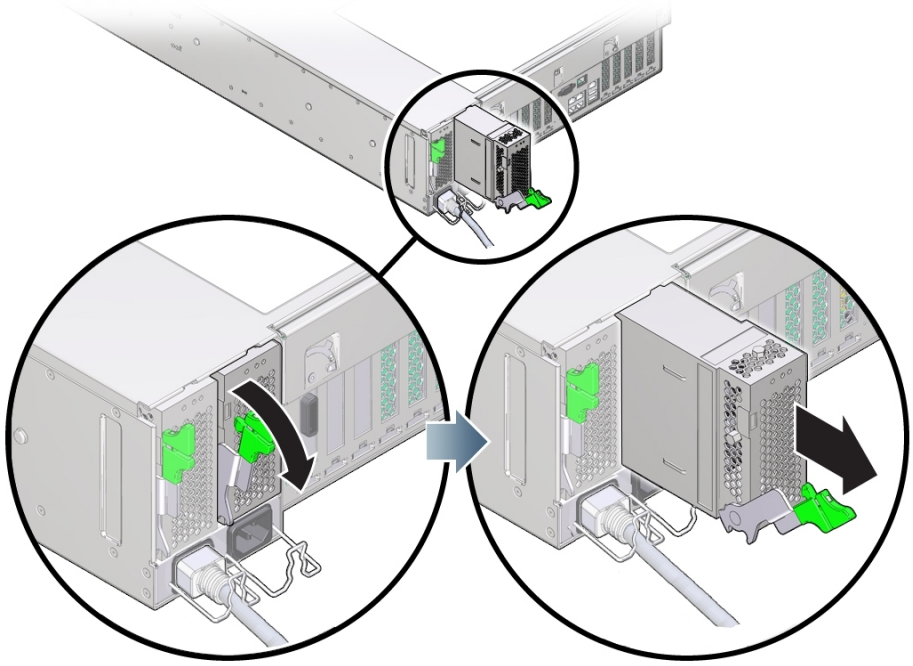

4. Rotate the release latch fully downward, to disengage the power supply from the internal power supply backplane connector.
5. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
6. Ensure that the replacement power supply handle is in the open position.
7. Align the replacement power supply with the empty power supply bay.

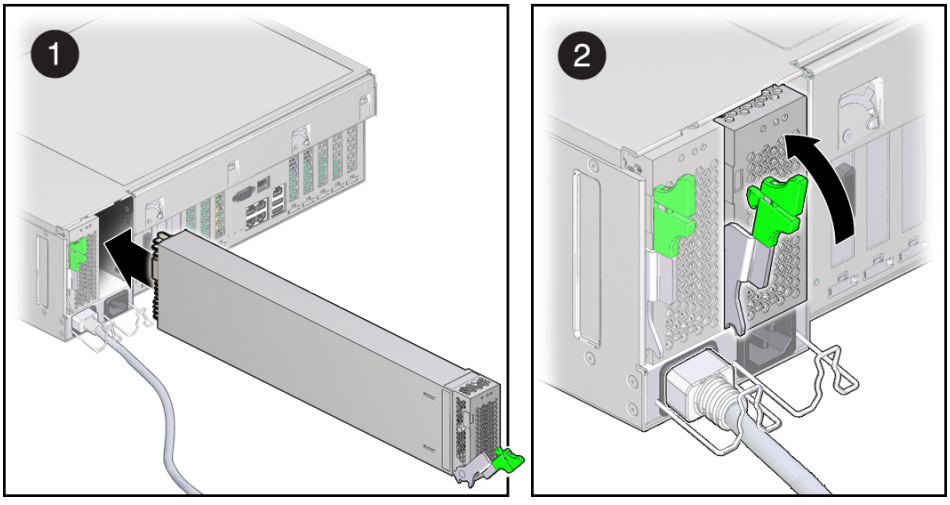

8. Slide the power supply into the bay [1] until it stops.

9. To seat the power supply, rotate the release handle upward [2] until it locks in place. This action of rotating the handle upward draws the power supply into the controller and engages it with the internal connector.

> **Note:**
>
> The latch at the hinge of the handle must engage the slot at the bottom of the power supply bay.

10. Connect the power cord to the power supply.
11. Verify that the green AC Present status indicator is lit.
12. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **PSU**, and verify that the status icon is green for the newly installed power supply. To view details about the new power supply, click its information icon .
13. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS4-4 Fan Module

> **Caution:**
>
> Operating a controller for an extended period of time with fans removed reduces the effectiveness of the cooling system. For this reason, the replacement fan should be unpacked in advance and ready to insert into the controller chassis as soon as the faulted fan is removed.

Fan modules cool the motherboard and its components. The fan modules are arranged in stacked pairs (a front row and a back row) to provide redundancy. If you remove a failed fan module, replace it immediately. Fan modules are hot-swappable and can be removed and installed while the storage controller is running without affecting other hardware capabilities. Use the following procedure to replace an Oracle ZFS Storage ZS4-4 fan module.

1. To flash the locator LED on the chassis to service, from the **Maintenance** BUI menu, select **Hardware**, and click the associated locator icon , or go to the service processor (SP) prompt and enter `set /SYS/LOCATE status=on`.
2. To locate the faulted component, from the **Maintenance** BUI menu, select **Hardware**, then select **Fan**, and click the row with an amber indicator.

   The component is highlighted in the diagram.
3. At the back of the controller, verify that cables have sufficient length and clearance to extend the controller from the rack.
4. Extend the controller from the rack as described in Extending the Controller from the Rack.
5. Lift the green handle at the top of the fan module to unlock the fan module, and pull the fan module straight up and out.

> **Caution:**
>
> Excessive movement or rocking of the fan module during removal can cause damage to the internal connector on the fan module board. When removing a fan module, do not rock it back and forth.

6. Position the replacement fan module so the connector on the bottom of the module is aligned with the connector on the inside of the slot and all keying and labeling are correctly positioned. The fan modules are keyed to ensure that they are installed in the correct orientation.

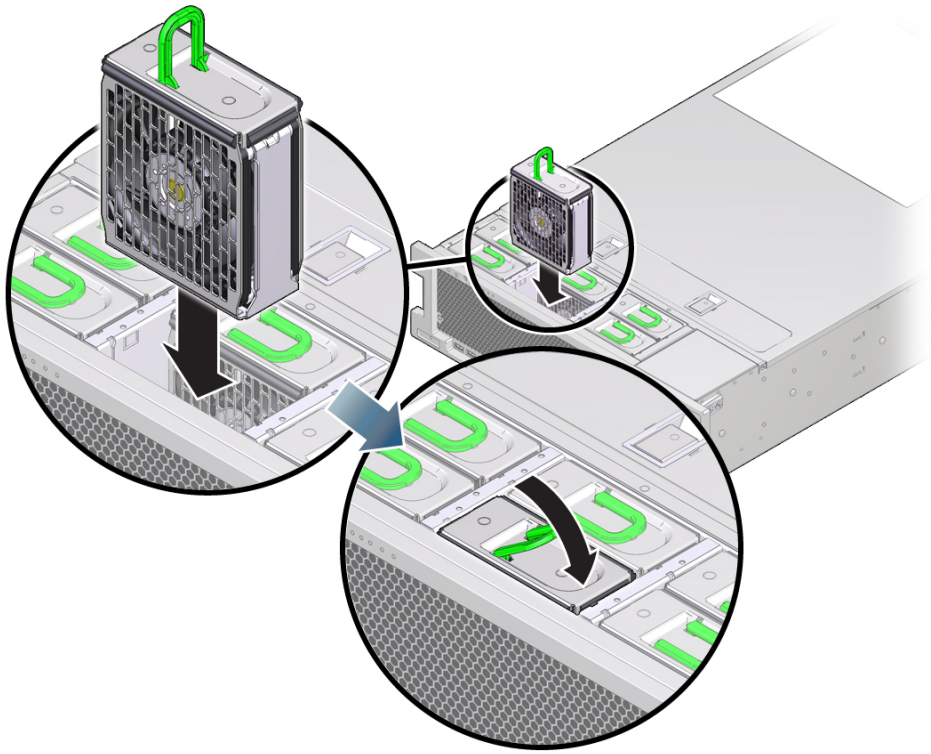

7. Slide the fan module into the slot until it stops.
8. Press downward on the top of the fan module, on the **Press Here to Latch** label, until the fan module is fully seated.
9. Replace the top cover on the chassis.
10. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.
11. Verify that the green OK indicator is lit.
12. Verify that the Top Fan indicator and the Service Action Required indicator on the front panel are not lit.
13. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **Fan**, and verify that the status icon is green for the newly installed fan module.
14. If a fault does not self-clear, see Repairing Active Problems.

Identifying a Faulted Oracle ZFS Storage ZS4-4 Memory Module

Replacing a DIMM is a cold-swap procedure, and you must remove power from the storage controller. Use the following procedure to identify a faulted Oracle ZFS Storage ZS4-4 memory module.

1. To identify a general memory fault, from the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon for the controller. Then click **DIMM** and note the ID of the faulted module, which has an amber icon. To highlight the location of the memory module in the controller picture, click on the module ID. To view details about the faulted module, click its information icon and then click **Active Problems**.
2. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
3. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

4. Extend the controller from the rack as described in Extending the Controller from the Rack.
5. Remove the top cover as described in Removing the Controller Top Cover.
6. To locate the faulted component, first press and hold the Fault Remind button to verify that the remind circuitry is usable. The Fault Remind button is located on the divider between cooling zone 1 and cooling zone 2. The Power LED, next to the button, is green when the remind circuitry is usable.

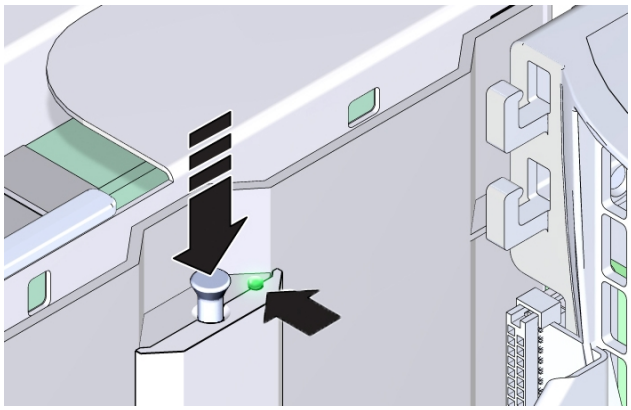

- If a memory riser card has failed, see Replacing an Oracle ZFS Storage ZS4-4 Memory Riser Card for removal and installation instructions.
- If a DIMM has failed, see Replacing an Oracle ZFS Storage ZS4-4 DIMM for removal and installation instructions.

- If a CPU has failed, the LEDs for both memory riser cards associated with the failed CPU turn on. The following example, shows the Fault indicators for memory riser cards, P0/MR0 and P0/MR1 are lit, as is the Fault indicator for CPU, P0.

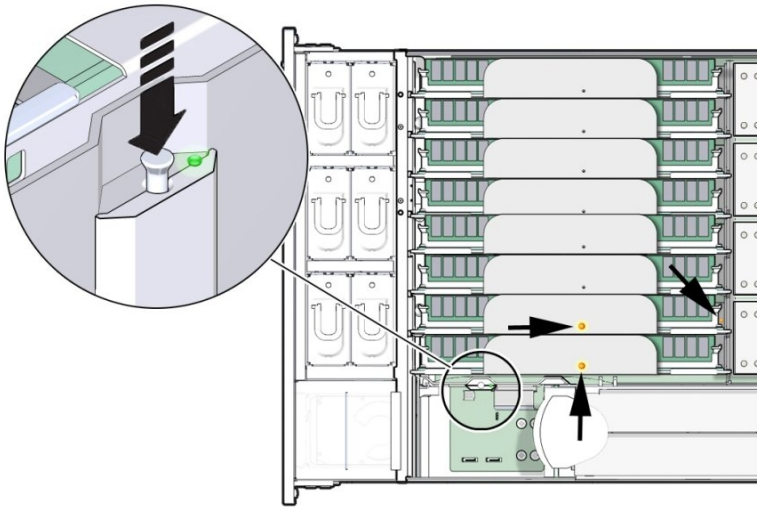

> **Caution:**
>
> A CPU is a field-replaceable unit (FRU) and must be replaced by a trained Oracle service technician.

Replacing an Oracle ZFS Storage ZS4-4 DIMM

Use the following procedure to replace an Oracle ZFS Storage ZS4-4 DIMM.

1. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. Remove the top cover as described in Removing the Controller Top Cover.
5. Remove the memory riser card with the lighted Fault indicators (P0/MR0 and P0/MR1), and locate the failed DIMM. Use the DIMM Fault Remind circuit to locate the failed DIMM. See Identifying a Faulted Oracle ZFS Storage ZS4-4 Memory Module.

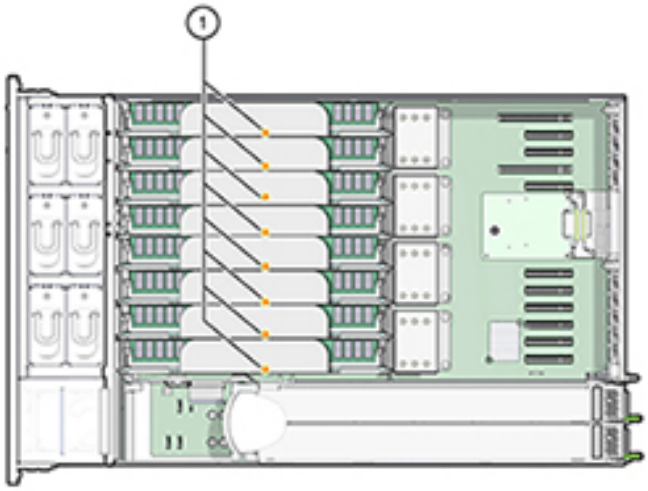

6. To remove the failed DIMM, rotate both DIMM slot ejector levers outward as far as they will go.

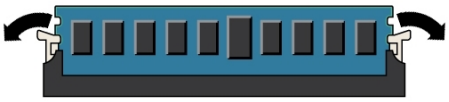

7. Carefully lift the DIMM straight up to remove it from the socket.

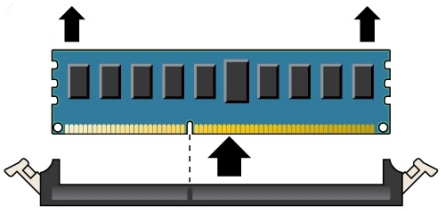

8. Repeat the previous steps for any other faulty DIMMs.
9. Ensure that the DIMM ejector levers at both ends of the DIMM slot are in a fully open position.
10. Align the replacement DIMM with the empty slot.

    The DIMM is keyed with a notch that must align with a protrusion in the DIMM slot. The keying ensures that the DIMM is installed correctly.
11. Gently and evenly push the DIMM into the slot until the ejector levers rise.

    The levers rise as the DIMM is pushed further into the slot.
12. Ensure that the levers have risen completely and have locked the DIMM in the slot.

> **Caution:**
>
> Open DIMM ejector levers can break off during the installation of the memory riser card. All populated and unpopulated DIMM slot ejector levers on the memory riser card must be in the fully closed position before installing the card in the controller. Ensure that all levers are closed and locked.

13. Replace the memory riser card. See Replacing an Oracle ZFS Storage ZS4-4 Memory Riser Card.
14. Replace the top cover on the chassis.
15. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.
16. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
17. Reconnect the power cords to the power supplies.
18. Verify that standby power is on, indicated by the Power/OK status indicator flashing on the front panel about two minutes after the power cords are plugged in.
19. Use a pen or other pointed object to press and release the recessed Power button on the storage controller front panel.

    The Power/OK status indicator next to the Power button lights and remains lit.
20. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **DIMM**, and verify that the status icon is green for the newly installed DIMM. To view details about the new DIMM, click its information icon .
21. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS4-4 Memory Riser Card

Use the following procedure to replace an Oracle ZFS Storage ZS4-4 memory riser card.

1. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. Remove the top cover as described in Removing the Controller Top Cover.
5. To remove the memory riser card, pull the handles upward to disengage the connector from the motherboard, and carefully lift the memory riser straight up and out of the controller. The handles act as levers against the sidewalls to extract the card connector from the connector on the motherboard.

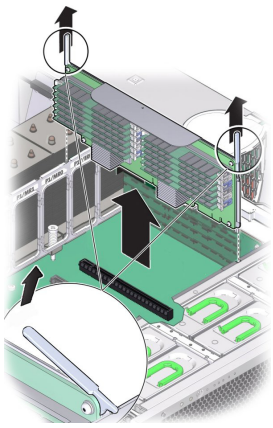

6. On the replacement memory riser card, ensure that all populated and unpopulated DIMM slot ejector levers are in the closed and locked position.

> **Caution:**
>
> Open DIMM ejector levers can break off during the installation of the memory riser card. All populated and unpopulated DIMM slot ejector levers on the memory riser card must be in the fully closed and locked position before installing the card in the controller.

7. In the controller, ensure that the memory riser card ejector levers are in the closed position.

> **Note:**
>
> The ejector levers are only used to remove a memory riser card and are not used to install a card.

8. Position the memory riser card over the slot on the motherboard.

   The DIMMs on the card must face to the left (when positioned at the front of the controller).

9. Lower the memory riser card into the controller and set it on the slot on the motherboard.

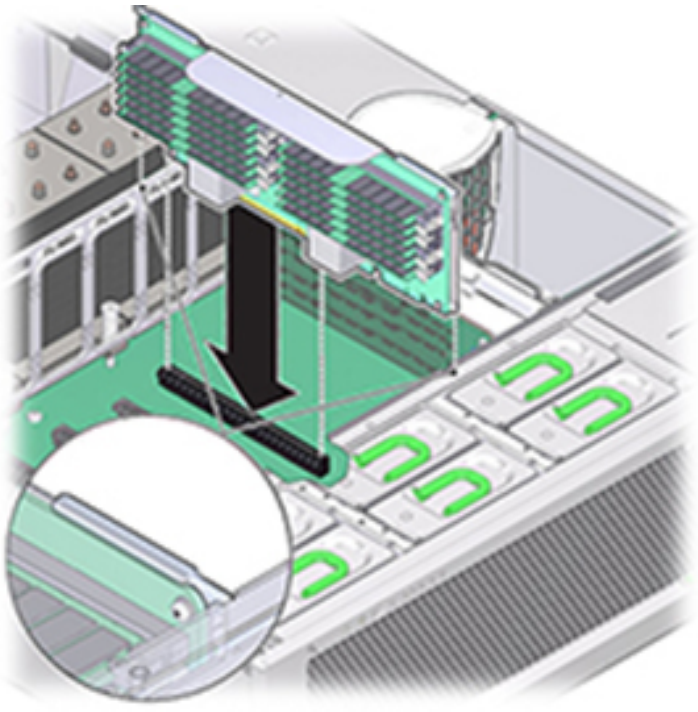

10. Ensure that the memory riser card connector is aligned with the slot.
11. Firmly push down on the metal bracket on top of the card, to seat the card within the connector on the motherboard.
12. Replace the top cover on the chassis.
13. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.
14. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
15. Reconnect the power cords to the power supplies.
16. Verify that standby power is on, indicated by the Power/OK status indicator flashing on the front panel about two minutes after the power cords are plugged in.
17. Use a pen or other pointed object to press and release the recessed Power button on the storage controller front panel.
18. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **DIMM**, and verify that the status icon is green for the newly installed memory riser card.
19. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS4-4 PCIe Card

Replacing an Oracle ZFS Storage ZS4-4 PCIe card is a cold-swap procedure, and you must remove power from the storage controller. If you are installing a newly released HBA, upgrade your system software before installing the HBA. All HBAs must be of the same type. Use the following procedure to replace an Oracle ZFS Storage ZS4-4 PCIe card.

1. From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, click **Slot**, and note the ID of the faulted PCIe card, which has an amber icon. To highlight the location of the card in the controller picture, click on the card ID. To view details about the faulted card, click its information icon and then click **Active Problems**.

   > **Caution:**
   >
   > This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

2. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
3. Disconnect the AC power cords from the rear panel of the storage controller.

   > **Caution:**
   >
   > Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

4. Extend the controller from the rack as described in Extending the Controller from the Rack.
5. Remove the top cover as described in Removing the Controller Top Cover.
6. Locate the PCIe card position in the storage controller.
7. To disengage the PCIe card retaining bar, push down on the bar, move it away from the controller back wall, and lift it to an upright position.

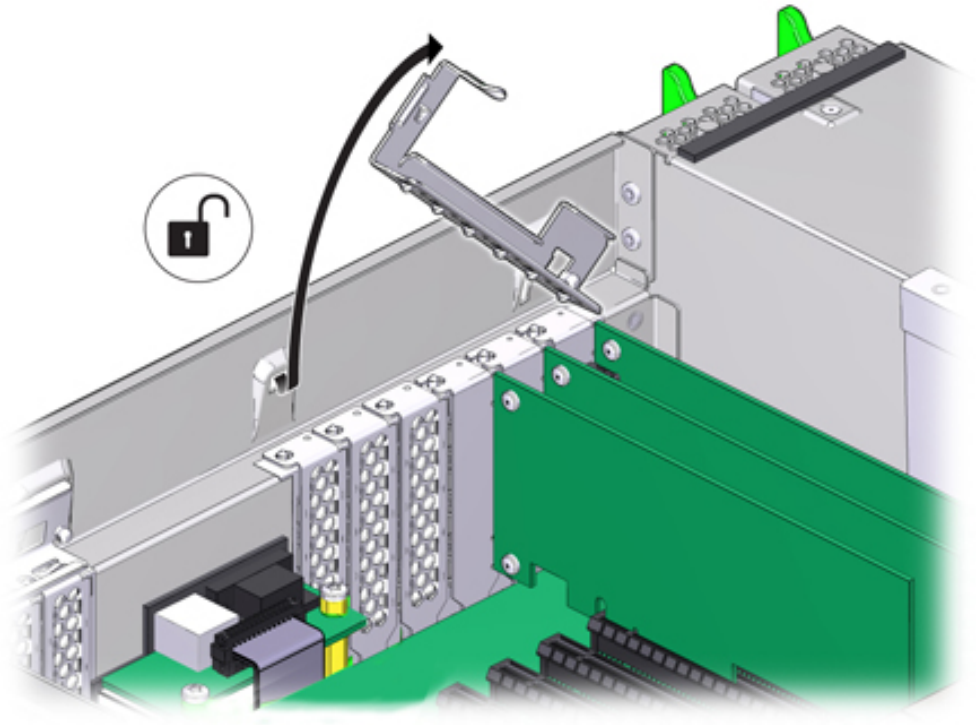

8. Carefully remove the PCIe card from the PCIe card slot.

> **Caution:**
>
> Whenever you remove a PCIe card, you should replace it with another PCIe card or a filler panel; otherwise, the controller might overheat due to improper airflow.

9. Ensure the PCIe card retaining bar is open.

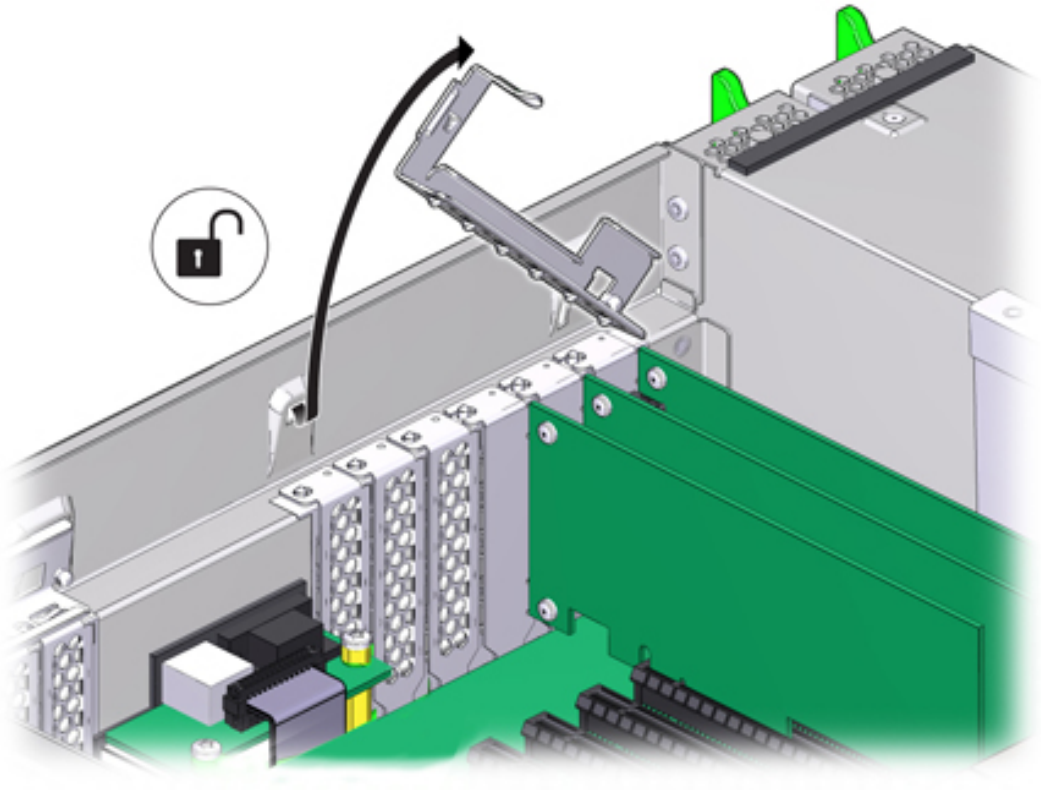

10. Install the replacement PCIe card into the PCIe card slot.

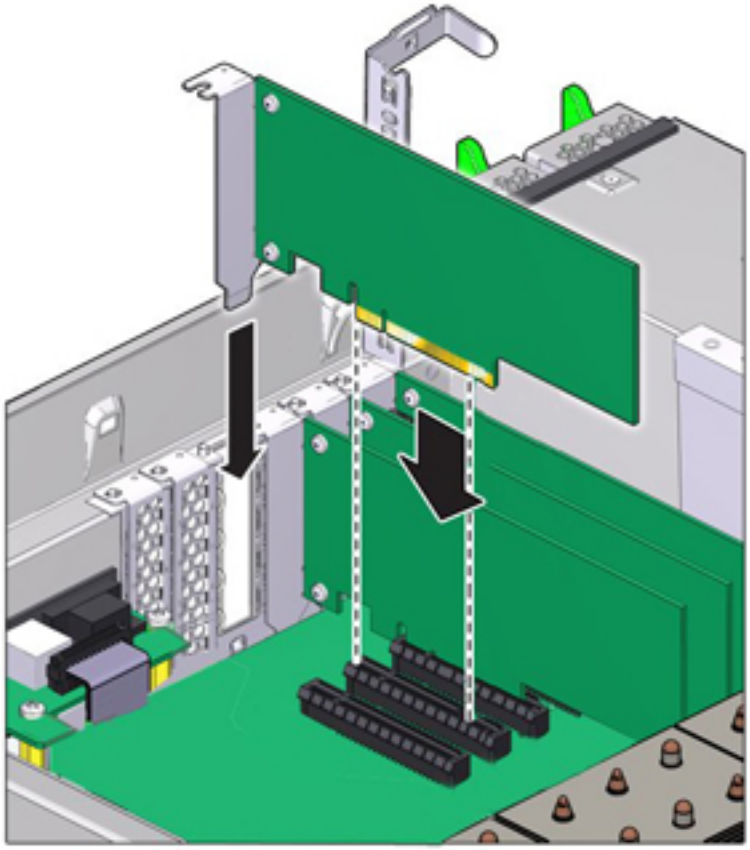

11. Return the retaining bar to its closed and locked position.
12. Replace the top cover on the chassis.
13. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.
14. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
15. Reconnect the power cords to the power supplies.
16. Verify that standby power is on, indicated by the Power/OK status indicator flashing on the front panel about two minutes after the power cords are plugged in.
17. Use a pen or other pointed object to press and release the recessed Power button on the storage controller front panel. The Power/OK status indicator next to the Power button lights and remains lit.
18. Connect data cables to the PCIe card, routing them through the cable management arm, if used.
19. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, click **Slot**, and verify that the status icon is green for the newly installed card. To view details about the new card, click its information icon .
20. If a fault does not self-clear, see Repairing Active Problems.

21. If you replaced an HBA, connect the disk shelves to it as described in Connecting to Attached Storage. If you installed a new HBA, install the disk shelf as described in Installation Prerequisites and Hardware Overviews in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*, and connect the expansion storage as described in Connecting to Attached Storage.

Replacing an Oracle ZFS Storage ZS4-4 Battery

Replacing an Oracle ZFS Storage ZS4-4 system battery is a cold-swap procedure, and you must remove power from the storage controller. You might need to replace the battery if the storage controller fails to maintain the proper time when powered off and disconnected from the network. You will need a small (No. 1 flat-blade) non-metallic screwdriver or equivalent. Use the following procedure to replace an Oracle ZFS Storage ZS4-4 system battery.

1. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. Remove the top cover as described in Removing the Controller Top Cover.

   The battery is located toward the back of the controller between PCIe slot 1 and the power supply side wall.

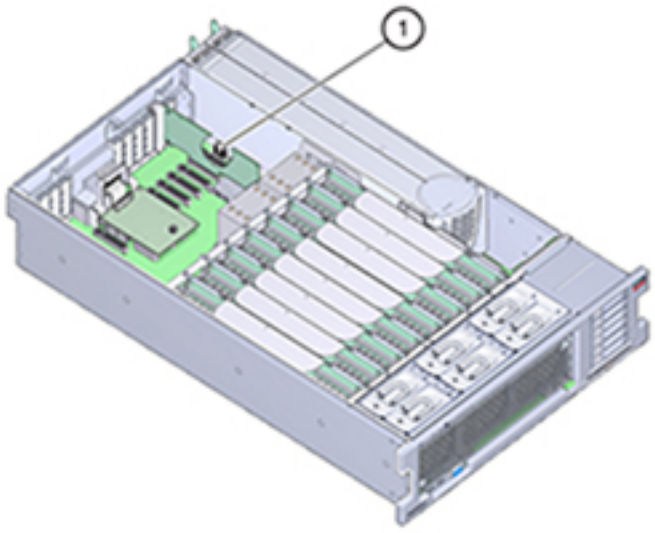

5. If necessary, remove the PCIe card in slot 1 to access the battery. See Replacing an Oracle ZFS Storage ZS4-4 PCIe Card.

6. Remove the battery from the battery holder by pushing the back of the battery, negative (-) side, toward the metal tab, on the positive (+) side, and lifting the battery up and out of the battery holder.

   Take care not to bend the tab when pushing against it.

   > **Caution:**
   >
   > Do not deform the metal tab on the positive (+) side of the battery. The metal tab maintains the positive battery connection and secures the battery in the holder.

7. Press the new battery into the battery holder with the positive (+) side facing the metal tab that holds it in place.

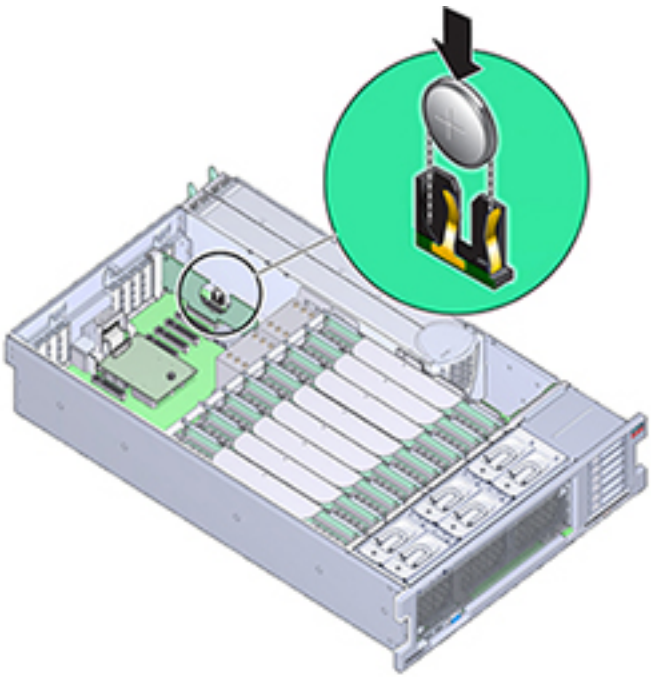

8. If you removed the PCIe card from slot 1, re-install it. See Replacing an Oracle ZFS Storage ZS4-4 PCIe Card.
9. Replace the top cover on the chassis.
10. Return the controller to the normal rack position by pushing the release tabs on the side of each rail while slowly pushing the storage controller into the rack.
11. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
12. Connect the power cords to the power supplies.
13. Verify that standby power is on, indicated by the Power/OK status indicator flashing on the front panel about two minutes after the power cords are plugged in.

14. Use a pen, or other pointed object, to press and release the recessed Power button on the storage controller front panel. The Power/OK status indicator next to the Power button lights and remains lit.
15. If you re-installed a PCIe card, connect the data cables to the PCIe card, routing them through the cable management arm, if used.
16. When the system has finished booting, log in, and set the time using the steps in the BUI Clock task. For more information, see Setting Clock Synchronization (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Oracle ZFS Storage ZS4-4 Controller Hardware Overview

To understand the Oracle ZFS Storage ZS4-4 controller, use the following topics:

- Oracle ZFS Storage ZS4-4 Base Configuration
- Oracle ZFS Storage ZS4-4 Front and Rear Panel Components
- Oracle ZFS Storage ZS4-4 PCIe I/O Cards
- Oracle ZFS Storage ZS4-4 Internal Components
- Oracle ZFS Storage ZS4-4 Attached Storage

Oracle ZFS Storage ZS4-4 Base Configuration

The Oracle ZFS Storage ZS4-4 controller can be configured as a standalone controller or clustered controllers to create a high-availability cluster configuration. The following table describes the base configuration.

Table 8-1 Oracle ZFS Storage ZS4-4 Controller Base Configuration

Component	Description
CPU	Four Intel Xeon 15-core, 2.8 GHz
Memory	1.5TB 16GB DDR3 LV RDIMM
Boot Disks	Two 2.5-inch SAS-2 HDDs
Read Flash	Up to four optional 2.5-inch SAS-2 SSDs
HBAs	Two 4x4-port SAS-2 (base configuration)
PCIe slots	11 (4 base configuration, 7 expansion slots)

Oracle ZFS Storage ZS4-4 Front and Rear Panel Components

Oracle ZFS Storage ZS4-4 Front Panel Components - The Oracle ZFS Storage ZS4-4 controller drive slots and front panel components are shown in the following figure.

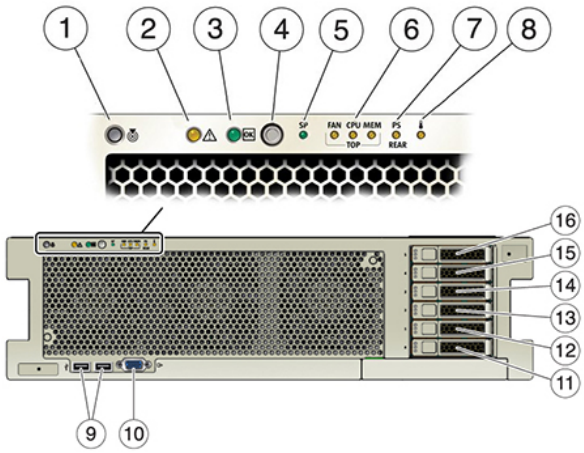

1 Locator LED/button (white)	**9** USB 2.0 connectors (2)
2 Service action required LED (amber)	**10** DB-15 video port
3 Power/OK LED (green)	**11** Boot drive 0 (required)
4 Power button	**12** Boot drive 1 (required)
5 Service Processor (SP) OK (green)/Fault (amber) LED	**13** Solid state drive 2 (optional)
6 Fan/CPU/Memory Service action required LED (amber)	**14** Solid state drive 3 (optional)
7 Power Supply (PS) Service action required LED (amber)	**15** Solid state drive 4 (optional)
8 Over temperature warning LED (amber)	**16** Solid state drive 5 (optional)

Oracle ZFS Storage ZS4-4 Rear Panel Components - The Oracle ZFS Storage ZS4-4 rear panel is shown in the following figure. Base configuration PCIe cards are not depicted in this illustration.

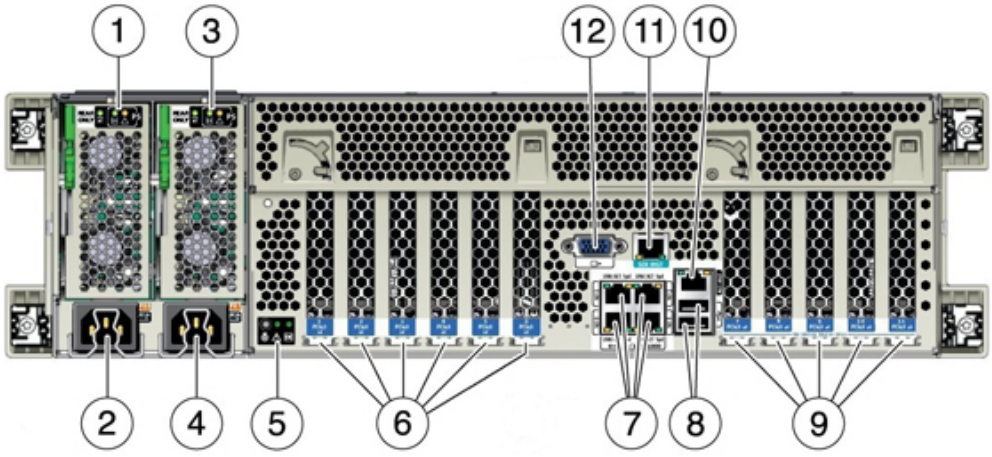

1 Power supply unit (PSU) 0 indicator panel	**7** Network (NET) 10 GbE ports: NET0NET3
2 PSU 0 AC inlet	**8** USB 2.0 connectors (2)
3 PSU 1 indicator panel	**9** PCIe card slots 7-11
4 PSU 1 AC inlet	**10** Network management (NET MGT) 10/100/1000 BASE-T Ethernet port
5 System status indicator panel	**11** Serial management (SER MGT) RJ-45 port
6 PCIe card slots 1-6	**12** DB-15 video port

Cluster Interface Card Compatibility - The Oracle ZFS Storage ZS4-4 controller can contain a Version 2 or Version 3 cluster interface card. For clustered controllers, ensure you are using the same version in each controller because these two versions are not compatible.

For information about how to connect cables to form a cluster, see Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

Ethernet Ports

The Oracle ZFS Storage ZS4-4 controller has four RJ-45 10-Gigabit Ethernet (10GbE) network connectors on the rear panel, labeled NET 0, NET 1, NET 2, and NET 3 (bottom left to top right), as shown in the following figure. Use these ports to connect the appliance to the network.

The LEDs located above the NET ports, labeled 2, 0, 3, 1 (left to right) are Link/Activity indicators.

LED	Status
OFF (1)	No Link
ON (0)	Link and no activity
Blink	Link and activity

> **Note:**
>
> Speed is not indicated for the NET ports.

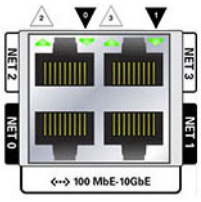

Network Management Port

The Oracle ZFS Storage ZS4-4 network management connector (NET MGT), shown in the following figure, is an RJ-45 10/100/1000 BASE-T Ethernet port and provides an alternate terminal interface to the service processor (SP) console.

Serial Management Port

The Oracle ZFS Storage ZS4-4 serial management connector (SER MGT), shown in the following figure, is an RJ-45 port and provides a terminal connection to the SP console.

4x4 SAS HBA

The Oracle ZFS Storage ZS4-4 4x4 SAS HBA provides connectivity to external disk shelves. The HBA logical ports are numbered 3-0, top to bottom, as shown in the following figure.

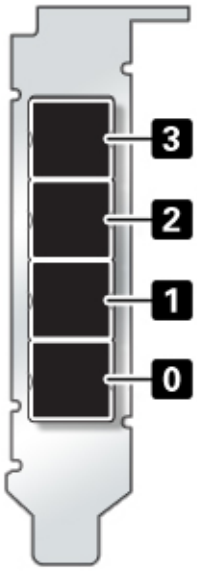

Oracle ZFS Storage ZS4-4 PCIe I/O Cards

The Oracle ZFS Storage ZS4-4 base configuration contains the following PCIe cards:

- One 8-port SAS-2 internal HBA (slot 2)
- Two 4-port (4x4) SAS external HBAs (slot 6 and slot 7)
- One cluster interface card (slot 4)

The following figure shows the PCIe I/O slot numbers.

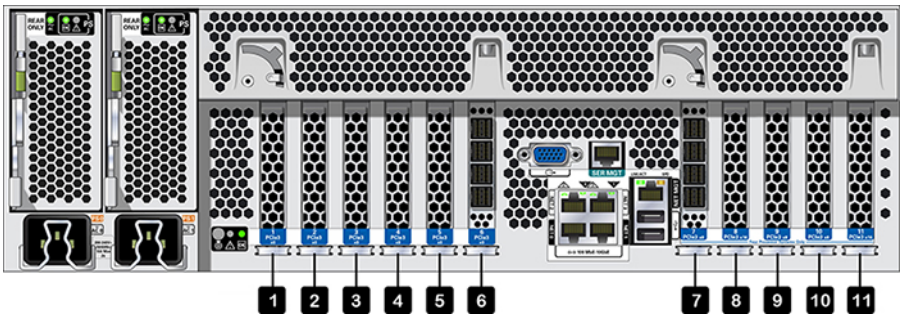

PCIe Slot Order - Install optional PCIe cards in the following order:

- Install additional 4x4 SAS HBAs in slot 9, then slot 3.
- Install InfiniBand CX3 HCAs into the first available client-option slot starting with slot 11, then slot 8, slot 5, slot 1, slot 10, slot 3, slot 9.
- Install 16Gb FC HBAs into the first available client-option slot starting with slot 11, then slot 8, slot 5, slot 1, slot 10, slot 3, slot 9.
- Install 10Gb Ethernet Optical NICs into the first available client-option slot starting with slot 11, then slot 8, slot 5, slot 1, slot 10, slot 3, slot 9.
- Install 10Gb Ethernet Copper NICs into the first available client-option slot starting with slot 11, then slot 8, slot 5, slot 1, slot 10, slot 3, slot 9.

Oracle ZFS Storage ZS4-4 PCIe Base and Optional Configurations - The following table describes the PCIe base and optional slot assignments for Oracle ZFS Storage ZS4-4 standalone and cluster configurations. PCIe slot numbering begins with slot 1.

See the legend for a description of the interconnect types and option codes shown in the Type column.

Slot	Description	Max	Type	Note
1	2-port InfiniBand CX3 HCA	4	A	Optional recommended front-end
1	2-port 10Gb Ethernet Optical NIC	4	C	Optional recommended front-end
1	2-port 10Gb Ethernet Copper NIC	4	D	Optional recommended front-end
1	16GB Dual Universal FC HBA	4	B	Optional FC target or initiator (Backup)
2	8-port SAS-2 Internal HBA	1	F	Base configuration
3	4-port (4x4) SAS External HBA	4	E	Optional back-end
3	2-port InfiniBand CX3 HCA	4	A	Optional recommended front-end
3	2-port 10Gb Ethernet Optical NIC	4	C	Optional recommended front-end
3	2-port 10Gb Ethernet Copper NIC	4	D	Optional recommended front-end

Slot	Description	Max	Type	Note
3	16GB Dual Universal FC HBA	4	B	Optional FC target or initiator (Backup)
4	Cluster Interface (second generation)	1	G	Base configuration
5	2-port InfiniBand CX3 HCA	4	A	Optional recommended front-end
5	2-port 10Gb Ethernet Optical NIC	4	C	Optional recommended front-end
5	2-port 10Gb Ethernet Copper NIC	4	D	Optional recommended front-end
5	16GB Dual Universal FC HBA	4	B	Optional FC target or initiator (Backup)
6	4-port (4x4) SAS External HBA	4	E	Base configuration
7	4-port (4x4) SAS External HBA	4	E	Base configuration
8	2-port InfiniBand CX3 HCA	4	A	Optional recommended front-end
8	2-port 10Gb Ethernet Optical NIC	4	C	Optional recommended front-end
8	2-port 10Gb Ethernet Copper NIC	4	D	Optional recommended front-end
8	16GB Dual Universal FC HBA	4	B	Optional FC target or initiator (Backup)
9	4-port (4x4) SAS External HBA	4	E	Optional back-end
9	2-port InfiniBand CX3 HCA	4	A	Optional recommended front-end
9	2-port 10Gb Ethernet Optical NIC	4	C	Optional recommended front-end
9	2-port 10Gb Ethernet Copper NIC	4	D	Optional recommended front-end
9	16GB Dual Universal FC HBA	4	B	Optional FC target or initiator (Backup)
10	2-port InfiniBand CX3 HCA	4	A	Optional recommended front-end
10	2-port 10Gb Ethernet Optical NIC	4	C	Optional recommended front-end
10	2-port 10Gb Ethernet Copper NIC	4	D	Optional recommended front-end
10	16GB Dual Universal FC HBA	4	B	Optional FC target or initiator (Backup)
11	2-port InfiniBand CX3 HCA	4	A	Optional recommended front-end
11	2-port 10Gb Ethernet Optical NIC	4	C	Optional recommended front-end
11	2-port 10Gb Ethernet Copper NIC	4	D	Optional recommended front-end
11	16GB Dual Universal FC HBA	4	B	Optional FC target or initiator (Backup)

The legend describes the interconnect types and option codes shown in the Type column.

A	InfiniBand QDR QSFP+ - QSFP Direct Copper Cable Connect - Optical Transceiver QSFP Short Range 40Gbs	E	Storage Array 4-port External SAS Mini-SAS HD - SFF-8644 Mini-SAS HD to SFF-8088 Mini-SAS Copper (used to connect Thebe2/3 to DE2) - SFF-8644 Mini-SAS HD to SFF-8644 Mini-SAS HD Copper (used to connect Thebe2/3 to DE3)

B	Fibre Channel 16Gb SFP+ - Optical Transceiver SFP Short Range 16Gbs	F	HDD 8-port Internal SAS-2 Mini-SAS - SFF-8087 Mini-SAS to Mini-SAS Copper
C	GbE NIC Multi-mode Fiber SFP+ 10GBase-SR/LR - Optical Transceiver SFP Short Range 10Gbs - Optical Transceiver SFP Long Range 10Gbs	G	Server Heartbeat 2-port RS-232 1-port 1GBase-T - Copper RJ-45 Serial RS-232 - Copper RJ-45 Unshielded Twisted Pair
D	GbE NIC UTP 10GBase-T - Copper RJ-45 Unshielded Twisted Pair		

Oracle ZFS Storage ZS4-4 Internal Components

The Oracle ZFS Storage ZS4-4 chassis contains both customer-replaceable units (CRUs) and field-replaceable units (FRUs) as shown in the following figure. FRUs must be replaced by trained Oracle service technicians.

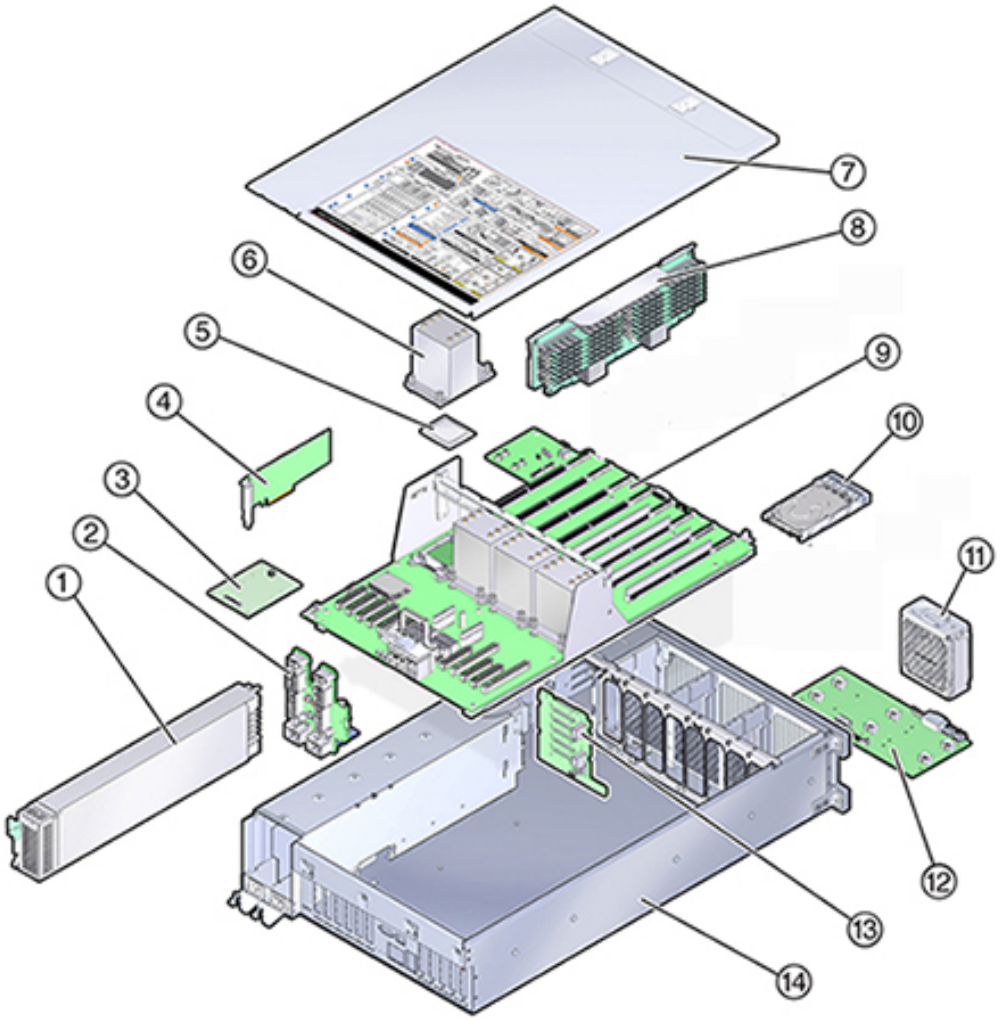

1 Power supplies (CRU)	**8** Memory riser card (CRU)
2 Power supply backplane (FRU)	**9** Motherboard (FRU)
3 SP card (FRU)	**10** System drive (CRU)
4 HBA/PCIe cards (CRU)	**11** Fan module (CRU)
5 CPU (FRU)	**12** Fan board (FRU)
6 Heatsink (FRU)	**13** Drive backplane (FRU)
7 Cover	**14** Chassis

Oracle ZFS Storage ZS4-4 System Drive - The Oracle ZFS Storage ZS4-4 controller has two system boot drives in slots 0 and 1, configured as a mirrored pair. Read-optimized cache devices can be installed in controller slots 2 through 5, or installed in disk shelf slots 20 through 23.

> **Note:**
>
> Read cache devices must be installed in either the controller or a disk shelf. Read cache devices cannot be installed in both controller and disk shelf slots at the same time. For software requirements and slot configuration rules, see Disk Shelf Configurations.

A filler panel must be installed in empty drive slots. The system drive LEDs are shown in the following figure.

1 Blue (Ready to remove)	**2** Service action required (amber)	**3** OK/Activity (green)

Oracle ZFS Storage ZS4-4 CPU and Memory - The Oracle ZFS Storage ZS4-4 controller has four Intel Xeon E7-8895 v2 15-core 2.8 GHz CPUs and eight memory riser cards as shown in the following figure. The memory configuration is 16GB DDR3 DIMMs to accommodate up to 1.5TB (ninety-six 16GB). All ZS4-4 DIMM risers are fully populated to accommodate this offering.

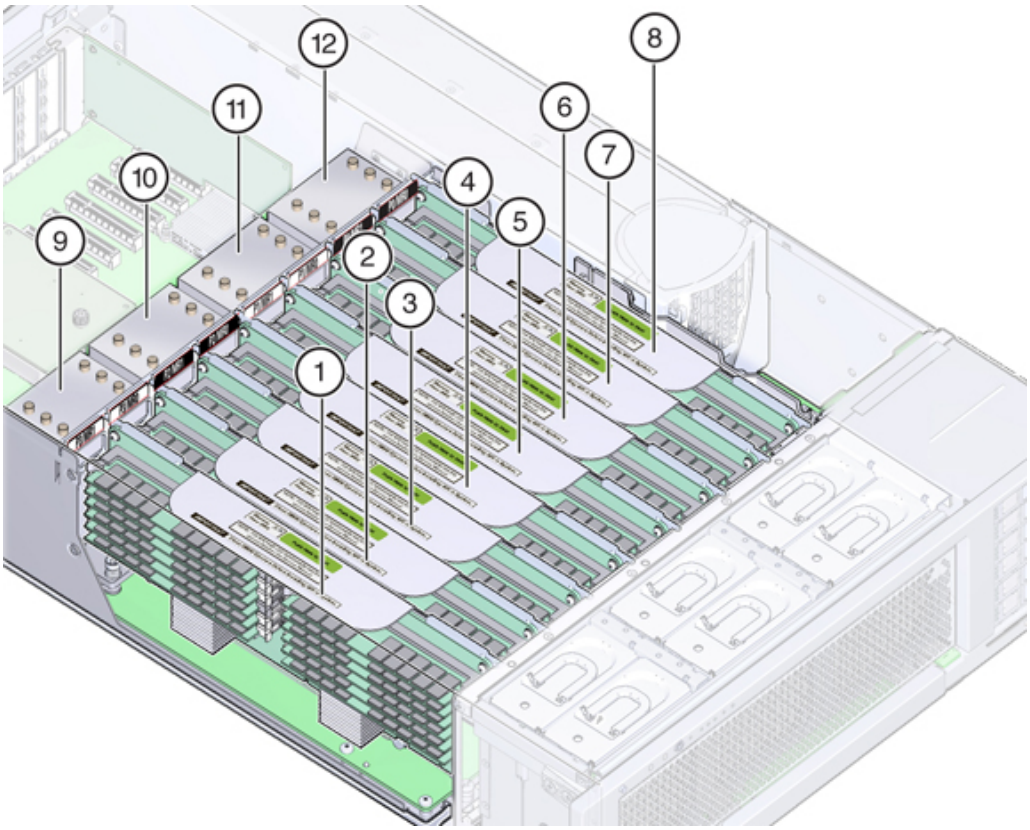

1 Memory riser card P3/MR1	**7** Memory riser card P0/MR1
2 Memory riser card P3/MR0	**8** Memory riser card P0/MR0
3 Memory riser card P2/MR1	**9** CPU P3
4 Memory riser card P2/MR0	**10** CPU P2
5 Memory riser card P1/MR1	**11** CPU P1
6 Memory riser card P1/MR0	**12** CPU P0

Each memory riser card contains twelve DIMM slots, four DDR3 channels, and two memory buffer ASICs. Each each memory buffer has two channels (A and B) and links to three DIMM slots per channel. Each memory buffer is connected to the processor's built-in memory controller by an SMI-2 link.

DIMM names in appliance logs and the **Maintenance: Hardware** view are displayed with the full name, such as `/SYS/MB/P0/D7`.

For more information about memory layout and procedures for replacing DIMMs, see Replacing an Oracle ZFS Storage ZS4-4 DIMM.

Oracle ZFS Storage ZS4-4 Cooling Subsystem - The Oracle ZFS Storage ZS4-4 internal components are cooled by air that is pulled in through the front of the controller and exhausted out the back of the controller. Cooling occurs in two areas of the chassis: the power supply area and the motherboard area.

The following figure shows the cooling zones and the approximate location of the temperature sensors. The accompanying legend table provides sensor NAC names and sensor motherboard designations.

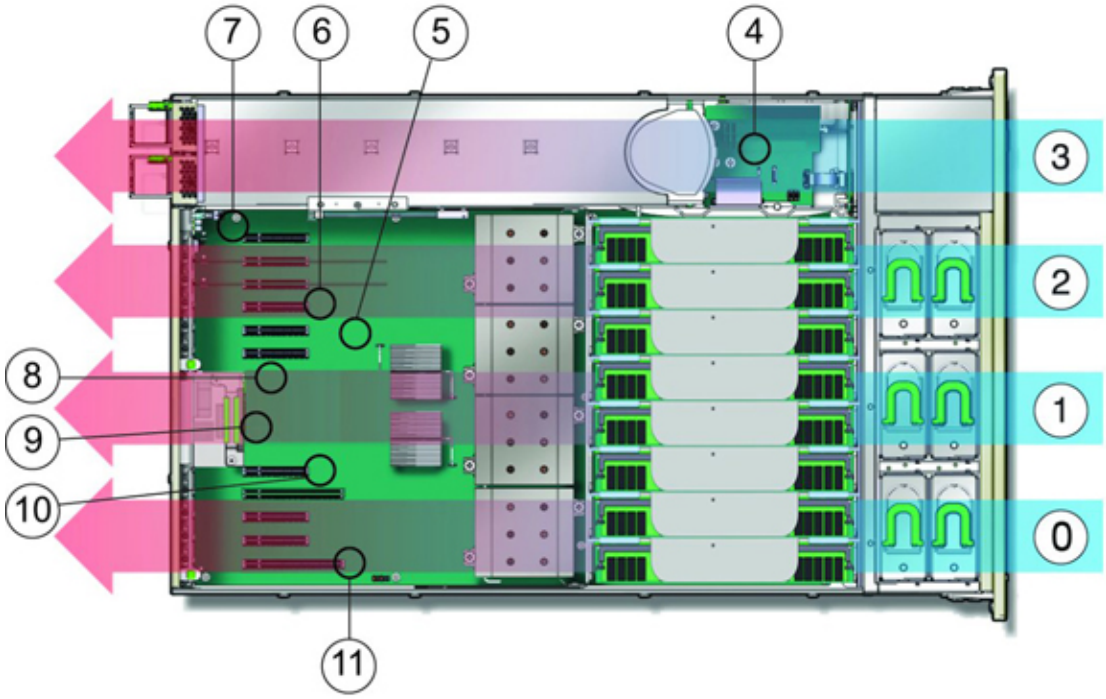

0 Cooling zone 0	**6** Temperature sensor TS_ZONE2 (U4505)
1 Cooling zone 1	**7** Temperature sensor TS_OUT (U4506)
2 Cooling zone 2	**8** Temperature sensor TS_TVL_1 (U4002)
3 Cooling zone 3 (power supply backplane area)	**9** Temperature sensor TS_TVL_0 (U4302)
4 Temperature sensor TS_PS (U4603)	**10** Temperature sensor TS_ZONE0_B (U4509)
5 Temperature sensor TS_ZONE1 (U4507)	**11** Temperature sensor TS_ZONE0_A (U4508)

Oracle ZFS Storage ZS4-4 Attached Storage

The Oracle ZFS Storage ZS4-4 standalone and cluster controller configurations allow up to 36 disk shelves, consisting of up to six chains of one to six disk shelves. Any combination of disk-only, or disk, log, and read cache devices, can be combined within the chain in any order. The Sun Disk Shelf is not supported. For more information, see Disk Shelf Configurations.

> **Note:**
>
> Read cache devices must be installed in either the controller or a disk shelf. Read cache devices cannot be installed in both controller and disk shelf slots at the same time. For software requirements and slot configuration rules, see Disk Shelf Configurations.

Related Topics

- Connecting to Attached Storage
- Disk Shelf Hardware Overview
- Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x

9

Servicing the Oracle ZFS Storage ZS3-4 Controller

Be sure to read sections Getting Started with Servicing the Appliance and Prerequisites for Servicing Hardware.

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see Repairing Active Problems.

Oracle Storage Drive Enclosure DE2-24s and Sun Disk Shelves can be used together with standalone and clustered controllers after certain hardware modifications. Contact Oracle Service to upgrade your Oracle ZFS Storage ZS3-4 controllers for disk shelf interoperability.

To service Oracle ZFS Storage ZS3-4 hardware components, use the following tasks.

- Replacing an Oracle ZFS Storage ZS3-4 HDD or SSD
- Replacing an Oracle ZFS Storage ZS3-4 Fan Module
- Replacing an Oracle ZFS Storage ZS3-4 Power Supply
- Replacing an Oracle ZFS Storage ZS3-4 DIMM
- Replacing an Oracle ZFS Storage ZS3-4 PCIe Card or Riser
- Replacing an Oracle ZFS Storage ZS3-4 Battery

To understand Oracle ZFS Storage ZS3-4, see the following topics:

- Oracle ZFS Storage ZS3-4 Base Configuration
- Oracle ZFS Storage ZS3-4 Specifications
- Oracle ZFS Storage ZS3-4 Front and Rear Panel Components
- Oracle ZFS Storage ZS3-4 Internal Components
- Oracle ZFS Storage ZS3-4 PCIe I/O Cards
- Oracle ZFS Storage ZS3-4 Attached Storage

Replacing an Oracle ZFS Storage ZS3-4 HDD or SSD

HDDs and SSDs are hot-swappable and can be removed and installed while the storage controller is running without affecting other hardware capabilities. Use the following procedure to replace an Oracle ZFS Storage ZS3-4 HDD or SSD.

> **Note:**
>
> If you have more than one failed drive, replace only one drive at a time. Removing more than one drive in quick succession causes a hardware/pool fault. Also, do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** BUI menu, select **System**, or navigate to `maintenance system updates` in the CLI.

1. Identify the faulted HDD or SSD.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, and open the detailed view of the controller. The faulted drive has either an amber or a blue-on-amber (ready to remove) status icon. Note the ID of the faulted drive. Click the entry for that drive in the table to highlight the location of the drive in the controller illustration. To view details about the faulted drive, click its information icon and then click **Active Problems**.
   - If you are physically at the system, the amber Service Action Required indicator on the HDD or SSD should be lit. Additionally, you can click the locator icon in the BUI to light the locator LED on the faulted drive.
2. Verify that the faulted drive is ready to be removed.

   > **Caution:**
   >
   > Removing the drive before it is ready to be removed can result in loss of data.

   Use any of the following methods to verify that the faulted drive is ready to be removed:
   - On the faulted drive, the Ready to Remove LED is lit.
   - In the BUI:
     - The status icon for the faulted drive in the detail view in **Maintenance: Hardware** has changed from amber to blue-on-amber.
     - The **Component Details** information window (click the information icon) for the faulted drive shows **Yes** as the value for the **Ready to remove** property.
   - In the CLI (and RESTful API), the value of the `readytoremove` property for the faulted drive is `true`:

     ```
     hostname:maintenance chassis-000 disk-001> get readytoremove
                     readytoremove = true
     ```
3. Push the release button on the HDD or SSD to open the latch.
4. Grasp the latch [2], and pull the drive out of the drive slot.

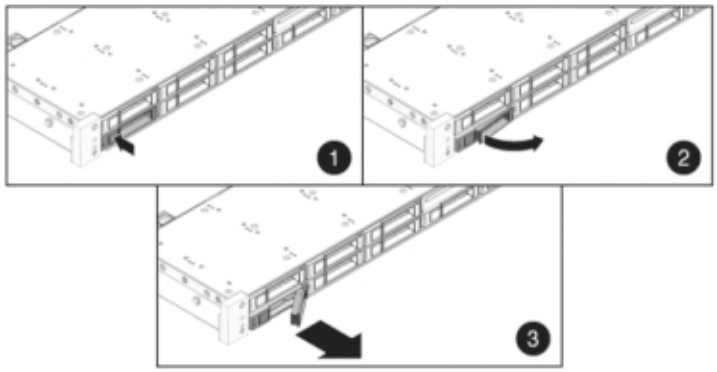

5. After at least 30 seconds, from the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon on the system controller to verify that the software has detected that the drive is not present.
6. Slide the replacement drive into the slot until it is fully seated.

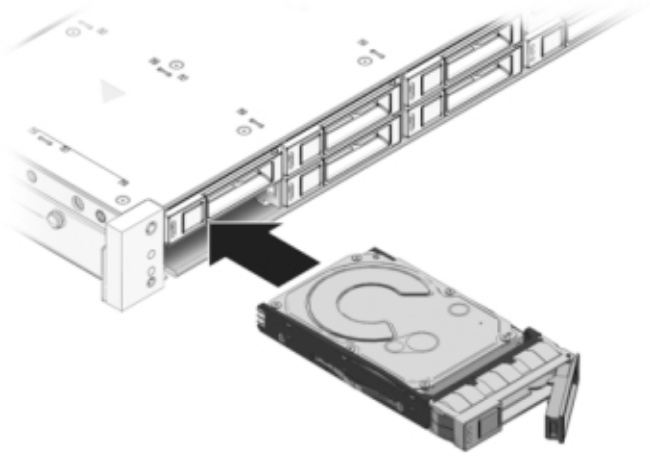

7. Close the latch to lock the drive in place.

   The Oracle ZFS Storage Appliance software automatically detects and configures the new drive.
8. From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, and verify that the status icon is green for the newly installed drive. To view details about the new drive, click its information icon .
9. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS3-4 Fan Module

> **Caution:**
>
> Operating a controller for an extended period of time with fans removed reduces the effectiveness of the cooling system. For this reason, the replacement fan should be unpacked in advance and ready to insert into the controller chassis as soon as the faulted fan is removed.

Fan modules are hot-swappable and can be removed and installed while the storage controller is running without affecting other hardware capabilities. Use the following procedure to replace an Oracle ZFS Storage ZS3-4 fan module.

1. To flash the locator LED on the chassis to service, from the **Maintenance** BUI menu, select **Hardware**, and click the associated locator icon, or go to the service processor (SP) prompt and enter `set /SYS/LOCATE status=on`.
2. To locate the faulted component, from the **Maintenance** BUI menu, select **Hardware**, then select **Fan**, and click the row with an amber indicator.

   The component is highlighted in the diagram.
3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. Lift the latch at the top of the fan module to unlock the fan module, and pull out the fan module.

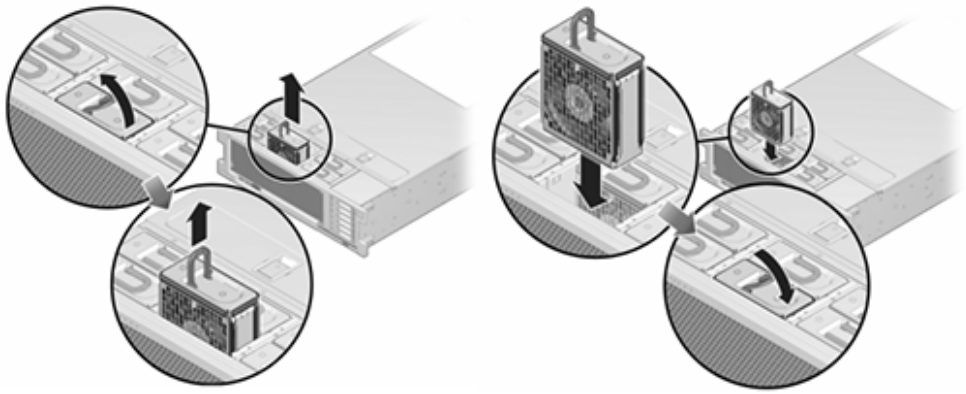

5. Unlock and insert the fan module.
6. Apply firm pressure to fully seat the fan module.
7. On the replaced fan module, verify that the Fan OK indicator is lit, and that the fault indicator is not lit.
8. Verify that the Top Fan status indicator, the Service Required status indicators, and the Locator status indicator/Locator button are not lit.

9. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller and then click **Fan** to verify that the online icon is green for the newly installed fan module.
10. If a fault does not self-clear, see Repairing Active Problems.
11. Push the release tabs on the side of each rail and slowly slide the storage controller into the rack.

Replacing an Oracle ZFS Storage ZS3-4 Power Supply

Oracle ZFS Storage ZS3-4 controllers are equipped with redundant hot-swappable power supplies. If a power supply fails and you do not have a replacement, leave the failed power supply installed to ensure proper air flow. A faulted power supply is indicated by an amber colored status LED. Use the following procedure to replace an Oracle ZFS Storage ZS3-4 power supply.

1. Gain access to the rear of the storage controller where the faulted power supply is located.
2. If a cable management arm (CMA) is installed, press and hold the CMA release tab and rotate the arm out of the way.

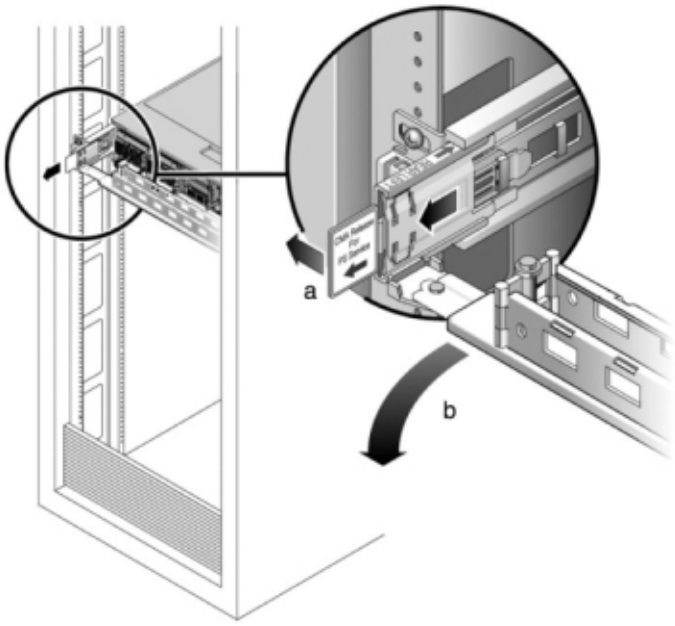

3. Disconnect the power cord from the faulted power supply.
4. Grasp the power supply handle and press the release latch to remove the power supply.

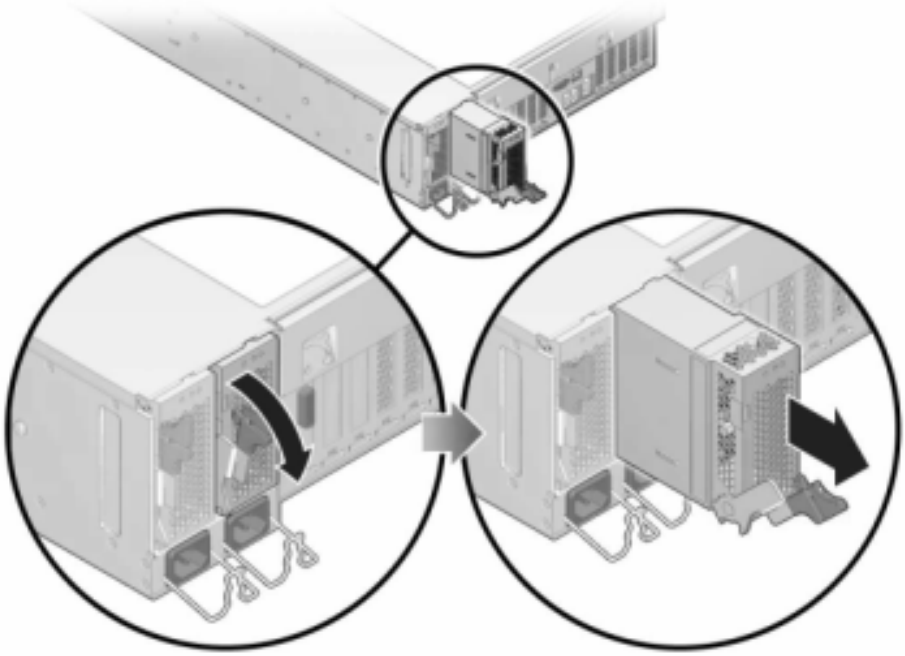

5. Align the replacement power supply with the empty power supply chassis bay.
6. Slide the power supply into the bay until it is fully seated.

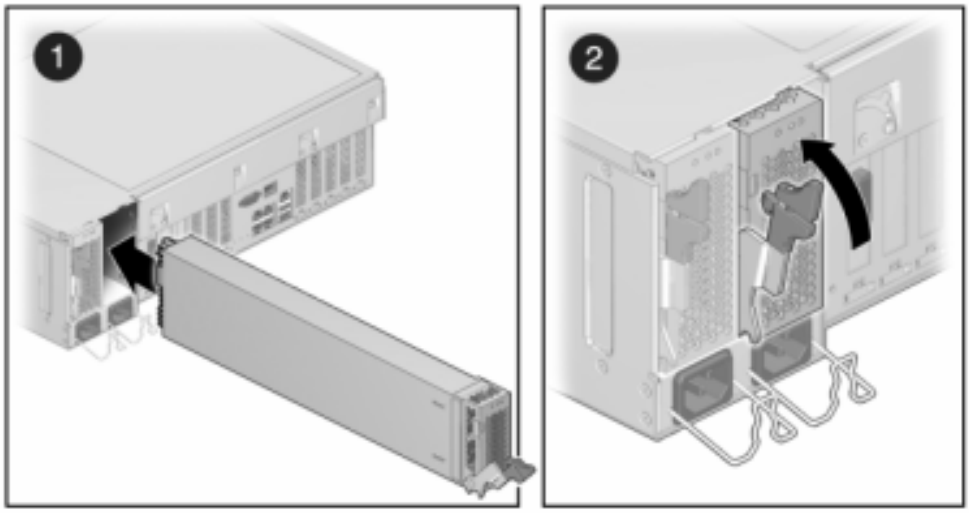

7. Connect the power cord to the power supply.
8. Verify that the green AC Present status indicator is lit.
9. Close the CMA, inserting the CMA into the rear left rail bracket.
10. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller and then click **PSU** to verify that the online icon is green for the newly installed power supply.
11. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS3-4 DIMM

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

To identify a specific memory module that has faulted, you must open the storage controller and use the amber status LEDs on the motherboard. To identify a general memory fault, from the **Maintenance** BUI menu, select **Hardware**, and click on the right-arrow icon on the controller. Then click **DIMM** to locate the faulted component, indicated by the warning icon .

1. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. Remove the top cover as described in Removing the Controller Top Cover.
5. To locate the DIMM you want to service, press the Fault Remind button on the storage controller.

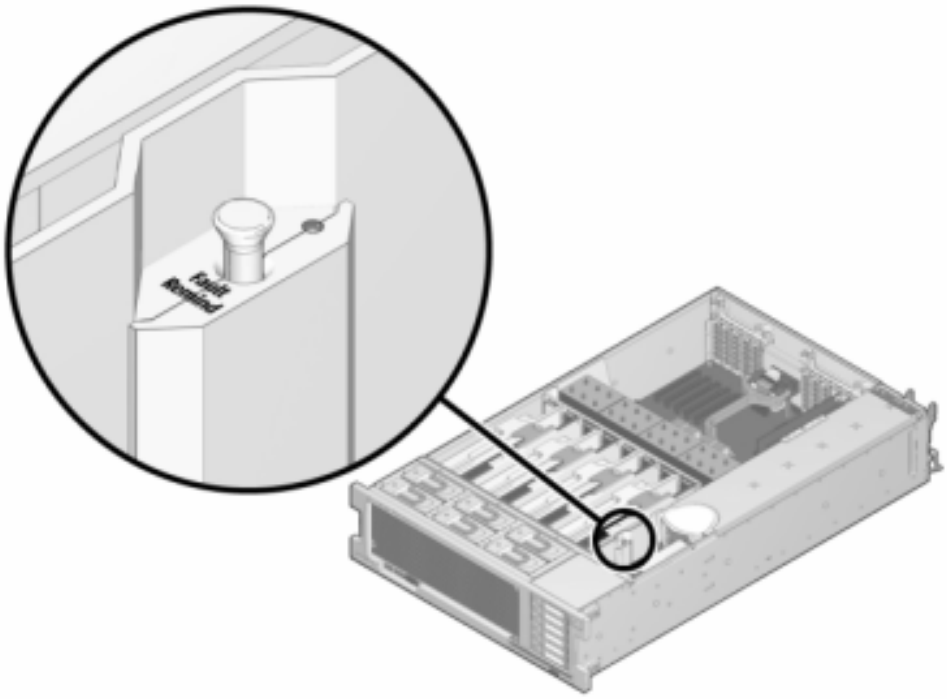

6. Identify the memory riser that hosts the faulted DIMM by the Service Required status indicator. Lift the memory riser straight up to remove it from the motherboard, and place it on an antistatic mat.

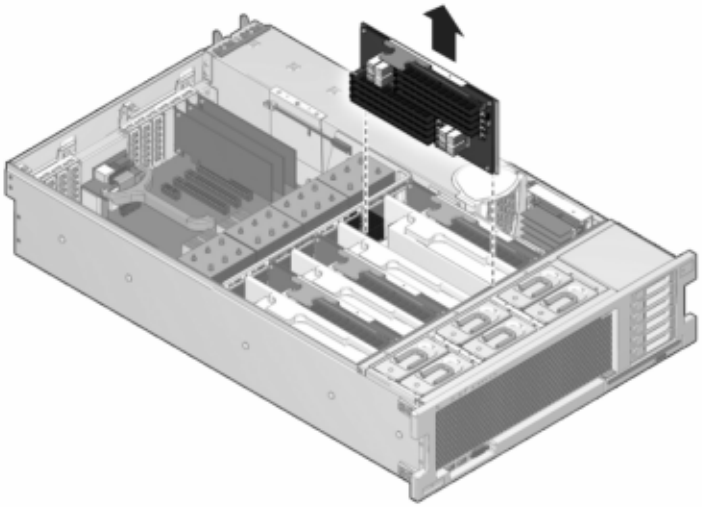

7. Rotate both DIMM slot ejectors outward as far as they will go and carefully lift the faulted DIMM straight up to remove it from the socket.

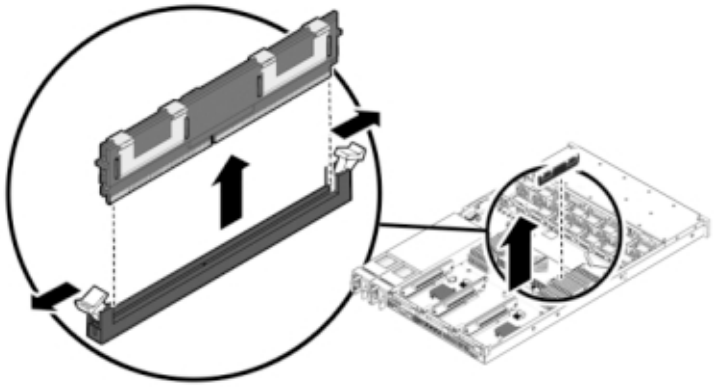

8. Line up the replacement DIMM with the connector, aligning the notch with the key to ensure that the component is oriented correctly.

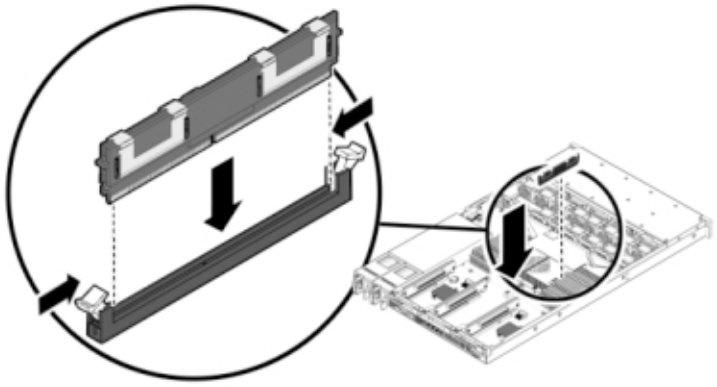

9. Push the DIMM into the connector until the ejector tabs lock the component in place.
10. To replace the cover:
    a. Push the memory riser module into the associated CPU memory riser slot until the riser module locks in place.

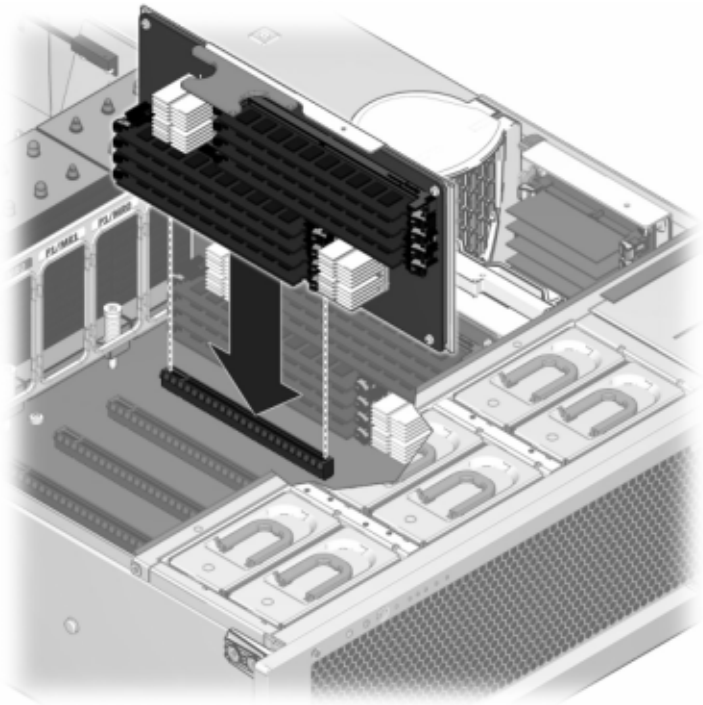

b. Place the top cover on the chassis so that it is forward of the rear of the storage controller by about an inch (2.5 cm).

c. Slide the top cover toward the rear of the chassis until it seats and press down on the cover with both hands until both latches engage.

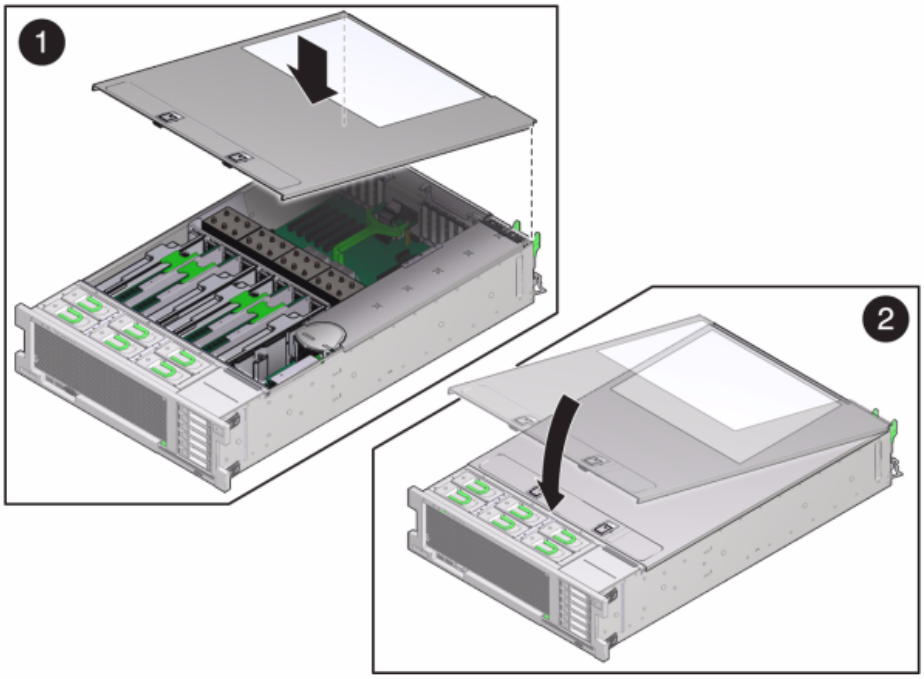

11. Push the release tabs on the side of each rail and slowly push the storage controller into the rack.

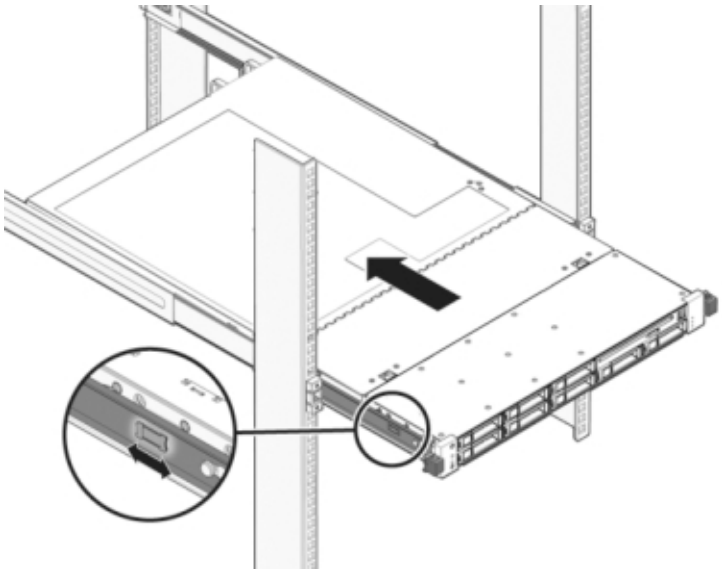

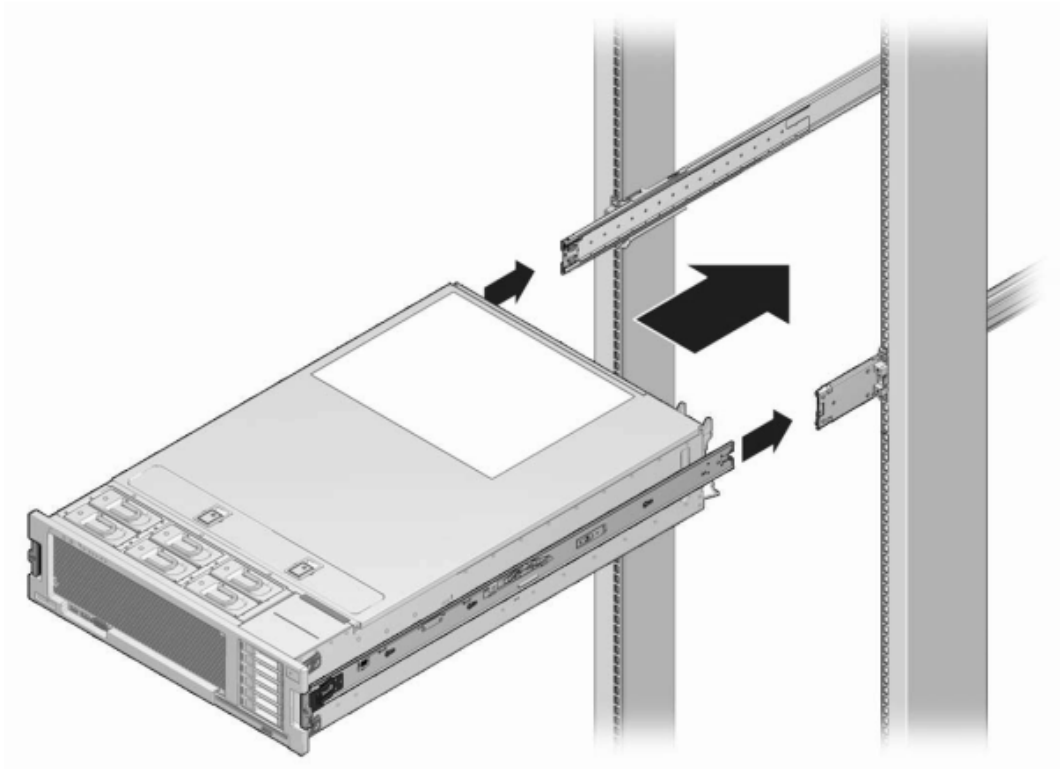

12. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
13. Connect the power cords to the power supplies.

14. Verify that standby power is on, indicated by the Power/OK status indicator flashing on the front panel about two minutes after the power cords are plugged in.
15. Use a pen or other pointed object to press and release the recessed Power button on the storage controller front panel.

    The Power/OK status indicator next to the Power button lights and remains lit.
16. From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon on the controller, and click **DIMM** to verify the status of the new component. The online icon should be green.
17. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS3-4 PCIe Card or Riser

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon on the controller, and then click **Slot** to locate the faulted component. All HBAs must be of the same type. Ensure that you upgrade your system software before installing a newly released HBA.

1. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the storage controller.

   > **Caution:**
   >
   > Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.
3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. Remove the top cover as described in Removing the Controller Top Cover.
5. Locate the PCIe card position in the storage controller, see Oracle ZFS Storage ZS3-4 PCIe I/O Cards.
6. To replace the PCIe card:
   a. Disengage the PCIe card slot crossbar from its locked position and rotate the crossbar into an upright position.

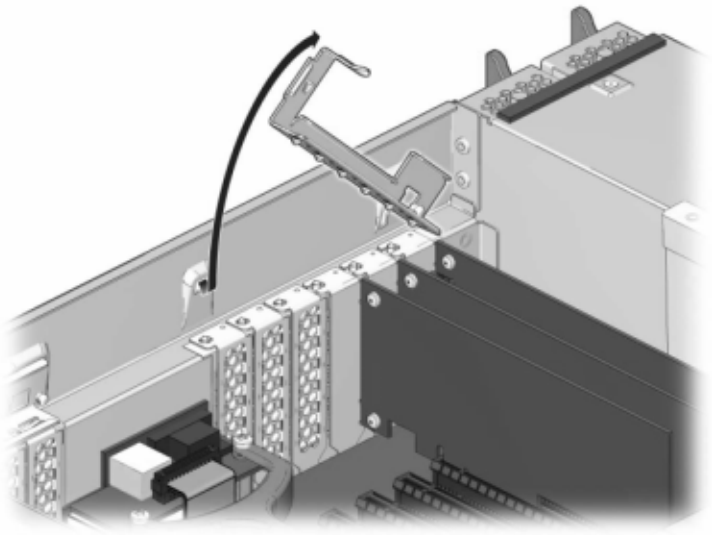

b. Remove the retaining screw that holds the PCIe card to the chassis.

c. Carefully remove the PCIe card from the connector and clean the slot with filtered, compressed air if necessary.

d. Install the replacement PCIe card into the PCIe card slot.

e. Install the retaining screw to hold the PCIe card to the chassis.

f. Return the crossbar to its closed and locked position.

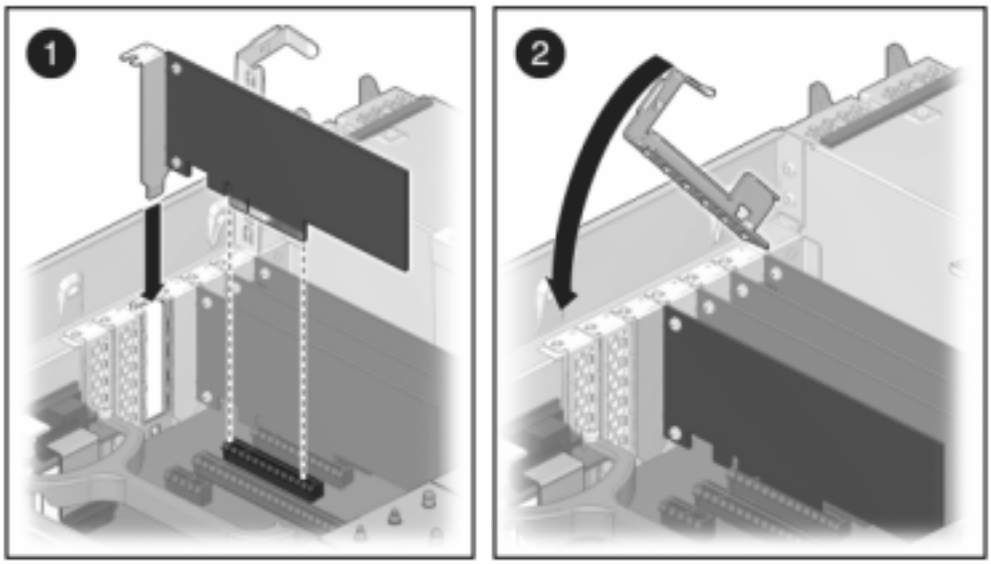

7. To install the top cover:

   a. Place the top cover on the chassis [1] so that it is forward of the rear of the storage controller by about an inch (2.5 cm).

   b. Slide the top cover toward the rear of the chassis [2] until it seats.

c. Press down on the cover with both hands until both latches engage.

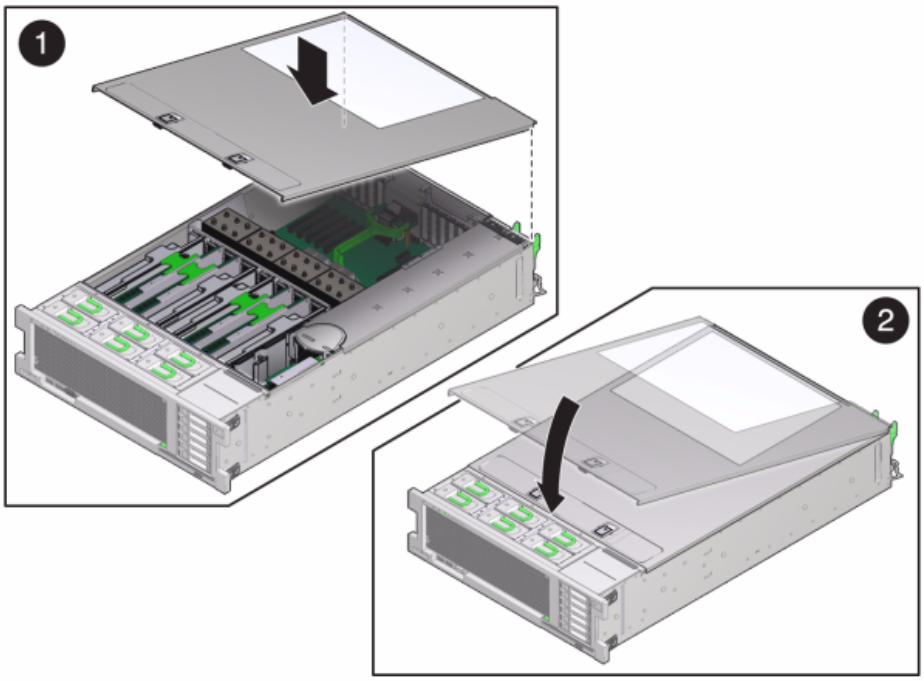

8. Push the release tabs on the side of each rail and slowly push the storage controller into the rack, making sure no cables obstruct the path of the controller.
9. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
10. Connect the power cords to the power supplies.
11. Verify that standby power is on, indicated by the Power/OK status indicator flashing on the front panel about two minutes after the power cords are plugged in.
12. Use a pen or other pointed object to press and release the recessed Power button on the storage controller front panel. The Power/OK status indicator next to the Power button lights and remains lit.
13. Connect data cables to the PCIe card, routing them through the cable management arm.
14. From the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon on the controller. Then click **Slot** to verify the status of the new component. The online icon should be green.
15. If a fault does not self-clear, see Repairing Active Problems.
16. Install the disk shelf as described in Installation Prerequisites and Hardware Overviews in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*, and connect the expansion storage as described in Connecting to Attached Storage.

Replacing an Oracle ZFS Storage ZS3-4 Battery

You might need to replace the battery if the Oracle ZFS Storage ZS3-4 controller fails to maintain the proper time when powered off and disconnected from the network. You will need a small (No. 1 flat-blade) non-metallic screwdriver or equivalent.

1. Shut down the controller using one of the power-off methods described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. Remove the top cover as described in Removing the Controller Top Cover.
5. Using a small, non-metallic screwdriver, press the latch and remove the battery from the motherboard.

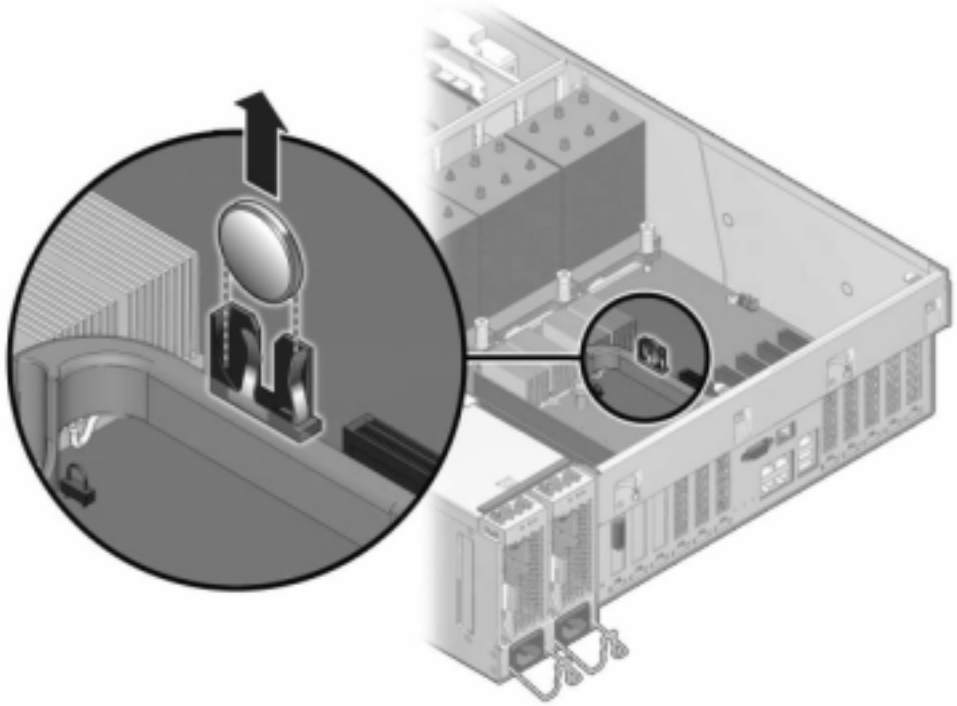

6. Press the replacement battery into the motherboard with the positive side (+) facing upward.
7. To install the top cover:

a. Place the top cover on the chassis [1] so that it is forward of the rear of the storage controller by about an inch (2.5 cm).

b. Slide the top cover toward the rear of the chassis [2] until it seats.

c. Press down on the cover with both hands until both latches engage.

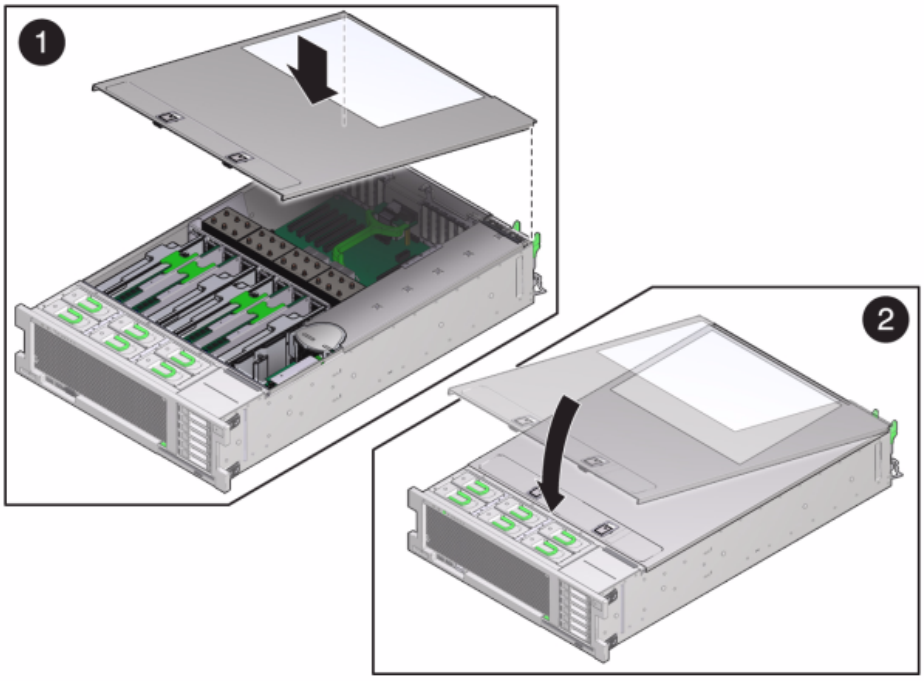

8. Push the release tabs on the side of each rail and slowly push the storage controller into the rack, making sure no cables obstruct the path of the controller.
9. Connect the power cords to the power supplies.
10. Verify that standby power is on, indicated by the Power/OK status indicator flashing on the front panel about two minutes after the power cords are plugged in.
11. Use a pen, or other pointed object, to press and release the recessed Power button on the storage controller front panel. The Power/OK status indicator next to the Power button lights and remains lit.
12. Connect data cables to the PCIe card, routing them through the cable management arm.
13. When the system has finished booting, log in, and set the time using the steps in the task Setting Clock Synchronization (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Oracle ZFS Storage ZS3-4 Controller Hardware Overview

To understand Oracle ZFS Storage ZS3-4, see the following topics:

- Oracle ZFS Storage ZS3-4 Base Configuration
- Oracle ZFS Storage ZS3-4 Specifications
- Oracle ZFS Storage ZS3-4 Front and Rear Panel Components

Oracle ZFS Storage ZS3-4 Base Configuration

The Oracle ZFS Storage ZS3-4 controller can be configured as a standalone controller or clustered controllers to create a high-availability cluster configuration. The following table describes the configuration options:

Table 9-1 Oracle ZFS Storage ZS3-4 Controller Features

Mktg Part Number	CPU	Memory	Read Cache SAS-2	Boot Drive SAS-2	HBA SAS-2	Software Version (min)
7105725	4x10-core, 2.40GHz	1TB (16GB DIMMs)	Four 1.6TB	Two 900GB	4X4-port	2013.1.0

Oracle ZFS Storage ZS3-4 Specifications

Oracle ZFS Storage ZS3-4 Physical Specifications - The Oracle ZFS Storage ZS3-4 controller chassis fits in a standard equipment rack, and occupies three rack units (3RU) in height. The chassis dimensions are as follows:

Table 9-2 Oracle ZFS Storage ZS3-4 Controller Dimensions

Dimension	Measurement	Dimension	Measurement
Height	13.3 cm/5.25 in	Depth	70.6 cm/27.8 in
Width	43.7 cm/17.19 in	Weight	16.36 kg/96 lbs

Oracle ZFS Storage ZS3-4 Electrical Specifications - The following list shows the electrical specifications for the controller.

> **Note:**
>
> The power dissipation numbers listed are the maximum rated power numbers for the power supply. The numbers are not a rating of the actual power consumption of the appliance.

Input

- Nominal frequencies: 50/60Hz
- AC operating range: 200-240 VAC
- Maximum current AC RMS: 12A @ 200 VAC

Power Dissipation

- Max power consumption: 1800 W

- Max heat output: 6143 BTU/hr
- Volt-Ampere rating: 1837 VA @ 240 VAC, 0.98 P.F.

Oracle ZFS Storage ZS3-4 Front and Rear Panel Components

Oracle ZFS Storage ZS3-4 Front Panel Components - The Oracle ZFS Storage ZS3-4 controller drive slots and front panel components are shown in the following figure.

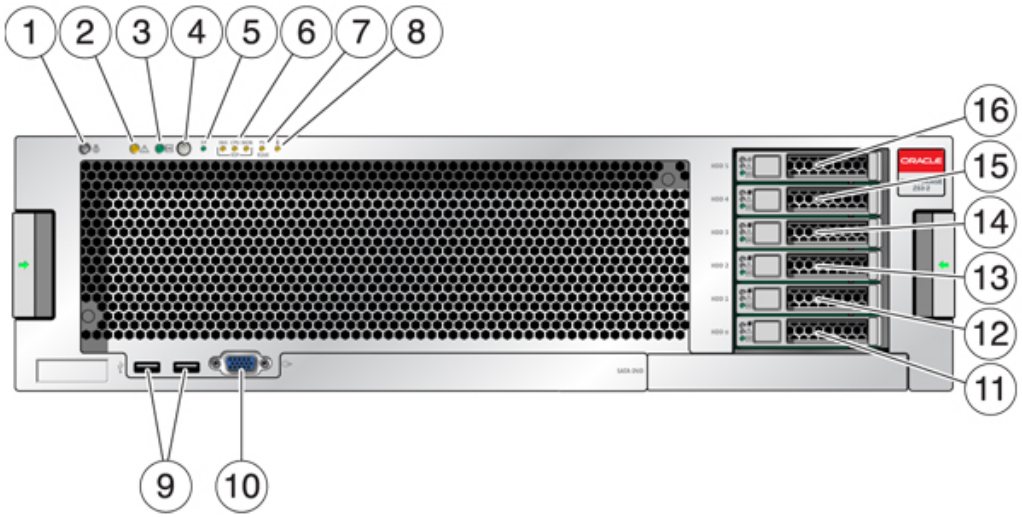

1 Locator LED and button (white)	**9** USB 2.0 Connectors
2 Service Required LED (amber)	**10** DB-15 video connector
3 Power/OK LED (green)	**11** Boot drive 0
4 Power button	**12** Boot drive 1 (required)
5 Service Processor (SP) OK LED (green)	**13** Solid state drive 2 (optional)
6 Fan/CPU/Memory Service Required LED	**14** Solid state drive 3 (optional)
7 Power Supply (PS) Service Required LED	**15** Solid state drive 4 (optional)
8 Over Temperature Warning LED	**16** Solid state drive 5 (optional)

Oracle ZFS Storage ZS3-4 Rear Panel Components - The following graphic shows the rear panel. Base configuration HBAs are not depicted in this illustration.

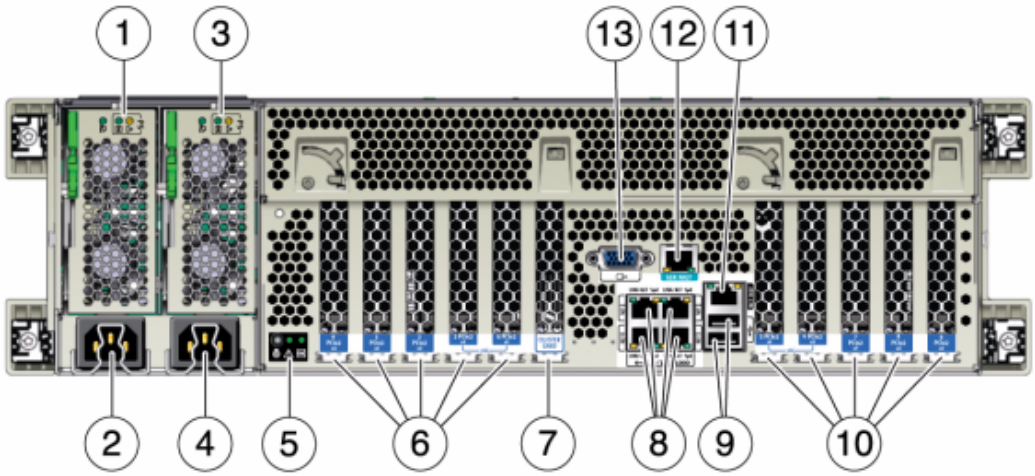

1 Power supply unit 0 status LEDs OK: green Power Supply Fail: amber AC OK: green	**8** Network (NET) 10/100/1000 ports: NET0-NET3
2 Power supply unit 0 AC inlet	**9** USB 2.0 ports
3 Power supply unit 1 status LEDs OK: green Power Supply Fail: amber AC OK: green	**10** PCIe slots 5-9
4 Power supply unit 1 AC inlet	**11** Network management (NET MGT) 10/100 BASE-T Ethernet port
5 System status LEDs Power: green Attention: amber Locate: white	**12** Serial management (SER MGT) port
6 PCIe slots 0-4	**13** DB-15 video connector
7 Cluster card slot	

Cluster Interface Card Compatibility - The Oracle ZFS Storage ZS3-4 controller contains a Version 2 cluster interface card. This card cannot be replaced with a Version 3 cluster interface card.

For information about how to connect cables to form a cluster, see Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

Ethernet Ports

There are four RJ-45 Gigabit Ethernet connectors (NET0, NET1, NET2, NET3) located on the Oracle ZFS Storage ZS3-4 motherboard that operate at 10/100/1000 Mbit/sec. These network interfaces must be configured before use.

Network Management Port

The Oracle ZFS Storage ZS3-4 network management connector (NET MGT) is an RJ-45 10/100 BASE-T Ethernet port and provides a LAN interface to the SP console.

Serial Management Port

The Oracle ZFS Storage ZS3-4 serial management connector (SER MGT) is an RJ-45 connector and provides a terminal connection to the SP console.

Oracle ZFS Storage ZS3-4 Internal Components

The components of the Oracle ZFS Storage ZS3-4 controller are shown in the following figure and identified in the legend.

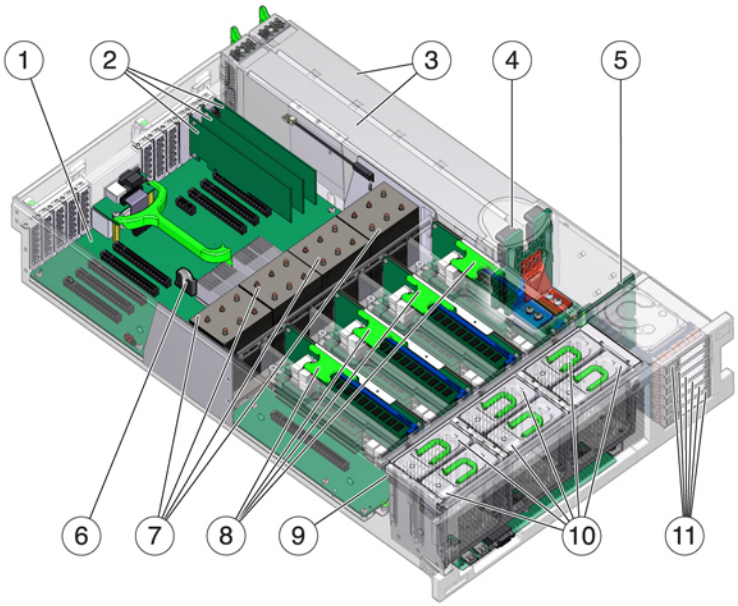

1 Motherboard	**7** CPUs and heatsinks
2 Low-profile PCIe cards	**8** Memory risers
3 Power supplies	**9** Fan board
4 Power supply backplane	**10** Fan modules
5 Drive backplane	**11** Boot drives and SSDs
6 System lithium battery	

Oracle ZFS Storage ZS3-4 Internal Boards - The Oracle ZFS Storage ZS3-4 controller chassis contains the following field-replaceable units (FRUs). FRUs are not customer-serviceable, and should only be replaced by trained Oracle service technicians.

- **Motherboard** - The motherboard includes CPU modules, slots for eight DIMM risers, memory control subsystems, and the service processor (SP) subsystem. The SP subsystem controls the host power and monitors host system events (power and environmental). The SP controller draws power from the host's 3.3V standby supply rail, which is available whenever the system is receiving AC input power, even when the system is turned off.
- **Power Distribution Board** - The power distribution board distributes main 12V power from the power supplies to the rest of the system. It is directly connected to the Vertical PDB card, and to the motherboard through a bus bar and ribbon cable. It also supports a top cover interlock ("kill") switch. In the controller, the power supplies connect to the power supply backplane which connects to the power distribution board.
- **Vertical PDB Card** - The vertical power distribution board, or Paddle Card serves as the interconnect between the power distribution board and the fan power boards, hard drive backplane, and I/O board.
- **Power Supply Backplane Card** - This board connects the power distribution board to power supplies 0 and 1.
- **Fan Power Boards** - The two fan power boards are FRUs and carry power to the controller fan modules. In addition, they contain fan module status LEDs and transfer I2C data for the fan modules.
- **Drive Backplane** - The six-drive backplane includes the connectors for the drives, as well as the interconnect for the I/O board, Power and Locator buttons, and system/component status LEDs. Each drive has an LED indicator for Power/Activity, Fault, and Locate.

Oracle ZFS Storage ZS3-4 System Drives - The Oracle ZFS Storage ZS3-4 controller has two system boot drives in slots 0 and 1, configured as a mirrored pair. Read-optimized cache devices can be installed in controller slots 2 through 5, or installed in Oracle Storage Drive Enclosure DE2-24 disk shelf slots.

> **Note:**
>
> Read cache devices must be installed in either the controller or an Oracle Storage Drive Enclosure DE2-24 disk shelf. Read cache devices cannot be installed in both controller and disk shelf slots at the same time. For software requirements and slot configuration rules, see Disk Shelf Configurations.

1 Ready to remove (blue)	**2** Service action required (amber)	**3** OK/Activity (green)

Oracle ZFS Storage ZS3-4 4x4 SAS-2 HBA - The 4x4 SAS-2 HBA, installed in the Oracle ZFS Storage ZS3-4 controller, provides connectivity to external Oracle Storage Drive Enclosure DE2 and Sun Disk Shelves. The HBA logical ports are numbered 3-0, top to bottom:

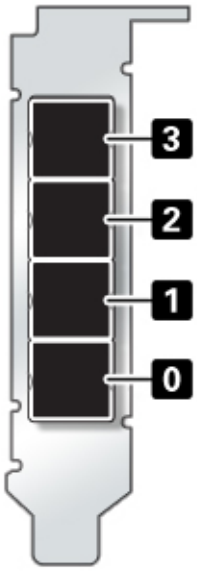

See Oracle ZFS Storage ZS3-4 PCIe I/O Cards for slot placement.

Oracle ZFS Storage ZS3-4 CPU and Memory - The Oracle ZFS Storage ZS3-4 controller supports 16GB DDR3 DIMMs installed on all eight risers, accommodating 1TB of memory.

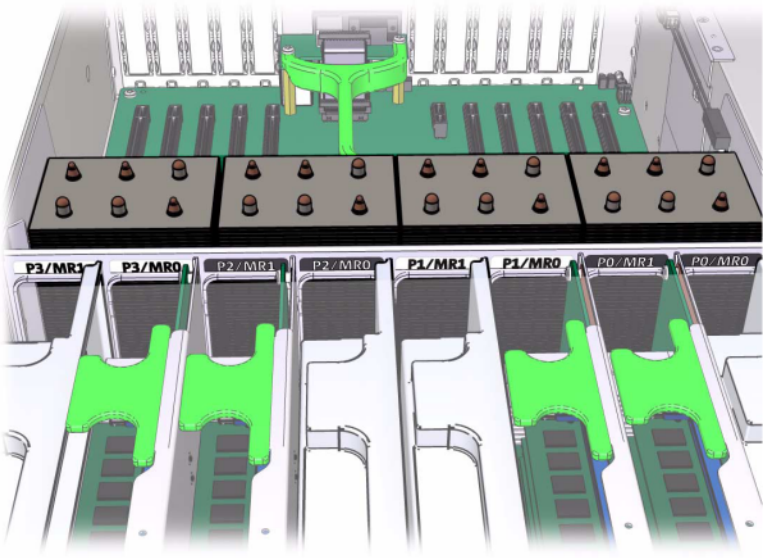

Refer to the service label on the cover for DIMM placement information. On every memory riser, slots D0, D2, D4, and D6 must be populated; optionally, slots D1, D3, D5, and D7 may be populated as a group on **all** installed memory risers. All DIMMs in the system *must be identical.*

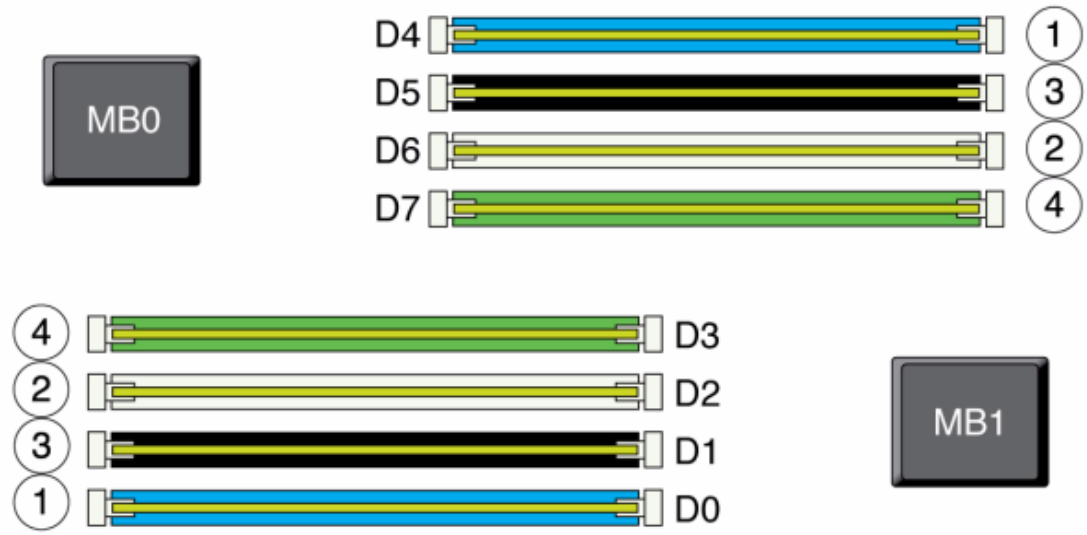

DIMM names in appliance logs and the **Maintenance: Hardware** view are displayed with the full name, such as `/SYS/MB/P0/D7`.

Oracle ZFS Storage ZS3-4 Fan Modules - The Fan Modules and Fan Module LEDs of the Oracle ZFS Storage ZS3-4 controller are shown in the following figure.

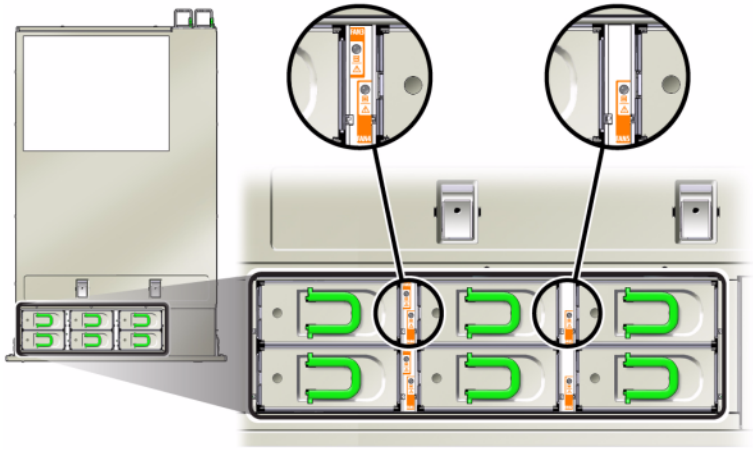

The following LEDs are lit when a fan module fault is detected:

- Front and rear Service Action Required LEDs
- Fan Module Service Action Required (TOP) LED on the front of the controller
- Fan Fault LED on or adjacent to the faulty fan module

The system Overtemp LED might light if a fan fault causes an increase in system operating temperature.

Oracle ZFS Storage ZS3-4 PCIe I/O Cards

Oracle ZFS Storage ZS3-4 PCIe Options - This table describes the PCIe base configuration and optional slot assignments for Oracle ZFS Storage ZS3-4 standalone and cluster configurations.

Table 9-3 Oracle ZFS Storage ZS3-4 Controller PCIe Options

Slot	Mfg Part Number	Description	Max	Note
0	7047852	8-port SAS-2 internal HBA	1	Base configuration
1	7067091	4-port (4x4) SAS-2 external HBA	4	Base configuration
2	7067091	4-port (4x4) SAS-2 external HBA	4	Additional optional back-end
2	7070006	2-port 10Gb Ethernet Copper NIC	4	Optional recommended front-end
2	375-3696-01	2-port InfiniBand CX2 HCA	4	Optional recommended front-end
2	7070195	4-port 1Gb Ethernet Copper UTP NIC	4	Optional recommended front-end
2	7051223	2-port 10Gb Ethernet Optical NIC	6	Optional recommended front-end
2	371-4325-02	2-port 8Gb FC HBA	4	Optional FC target or initiator (backup)
2	7023303	2-port 16Gb FC HBA	4	Optional FC target or initiator (backup)
3	7070006	2-port 10Gb Ethernet Copper NIC	4	Optional recommended front-end
3	375-3696-01	2-port InfiniBand CX2 HCA	4	Optional recommended front-end
3	7070195	4-port 1Gb Ethernet Copper UTP NIC	4	Optional recommended front-end
3	7051223	2-port 10Gb Ethernet Optical NIC	6	Optional recommended front-end
3	371-4325-02	2-port 8Gb FC HBA	4	Optional FC target or initiator (backup)

Table 9-3 (Cont.) Oracle ZFS Storage ZS3-4 Controller PCIe Options

Slot	Mfg Part Number	Description	Max	Note
3	7023303	2-port 16Gb FC HBA	4	Optional FC target or initiator (backup)
4	7070006	2-port 10Gb Ethernet Copper NIC	4	Optional recommended front-end
4	375-3696-01	2-port InfiniBand CX2 HCA	4	Optional recommended front-end
4	7070195	4-port 1Gb Ethernet Copper UTP NIC	4	Optional recommended front-end
4	7051223	2-port 10Gb Ethernet Optical NIC	6	Optional recommended front-end
4	371-4325-02	2-port 8Gb FC HBA	4	Optional FC target or initiator (backup)
4	7023303	2-port 16Gb FC HBA	4	Optional FC target or initiator (backup)
C	511-1496-05	Cluster Controller 200	1	Cluster base configuration
5	7070006	2-port 10Gb Ethernet Copper NIC	4	Optional recommended front-end
5	375-3696-01	2-port InfiniBand CX2 HCA	4	Optional recommended front-end
5	7070195	4-port 1GbE UTP Ethernet	4	Optional recommended front-end
5	7051223	2-port 10Gb Ethernet Optical NIC	6	Optional recommended front-end
5	371-4325-02	2-port 8Gb FC HBA	4	Optional FC target or initiator (backup)
5	7023303	2-port 16Gb FC HBA	4	Optional FC target or initiator (backup)
6	7070006	2-port 10Gb Ethernet Copper NIC	4	Optional recommended front-end
6	375-3696-01	2-port InfiniBand CX2 HCA	4	Optional recommended front-end

Table 9-3 (Cont.) Oracle ZFS Storage ZS3-4 Controller PCIe Options

Slot	Mfg Part Number	Description	Max	Note
6	7070195	4-port 1Gb Ethernet Copper UTP NIC	4	Optional recommended front-end
6	7051223	2-port 10Gb Ethernet Optical NIC	6	Optional recommended front-end
6	371-4325-02	2-port 8Gb FC HBA	4	Optional FC target or initiator (backup)
6	7023303	2-port 16Gb FC HBA	4	Optional FC target or initiator (backup)
7	7067091	4-port (4x4) external HBA	4	Additional optional back-end
7	7070006	2-port 10Gb Ethernet Copper NIC	4	Optional recommended front-end
7	371-4325-02	2-port 8Gb FC HBA	4	Optional recommended front-end
7	375-3696-01	2-port InfiniBand CX2 HCA	4	Optional recommended front-end
7	7051223	2-port 10Gb Ethernet Optical NIC	6	Optional recommended front-end
7	7070195	4-port 1Gb Ethernet Copper UTP NIC	4	Optional recommended front-end
7	371-4325-02	2-port 8Gb FC HBA	4	Optional FC target or initiator (backup)
7	7023303	2-port 16Gb FC HBA	4	Optional FC target or initiator (backup)
8	7067091	4-port (4x4) SAS-2 external HBA	4	Base configuration
9	371-4325-02	2-port 8Gb FC HBA	4	Optional FC target or initiator (backup)
9	7070006	2-port 10Gb Ethernet Copper NIC	4	Optional recommended front-end
9	7070195	4-port 1Gb Ethernet Copper UTP NIC	4	Optional recommended front-end

Table 9-3 (Cont.) Oracle ZFS Storage ZS3-4 Controller PCIe Options

Slot	Mfg Part Number	Description	Max	Note
9	7051223	2-port 10Gb Ethernet Optical NIC	6	Optional recommended front-end
9	7023303	2-port 16Gb FC HBA	4	Optional FC target or initiator (backup)

Oracle ZFS Storage ZS3-4 NIC/HBA Options - This table describes NIC/HBA PCIe card options for Oracle ZFS Storage ZS3-4 standalone and cluster configurations.

Table 9-4 Oracle ZFS Storage ZS3-4 Controller NIC/HBA Options

Mktg Part Number	Description
SG-SAS6-INT-Z	8-port 6Gb/s SAS-2 internal HBA
SG-XPCIE2FC-QF8-Z	2-port 8Gb FC HBA
7103791	4-port (4x4) SAS-2 6Gb/s external HBA
7101674	2-port 16Gb FC HBA
7100477	4-port 1Gb Ethernet Copper UTP
7100488	2-port 10Gb Ethernet Copper Base-T
X4242A	2-port InfiniBand CX2 HCA
X1109A-Z	2-port 10GbE SFP+ NIC
X2129A	Transceiver 850NM, 1/10GPS, Short Reach, SFP
X5562A-Z	Transceiver 10GbE/1GbE, Long Reach, SFP

Oracle ZFS Storage ZS3-4 PCIe Slot Order - When installing additional PCIe cards to your base configuration, cards must be added in a specific order.

> **Note:**
>
> Slot 0 is always populated with a SAS-2 internal HBA. Slots 1 and 8 are reserved for 4x4 SAS-2 external HBAs.

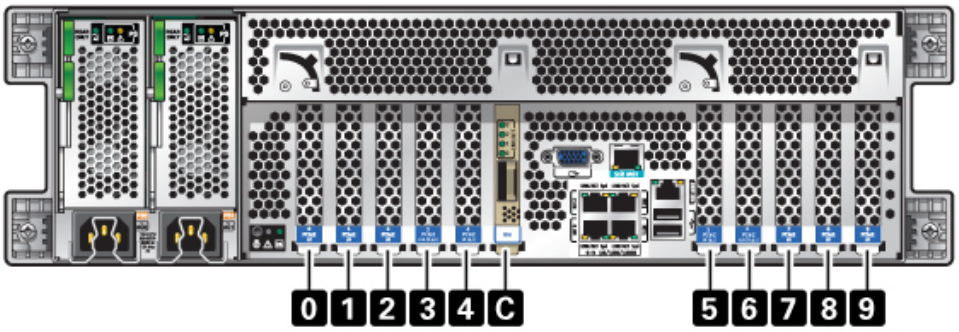

Install optional PCIe cards in the following order:

Table 9-5 Oracle ZFS Storage ZS3-4 Controller Optional PCIe Cards and Slot Order

PCIe Card	Slot Order
1 4-port (4x4) SAS-2 external HBA	Slot 2, 7
2 8Gb FC HBA	Slot 9
.	Additional FC HBAs are installed last
3 InfiniBand HCAs	Slot 3, 6, 4, 5, 2, 7
4 2-port 10Gb Ethernet NIC	Slot 3, 6, 4, 5, 2, 7, 9
.	Note: If you are adding a 2-port 10Gb Ethernet Optical NIC and 2-port 10Gb Ethernet Copper NIC, the 10Gb Ethernet Optical NIC has slot priority.
5 4-port 1Gb Ethernet NIC	Slot 3, 6, 4, 5, 2, 7, 9
6 16Gb FC HBA	Slot 3, 6, 4, 5, 2, 7, 9
7 Remaining 8Gb FC HBAs	Slot 3, 6, 4, 5, 2, 7 (slot 9, see step 2)

Oracle ZFS Storage ZS3-4 Attached Storage

The Oracle ZFS Storage ZS3-4 standalone and cluster controller configurations allow up to 36 disk shelves, consisting of up to six chains of one to six disk shelves. Any combination of disk-only, or disk, log, and read cache devices can be combined within the chain in any order. You can also attach mixed disk shelf types (Oracle Storage Drive Enclosure DE2 family and Sun Disk Shelves) behind the same controllers, but each chain must contain only the same disk shelf type. Directly connecting different disk shelf types is not supported.

> **Note:**
>
> Read cache devices must be installed in either the Oracle ZFS Storage ZS3-4 controller or an Oracle Storage Drive Enclosure DE2-24 disk shelf. Read cache devices cannot be installed in both controller and Oracle Storage Drive Enclosure DE2 disk shelf slots at the same time. For software requirements and slot configuration rules, see Disk Shelf Configurations.

Related Topics

- Connecting to Attached Storage
- Disk Shelf Hardware Overview

10

Servicing the Oracle ZFS Storage ZS3-2 Controller

Be sure to read sections Getting Started with Servicing the Appliance and Prerequisites for Servicing Hardware.

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see Repairing Active Problems.

Oracle Storage Drive Enclosure DE2-24 disk shelves and Sun Disk Shelves can be used together with standalone and clustered controllers after certain hardware modifications. Contact Oracle Service to upgrade your Oracle ZFS Storage ZS3-2 controllers for disk shelf interoperability.

To service Oracle ZFS Storage ZS3-2 hardware components, use the following tasks.

- Installing Optional PCIe Slot Cards in Order
- Disconnecting the RJ-45 Cable from the Cluster Serial Port
- Replacing an Oracle ZFS Storage ZS3-2 HDD or SSD
- Replacing an Oracle ZFS Storage ZS3-2 Power Supply
- Replacing an Oracle ZFS Storage ZS3-2 Fan Module
- Replacing an Oracle ZFS Storage ZS3-2 DIMM
- Replacing an Oracle ZFS Storage ZS3-2 PCIe Card or Riser
- Replacing an Oracle ZFS Storage ZS3-2 Flash Drive
- Replacing an Oracle ZFS Storage ZS3-2 Battery

To understand the Oracle ZFS Storage ZS3-2 hardware, use the following topics.

- Oracle ZFS Storage ZS3-2 Base Configuration
- Oracle ZFS Storage ZS3-2 Front and Rear Panel Components
- Oracle ZFS Storage ZS3-2 Specifications
- Oracle ZFS Storage ZS3-2 Internal Components
- Oracle ZFS Storage ZS3-2 PCIe I/O Cards
- Oracle ZFS Storage ZS3-2 Optional Cable Management Arm
- Oracle ZFS Storage ZS3-2 Attached Storage

Installing Optional PCIe Slot Cards in Order

Install Oracle ZFS Storage ZS3-2 optional PCIe cards in the following order:

1. Install optional second 4x4 SAS-2 HBA (7067091) in slot 3.

> **Note:**
>
> For the remainder of PCIe option cards, always start installing in slot 6, and then work down to slot 3 in this order:

2. Install InfiniBand CX2 HCA (375-3696-01) in slot 6 and continue until all InfiniBand options are installed and or,
3. Install 10Gb Ethernet Optical NICs (7051223) and or,
4. Install 10Gb Ethernet Copper NICs (7070006) and or,
5. Install 8Gb FC HBAs (371-4325-02) or 16Gb FCs HBA (7023303) and or,
6. Install 1Gb Ethernet Copper NICs (7070195).

Disconnecting the RJ-45 Cable from the Cluster Serial Port

1. Using your index finger, press down fully on the RJ-45 release tab. Be sure the tab is fully disengaged from the port.

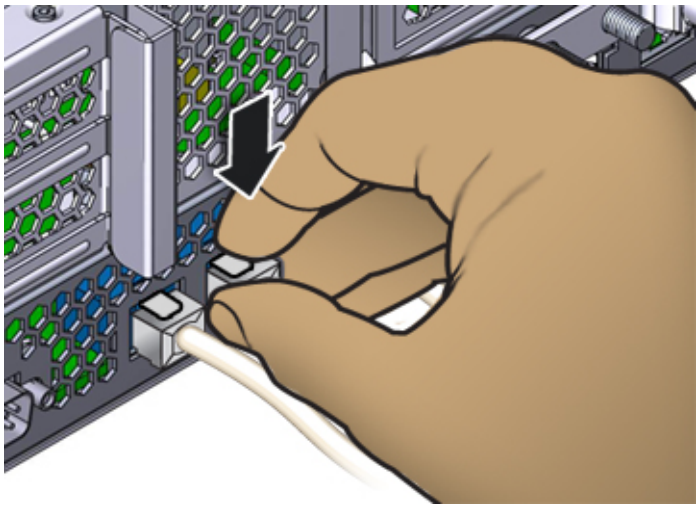

2. Using your thumb and middle finger, apply a slight downward pressure while pulling the plug out of the port. Do not pull the plug upward or pinch the release tab with your fingers below the plug, as shown in the following figure.

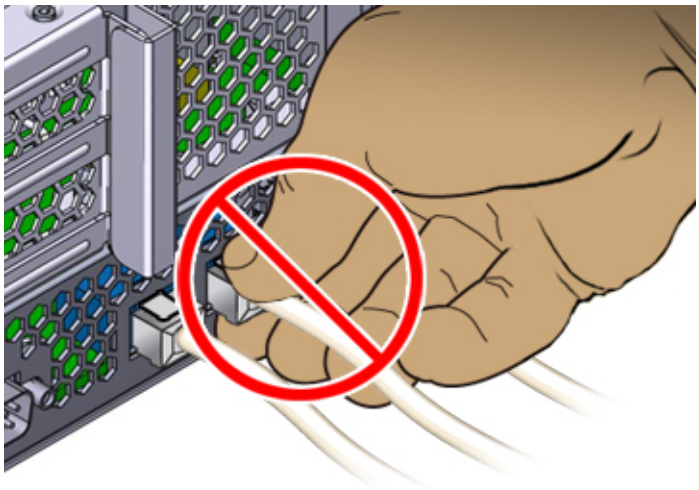

Replacing an Oracle ZFS Storage ZS3-2 HDD or SSD

HDDs and SSDs are hot-swappable and can be removed and installed while the storage controller is running without affecting other hardware capabilities. Use the following procedure to replace an Oracle ZFS Storage ZS3-2 HDD or SSD.

> **Note:**
>
> If you have more than one failed drive, replace only one drive at a time. Removing more than one drive in quick succession causes a hardware/pool fault. Also, do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** BUI menu, select **System**, or navigate to `maintenance system updates` in the CLI.

1. Identify the faulted HDD or SSD.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, and open the detailed view of the controller. The faulted drive has either an amber or a blue-on-amber (ready to remove) status icon. Note the ID of the faulted drive. Click the entry for that drive in the table to highlight the location of the drive in the controller illustration. To view details about the faulted drive, click its information icon and then click **Active Problems**.
   - If you are physically at the system, the amber Service Action Required indicator on the HDD or SSD should be lit. Additionally, you can click the locator icon in the BUI to light the locator LED on the faulted drive.
2. Verify that the faulted drive is ready to be removed.

> **Caution:**
>
> Removing the drive before it is ready to be removed can result in loss of data.

Use any of the following methods to verify that the faulted drive is ready to be removed:

- In the BUI:
  - The status icon for the faulted drive in the detail view in **Maintenance: Hardware** has changed from amber to blue-on-amber.
  - The **Component Details** information window (click the information icon) for the faulted drive shows **Yes** as the value for the **Ready to remove** property.
- In the CLI (and RESTful API), the value of the `readytoremove` property for the faulted drive is `true`:

```
hostname:maintenance chassis-000 disk-001> get readytoremove
                readytoremove = true
```

3. At the front of the controller, pinch the air filter's left and right release levers inward [1].
4. Pivot the air filter forward and lift if out of the chassis [2].

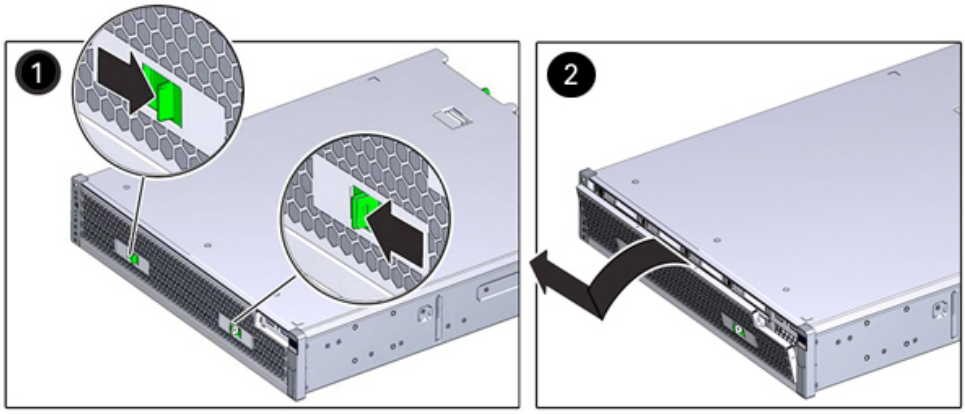

5. Press the release button of the failed drive, and pivot the release lever open to the right.

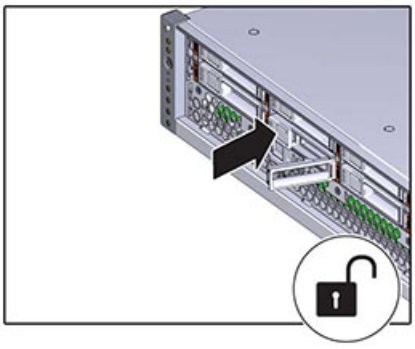

6. Grasp the release lever and pull the drive out of the drive slot.

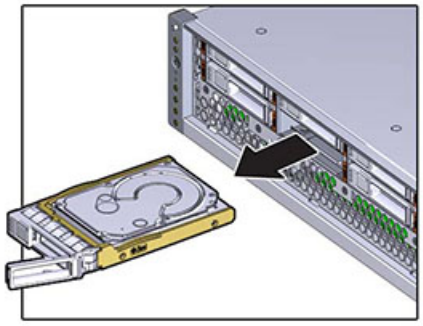

7. After at least 30 seconds, from the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon on the system controller to verify that the software has detected that the drive is not present.
8. Align the replacement drive with the drive slot.
9. Press on the release button to slide the drive into the chassis until the release lever moves slightly inward.

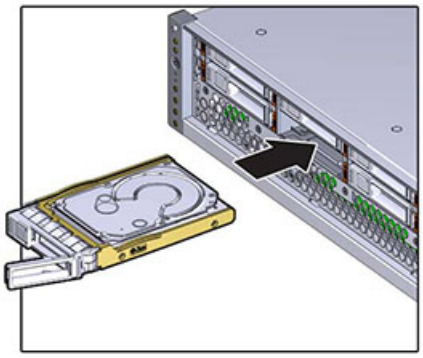

10. Press the release lever fully closed with a click, securing the drive into the drive slot.

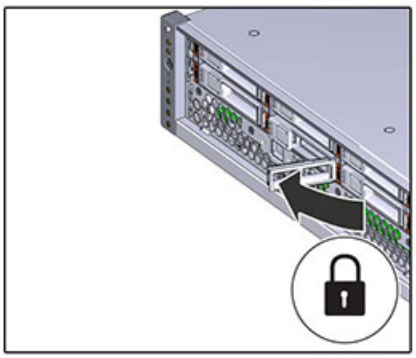

11. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, and then click **Disk** to verify that the online icon is green for the newly installed drive.
12. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS3-2 Power Supply

Oracle ZFS Storage ZS3-2 controllers are equipped with redundant hot-swappable power supplies. If a power supply fails and you do not have a replacement, leave the failed power supply installed to ensure proper air flow. A faulted power supply is indicated by an amber colored status LED. Use the following procedure to replace an Oracle ZFS Storage ZS3-2 power supply.

1. Gain access to the rear of the storage controller where the faulted power supply is located.
2. If a cable management arm (CMA) is installed, press and hold the CMA release tab and rotate the arm out of the way.

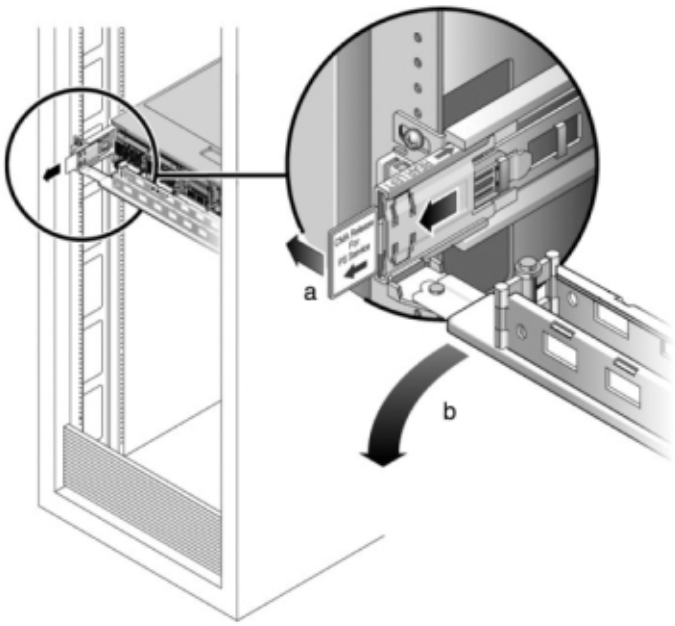

3. Disconnect the power cord from the faulted power supply.
4. Press the release tab to the left and pull on the handle.
5. Continue to pull on the handle to slide the power supply out of the chassis.

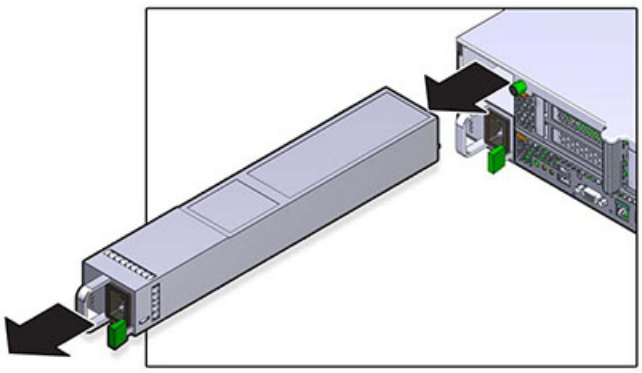

6. Align the replacement power supply with the empty power supply chassis bay.
7. Slide the power supply into the bay until the power supply seats and the release tab clicks.

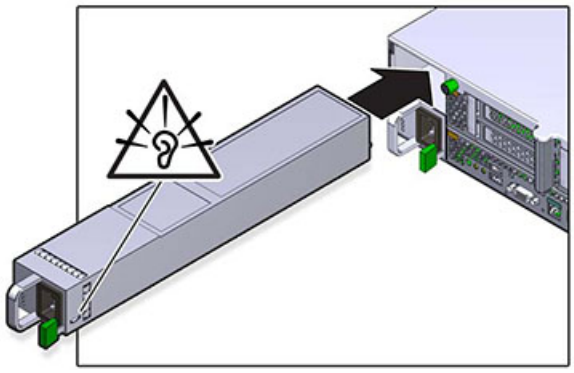

8. Push the release tabs on the side of each rail, and slowly push the storage controller into the rack, making sure no cables obstruct the path of the controller.
9. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
10. Connect the power cords to the power supplies. Verify that standby power is on, indicated by the Power/OK and SP status indicators flashing green on the front panel.
11. After the SP indicator is solid, use a pen or other pointed object to press and release the recessed Power button on the front panel. The Power/OK status indicator next to the Power button lights and remains lit.
12. Close the CMA, inserting the CMA into the rear left rail bracket.
13. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, and then click **PSU** to verify that the online icon is green for the newly installed power supply.
14. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS3-2 Fan Module

Removing a fan module is a cold-service operation. You *must* power off the appliance before you remove the fan module. There will be a loss of access to the storage unless the system is in a clustered configuration.

The status of each fan module is represented by a single bi-color LED. The LEDs are located on the motherboard near each fan module, and are visible from the rear panel, by looking through the grille below the risers. Use the following procedure to replace an Oracle ZFS Storage ZS3-2 fan module.

1. To flash the locator LED on the chassis to service, from the **Maintenance** BUI menu, select **Hardware**, and click the associated locator icon , or go to the service processor (SP) prompt and enter `set /SYS/LOCATE status=on`.

2. To locate the faulted component, from the **Maintenance** BUI menu, select **Hardware**, then select **Fan**, and click the row with an amber indicator.

   The component is highlighted in the diagram.

3. Power off the storage controller as described in Powering Off the Controller.
4. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

5. Extend the controller from the rack as described in Extending the Controller from the Rack.
6. Remove the top cover as described in Removing the Controller Top Cover.
7. Grasp the green band of the fan module to be removed with your forefinger and press down on the fan module with your thumb.
8. Pull the band straight up and lift the fan module out of the chassis.

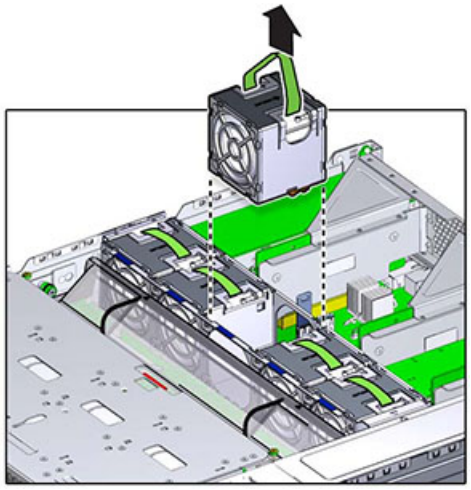

9. Align the replacement fan module into the fan tray slot. The green band is on the top of the fan module and the arrow points to the rear of the chassis.

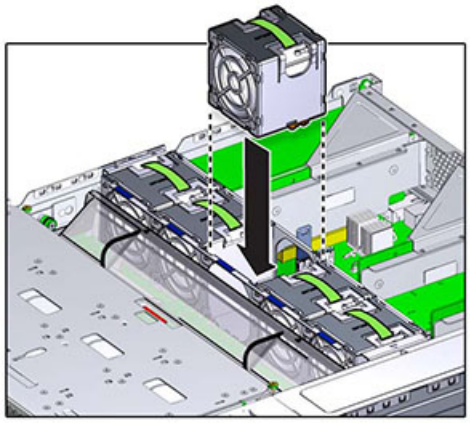

10. Lower the fan module into the fan tray and press down until the module clicks securely in the fan tray.

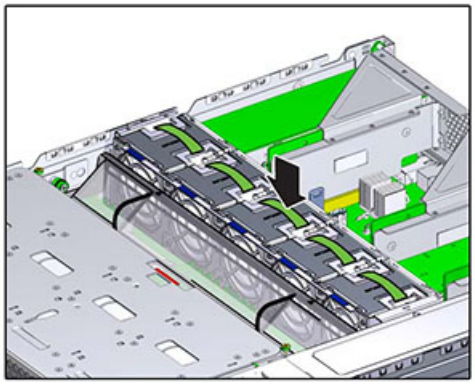

11. Replace the top cover and tighten the two captive screws.
12. Push the release tabs on the side of each rail and slowly push the storage controller into the rack, making sure no cables obstruct the path of the controller.
13. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
14. Connect the power cords to the power supplies. Verify that standby power is on, indicated by the Power/OK and SP status indicators flashing green on the front panel.
15. After the SP status indicator is solid green, use a pen or other pointed object to press and release the recessed Power button on the front panel. The Power/OK status indicator next to the Power button lights and remains lit.
16. Verify that the Fan OK status indicator is lit, and that the fault status indicator on the replaced fan module is dim.
17. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, and then click **Fan** to verify that the online icon is green for the newly installed fan module.
18. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS3-2 DIMM

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

Removing a DIMM is a cold-service operation. You *must* power off the appliance before you remove the DIMM. There will be a loss of access to the storage unless the system is in a clustered configuration.

1. To identify a general memory fault, from the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon for the controller. Then click **DIMM** to locate the faulted component, indicated by the warning icon . To identify a specific memory module that has faulted, you can open the storage controller and use the amber status LEDs on the motherboard.
2. Power off the storage controller as described in Powering Off the Controller.
3. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

4. Extend the controller from the rack as described in Extending the Controller from the Rack.
5. Remove the top cover as described in Removing the Controller Top Cover.
6. Disconnect the cables attached to the drive backplane, and the cable attached to the USB board.
7. Fully loosen the four screws securing the drive cage to the chassis [1].
8. Raise the drive cage to the vertical position [2].

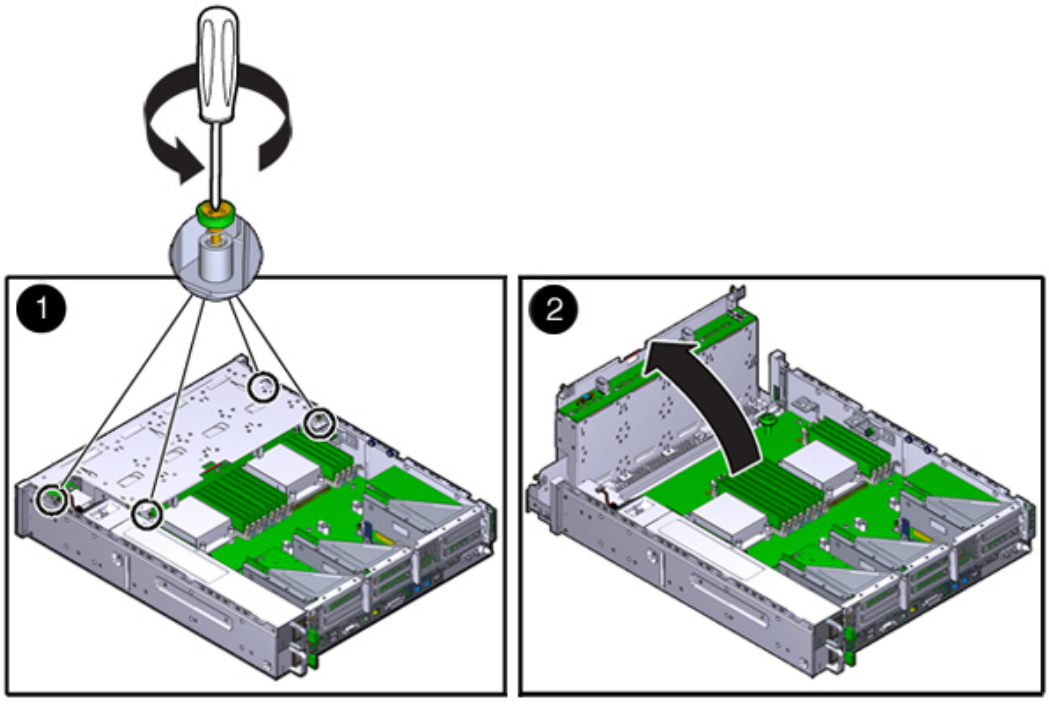

9. Fully loosen the thumbscrews that secure the air duct.

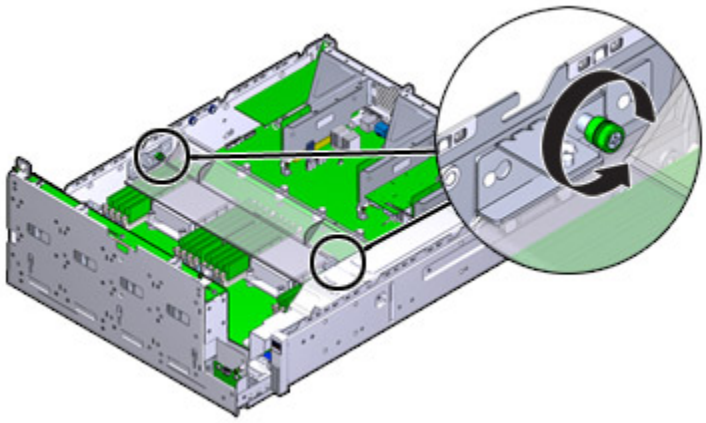

10. Slide the left (L) thumbscrew bracket forward and off of the mounting pins.
11. Slide the right (R) thumbscrew bracket rearward and off of the mounting pins. Set the thumbscrew bracket aside.
12. Lift the air duct straight up and off the alignment pins. Set the air duct aside.

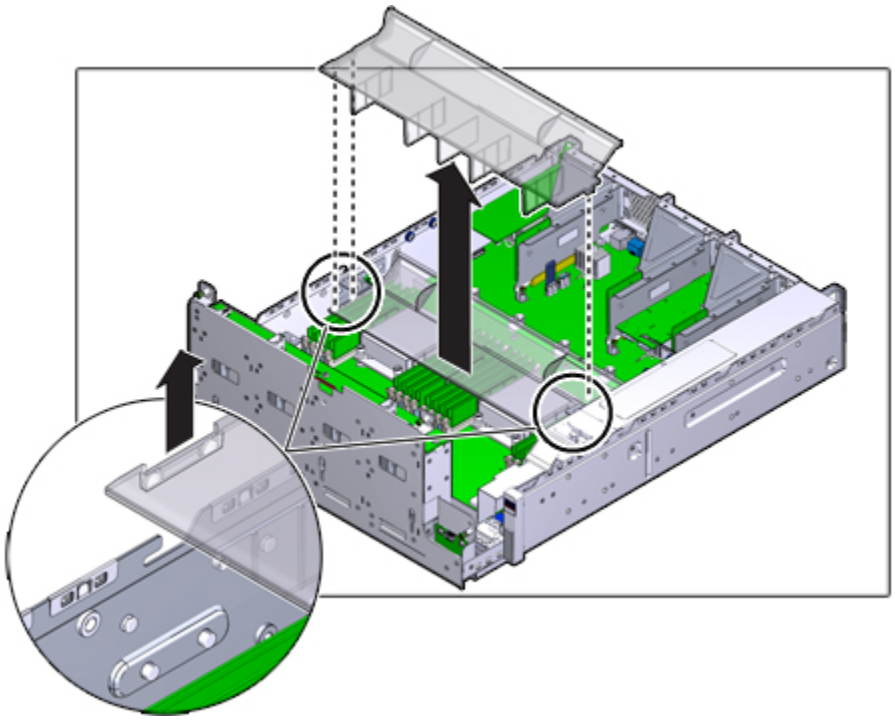

13. Press the button on the midplane. An amber LED indicates the faulty DIMM.
14. Press down and out on the release levers at both ends of the DIMM slot.
15. Lift the DIMM up and out of the slot. Set the DIMM aside on an antistatic mat.

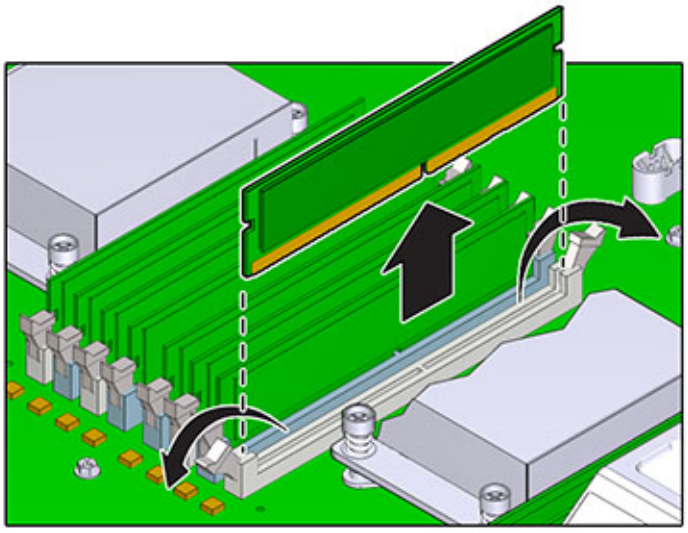

16. Align the DIMM to the location where it installs into the slot. Ensure that the notch in the DIMM lines up with the key in the slot.

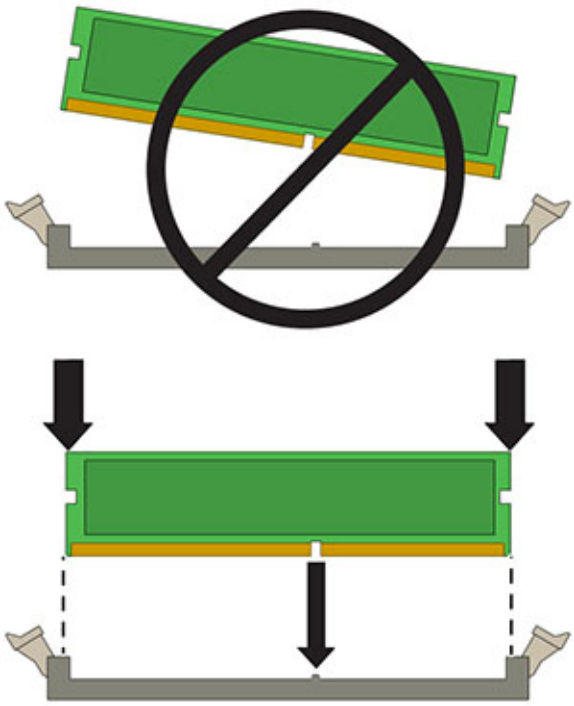

17. Insert the DIMM into the slot, pressing firmly so that both release levers click closed.

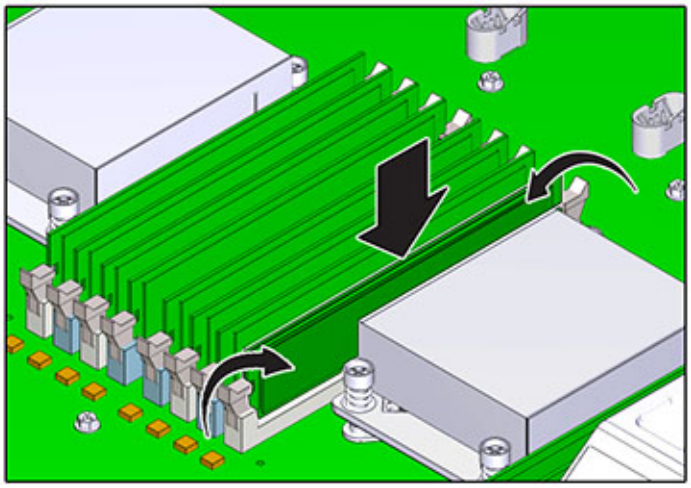

18. Replace and secure the air duct.
19. Lower the drive cage to the horizontal position and with the cage slightly elevated, reconnect the power cable. Tighten the four screws. Reconnect the cables to the drive backplane and the cable attached to the USB board.
20. Replace the top cover and tighten the two captive screws.
21. Push the release tabs on the side of each rail and slowly push the storage controller into the rack, making sure no cables obstruct the path of the controller.

22. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
23. Connect the power cords to the power supplies. Verify that standby power is on, indicated by the Power/OK and SP status indicators flashing green on the front panel.
24. After the SP status indicator is solid green, use a pen or other pointed object to press and release the recessed Power button on the front panel. The Power/OK status indicator next to the Power button lights and remains lit.
25. From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, and then click **DIMM** to verify that the online icon is green for the newly installed DIMM.
26. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS3-2 PCIe Card or Riser

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

Removing a PCIe card is a cold-service operation. You *must* power off the appliance before you remove the card. There will be a loss of access to the storage unless the system is in a clustered configuration.

1. To locate a faulted component, from the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller, and then click **Slot**.
2. Power off the storage controller as described in Powering Off the Controller.
3. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

4. Extend the controller from the rack as described in Extending the Controller from the Rack.
5. Remove the top cover as described in Removing the Controller Top Cover.
6. Locate the PCIe card position in the storage controller, see Oracle ZFS Storage ZS3-2 PCIe I/O Cards.
7. Fully loosen the two captive screws securing the riser to the motherboard.

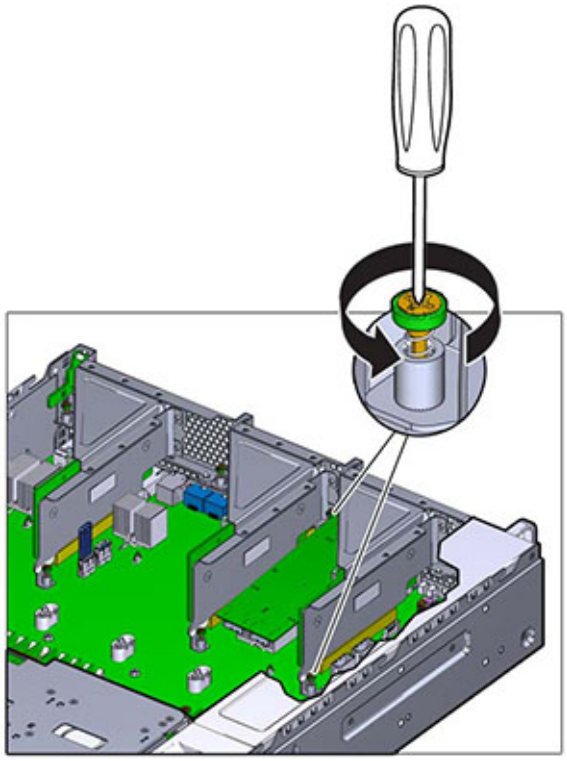

8. Lift the riser straight up and out of its socket on the motherboard.

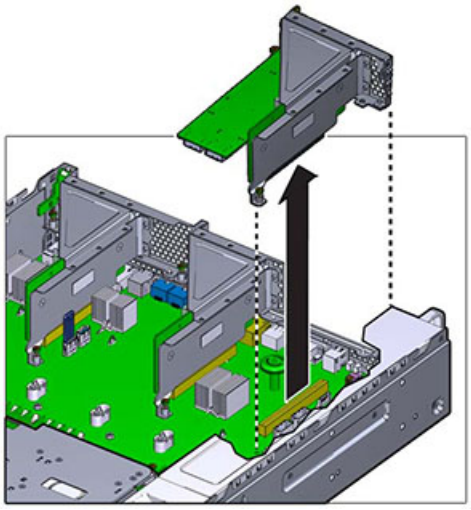

9. Disconnect any internal cables attached to the PCIe cards installed in the riser.
10. Swing the retainer to its fully open (120 degrees) position.

    Note that the riser 3 latch is slightly different than the latch for riser 1 and 2.

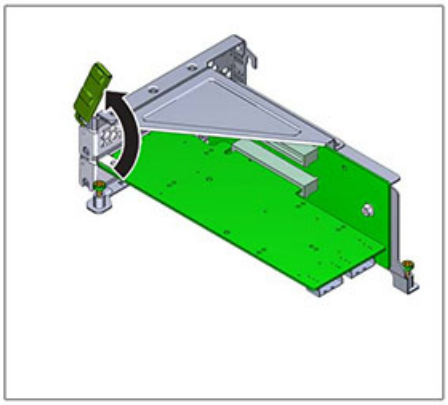

11. Lift the PCIe card off of the bracket alignment pin and out of the socket.

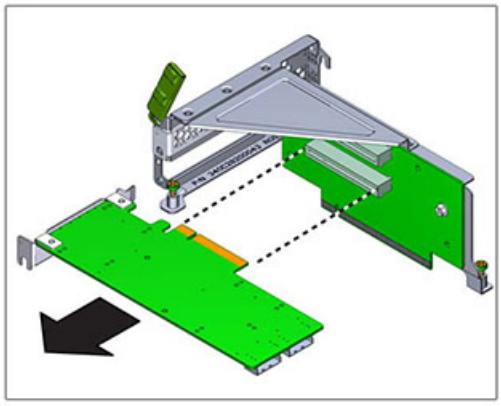

12. Position the replacement PCIe card with where it will install into the riser. The PCIe card bracket mates with the alignment pin and the socket of the riser.

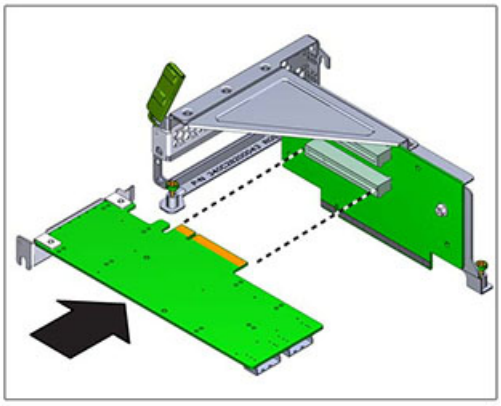

13. Press the PCIe card into the socket.
14. Swing the retainer to its fully closed position. If there is resistance, check the alignment of the PCIe card bracket and try again.

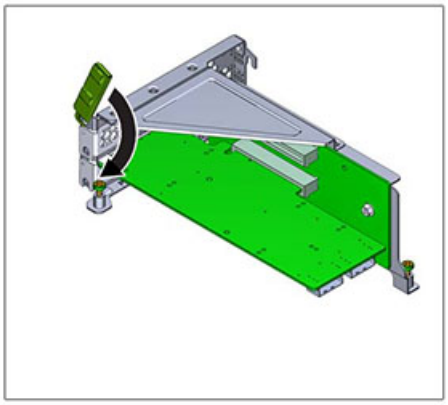

15. Align the riser to the location where it installs into the chassis.
16. Reconnect any previously removed internal cables to the PCIe card installed in the riser, if doing so will ease installation of the riser.
17. Lower the riser card onto the motherboard and press the card edge connector securely into the socket. The bracket of the riser (riser 2 and 3) overlaps the bracket of the adjacent riser.

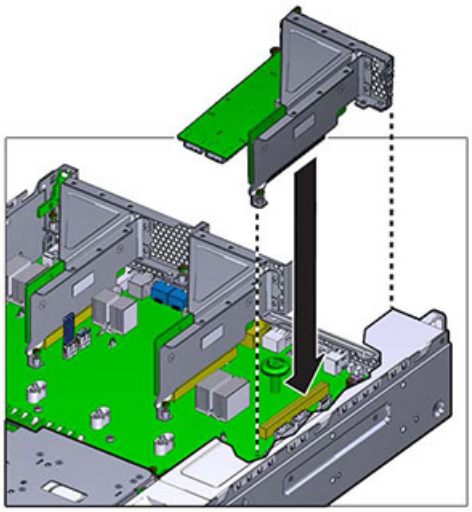

18. Tighten the two captive screws that secure the riser to the motherboard.
19. Replace the top cover and tighten the two captive screws.
20. Push the release tabs on the side of each rail and slowly push the storage controller into the rack, making sure no cables obstruct the path of the controller.
21. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
22. Connect the power cords to the power supplies. Verify that standby power is on, indicated by the Power/OK and SP status indicators flashing green on the front panel.
23. After the SP status indicator is solid green, use a pen or other pointed object to press and release the recessed Power button on the front panel. The Power/OK status indicator next to the Power button lights and remains lit.

24. From the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon for the controller. Then click **Slot** to verify the status of the new component. The online icon should be green.
25. If a fault does not self-clear, see Repairing Active Problems.

Replacing an Oracle ZFS Storage ZS3-2 Flash Drive

Removing the USB flash drive is a cold-service operation. You *must* power off the appliance before you remove the USB. There will be a loss of access to the storage unless the system is in a clustered configuration.

The USB flash drive is located in the forward-most USB socket at the center rear of the motherboard, labeled USB 1, near where Riser 2 installs. Use the following procedure to replace an Oracle ZFS Storage ZS3-2 flash drive.

1. Power off the storage controller as described in Powering Off the Controller.
2. Extend the controller from the rack as described in Extending the Controller from the Rack.
3. Remove the top cover as described in Removing the Controller Top Cover.
4. If a PCIe card installed in Riser 2 prevents you from accessing the USB drive, remove Riser 2.
5. Grasp the USB drive and pull it straight up and out of the chassis.

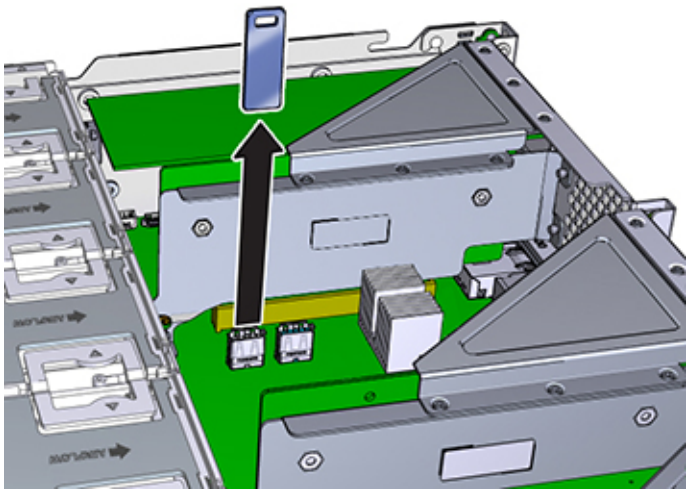

6. Position the replacement USB drive to where it installs onto the motherboard. The USB drive installs into the USB socket that is most forward. The socket is labeled USB 1.
7. Gently insert the USB drive into the USB socket.
8. If you removed Riser 2 to access the USB drive, reinstall Riser 2.
9. Replace the top cover and tighten the two captive screws.
10. Push the release tabs on the side of each rail, and slowly push the storage controller into the rack, making sure no cables obstruct the path of the controller.

11. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
12. Connect the power cords to the power supplies. Verify that standby power is on, indicated by the Power/OK and SP status indicators flashing green on the front panel.
13. After the SP status indicator is solid green, use a pen or other pointed object to press and release the recessed Power button on the front panel. The Power/OK status indicator next to the Power button lights and remains lit.

Replacing an Oracle ZFS Storage ZS3-2 Battery

Removing the battery is a cold-service operation. You *must* power off the appliance before you remove the battery. There will be a loss of access to the storage unless the system is in a clustered configuration.

You might need to replace the battery if the storage controller fails to maintain the proper time when powered off and disconnected from the network. You will need a small (No. 1 flat-blade) non-metallic screwdriver or equivalent. Use the following procedure to replace an Oracle ZFS Storage ZS3-2 battery.

1. Power off the storage controller as described in Powering Off the Controller.
2. Extend the controller from the rack as described in Extending the Controller from the Rack.
3. Remove the top cover as described in Removing the Controller Top Cover.
4. Remove Riser 1.
5. Press the tab of the battery receptacle forward to release the battery. The battery pops up.

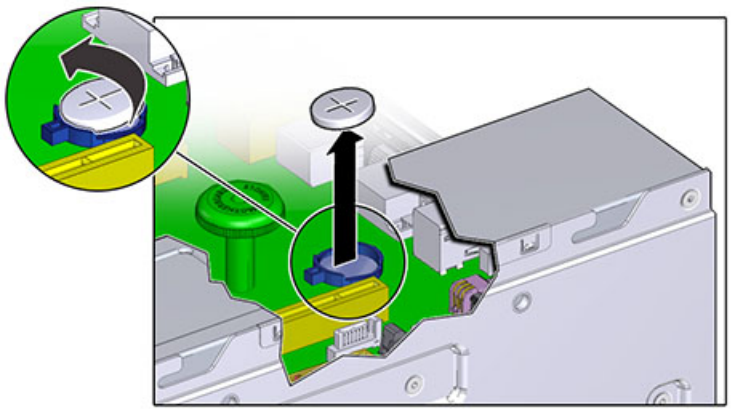

6. Lift the battery out of the receptacle and set it aside.
7. Position the battery where it will install into the receptacle, with the positive (+) side of the battery facing up.
8. Press the battery into the receptacle until you hear or feel a click.

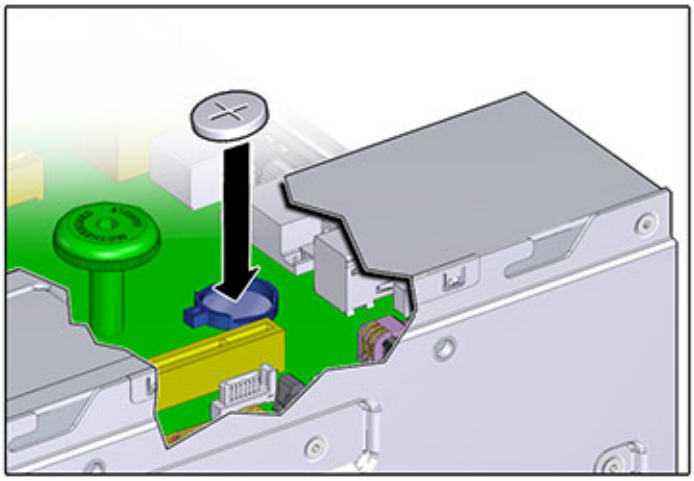

9. Install Riser 1.
10. Replace the top cover and tighten the two captive screws.
11. Push the release tabs on the side of each rail, and slowly push the storage controller into the rack, making sure no cables obstruct the path of the controller.
12. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
13. Connect the power cords to the power supplies. Verify that standby power is on, indicated by the Power/OK and SP status indicators flashing green on the front panel.
14. After the SP status indicator is solid green, use a pen or other pointed object to press and release the recessed Power button on the front panel. The Power/OK status indicator next to the Power button lights and remains lit.
15. When the system has finished booting, log in, and set the time using the steps in the task Setting Clock Synchronization (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Oracle ZFS Storage ZS3-2 Controller Hardware Overview

To understand the Oracle ZFS Storage ZS3-2 hardware, use the following topics.

- Oracle ZFS Storage ZS3-2 Base Configuration
- Oracle ZFS Storage ZS3-2 Front and Rear Panel Components
- Oracle ZFS Storage ZS3-2 Specifications
- Oracle ZFS Storage ZS3-2 Internal Components
- Oracle ZFS Storage ZS3-2 PCIe I/O Cards
- Oracle ZFS Storage ZS3-2 Optional Cable Management Arm
- Oracle ZFS Storage ZS3-2 Attached Storage

Oracle ZFS Storage ZS3-2 Base Configuration

The Oracle ZFS Storage ZS3-2 controller is an enterprise-class, rackmount x86 system powered by the Intel Xeon processor. It provides high performance and room for growth with expandable PCIe slots and 16 DIMM slots in a compact 2U footprint.

Table 10-1 Oracle ZFS Storage ZS3-2 Controller Features

CPU	Memory	Read Cache	Boot Drive	PCIe Slots	HBA
2x8-core, 2.1GHz	256GB 16x16GB 512GB 16x32GB	1-4 SAS-2 SSDs	2x2.5-inch SAS-2	2 dedicated; 4 available	4-port (4x4) SAS-2 6Gb/s external

Oracle ZFS Storage ZS3-2 Front and Rear Panel Components

Oracle ZFS Storage ZS3-2 Front Panel Components - The Oracle ZFS Storage ZS3-2 controller drive slots and front panel components are shown in the following figure.

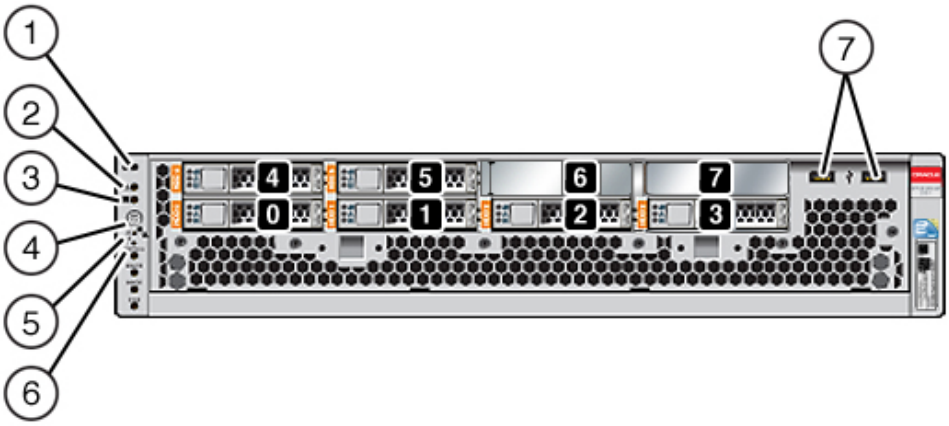

Drive slots 0 and 1 have two mirrored boot drives. Up to four read-optimized cache solid state drives (SSDs) fill slots 2 through 5, in order. Slots 6 and 7 are empty and must contain drive fillers.

1 Locator LED/button (white)	**5** Service Processor OK LED (green)
2 Service Action Required LED (amber)	**6** Fan/CPU/Memory Service Required LED
3 Power OK LED (green)	**7** USB 2.0 ports
4 Power button	

> **Note:**
>
> The LEDs below the Fan/CPU/Memory Service Required LED are not currently used.

Oracle ZFS Storage ZS3-2 Rear Panel Components - The Oracle ZFS Storage ZS3-2 controller PCIe slots and rear panel components are shown in the following figure.

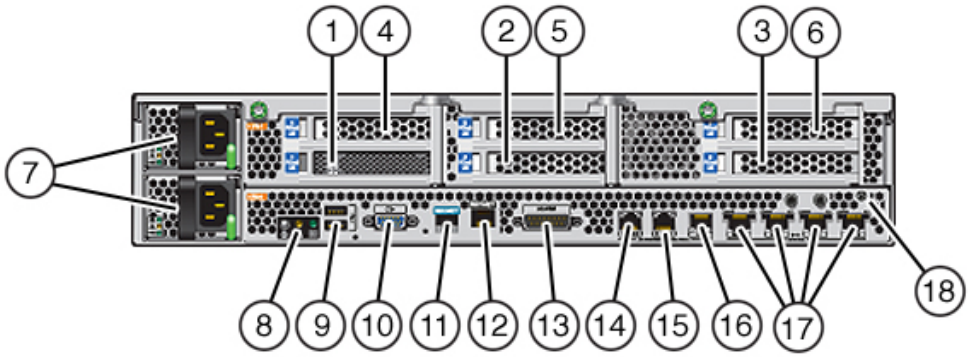

1 8-port SAS-2 internal HBA (slot 1)	**7** AC power supplies PS1 (top), PS0 (bottom)	**13** Alarm port, DB-15 connector (not used for the ZS3-2)
2 4x4 SAS HBA (slot 2)	**8** System status LEDs	**14-16** Cluster I/O ports
3 PCIe slot 3	**9** USB 2.0 ports	**17** 10-Gbit Ethernet ports
4 PCIe slot 4	**10** SP 15-pin VGA video port	**18** Chassis ground post
5 PCIe slot 5	**11** Serial management port	
6 PCIe slot 6	**12** Network management (NET MGT) 10/100 BASE-T Ethernet port	

> **Note:**
>
> The three Cluster I/O ports (0, 1, and GigE) are reserved for cluster interconnection only.

Cluster Interface Compatibility - The Oracle ZFS Storage ZS3-2 controller contains embedded support for Version 2 clustering ports. It does not have a cluster interface card nor can one be installed in a PCIe slot.

Ethernet Ports

The Oracle ZFS Storage ZS3-2 controller has four RJ-45 10-Gigabit Ethernet (10GbE) network connectors, labeled NET 3, NET 2, NET 1, and NET 0 (left to right) on the

rear panel. The ports operate at 100 Mbits/sec, 1000 Mbits/sec, or 10-Gbits/sec. Use these ports to connect the appliance to the network.

Table 10-2 Oracle ZFS Storage ZS3-2 Ethernet Port LEDs Description

Connection Type	EEE Terminology	Speed LED Color	Transfer Rate
Fast Ethernet	100BASE-TX	Off	100 Mbits/sec
Gigabit Ethernet	1000BASE-T	Amber	1000 Mbits/sec
10 Gigabit Ethernet	10GBASE-T	Green	10000 Mbits/sec

The LEDs located above each NET port are Link/Activity (left) and Speed (right) indicators for each port.

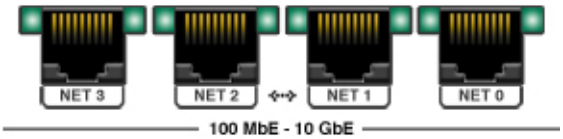

Network Management Port

The Oracle ZFS Storage ZS3-2 network management connector (NET MGT) is an RJ-45 10/100 BASE-T Ethernet port and provides an alternate terminal interface to the SP console.

Serial Management Port

The Oracle ZFS Storage ZS3-2 serial management connector (SER MGT) is an RJ-45 port and provides a terminal connection to the service processor (SP) console.

Cluster I/O Ports

Two cluster serial ports (0 and 1) and one Ethernet port provide communication between two Oracle ZFS Storage ZS3-2 controllers to form a cluster configuration. For information about

how to connect cables to form a cluster, see Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

> **Caution:**
>
> When disconnecting an RJ-45 cable from a cluster serial port (0 and 1), use extreme care not to damage the internal RJ-45 receptacle.

Oracle ZFS Storage ZS3-2 Controller Cluster I/O Ports: Serial 0, Serial 1, Ethernet

4x4 SAS HBA

The 4x4 SAS HBA, installed in PCIe slot 2 of the Oracle ZFS Storage ZS3-2 controller, provides connectivity to external disk shelves. The HBA logical ports are numbered 0-3, as shown in the following figure.

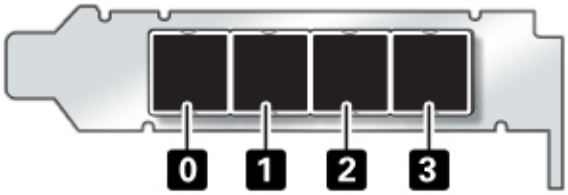

For controllers with an additional 4x4 SAS HBA, see Oracle ZFS Storage ZS3-2 PCIe I/O Cards for slot placement.

Oracle ZFS Storage ZS3-2 Specifications

Oracle ZFS Storage ZS3-2 Physical Specifications - The Oracle ZFS Storage ZS3-2 controller 2U chassis dimensions are as follows:

Table 10-3 Oracle ZFS Storage ZS3-2 Controller Physical Specifications

Dimension	Measurement	Dimension	Measurement
Height	3.44 in./87.4 mm	Depth	20.25 in./514 mm
Width	17.52 in./445 mm	Weight	41.23 lb/18.70 kg

Oracle ZFS Storage ZS3-2 Electrical Specifications - The Oracle ZFS Storage ZS3-2 controller electrical specifications are listed below. The power dissipation

numbers listed are the *maximum* rated power numbers for the power supply. The numbers are not a rating of the actual power consumption of the appliance.

Connectors

- Two C13 connectors which work on 110-220v outlets

Input

- Nominal frequencies: 50/60Hz (47 to 63 Hz range)
- Nominal voltage range: 100 to 120/200 to 240 VAC
- Maximum current AC RMS: 6.8 A @ 100-120 V / 3.4 A @ 200-240 V
- AC operating range: 90-264 VAC

Output

- 3.3 VDC STBY: 3.0A
- +12 VDC: 86.7A

Power Dissipation

- Max power consumption: 890 W max
- Max heat output: 3026 BTU/hr
- Volt-Ampere rating: 908 VA @ 240 VAC, 0.98P.F.

Oracle ZFS Storage ZS3-2 Acoustic Noise Emissions - In compliance with the requirements defined in ISO 7779, the workplace-dependent noise level of this product is less than 70 db (A).

Oracle ZFS Storage ZS3-2 Internal Components

The Oracle ZFS Storage ZS3-2 chassis has the following field-replaceable units (FRUs). FRUs are not customer-serviceable and must be replaced by trained Oracle service technicians.

- **PCIe Risers** - There are three risers per system, each attached to the rear of the motherboard. Each riser supports two PCIe cards (which are customer-replaceable units (CRUs)).
- **Motherboard** - The motherboard includes CPU modules, slots for 16 DIMMs, memory control subsystems, and the service processor (SP) subsystem. The SP subsystem controls the host power and monitors host system events (power and environmental). The SP controller draws power from the host 3.3V standby supply rail, which is available whenever the system is receiving AC input power, even when the system is turned off.
- **Power Distribution Board** - The power distribution board distributes main 12V power from the power supplies to the rest of the controller. It is directly connected to the connector break out board and to the motherboard through a bus bar and ribbon cable. It also supports a top cover interlock *kill* switch. The power supplies connect directly to the power distribution board.
- **Storage Drive Backplane** - The storage drive backplane includes the connectors for the storage drives, as well as the interconnect for the I/O board, power and locator buttons, and system/component status LEDs. The system has an 8-disk backplane. Each drive has an LED indicator for Power/Activity, Fault, and Locate.

Oracle ZFS Storage ZS3-2 Motherboard, Memory, and PCIe Cards - The Oracle ZFS Storage ZS3-2 controller motherboard, memory, and PCIe components are described in the following figure and legend.

> **Note:**
>
> Field-replaceable components (FRUs) are not customer-serviceable, and must be replaced by trained Oracle service technicians.

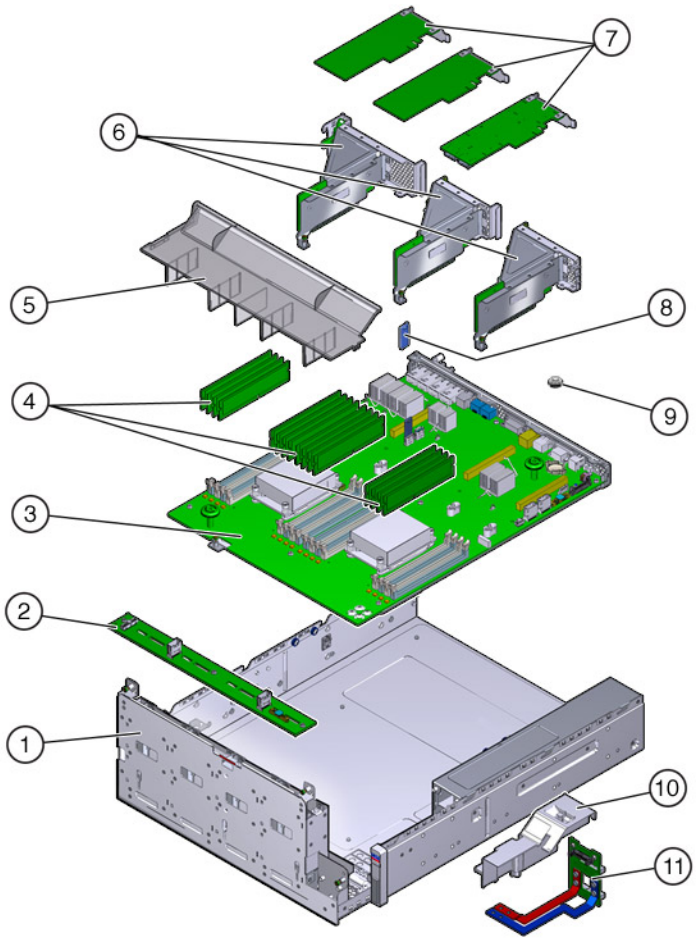

1 Drive cage (FRU)	**5** Airduct	**9** Battery
2 Drive backplane (FRU)	**6** PCIe Risers (FRU)	**10** PDB duct
3 Motherboard (FRU)	**7** PCIe cards	**11** Power distribution board (FRU)
4 DIMMs	**8** USB flash drive	

Oracle ZFS Storage ZS3-2 CPU and Memory - The Oracle ZFS Storage ZS3-2 controller motherboard has 16 slots in two groups that hold industry-standard DDR3 low voltage (LV) DIMMS.

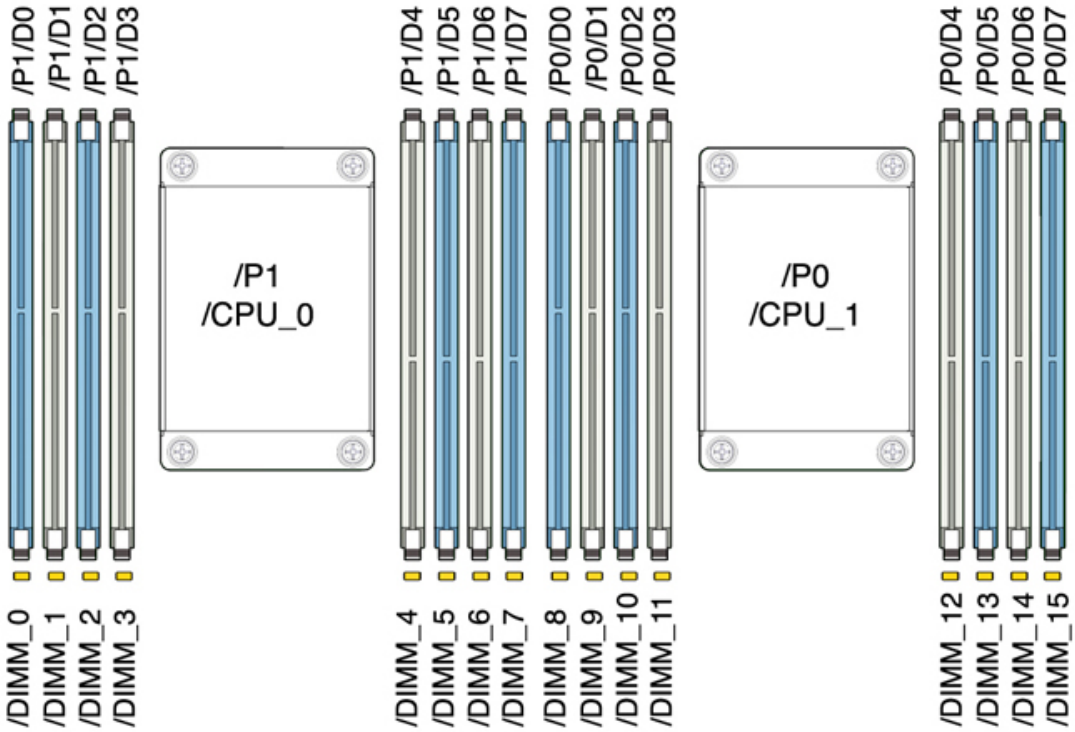

Note:

All sockets must be occupied by identical DDR3 DIMMs.

Table 10-4 Oracle ZFS Storage ZS3-2 Controller CPU Descriptions

Capacity	CPU 0	CPU 1
256GB	D0, D2, D5, D7 (blue)	D0, D2, D5, D7 (blue)
.	D1, D3, D4, D6 (white)	D1, D3, D4, D6 (white)

The Oracle ZFS Storage ZS3-2 controller replaceable memory components and part numbers are listed next.

Table 10-5 Oracle ZFS Storage ZS3-2 Controller Replaceable Memory Components

Component	Description	FRU/CRU	Part Number
CPU	Intel E5-2658, 2.1G, 8-core	FRU	7019701
Memory	DIMM, 16GB,DDR3,1600,2Rx4,1.35V	CRU	7041603
Memory	DIMM, 32GB,DDR3,1066,4Rx4,1.35V	CRU	7055964

Oracle ZFS Storage ZS3-2 Storage, Power, and Fan Components - The Oracle ZFS Storage ZS3-2 controller's internal storage, power, and cooling components are described in the following figure and legend. A component identified as a field-replaceable unit (FRU) must be replaced by trained Oracle service technicians.

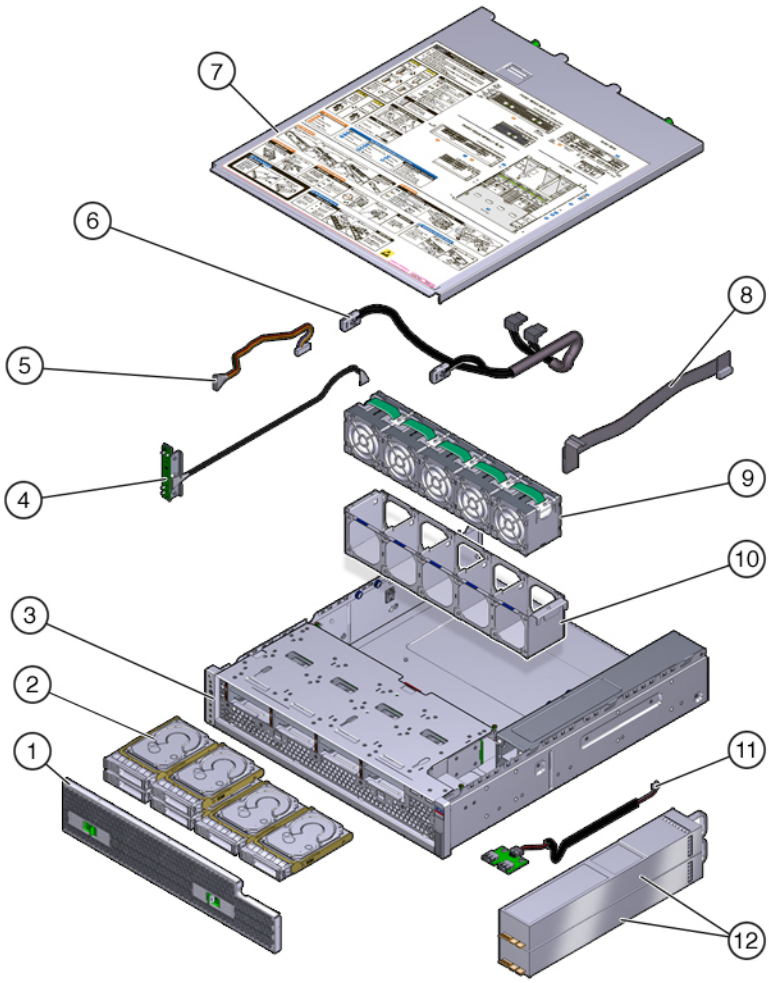

1 Air Filter	**5** Drive power cable (FRU)	**9** Fan modules
2 Drives	**6** Drive signal cable (FRU)	**10** Fan tray
3 Chassis (FRU)	**7** Top cover	**11** USB board (FRU)
4 LED board (FRU)	**8** PDB signal cable (FRU)	**12** Power supplies

Oracle ZFS Storage ZS3-2 Internal Cables - The Oracle ZFS Storage ZS3-2 controller contains the following field-replaceable unit (FRU) internal cables. FRUs are not customer-serviceable, and must be replaced by trained Oracle service technicians.

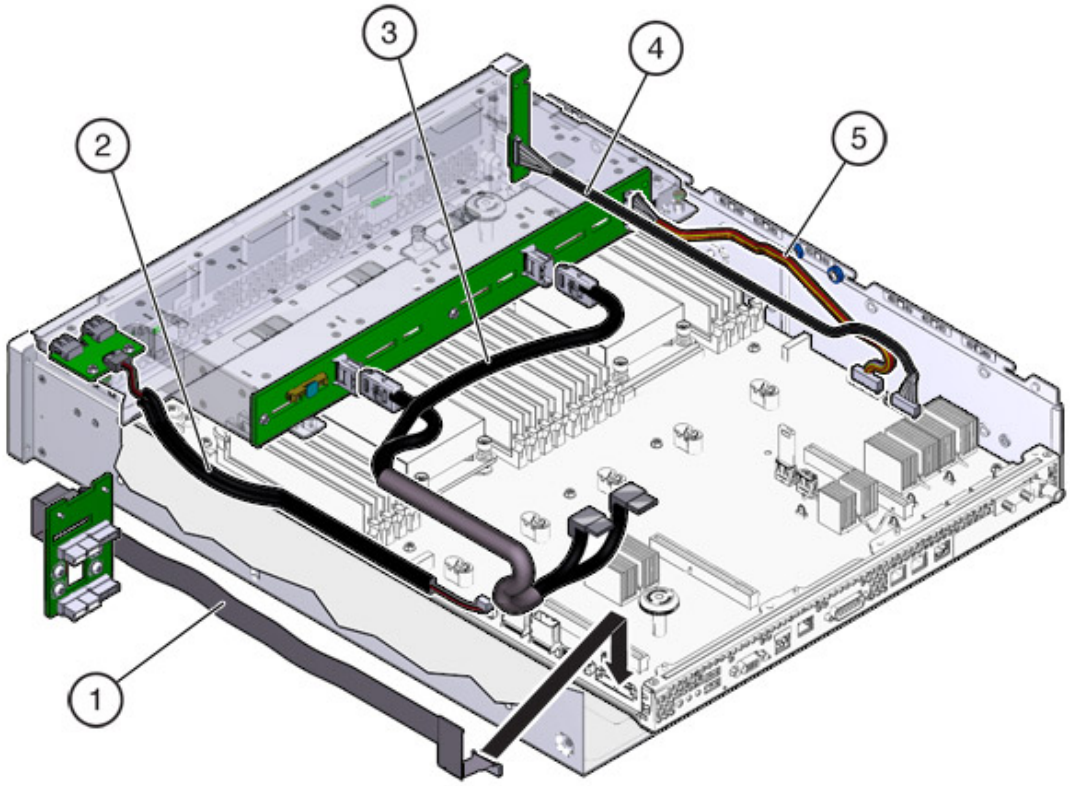

1 PDB signal cable	**3** Drive signal cable	**5** Drive power cable
2 USB board cable	**4** LED board cable	

Oracle ZFS Storage ZS3-2 PCIe I/O Cards

Oracle ZFS Storage ZS3-2 PCIe I/O Options - The following table describes the supported PCIe configuration options for the Oracle ZFS Storage ZS3-2 controller. Slots 1 and 2 are reserved for internal and external HBAs, as shown in the following table. When adding PCIe cards, populate from higher order slots (6) first toward lower order slots.

> **Note:**
>
> Any empty PCIe slots must have a filler panel installed.

Table 10-6 Oracle ZFS Storage ZS3-2 Controller PCIe Configuration Options

Slot	Mfg Part Number	Description	Max	Note
1	7047852	8-port SAS-2 internal HBA	1	Base configuration
2	7067091	4-port (4x4) SAS external HBA	2	Base configuration
3	7067091	4-port (4x4) SAS external HBA	2	Second 4x4-port SAS external HBA

Table 10-6 (Cont.) Oracle ZFS Storage ZS3-2 Controller PCIe Configuration Options

Slot	Mfg Part Number	Description	Max	Note
3	7070195	4-port 1Gb Ethernet Copper UTP NIC	4	Optional recommended front-end
3	371-4325-02	2-port 8Gb FC HBA	4	Optional FC target or initiator (backup)
3	7023303	2-port 16Gb FC HBA	4	Optional FC target or initiator (backup)
3	375-3696-01	2-port InfiniBand CX2 HCA	4	Optional recommended front-end
3	7051223	2-port 10Gb Ethernet Optical NIC	4	Optional recommended front-end
3	7070006	2-port 10Gb Ethernet Copper Base-T NIC	4	Optional recommended front-end
4-6	7070195	4-port 1Gb Ethernet Copper UTP NIC	4	Optional recommended front-end
4-6	371-4325-02	2-port 8Gb FC HBA	4	Optional FC target or initiator (backup)
4-6	7023303	2-port 16Gb FC HBA	4	Optional FC target or initiator (backup)
4-6	375-3696-01	2-port InfiniBand CX2 HCA	4	Optional recommended front-end
4-6	7051223	2-port 10Gb Ethernet Optical NIC	4	Optional recommended front-end
4-6	7070006	2-port 10Gb Ethernet Copper Base-T NIC	4	Optional recommended front-end

Oracle ZFS Storage ZS3-2 NIC/HBA Options - The following table describes the NIC/HBA options for the Oracle ZFS Storage ZS3-2 controller.

Table 10-7 Oracle ZFS Storage ZS3-2 Controller NIC/HBA Options

Mktg Part Number	Description
SG-SAS6-INT-Z	8-port 6Gb/s SAS-2 internal HBA
SG-XPCIE2FC-QF8-Z	2-port 8Gb FC HBA
7103791	4-port (4x4) SAS-3 12 Gb/s external HBA
7110372	4-port (4x4) SAS-2 6Gb/s external HBA
7101674	2-port 16Gb FC HBA
X1109A-Z	2-port 10GbE SFP+ NIC

Table 10-7 (Cont.) Oracle ZFS Storage ZS3-2 Controller NIC/HBA Options

Mktg Part Number	Description
X4242A	2-port InfiniBand CX2 HCA
7100477	4-port 1Gb Ethernet Copper UTP NIC
7100488	2-port 10Gb Ethernet Copper Base-T NIC
X2129A	XCVR 850NM, 1/10GPS, SFP, short reach
X5562A-Z	10GbE/1GbE SFP+ Transceiver, long reach

Oracle ZFS Storage ZS3-2 PCIe Riser Configuration - The three risers are labeled Riser 1, Riser 2, and Riser 3. Though similar, the risers are not interchangeable. Riser 1 installs at the left rear of the chassis, Riser 2 at the center rear, and Riser 3 at the right rear of the chassis. Each riser can accommodate two PCIe cards:

- Riser 1 contains slots 1 and 4
- Riser 2 contains slots 2 and 5
- Riser 3 contains slots 3 and 6

Oracle ZFS Storage ZS3-2 Optional Cable Management Arm

The following figure identifies the components of the second-generation cable management arm (CMA). See the CMA installation instructions.

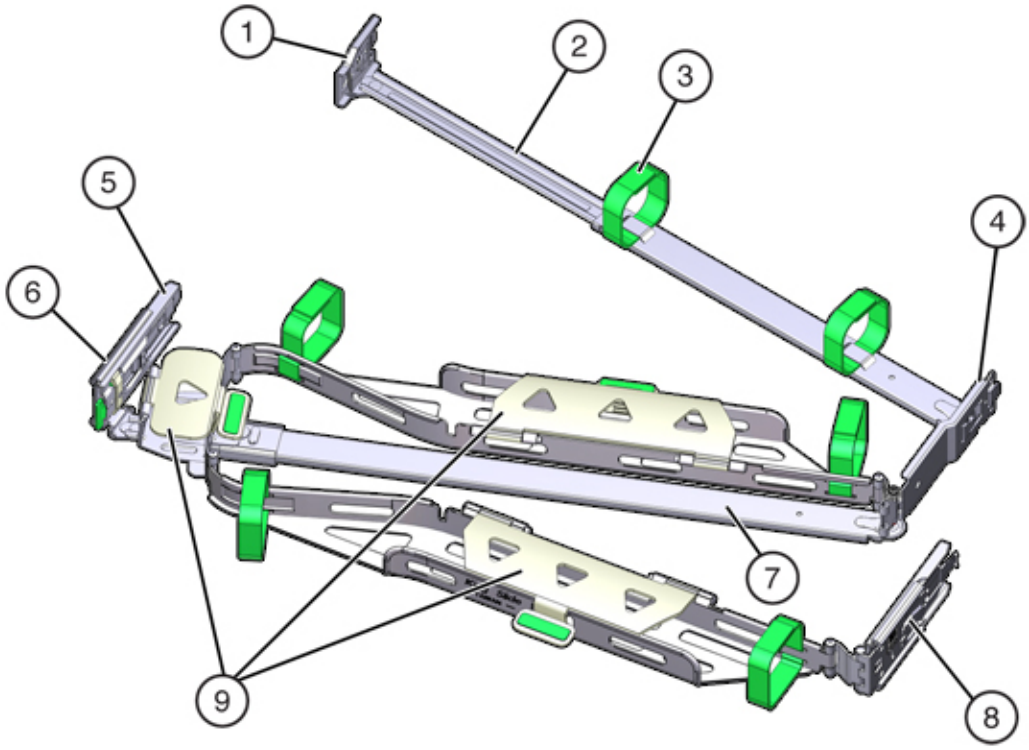

1 Connector A	**6** Connector D
2 Front slide bar	**7** Slide-rail latching bracket (used with connector D)
3 Velcro straps (6)	**8** Rear slide bar
4 Connector B	**9** Cable covers
5 Connector C	

Oracle ZFS Storage ZS3-2 Attached Storage

The Oracle ZFS Storage ZS3-2 controller connects to external storage through a 4-port (4x4) SAS HBA. You can attach from one to 16 HDD-only or HDD/SSD Log cache disk shelves to the controller. You can also attach mixed disk shelf types behind the same controllers, but each chain must contain only the same disk shelf type. Directly connecting different disk shelf types is not supported. For more information, see Connecting to Attached Storage. See also Disk Shelf Hardware Overview for component specifications and diagrams.

11

Servicing the Sun ZFS Storage 7x20 Controller

Be sure to read sections Getting Started with Servicing the Appliance and Prerequisites for Servicing Hardware.

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see Repairing Active Problems.

Oracle Storage Drive Enclosure DE2-24 disk shelves and Sun Disk Shelves can be used together with standalone and clustered controllers after certain hardware modifications. Contact Oracle Service to upgrade your Sun ZFS Storage 7x20 controllers for disk shelf interoperability.

To service Sun ZFS Storage 7x20 hardware components, use the following tasks:

- Replacing a Sun ZFS Storage 7x20 HDD or SSD
- Replacing a Sun ZFS Storage 7x20 Fan Module
- Replacing a Sun ZFS Storage 7x20 Power Supply
- Replacing a Sun ZFS Storage 7x20 DIMM
- Replacing a Sun ZFS Storage 7x20 Controller PCIe Card or Riser
- Replacing a Sun ZFS Storage 7x20 Battery

To understand the Sun ZFS Storage 7x20 hardware, see the following topics:

- Sun ZFS Storage 7420 Controller Hardware Overview
- Sun ZFS Storage 7320 Controller Hardware Overview
- Sun ZFS Storage 7120 Controller Hardware Overview

Replacing a Sun ZFS Storage 7x20 HDD or SSD

Sun ZFS Storage 7x20 HDDs and SSDs are hot-swappable and can be removed and installed while the storage controller is running without affecting other hardware capabilities. Use the following procedure to replace a Sun ZFS Storage 7x20 HDD or SSD.

> **Note:**
>
> If you have more than one failed drive, replace only one drive at a time. Removing more than one drive in quick succession causes a hardware/pool fault. Also, do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** BUI menu, select **System**, or navigate to `maintenance system updates` in the CLI.

1. Identify the faulted HDD or SSD.

- If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, and open the detailed view of the controller. The faulted drive has either an amber or a blue-on-amber (ready to remove) status icon. Note the ID of the faulted drive. Click the entry for that drive in the table to highlight the location of the drive in the controller illustration. To view details about the faulted drive, click its information icon and then click **Active Problems**.
- If you are physically at the system, the amber Service Action Required indicator on the HDD or SSD should be lit. Additionally, you can click the locator icon in the BUI to light the locator LED on the faulted drive.

2. Verify that the faulted drive is ready to be removed.

> **Caution:**
>
> Removing the drive before it is ready to be removed can result in loss of data.

Use any of the following methods to verify that the faulted drive is ready to be removed:

- On the faulted drive, the Ready to Remove LED is lit as shown in Sun ZFS Storage 7420 System Drives.
- In the BUI:
  - The status icon for the faulted drive in the detail view in **Maintenance: Hardware** has changed from amber to blue-on-amber.
  - The **Component Details** information window (click the information icon) for the faulted drive shows **Yes** as the value for the **Ready to remove** property.
- In the CLI (and RESTful API), the value of the `readytoremove` property for the faulted drive is `true`:

```
hostname:maintenance chassis-000 disk-001> get readytoremove
                readytoremove = true
```

3. Push the release button on the HDD or SSD to open the latch.
4. Grasp the latch [2], and pull the drive out of the drive slot.

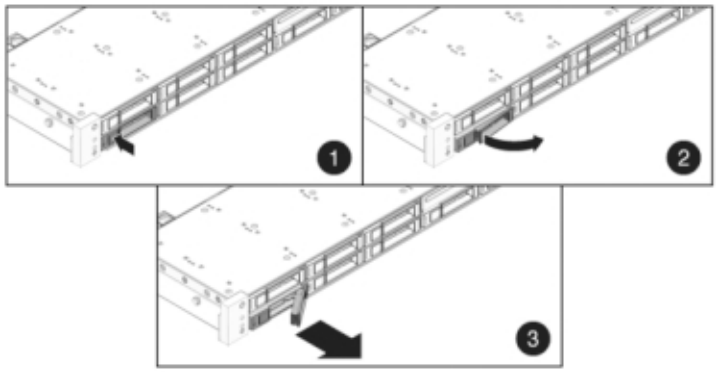

5. After at least 30 seconds, from the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon on the system controller to verify that the software has detected that the drive is not present.
6. Slide the replacement drive into the slot until it is fully seated.

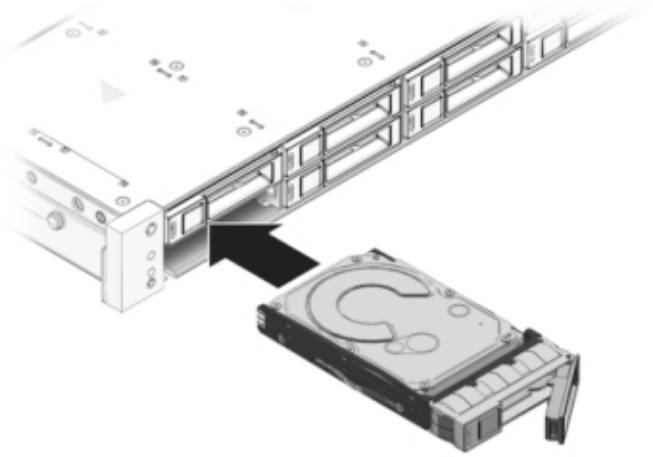

7. Close the latch to lock the drive in place.
8. From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon for the controller or disk shelf, and then click **Disk** to verify that the online icon is green for the newly installed drive.
9. If a fault does not self-clear, seeRepairing Active Problems .

Replacing a Sun ZFS Storage 7x20 Fan Module

> **Caution:**
>
> Operating a controller for over 60 seconds with fans removed reduces the effectiveness of the cooling system. For this reason, the replacement fan should be unpacked in advance and ready to insert into the controller chassis as soon as the faulted fan is removed.

Fan modules are hot-swappable and can be removed and installed while the storage controller is running without affecting other hardware capabilities.

Sun ZFS Storage 7120 or Sun ZFS Storage 7320: The fan modules and status indicators are hidden under a fan door in the Sun ZFS Storage 7120 and Sun ZFS Storage 7320 storage controllers. Components may differ slightly between the Sun ZFS Storage 7120 and Sun ZFS Storage 7320; however, service procedures for each are identical. The illustration shows the Sun ZFS Storage 7320 controller.

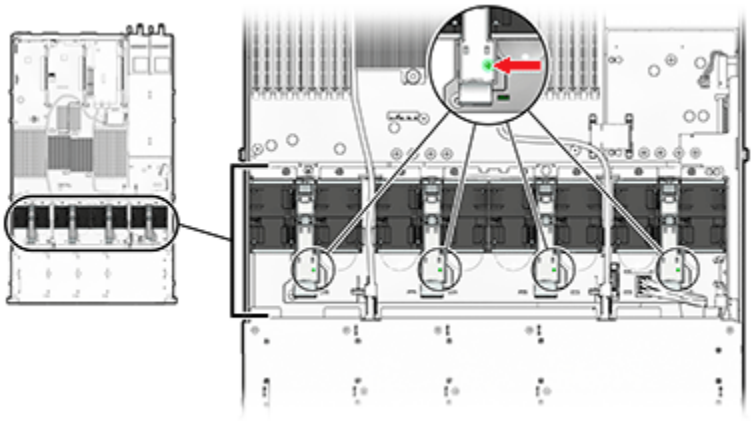

Caution:

The storage controller might overheat and shut down if the door is left open for more than 60 seconds.

Sun ZFS Storage 7420: The following illustration shows the fan modules in the Sun ZFS Storage 7420 storage controller.

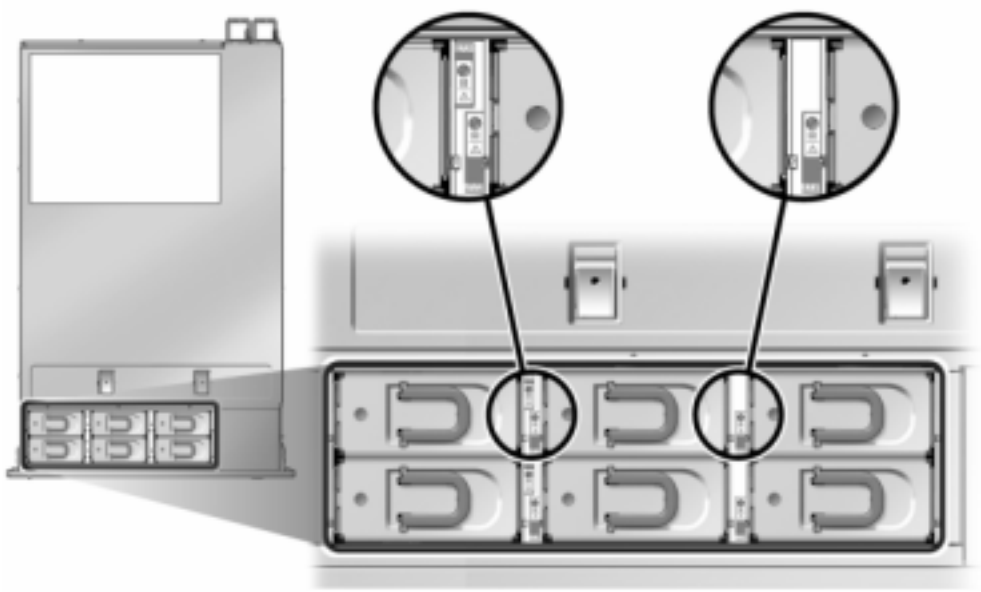

1. To flash the locator LED on the chassis to service, from the **Maintenance** BUI menu, select **Hardware**, and click the associated locator icon, or go to the service processor (SP) prompt and enter `set /SYS/LOCATE status=on`.
2. To locate the faulted component, from the **Maintenance** BUI menu, select **Hardware**, then select **Fan**, and click the row with an amber indicator.

   The component is highlighted in the diagram.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. **Sun ZFS Storage 7120 or Sun ZFS Storage 7320:** To replace the fan module:
   a. Open the fan module door by unlatching its release tabs.

> **Caution:**
>
> The storage controller might overheat and shut down if the door is left open for more than 60 seconds.

   b. If the system is still powered on, identify the fan module with the lit fault indicator.
   c. Using thumb and forefinger, squeeze the tabs together to pull the fan module up and out.

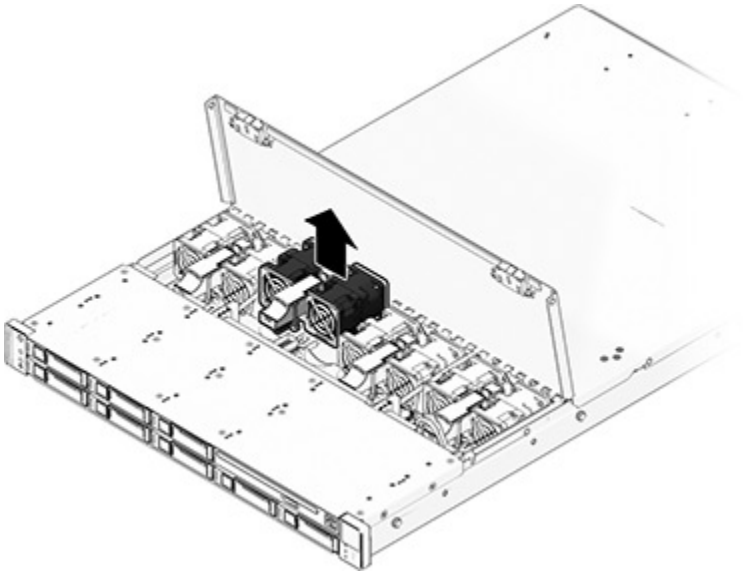

   d. Insert the replacement fan module into the slot and apply firm pressure to fully seat it.

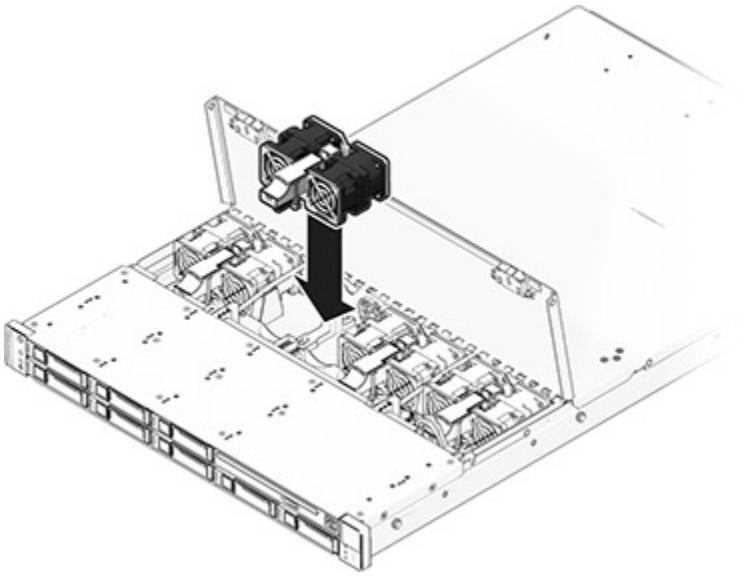

e. On the replaced fan module, verify that the Fan OK indicator is lit, and that the fault indicator is not lit.

f. Close the door immediately to maintain airflow in the storage controller.

5. **Sun ZFS Storage 7420:** To replace the fan module:

a. Lift the latch at the top of the fan module to unlock the fan module, and pull the fan module out.

> **Caution:**
>
> The fan module must be replaced within 60 seconds because the storage controller might overheat and shut down.

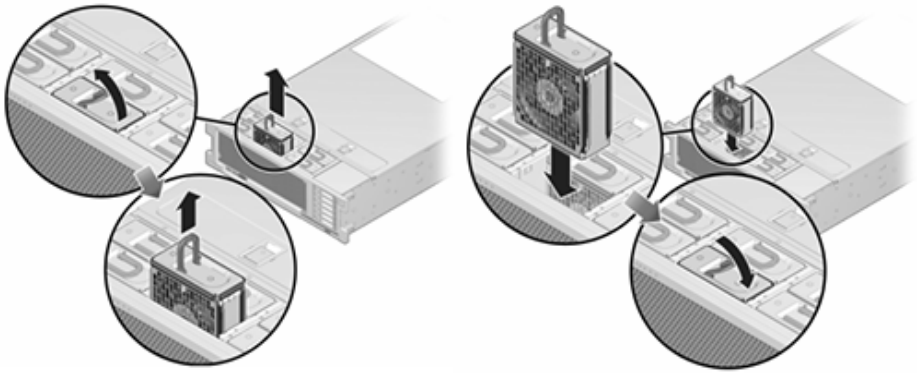

b. Unlock and insert the Sun ZFS Storage 7420 fan module into the slot and apply firm pressure to fully seat it.

c. On the replaced fan module, verify that the Fan OK indicator is lit, and that the fault indicator is not lit.

6. Verify that the controller's Top Fan status indicator, Service Required status indicators, and the Locator status indicator/Locator button are not lit.
7. Push the release tabs on the side of each rail and slowly slide the storage controller into the rack.
8. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, and then click **Fan** to verify that the online icon is green for the newly installed fan module.
9. If a fault does not self-clear, see Repairing Active Problems.

Replacing a Sun ZFS Storage 7x20 Power Supply

Sun ZFS Storage 7x20 controllers are equipped with redundant hot-swappable power supplies. If a power supply fails and you do not have a replacement, leave the failed power supply installed to ensure proper air flow. A faulted power supply is indicated by an amber-colored status LED. Use the following procedure to replace a Sun ZFS Storage 7x20 power supply.

1. Gain access to the rear of the storage controller where the faulted power supply is located.
2. If a cable management arm (CMA) is installed, press and hold the CMA release tab and rotate the arm out of the way.

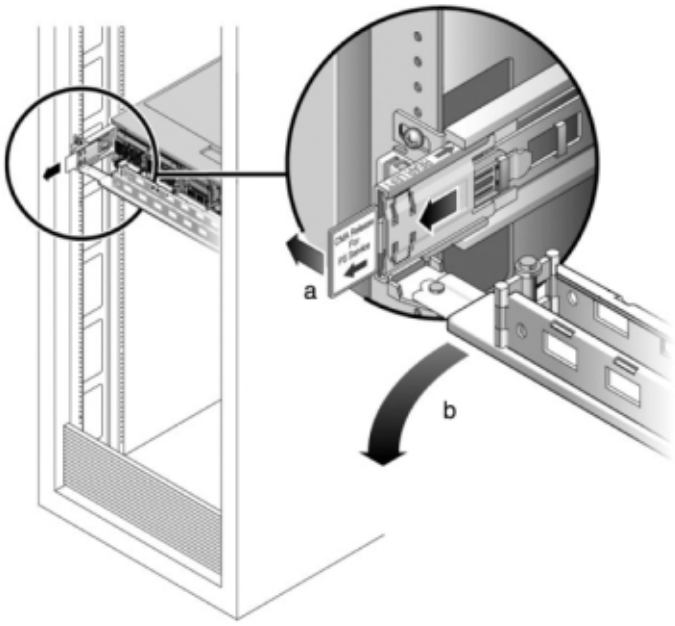

3. Disconnect the power cord from the faulted power supply.
4. Shut down the controller using one of the power-off methods described in Powering Off the Controller.

   Sun ZFS Storage 7120 or Sun ZFS Storage 7320: Release the latch, then remove the power supply. Components may differ slightly between the Sun ZFS Storage 7120 and Sun ZFS Storage 7320; however, service procedures for each are identical. The illustration shows the Sun ZFS Storage 7320 controller.

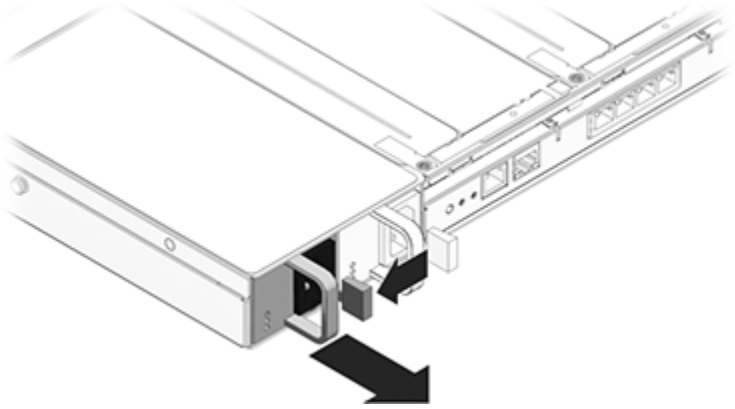

Sun ZFS Storage 7420: Grasp the power supply handle and press the release latch to remove the power supply.

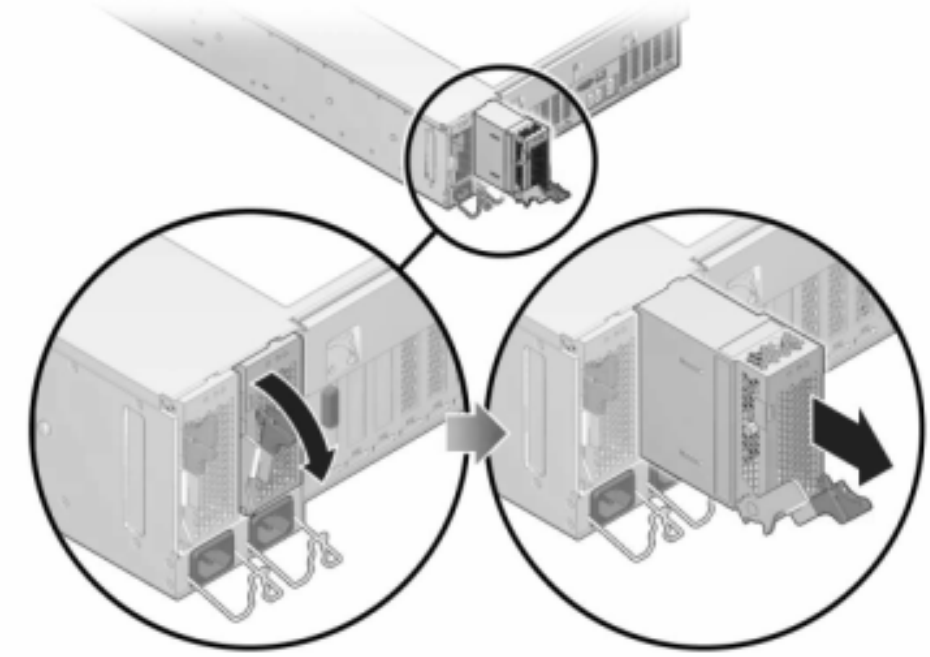

5. Align the replacement power supply with the empty power supply chassis bay.
6. Slide the power supply into the bay until it is fully seated. The following figure shows the Sun ZFS Storage 7420 power supply.

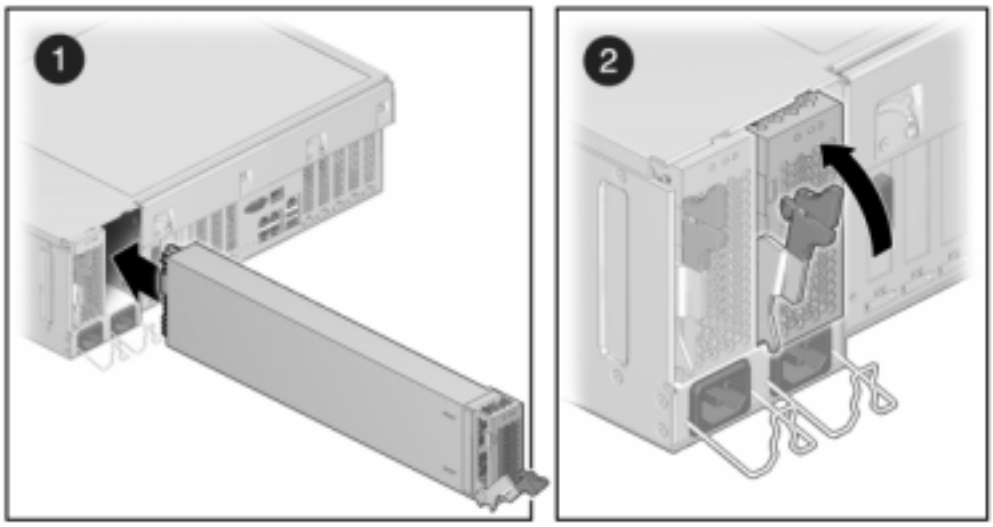

7. Connect the power cord to the power supply.
8. Verify that the green AC Present status indicator is lit.
9. Close the CMA, inserting the CMA into the rear left rail bracket.
10. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, and then click **PSU** to verify that the online icon is green for the newly installed power supply.
11. If a fault does not self-clear, seeRepairing Active Problems .

Replacing a Sun ZFS Storage 7x20 DIMM

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components.

To identify a specific memory module that has faulted, you must open the storage controller and use the amber status LEDs on the motherboard. To identify a general memory fault, from the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon on the controller. Then click **DIMM** to locate the faulted component, indicated by the warning icon .

1. Power off the storage controller as described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. **Sun ZFS Storage 7120 or Sun ZFS Storage 7320:** Components may differ slightly between the Sun ZFS Storage 7120 and Sun ZFS Storage 7320; however, service procedures for each are identical. The illustration shows the Sun ZFS Storage 7320 controller. To remove the top cover:
   a. Unlatch the fan module door, pull the two release tabs back, rotate the fan door to the open position, and hold it there.
   b. Press the top cover release button and slide the top cover to the rear about a half-inch (1.3 cm).
   c. Lift up and remove the top cover.
   d. Also remove the air baffle by pressing the air baffle connectors outward and lifting the air baffle up and out of the controller.

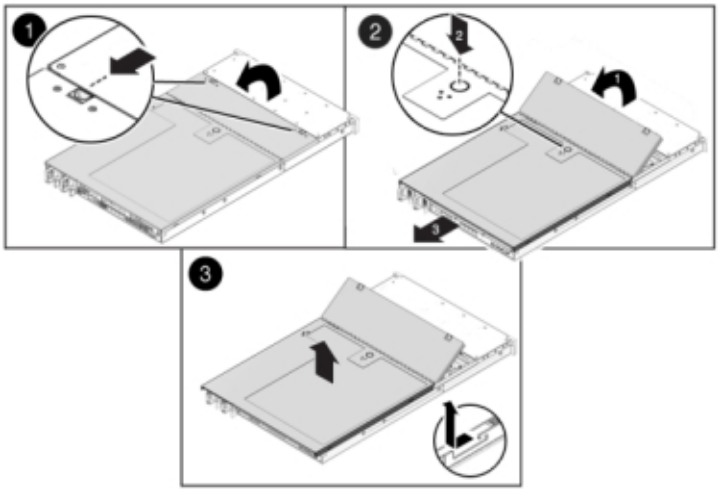

5. **Sun ZFS Storage 7420:** To remove the top cover:
   a. Simultaneously lift both cover latches in an upward motion.

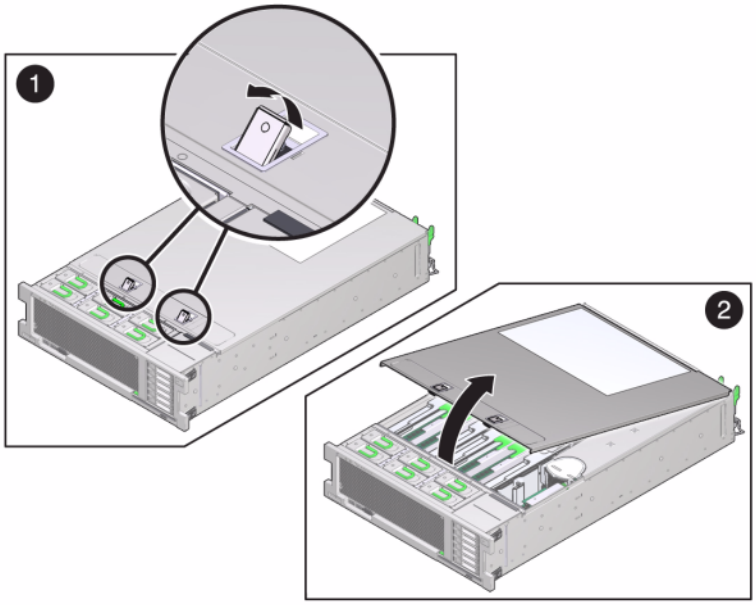

b. Lift up and remove the top cover.

6. To locate the DIMM you want to service, press the Fault Remind Button on the storage controller.

   The following illustration shows the Fault remind button on the Sun ZFS Storage 7120 controller.

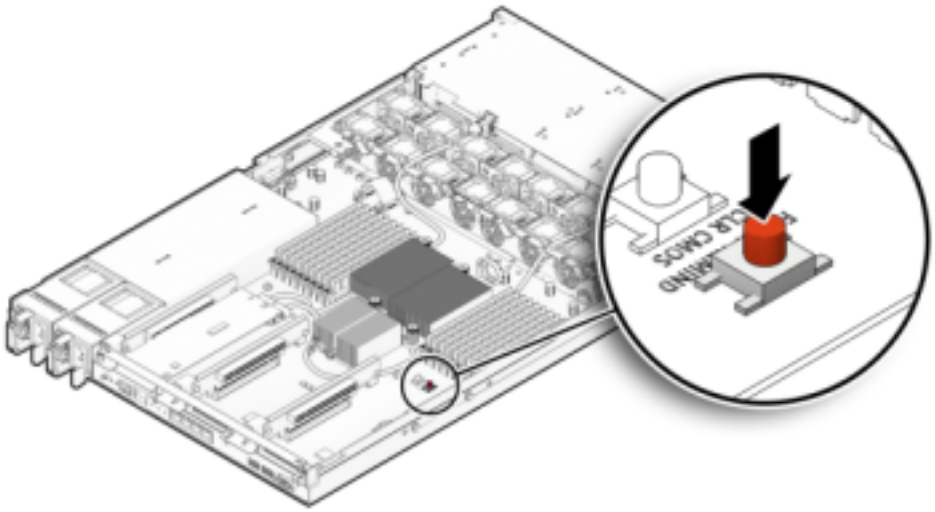

   The following illustration shows the Fault remind button on the Sun ZFS Storage 7420 controller.

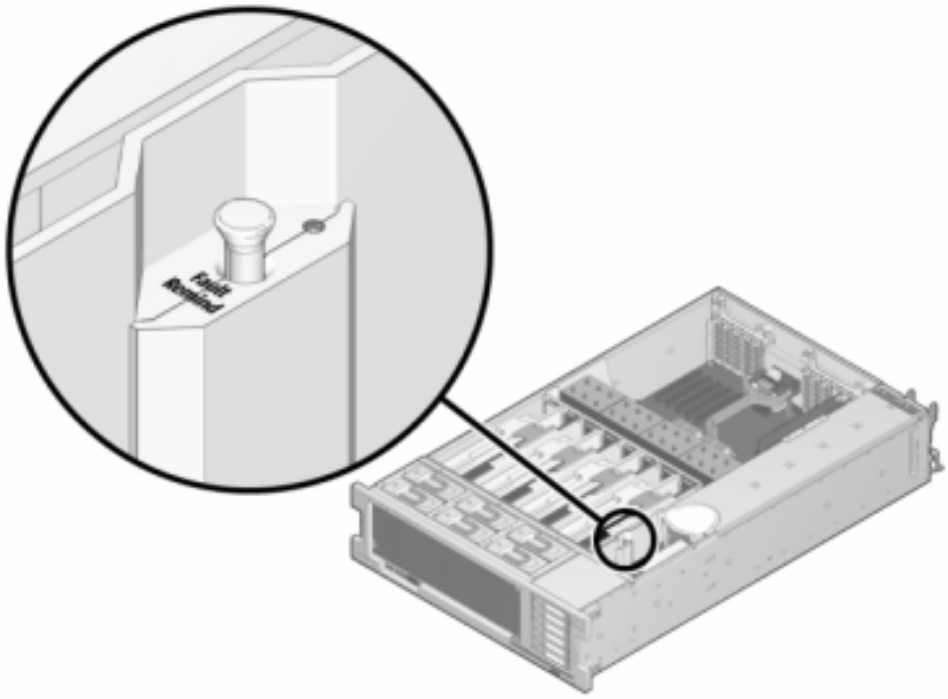

7. **Sun ZFS Storage 7420:** Identify the memory riser that hosts the faulted DIMM by the Service Required status indicator. Lift the memory riser straight up to remove it from the motherboard, and place it on an antistatic mat.

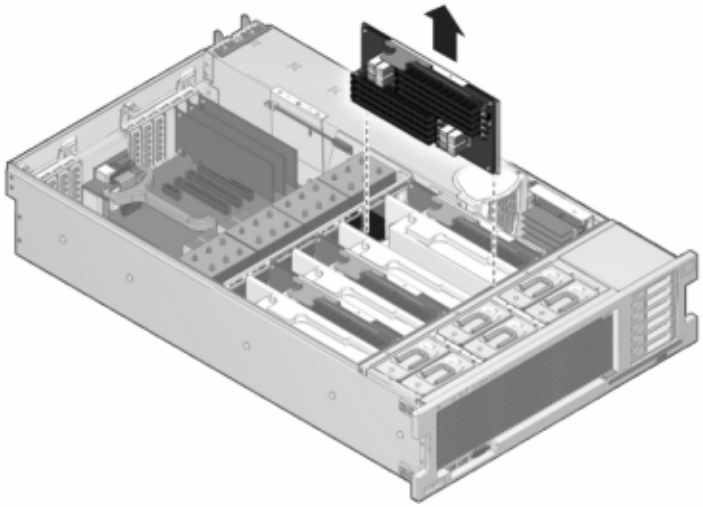

8. Rotate both DIMM slot ejectors outward as far as they will go, and carefully lift the faulted DIMM straight up to remove it from the socket.

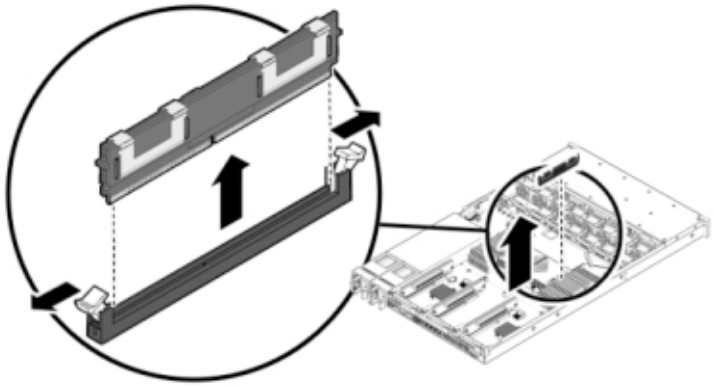

9. Line up the replacement DIMM with the connector, aligning the notch with the key to ensure that the component is oriented correctly.

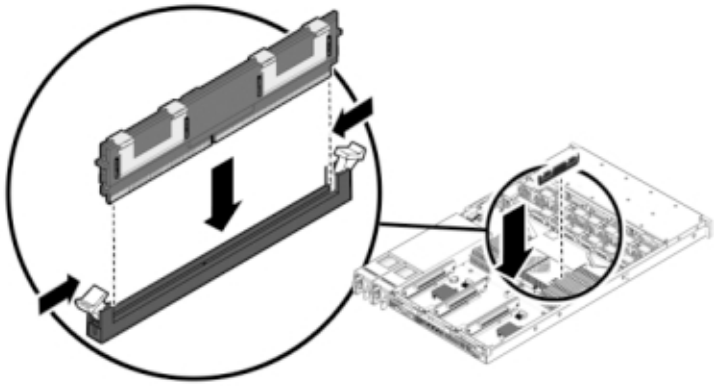

10. Push the DIMM into the connector until the ejector tabs lock the component in place.
11. **Sun ZFS Storage 7120 or Sun ZFS Storage 7320:** Components may differ slightly between the Sun ZFS Storage 7120 and Sun ZFS Storage 7320; however, service procedures for each are identical. The illustration shows the Sun ZFS Storage 7320. To replace the cover:
    a. Place the top cover on the chassis so that it hangs over the rear of the storage controller by about an inch (2.5 cm).
    b. Slide the top cover forward until it seats.
    c. Close the fan cover and engage the fan cover latches. The cover must be completely closed for the storage controller to power on.

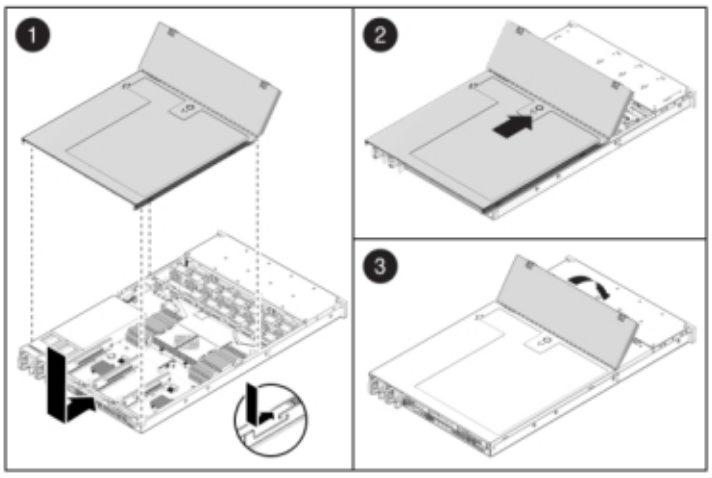

12. **Sun ZFS Storage 7420:** To replace the cover:

    a. Push the memory riser module into the associated CPU memory riser slot until the riser module locks in place.

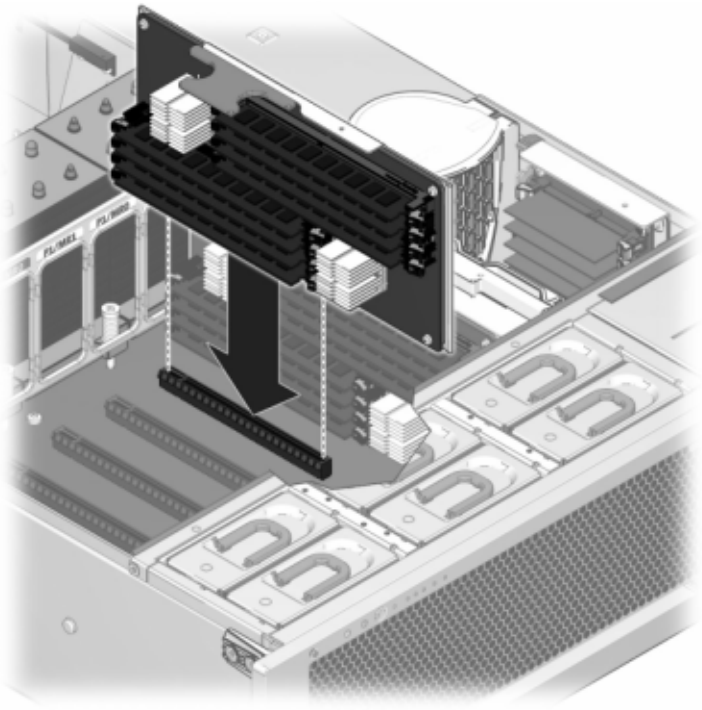

    b. Place the top cover on the chassis [1] so that it is forward of the rear of the storage controller by about an inch (2.5 cm).

    c. Slide the top cover toward the rear of the chassis [2] until it seats, and press down on the cover with both hands until both latches engage.

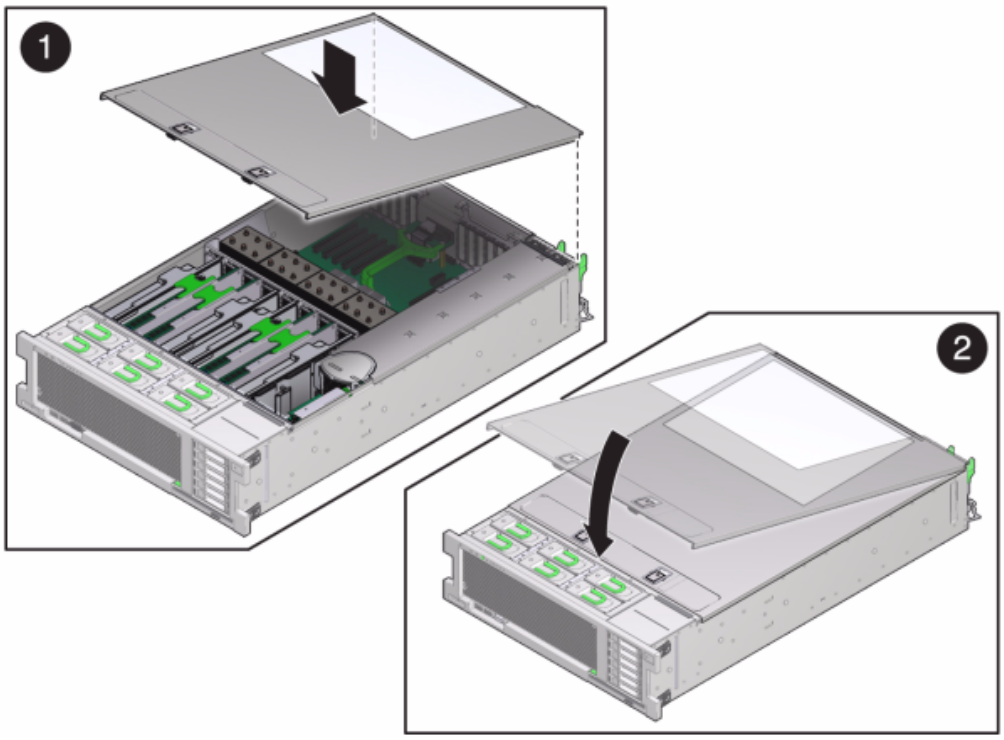

13. Push the release tabs on the side of each rail, and slowly push the storage controller into the rack.

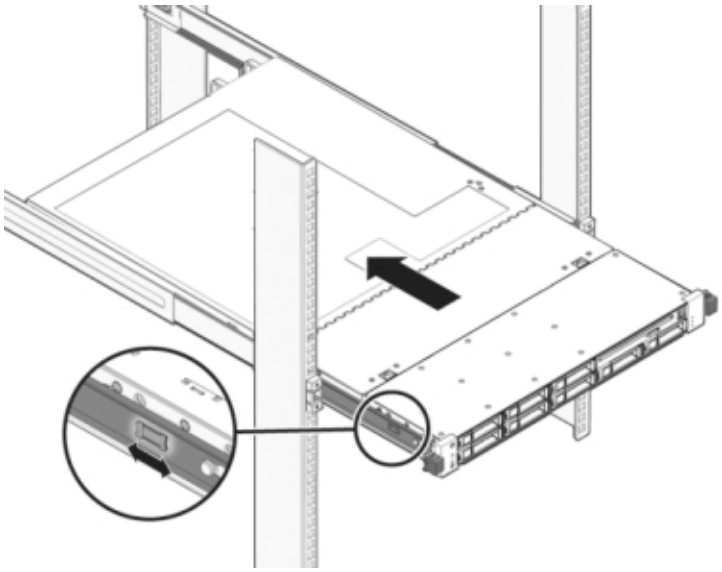

The following image shows the Sun ZFS Storage 7420 chassis.

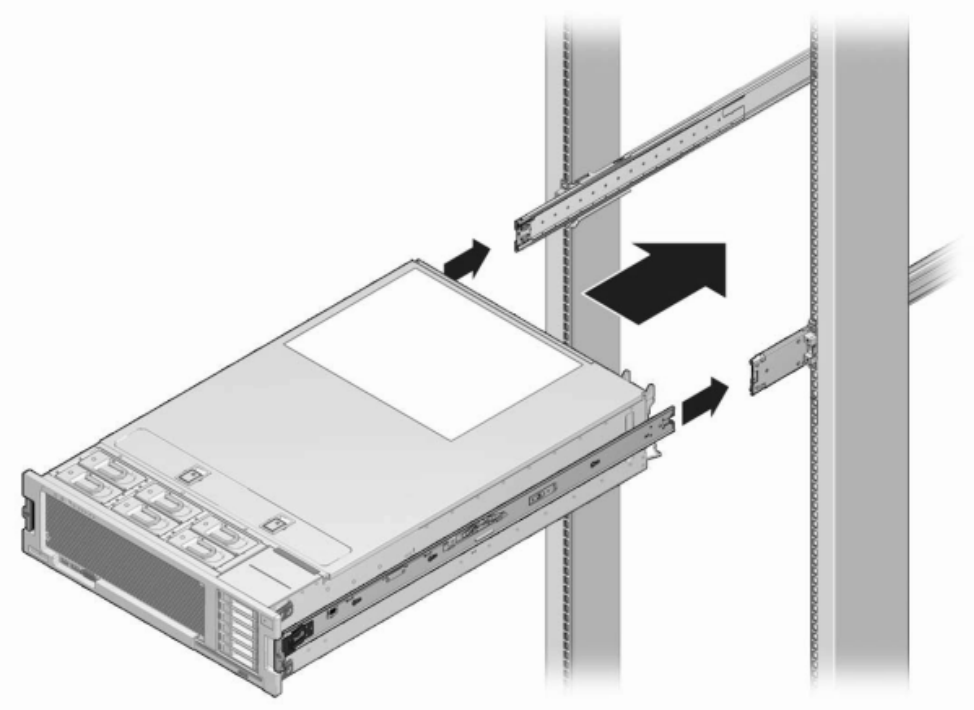

14. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
15. Connect the power cords to the power supplies.
16. Verify that standby power is on, indicated by the Power/OK status indicator flashing on the front panel about two minutes after the power cords are plugged in.
17. Use a pen or other pointed object to press and release the recessed Power button on the storage controller front panel.

    The Power/OK status indicator next to the Power button lights and remains lit.
18. From the **Maintenance** BUI menu, select **Hardware**. Click the right-arrow icon for the controller, and then click **DIMM** to verify that the online icon is green for the newly installed DIMM.
19. If a fault does not self-clear, see Repairing Active Problems.

Replacing a Sun ZFS Storage 7x20 Controller PCIe Card or Riser

> **Caution:**
>
> This procedure requires that you handle components that are sensitive to static discharge, which can cause the component to fail. To avoid damage, wear an antistatic wrist strap and use an antistatic mat when handling components. Note that the Sun ZFS Storage 7120 Sun Flash Accelerator F20 card is a field-replaceable unit (FRU) and *must* be replaced by an Oracle service representative.

From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon on the controller, and then click **Slot** to locate the faulted component. All HBAs must be of the same type. Ensure that you upgrade your system software before installing a newly released HBA.

1. Power off the storage controller as described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. **Sun ZFS Storage 7120 or Sun ZFS Storage 7320:** To remove the top cover:
   a. Unlatch the fan module door, pull the two release tabs back, rotate the fan door to the open position, and hold it there.
   b. Press the top cover release button, and slide the top cover to the rear about a half-inch (1.3 cm).
   c. Lift up and remove the top cover.
5. **Sun ZFS Storage 7420:** To remove the top cover:
   a. Simultaneously lift both cover latches in an upward motion.

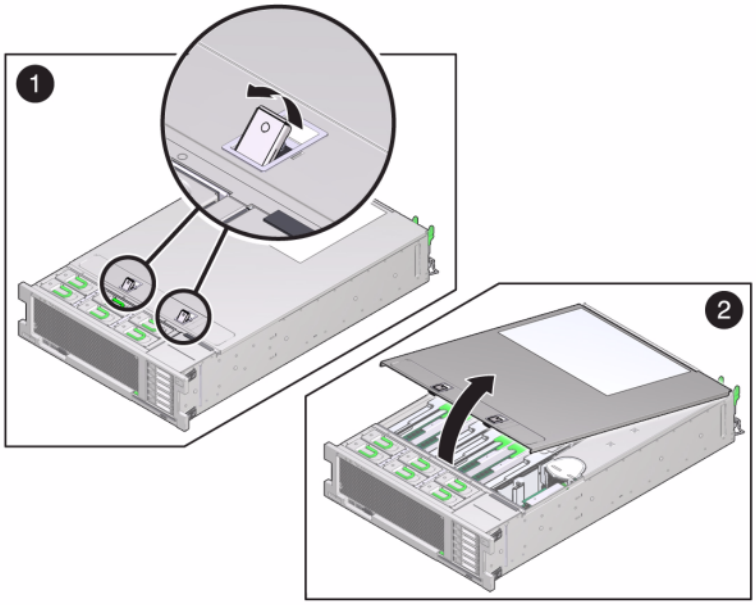

b. Lift up and remove the top cover.

6. Locate the PCIe card position in the storage controller, see Sun ZFS Storage 7320 Configuration Options for the Sun ZFS Storage 7320, the Sun ZFS Storage 7120 PCIe Options for the Sun ZFS Storage 7120, or Sun ZFS Storage 7420 PCIe I/O Cards for the Sun ZFS Storage 7420.

7. **Sun ZFS Storage 7120 or Sun ZFS Storage 7320:** To replace the PCIe card:

   a. Disconnect any data cables connected to the cards on the PCIe riser you want to replace. Label the cables for proper connection later.

   b. Loosen the two captive Phillips screws on the end of the rear panel crossbar, and lift the crossbar up and back to remove it.

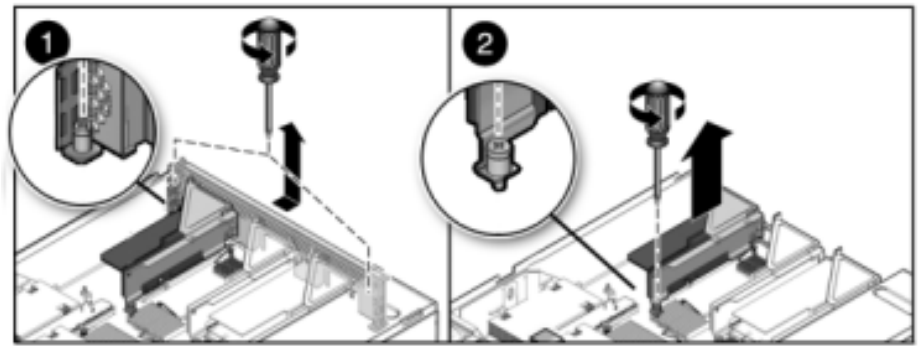

   c. Loosen the captive retaining screw holding the front end of the riser, and the Phillips screw on the end of the riser.

   d. Lift the riser up to remove it from the storage controller.

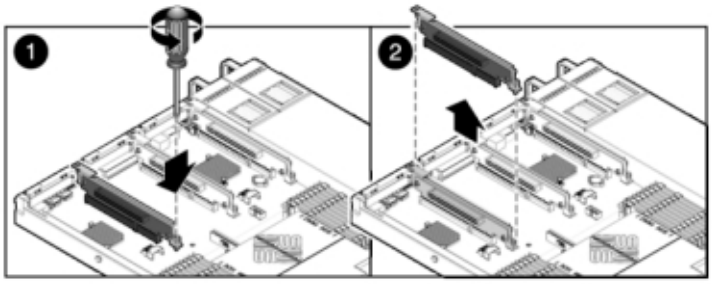

e. Carefully remove the PCIe card from the riser board connector, and clean the slot with filtered, compressed air if necessary.

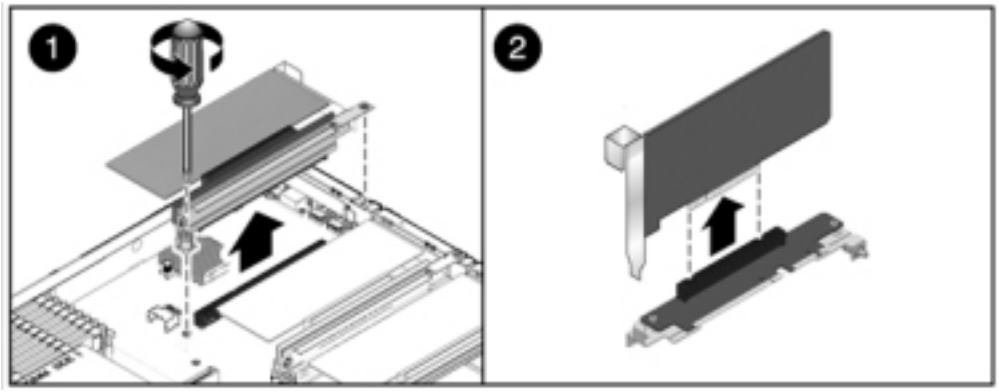

f. Seat the replacement PCIe card in the slot of the riser, and connect the cables.

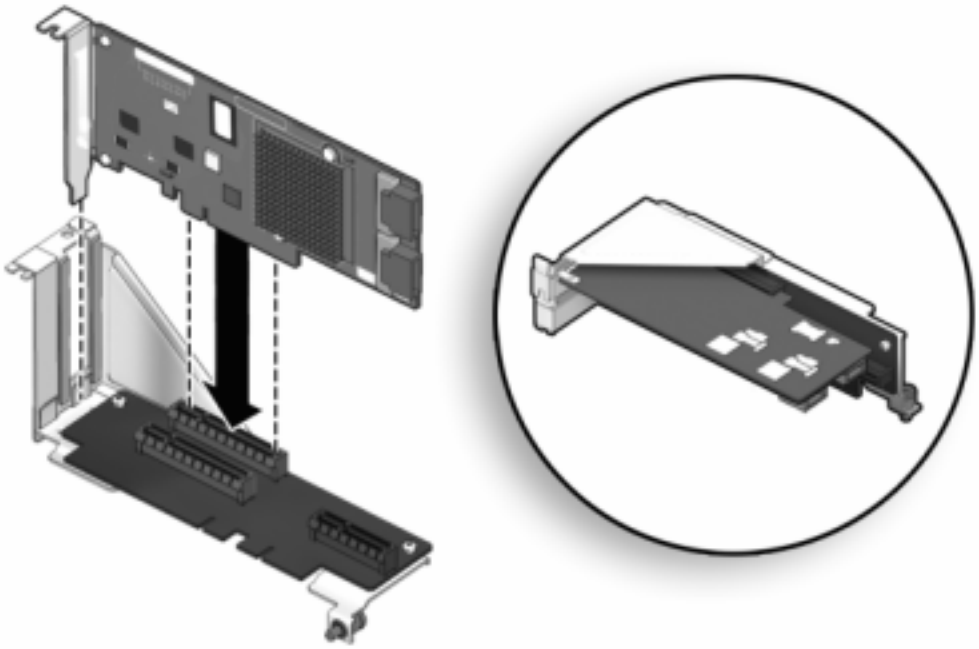

g. Align the riser, together with any attached PCIe cards, with the intended location on the motherboard, and carefully insert it into its slot.

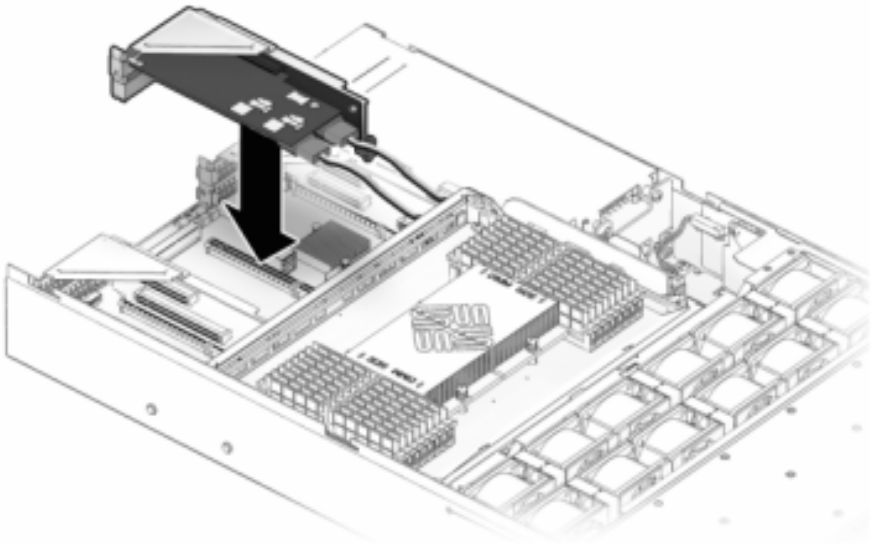

h. Slide the back of the riser into the motherboard rear panel stiffener.

i. Tighten the screw that secures the riser to the motherboard.

j. Replace the rear panel PCI crossbar by sliding it down over the PCIe risers, ensuring the crossbar is secured with two captive Phillips screws.

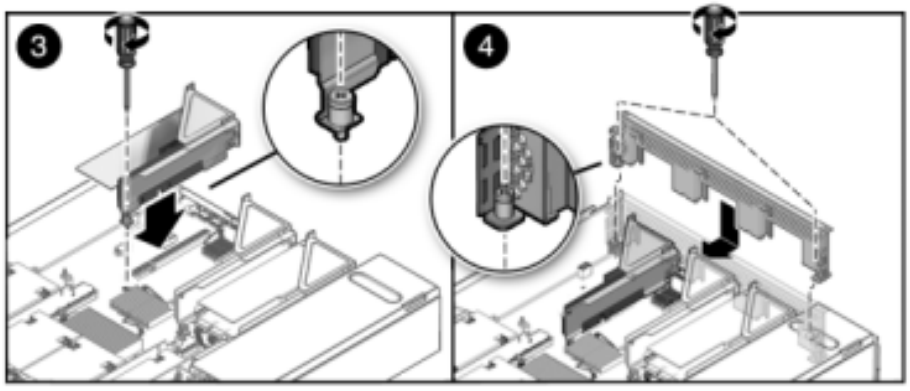

8. **Sun ZFS Storage 7420:** To replace the PCIe card:

   a. Disengage the PCIe card slot crossbar from its locked position, and rotate the crossbar into an upright position.

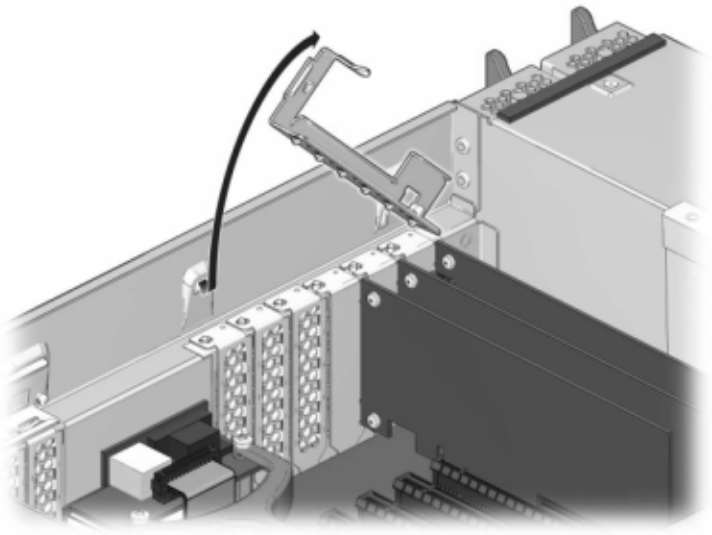

b. Remove the retaining screw that holds the PCIe card to the chassis.

c. Carefully remove the PCIe card from the connector, and clean the slot with filtered, compressed air if necessary.

d. Install the replacement PCIe card into the PCIe card slot.

e. Install the retaining screw to hold the PCIe card to the chassis.

f. Return the crossbar to its closed and locked position.

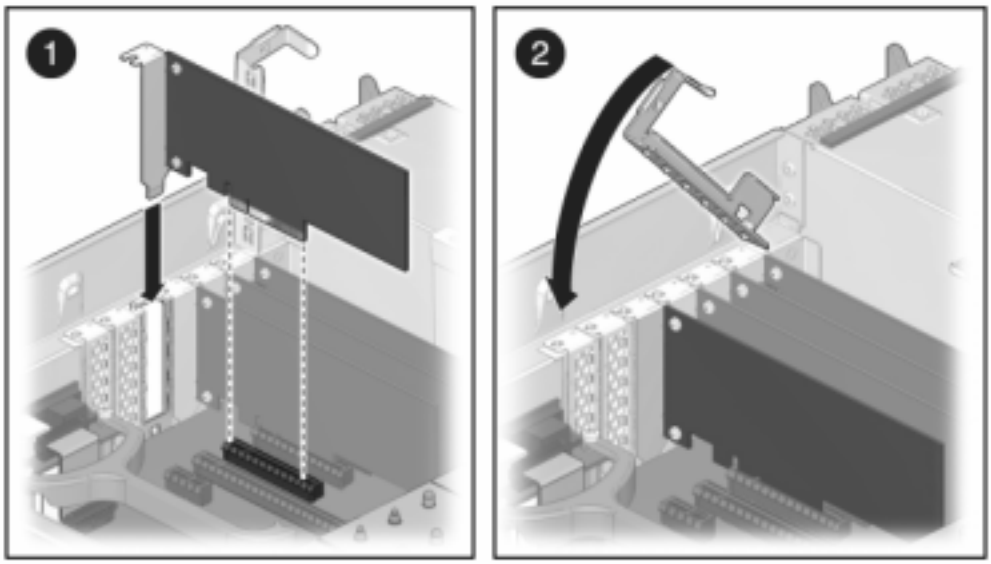

9. **Sun ZFS Storage 7120 or Sun ZFS Storage 7320:** Components may differ slightly between the Sun ZFS Storage 7120 and Sun ZFS Storage 7320; however, service procedures for each are identical. The illustration shows the Sun ZFS Storage 7320 controller. To install the top cover:

a. Place the top cover on the chassis so that it hangs over the rear of the storage controller by about an inch (2.5 cm), then slide the top cover forward until it seats.

b. Close the fan cover and engage the fan cover latches. The cover must be completely closed for the storage controller to power on.

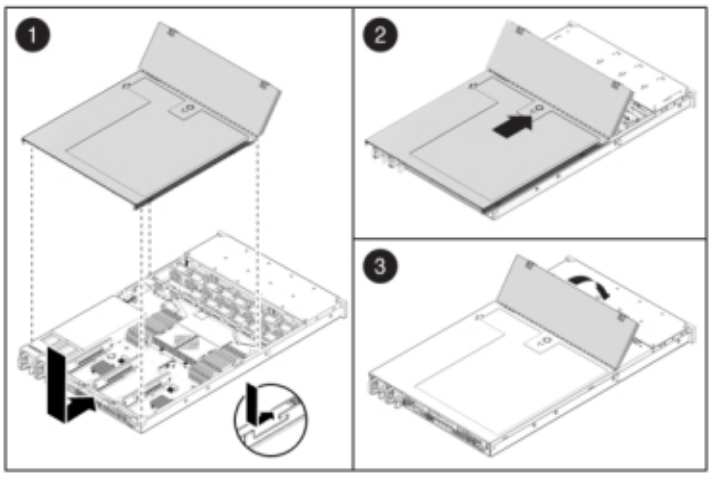

10. **Sun ZFS Storage 7420:** To install the top cover:

a. Place the top cover on the chassis [1] so that it is forward of the rear of the storage controller by about an inch (2.5 cm).

b. Slide the top cover toward the rear of the chassis [2] until it seats.

c. Press down on the cover with both hands until both latches engage.

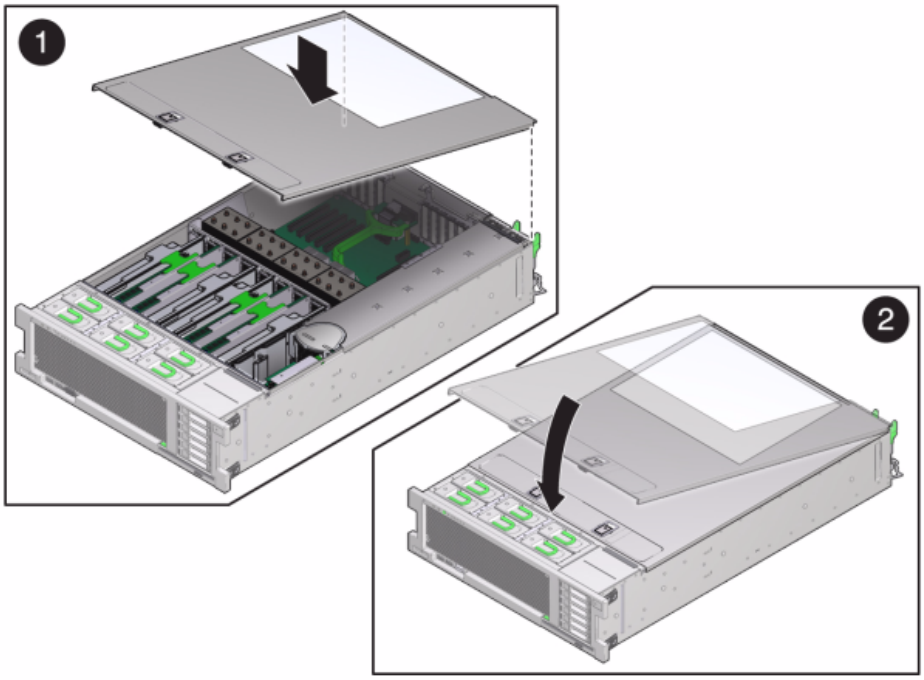

11. Push the release tabs on the side of each rail, and slowly push the storage controller into the rack, making sure no cables obstruct the path of the controller.
12. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
13. Connect the power cords to the power supplies.
14. Verify that standby power is on, indicated by the Power/OK status indicator flashing on the front panel about two minutes after the power cords are plugged in.
15. Use a pen or other pointed object to press and release the recessed Power button on the storage controller front panel. The Power/OK status indicator next to the Power button lights and remains lit.
16. Connect data cables to the PCIe card, routing them through the cable management arm.
17. From the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon on the controller. Then click **Slot** to verify the status of the new component. The online icon should be green.
18. If a fault does not self-clear, see Repairing Active Problems.
19. Install the disk shelf as described in Installing Disk Shelves in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*, and connect the expansion storage as described in Connecting to Attached Storage.

Replacing a Sun ZFS Storage 7x20 Battery

You might need to replace the battery if the Sun ZFS Storage 7x20 controller fails to maintain the proper time when powered off and disconnected from the network. You will need a small (No.1 flat-blade) non-metallic screwdriver or equivalent. Use the following procedure to replace a Sun ZFS Storage 7x20 battery.

1. Power off the storage controller as described in Powering Off the Controller.
2. Disconnect the AC power cords from the rear panel of the storage controller.

> **Caution:**
>
> Because 3.3 VDC standby power is always present in the system, you must unplug the power cords before accessing any cold-serviceable components.

3. Extend the controller from the rack as described in Extending the Controller from the Rack.
4. **Sun ZFS Storage 7120 or Sun ZFS Storage 7320:** To remove the top cover:
   a. Unlatch the fan module door, pull the two release tabs back, rotate the fan door to the open position, and hold it there.
   b. Press the top cover release button, and slide the top cover to the rear about a half-inch (1.3 cm).
   c. Lift up and remove the top cover.
5. **Sun ZFS Storage 7420:** To remove the top cover:
   a. Simultaneously lift both cover latches in an upward motion.

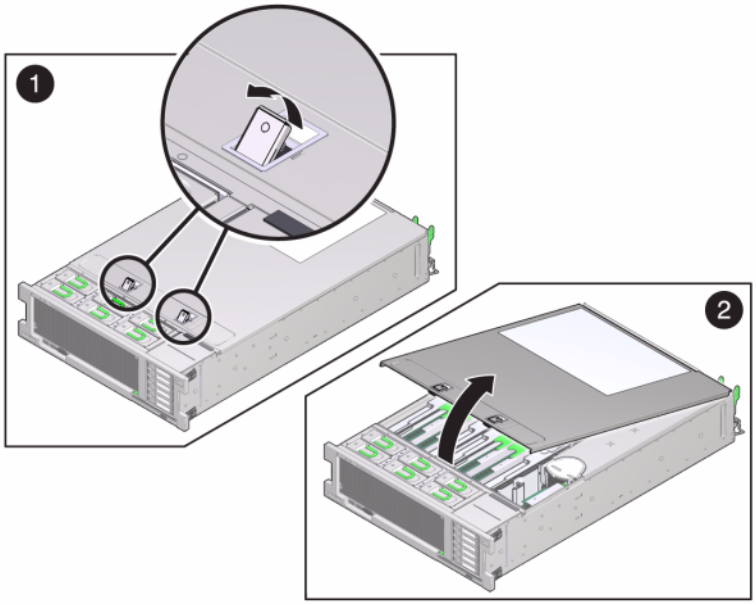

b. Lift up and remove the top cover.

6. Using a small, non-metallic screwdriver, press the latch, and remove the battery from the motherboard. The Sun ZFS Storage 7420 battery is shown here.

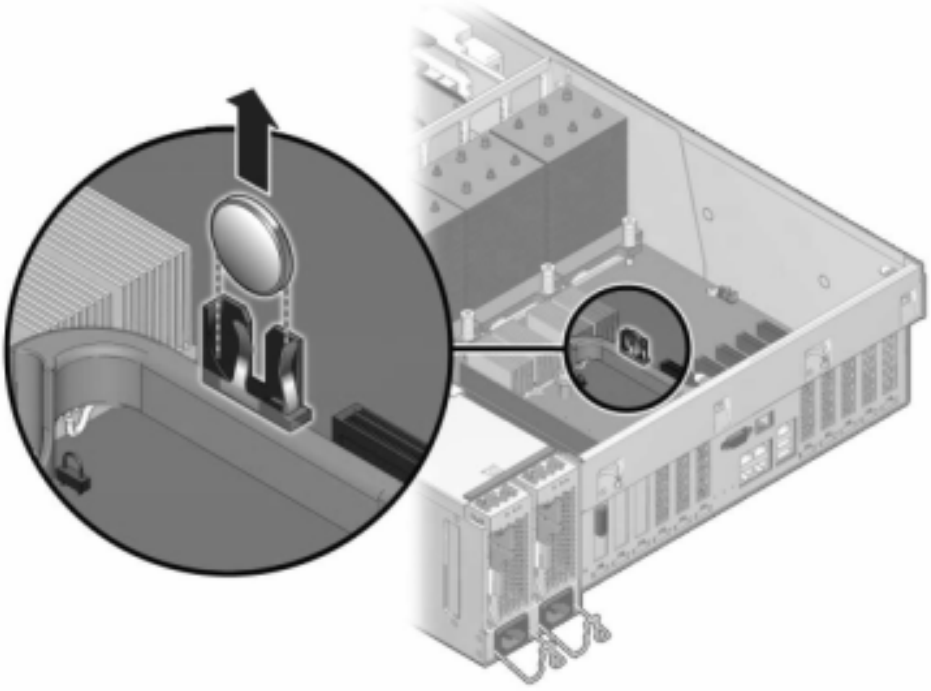

The following figure shows the Sun ZFS Storage 7120 battery.

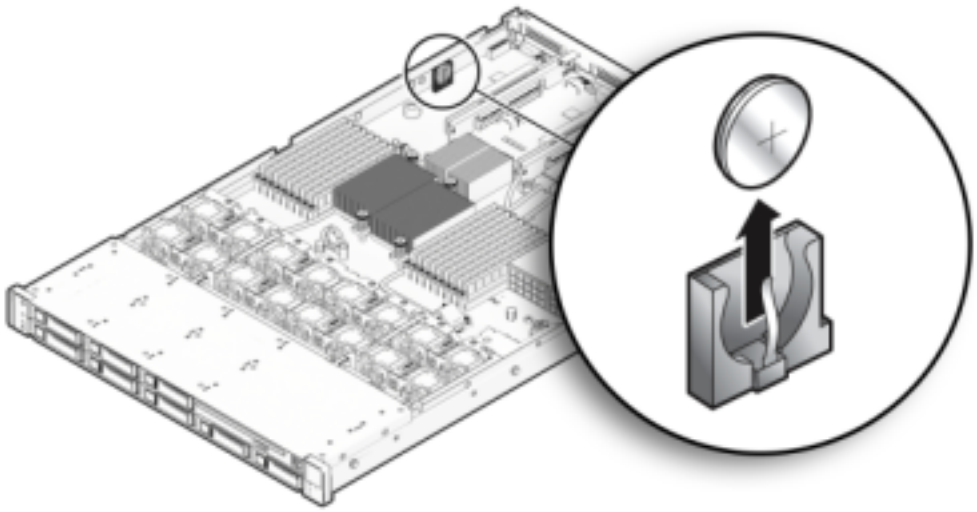

7. Press the replacement battery into the motherboard with the positive side (+) facing upward.
8. **Sun ZFS Storage 7120 or Sun ZFS Storage 7320:** Components may differ slightly between the Sun ZFS Storage 7120 and Sun ZFS Storage 7320; however, service procedures for each are identical. The illustration shows the Sun ZFS Storage 7320. To install the top cover:
   a. Place the top cover on the chassis so that it hangs over the rear of the storage controller by about an inch (2.5 cm), then slide the top cover forward until it seats.
   b. Close the fan cover and engage the fan cover latches. The cover must be completely closed for the storage controller to power on.

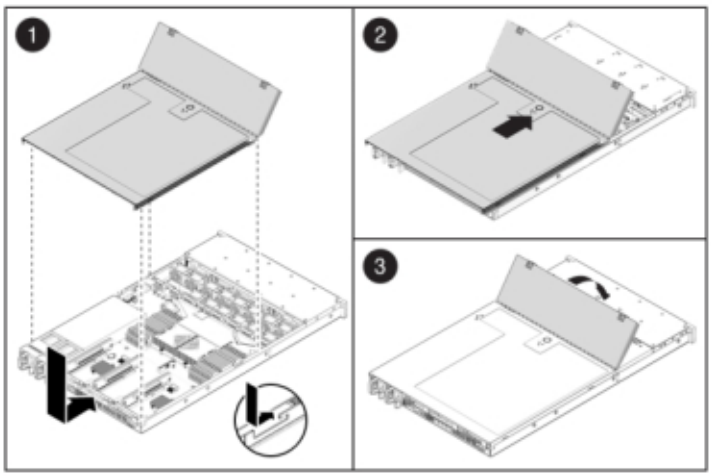

9. **Sun ZFS Storage 7420:** To install the top cover:
   a. Place the top cover on the chassis [1] so that it is forward of the rear of the storage controller by about an inch (2.5 cm).
   b. Slide the top cover toward the rear of the chassis [2] until it seats.
   c. Press down on the cover with both hands until both latches engage.

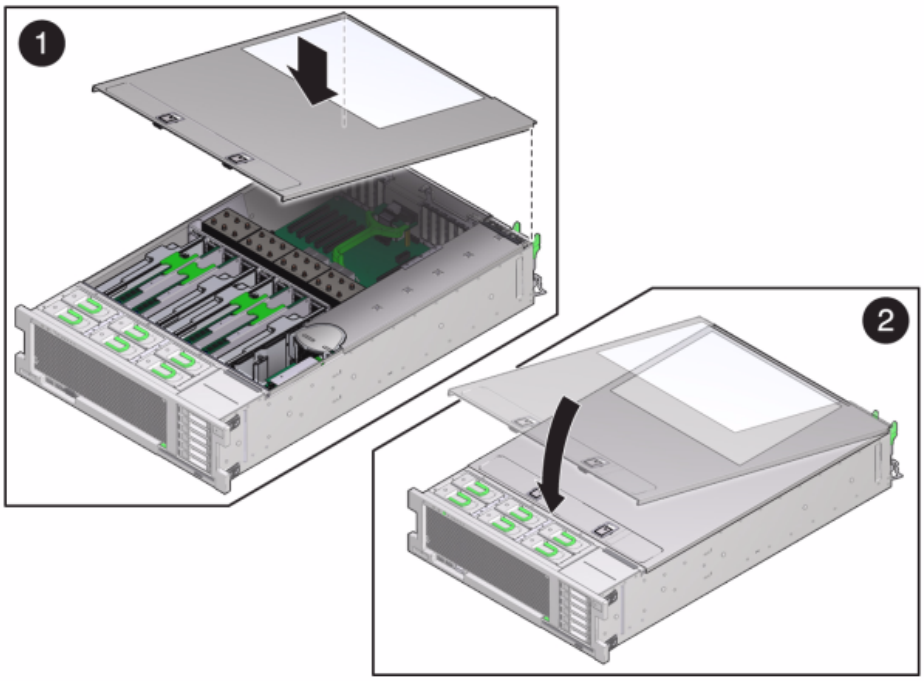

10. Push the release tabs on the side of each rail, and slowly push the storage controller into the rack, making sure no cables obstruct the path of the controller.
11. Reconnect all cables to their original ports, referring to the labels previously added as a guide. For cabling details, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.
12. Connect the power cords to the power supplies.
13. Verify that standby power is on, indicated by the Power/OK status indicator flashing on the front panel about two minutes after the power cords are plugged in.
14. Use a pen, or other pointed object, to press and release the recessed Power button on the storage controller front panel. The Power/OK status indicator next to the Power button lights and remains lit.
15. Connect data cables to the PCIe card, routing them through the cable management arm.
16. When the system has finished booting, log in, and set the time using the steps in the task Setting Clock Synchronization (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Sun ZFS Storage 7420 Controller Hardware Overview

This section describes the internal and external components of the Sun ZFS Storage 7420 controller. See the following for details:

- Sun ZFS Storage 7420 Configuration Options
- Sun ZFS Storage 7420 Front and Back Panel Components
- Sun ZFS Storage 7420 Specifications
- Sun ZFS Storage 7420 Internal Components

- Sun ZFS Storage 7420 PCIe I/O Cards
- Sun ZFS Storage 7420 Attached Storage

Sun ZFS Storage 7420 Configuration Options

The Sun ZFS Storage 7420 consists of a standalone storage controller, or two storage controllers in a high-availability cluster configuration, and 1 to 36 disk shelves. The following table describes the Sun ZFS Storage 7420 configuration options:

> **Note:**
>
> Sun ZFS Storage 7420 M2 controllers cannot be clustered with Sun ZFS Storage 7420 controllers.

Table 11-1 Sun ZFS Storage 7420 Controller Configuration Options

Product ID	Mktg Part Number	CPU	DIMMs	Read Cache	Boot Drive	Mfg Part Number
7420 M2	7107089	4x8-core, 2.0GHz	8GB, 16GB	1.6TB SAS	900GB SAS	7075466
7420 M2	7107090	4x10-core, 2.40GHz	8GB, 16GB	1.6TB SAS	900GB SAS	7075470
7420	7100566	4x8-core, 2.0GHz	8GB, 16GB	512GB SATA	500GB SATA	7014572
7420	7100568	4x10-core, 2.40GHz	8GB, 16GB	512GB SATA	500GB SATA	7014573

You can identify your configuration by looking at the product ID on the BUI **Maintenance** screen or by using the CLI `configuration version show` command. You can also check the physical label on the boot drive, as shown in the boot drive illustration in section Sun ZFS Storage 7320 Front and Back Panel Components.

Sun ZFS Storage 7420 Front and Back Panel Components

Sun ZFS Storage 7420 Front Panel Components - The Sun ZFS Storage 7420 controller drive slots and front panel components are shown in the following figure.

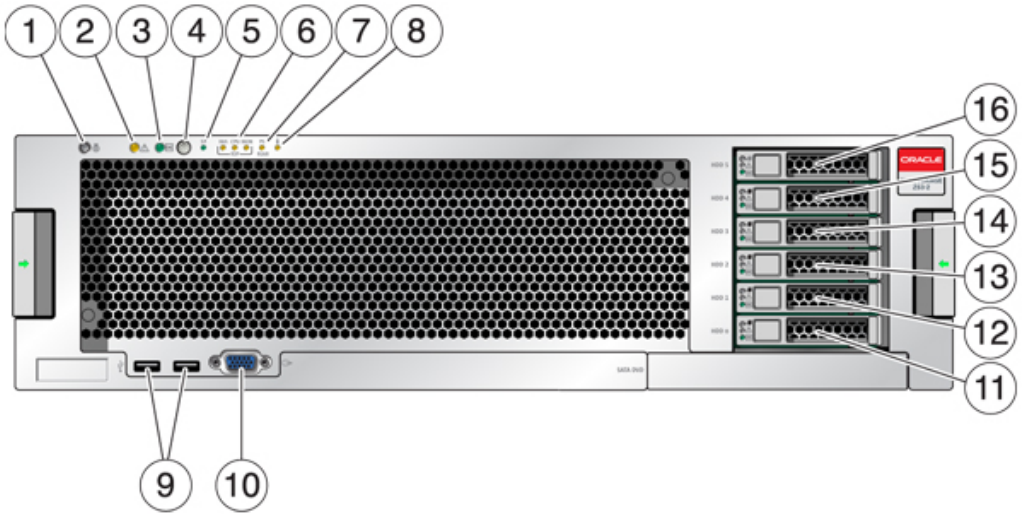

1 Locator LED and button (white)	**9** USB 2.0 Connectors
2 Service Required LED (amber)	**10** DB-15 video connector
3 Power/OK LED (green)	**11** Boot drive 0 (mirrored)
4 Power button	**12** Boot drive 1 (mirrored)
5 Service Processor (SP) OK LED (green)	**13** SSD 2 (optional)
6 Fan/CPU/Memory Service Required LED	**14** SSD 3 (optional)
7 Power Supply (PS) Service Required LED	**15** SSD 4 (optional)
8 Over Temperature Warning LED	**16** SSD 5 (optional)

Sun ZFS Storage 7420 Back Panel Components - The following graphic shows the rear panel of the controller. Base configuration HBAs are not depicted in this illustration.

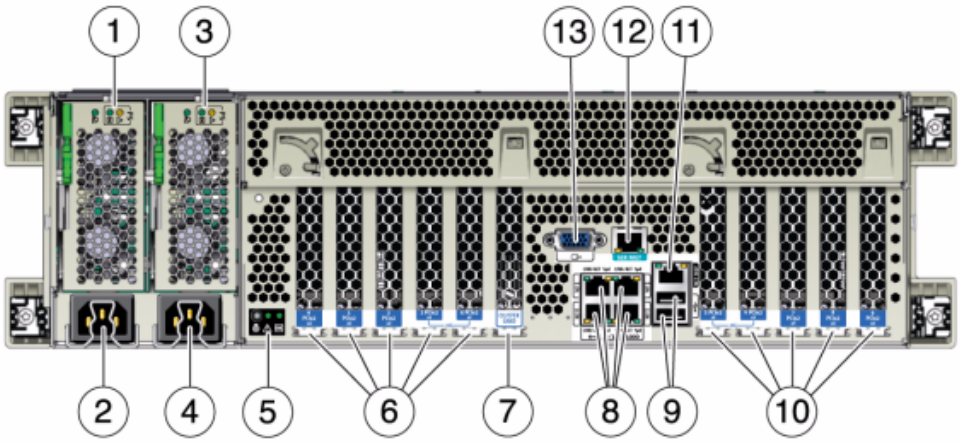

1 Power supply unit 0 status LEDs OK: green Power Supply Fail: amber AC OK: green	**8** Network (NET) 10/100/1000 ports: NET0-NET3
2 Power supply unit 0 AC inlet	**9** USB 2.0 ports
3 Power supply unit 1 status LEDs OK: green Power Supply Fail: amber AC OK: green	**10** PCIe slots 5-9
4 Power supply unit 1 AC inlet	**11** Network management (NET MGT) 10/100 BASE-T Ethernet port
5 System status LEDs Power: green Attention: amber Locate: white	**12** Serial management (SER MGT) port
6 PCIe slots 0-4	**13** DB-15 video connector
7 Cluster card slot	

Cluster Interface Card Compatibility - The Sun ZFS Storage 7420 controller contains a Version 2 cluster interface card. This card cannot be replaced with a Version 3 cluster interface card.

For information about how to connect cables to form a cluster, see Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

Ethernet Ports

There are four RJ-45 Gigabit Ethernet connectors (NET0, NET1, NET2, NET3) located on the Sun ZFS Storage 7420 motherboard that operate at 10/100/1000 Mbit/sec. These network interfaces must be configured before use.

Network Management Port

The Sun ZFS Storage 7420 network management connector (NET MGT) is an RJ-45 10/100 BASE-T Ethernet port and provides a LAN interface to the SP console.

Serial Management Port

The Sun ZFS Storage 7420 serial management connector (SER MGT) is an RJ-45 connector and provides a terminal connection to the SP console.

Sun ZFS Storage 7420 Specifications

Sun ZFS Storage 7420 Physical Specifications - The 3U chassis form factor dimensions are as follows:

Table 11-2 Sun ZFS Storage 7420 Controller Dimensions

Dimension	Measurement	Dimension	Measurement
Height	13.3 cm/5.25 in	Depth	70.6 cm/27.8 in
Width	43.7 cm/17.19 in	Weight	16.36 kg/96 lbs

Sun ZFS Storage 7420 Electrical Specifications - The following list shows the electrical specifications for Sun ZFS Storage 7420 controllers.

> **Note:**
>
> The power dissipation numbers listed are the maximum rated for the power supply and are not a rating of the actual power consumption of the appliance.

Input

- Nominal frequencies: 50/60Hz
- AC operating range: 100-127 VAC for 2 CPUs and 200-240 VAC for 2 or 4 CPUs
- Maximum current AC RMS: 12A @ 100 VAC / 12A @ 200 VAC

Power Dissipation

- Max power consumption: 1800 W
- Max heat output: 6143 BTU/hr
- Volt-Ampere rating: 1837 VA @ 240 VAC, 0.98 P.F.

Sun ZFS Storage 7420 Internal Components

The Sun ZFS Storage 7420 controller components are shown in the following figure and identified in the table.

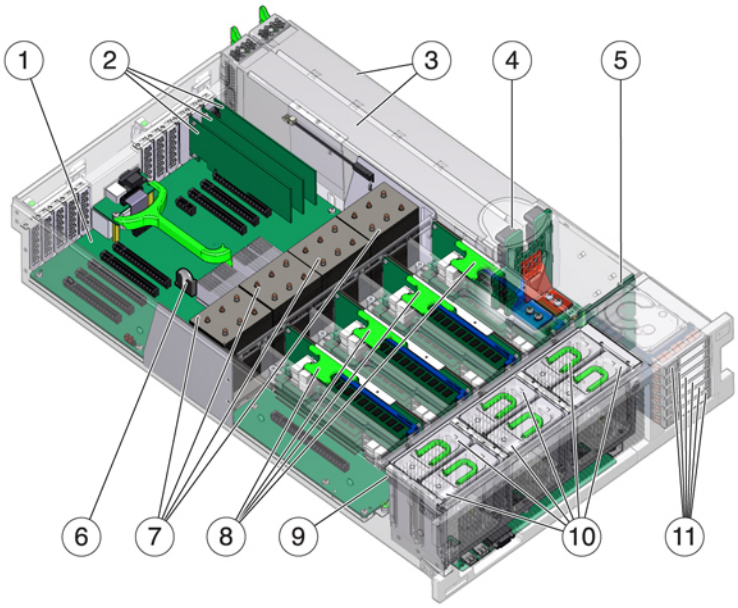

1 Motherboard	**7** CPUs and heatsinks
2 Low-profile PCIe cards	**8** Memory risers
3 Power supplies	**9** Fan board
4 Power supply backplane	**10** Fan modules
5 Drive backplane	**11** Boot drives and SSDs
6 System lithium battery	

Sun ZFS Storage 7420 Internal Boards - The Sun ZFS Storage 7420 controller chassis has the following boards installed. Field-replaceable units (FRUs) should only be replaced by trained Oracle service technicians.

- **Motherboard:** The motherboard is a FRU and includes CPU modules, slots for eight DIMM risers, memory control subsystems, and the service processor (SP) subsystem. The SP subsystem controls the host power and monitors host system events (power and environmental). The SP controller draws power from the host's 3.3V standby supply rail, which is available whenever the system is receiving AC input power, even when the system is turned off.
- **Power Distribution Board:** The power distribution board is a FRU and distributes main 12V power from the power supplies to the rest of the system. It is directly connected to the Vertical PDB card, and to the motherboard through a bus bar and ribbon cable. It also supports a top cover interlock ("kill") switch. In the storage controller, the power supplies connect to the power supply backplane which connects to the power distribution board.
- **Vertical PDB Card:** The vertical power distribution board, or Paddle Card is a FRU and serves as the interconnect between the power distribution board and the fan power boards, hard drive backplane, and I/O board.
- **Power Supply Backplane Card:** This board connects the power distribution board to power supplies 0 and 1.

- **Fan Power Boards:** The two fan power boards are FRUs and carry power to the storage controller fan modules. In addition, they contain fan module status LEDs and transfer I2C data for the fan modules.
- **Drive Backplane:** The six-drive backplane is a FRU and includes the connectors for the drives, as well as the interconnect for the I/O board, Power and Locator buttons, and system/component status LEDs. Each drive has an LED indicator for Power/Activity, Fault, and Locate.

Sun ZFS Storage 7420 System Drives - The Sun ZFS Storage 7420 M2 controller has two system boot drives in slots 0 and 1, configured as a mirrored pair. Up to four read-optimized cache solid state drives (SSDs) may fill slots 2 through 5, in order. In a Sun ZFS Storage 7420 M2 clustered configuration, the number of SSDs installed in each controller can vary.

Sun ZFS Storage 7420 controllers have two system boot drives in slots 0 and 1, configured as a mirrored pair. Up to four read-optimized cache SSDs, may fill slots 2 through 5, in order. In a Sun ZFS Storage 7420 clustered configuration, the number of SSDs installed in both controllers must match.

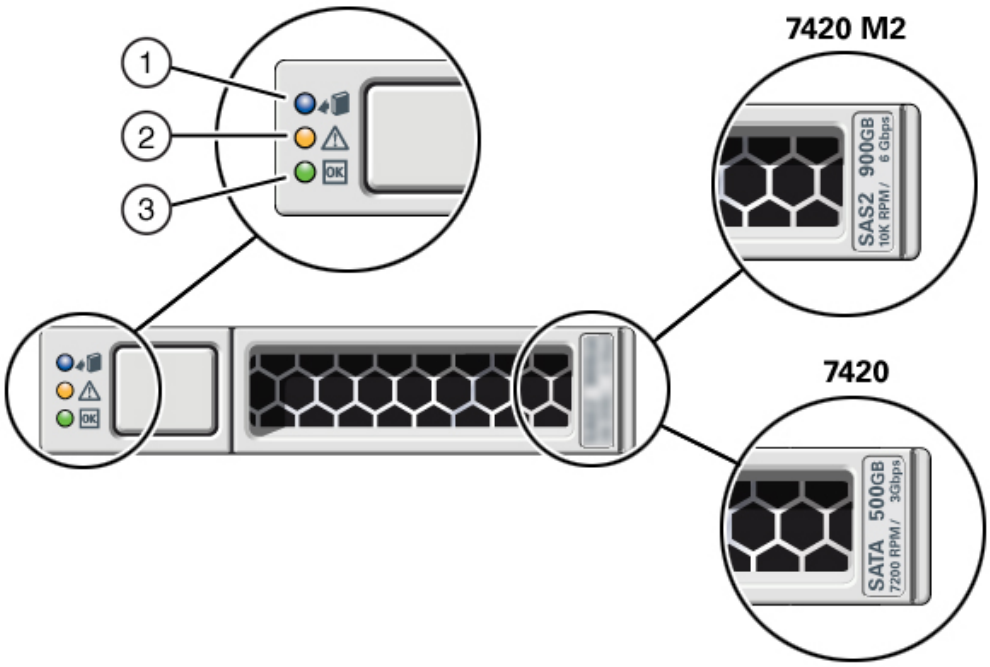

1 Ready to remove (blue)	**2** Service action required (amber)	**3** OK/Activity (green)

The Sun ZFS Storage 7420 controller supports the following configurations:

- Two, four, or eight 8GB DDR3 DIMMs installed on each riser, accommodating 128GB, 256GB, or 512GB of memory for 2.0Ghz CPUs.
- Four or eight 8GB DDR3 DIMMs installed on each riser, accommodating 256GB or 512GB of memory for 2.0GHz and 2.4Ghz CPUs.

- Four or eight 16GB DDR3 DIMMs installed on each riser, accommodating 512GB or 1TB of memory for 2.4Ghz CPUs.

> **Note:**
>
> For proper cooling, memory risers corresponding to unpopulated CPU sockets must have fillers installed.

Sun ZFS Storage 7420 CPU and Memory - Previous Sun ZFS Storage 7420 controller configurations support two or four (1.86GHz or 2.00GHz) CPUs, with two memory risers required by each CPU. Four or eight 4GB or 8GB DDR3 DIMMs are installed on each riser, accommodating up to 256GB of memory for two CPUs, or up to 512GB for four CPUs.

Sun ZFS Storage 7420 Controller Memory Risers

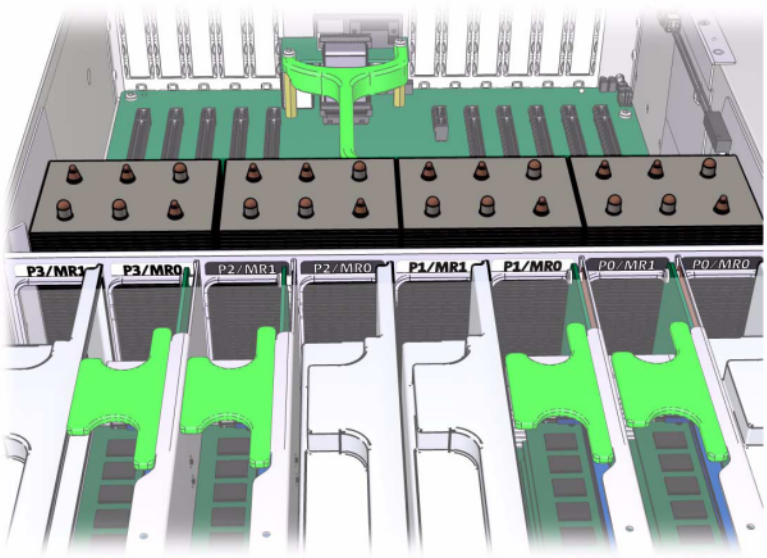

Refer to the service label on the cover for DIMM placement information. On every memory riser, slots D0, D2, D4, and D6 must be populated; optionally, slots D1, D3, D5, and D7 may be populated as a group on *all* installed memory risers. All DIMMs in the system *must be identical.*

Sun ZFS Storage 7420 Controller CPU and Memory

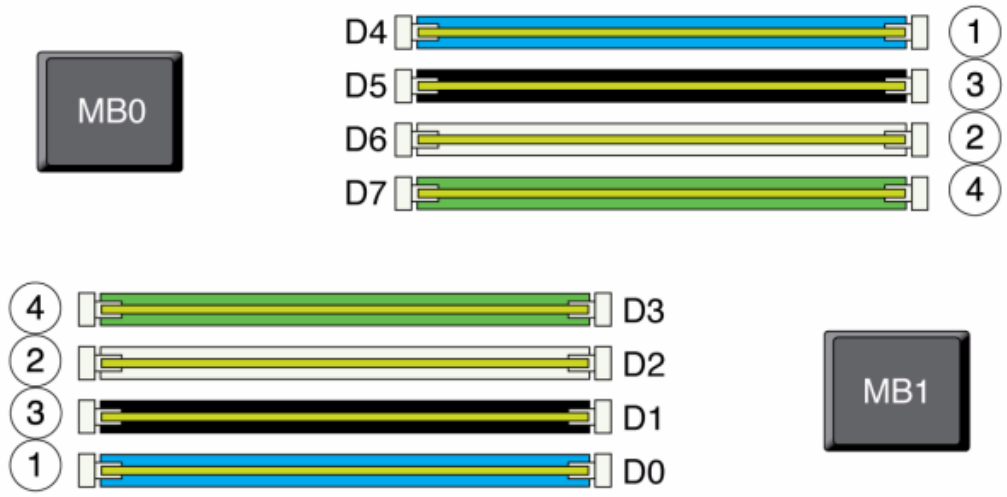

DIMM names in appliance logs and the **Maintenance: Hardware** view are displayed with the full name, such as `/SYS/MB/P0/D7`.

Sun ZFS Storage 7420 Fan Modules - The Fan Modules and Fan Module LEDs of the Sun ZFS Storage 7420 controller are shown in the following figure. The following LEDs are lit when a fan module fault is detected:

Sun ZFS Storage 7420 Controller Fan Module LEDs

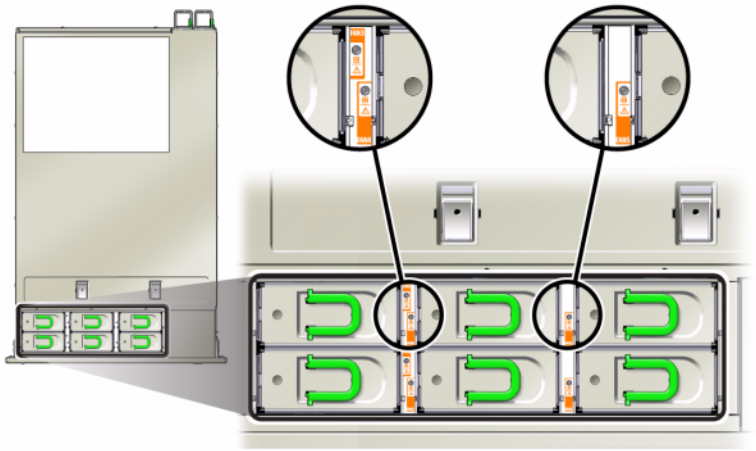

- Front and rear Service Action Required LEDs
- Fan Module Service Action Required (TOP) LED on the front of the controller
- Fan Fault LED on or adjacent to the faulty fan module

The system Overtemp LED might light if a fan fault causes an increase in system operating temperature.

Sun ZFS Storage 7420 PCIe I/O Cards

Sun ZFS Storage 7420 PCIe I/O Options - The Sun ZFS Storage 7420 controller has ten PCIe slots. This table describes base and optional PCIe card slots for standalone and cluster configurations.

> **Note:**
>
> The Sun ZFS Storage 7420 M2 reserves slot 0 for a SAS-2 internal HBA, and slots 1 and 8 for two SAS-2 external HBAs.

Table 11-3 Sun ZFS Storage 7420 Controller PCIe Options

Slot	Mfg Part Number	Description	Max	Note
0	7054001	8-port 6Gb/s SAS-2 internal HBA	1	Slot 0 of the 7420 M2 is reserved for the SAS-2 internal HBA
0	371-4325-02	8Gb 2-port FC HBA	6	Optional FC target or initiator (Backup) for the 7420 controller
0	7051223	2-port optical 10GbE NIC	6	Optional recommended front-end for the 7420 controller
0	7054739	4-port copper NIC	6	Optional recommended front-end for the 7420 controller
1	F375-3609-03 (7059331 for EU countries)	2-port SAS external HBA	6	Base configuration (2 minimum)
2	F375-3609-03 (7059331 for EU countries)	2-port SAS external HBA	6	Additional optional back-end
2	7054739	4-port copper NIC	6	Optional recommended front-end
2	371-4325-02	8Gb DP FC HBA	6	Optional FC target or initiator (Backup)
2	375-3606-03	CX1 InfiniBand HCA	4	Optional recommended front-end
2	7051223	2-port optical 10GE NIC	6	Optional recommended front-end
2	375-3696-01	CX2 InfiniBand HCA	4	Optional recommended front-end

Table 11-3 (Cont.) Sun ZFS Storage 7420 Controller PCIe Options

Slot	Mfg Part Number	Description	Max	Note
3	F375-3609-03 (7059331 for EU countries)	2-port SAS external HBA	6	Additional optional back-end
3	7054739	4-port copper NIC	6	Optional recommended front-end
3	371-4325-02	8Gb 2-port FC HBA	6	Optional FC target or initiator (Backup)
3	375-3606-03	CX1 InfiniBand HCA	4	Optional recommended front-end
3	7051223	2-port optical 10GE NIC	6	Optional recommended front-end
3	375-3696-01	CX2 InfiniBand HCA	4	Optional recommended front-end
4	7054739	4-port copper NIC	6	Optional recommended front-end
4	375-3606-03	CX1 InfiniBand HCA	4	Optional recommended front-end
4	7051223	2-port optical 10GE NIC	6	Optional recommended front-end
4	371-4325-02	8Gb 2-port FC HBA	6	Optional FC target or initiator (Backup)
4	375-3696-01	CX2 InfiniBand HCA	4	Optional recommended front-end
C	7056175	Cluster Controller 200	1	Cluster base configuration
5	7054739	4-port copper NIC	6	Optional recommended front-end
5	375-3606-03	CX1 InfiniBand HCA	4	Optional recommended front-end
5	7051223	2-port optical 10GE NIC	6	Optional recommended front-end
5	371-4325-02	8Gb 2-port FC HBA	6	Optional FC target or initiator (Backup)

Table 11-3 (Cont.) Sun ZFS Storage 7420 Controller PCIe Options

Slot	Mfg Part Number	Description	Max	Note
5	375-3696-01	CX2 InfiniBand HCA	4	Optional recommended front-end
6	F375-3609-03 (7059331 for EU countries)	2-port SAS external HBA	6	Additional optional back-end
6	7054739	4-port copper NIC	6	Optional recommended front-end
6	371-4325-02	8Gb 2-port FC HBA	6	Optional FC target or initiator (Backup)
6	375-3606-03	CX1 InfiniBand HCA	4	Optional recommended front-end
6	7051223	2-port optical 10GE NIC	6	Optional recommended front-end
6	375-3696-01	CX2 InfiniBand HCA	4	Optional recommended front-end
7	F375-3609-03 (7059331 for EU countries)	DP SAS external HBA	6	Additional optional back-end
7	7054739	4-port copper NIC	6	Optional recommended front-end
7	371-4325-02	8Gb 2-port FC HBA	6	Optional FC target or initiator (Backup)
7	375-3606-03	CX1 InfiniBand HCA	4	Optional recommended front-end
7	7051223	2-port optical 10GE NIC	6	Optional recommended front-end
7	375-3696-01	CX2 InfiniBand HCA	4	Optional recommended front-end
8	F375-3609-03 (7059331 for EU countries)	2-port SAS external HBA	6	Base configuration (2 minimum)
9	371-4325-02	8Gb 2-port FC HBA	6	Optional FC target or initiator (Backup)
9	7051223	2-port optical 10GbE NIC	6	Optional recommended front-end

Table 11-3 (Cont.) Sun ZFS Storage 7420 Controller PCIe Options

Slot	Mfg Part Number	Description	Max	Note
9	7054739	4-port copper NIC	6	Optional recommended front-end

Sun ZFS Storage 7420 NIC/HBA Options - This table describes NIC/HBA PCIe card options for standalone and cluster Sun ZFS Storage 7420 configurations.

Table 11-4 Sun ZFS Storage 7420 Controller NIC/HBA Options

Mktg Part Number	Description	Mfg Part Number
SGX-SAS6-INT-Z	8-port 6Gb/s SAS-2 internal HBA	7054001
SG-XPCIE2FC-QF8-Z	2-port 8Gb FC HBA	371-4325-02
SG-XPCIESAS-GEN2-Z	2-port SAS (x4) external HBA	F375-3609-03
7105394	2-port SAS (x4) external HBA for EU countries	7059331
X4446A-Z-N	4-port GigE UTP	7054739
X4242A	2-port CX2 4xQDR, HCA	594-6776-01
X1109A-Z	2-port 10GbE SFP+ NIC	7051223
X2129A	XCVRm 850NM, 1/10GPS, Short Reach, SFP	7015839
X5562A-Z	10GbE/1GbE SFP+ Transceiver, LR	594-6689-01

Sun ZFS Storage 7420 PCIe Slot Order - When installing additional PCIe cards to your base Sun ZFS Storage 7420 configuration, cards must be added in a specific order.

> **Note:**
>
> The Sun ZFS Storage 7420 M2 reserves slot 0 for a SAS-2 internal HBA, and slots 1 and 8 for two SAS-2 external HBAs (see Sun ZFS Storage 7420 PCIe I/O Cards).

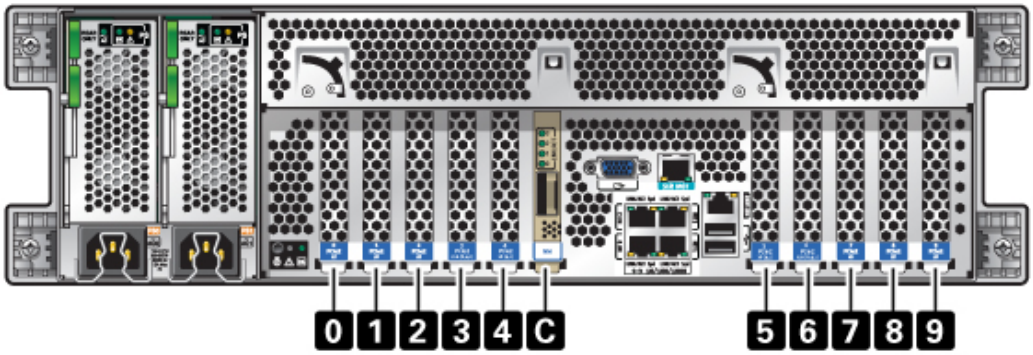

Install optional PCIe cards in the following order:

Table 11-5 Sun ZFS Storage 7420 Controller Optional PCIe Cards and Slot Order

PCIe Card	Slot Order
1. SAS-2 external HBA	Slot 2, 7, 3, 6
2. 8Gb FC HBA	Slot 9 Additional FC HBAs are installed last
3. InfiniBand HCAs	Slot 7, 2, 6, 3, 5, 4
4. 10Gb Ethernet	Slot 7, 2, 6, 3, 5, 4, 0 (not available for 7420 M2), 9
5. Quad 1Gb Ethernet	Slot 7, 2, 6, 3, 5, 4, 0 (not available for 7420 M2), 9
6. Remaining 8Gb FC HBAs	Slot 7, 2, 6, 3, 5, 4, 0 (not available for 7420 M2)

Sun ZFS Storage 7420 Attached Storage

Primary storage for Sun ZFS Storage 7420 controllers is provided by one to six chains of attached storage. Each chain consists of up to six HDD-only or SSD-capable disk shelves, combined in any order. For more information, see Connecting to Attached Storage. See also Disk Shelf Hardware Overview for component specifications and diagrams.

Sun ZFS Storage 7320 Controller Hardware Overview

This section describes the internal and external components of the Sun ZFS Storage 7320 controller. See the following for details:

- Sun ZFS Storage 7320 Configuration Options
- Sun ZFS Storage 7320 Front and Back Panel Components
- Sun ZFS Storage 7320 Electrical Specifications
- Sun ZFS Storage 7320 Internal Components
- Sun ZFS Storage 7320 Attached Storage

Sun ZFS Storage 7320 Configuration Options

Sun ZFS Storage 7320 Chassis Overview - The Sun ZFS Storage 7320 consists of either a standalone storage controller or two storage controllers in a high availability cluster configuration. Both the standalone and clustered configuration support one to six disk shelves.

The Sun ZFS Storage 7320 controller base configuration includes two CPUs, built-in 4 x 1Gb/s front-end GigE ports, redundant power supplies, NIC options for expanded front-end support, tape backup, InfiniBand, and dual port SAS HBA for storage expansion.

The CPUs are Intel Xeon 5600 series, 2.40GHz, 80W, with 4 core processors. The standard memory configuration is 96GB, 6 x 8GB DDR3-1333 low voltage (LV) DIMMs per CPU. Memory can be upgraded to 144GB using 9 x 8GB DDR3-1333 LV DIMMs per CPU (for a total of 18 x 8GB for two CPUs). Earlier versions of the Sun ZFS Storage 7320 controller included 24GB (base), 48GB, or 72GB memory options. The clustered configuration simply uses two controllers and a cluster card in each controller for a heartbeat connection between the controllers.

All user-accessible storage is provided by one to six disk shelves that are external to the controller(s). RAID functions are managed by software. Solid state SAS-1 drives (Sun ZFS Storage 7320 SAS-2) are used for a high-performance write cache (known as LogZilla) or ZFS intent log (ZIL) devices, and are used in place of up to 4 of the 24 drives in a disk shelf. The remaining 20 drives are available for storage.

The Sun ZFS Storage 7320 is a SAS-2 (Serial Attached SCSI 2.0) device that consists of an HBA, disk shelf, and disks. The SAS-2 storage fabric supports a greater number of targets, greater bandwidth, higher reliability, and bigger scale than the SAS-1 fabric.

Standalone Controller Configuration - The standalone controller base configuration is 96GB RAM, 2x2.4GHz Quad-Core processors, one external SAS HBA, and four 10/100/1000 Ethernet ports.

The following table describes base configurations for the Sun ZFS Storage 7320 controller.

Table 11-6 Sun ZFS Storage 7320 Controller Base Configurations

Mktg Part Number	Description	Mfg Part Number
TA7320-24A	S7320, 2xCPU, 24GB, Single	597-1060-01
7104054	S7320, 2xCPU, 96GB, Single	7045900
TA7320-24A-HA	S7320, 2xCPU, 24GB, Cluster	597-1061-01
7104055	S7320, 2xCPU, 96GB, Cluster	7045903

Following are the PCIe configuration options for a standalone controller. All PCIe cards are low profile and must be fitted with low-profile mounting brackets.

Table 11-7 Sun ZFS Storage 7320 Standalone Controller PCIe Configuration Options

Slot	Type	Part Number	Vendor Part	Description	Note
0	PCIe	375-3617-01	Intel Niantic	DP Optical 10GE NIC	Optional Recommended Front-end
0	PCIe	375-3696-01	Mellanox	InfiniBand HCA	Optional Recommended Front-end
0	PCIe	375-3606-03	MHJH29-XTC	InfiniBand HCA	Optional Recommended Front-end
0	PCIe	375-3481-01	Intel EXPI9404PT	QP Copper NIC	Optional Recommended Front-end
0	PCIe	371-4325-01	QLogic	8Gb DP FC HBA	Optional FC Target or Initiator (Backup)
1	PCIe	375-3617-01	Intel Niantic	DP Optical 10GE NIC	Optional Recommended Front-end
1	PCIe	375-3696-01	Mellanox	InfiniBand HCA	Optional Recommended Front-end
1	PCIe	375-3606-03	MHJH29-XTC	InfiniBand HCA	Optional Recommended Front-end
1	PCIe	375-3481-01	Intel EXPI9404PT	QP Copper NIC	Optional Recommended Front-end
1	PCIe	371-4325-01	QLogic	8Gb DP FC HBA	Optional FC Target or Initiator (Backup)
2	PCIe	375-3609-03	Sun Thebe	External SAS HBA	Base Configuration

Sun ZFS Storage 7320 Cluster Configurations - The Sun ZFS Storage 7320 cluster base configuration is 96GB RAM, 2x2.4GHz Quad-Core processors, one external SAS HBA, and four 10/100/1000 Ethernet ports, and a Cluster card. The Sun Storage 7420C Cluster Upgrade Kit (XOPT 594-4680-01) contains two cluster cards with cables for converting two Sun ZFS Storage 7320 or two Sun ZFS Storage 7420 controllers to a cluster.

The following options are available for clustered storage controllers.

> **Note:**
>
> When you cluster a Sun ZFS Storage 7320, you must identically configure the cards in both of the clustered storage controllers, and you must identically configure all optional NIC/HBA cards used in the clustered storage controllers in both chassis.

Table 11-8 Sun ZFS Storage 7320 Clustered Controllers Configurations

Slot	Type	Part Number	Vendor Part	Description	Note
0	PCIe	375-3617-01	Intel Niantic	DP Optical 10GE NIC	Optional Recommended Front-end
0	PCIe	375-3696-01	Mellanox	InfiniBand HCA	Optional Recommended Front-end

Table 11-8 (Cont.) Sun ZFS Storage 7320 Clustered Controllers Configurations

Slot	Type	Part Number	Vendor Part	Description	Note
0	PCIe	375-3606-03	MHJH29-XTC	InfiniBand HCA	Optional Recommended Front-end
0	PCIe	375-3481-01	Intel EXPI9404PT	QP Copper NIC	Optional Recommended Front-end
0	PCIe	371-4325-01	QLogic	8Gb DP FC HBA	Optional FC Target or Initiator (Backup)
1	PCIe	542-0298-01	Sun	Fishworks Cluster Card 2	Cluster Base Configuration
2	PCIe	375-3609-03	Sun Thebe	External SAS HBA	Cluster Base Configuration

Sun ZFS Storage 7320 Front and Back Panel Components

Sun ZFS Storage 7320 Front Panel Components - The following figure and legend identify the front panel LEDs.

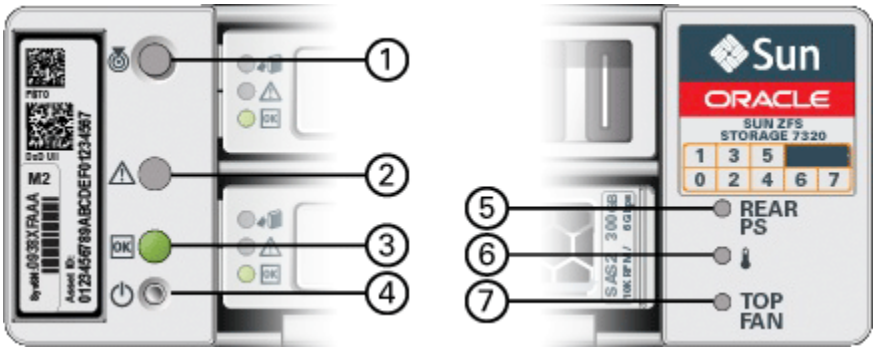

1 Locate Button/LED	**5** Rear Power Supply
2 Service Required LED (amber)	**6** System Overtemperature LED
3 Power/OK LED (green)	**7** Top Fan
4 Power Button	

The following figure and legend identify the Sun ZFS Storage 7320 front panel drive locations. Two mirrored hard disk drives (HDDs) that store the operating system reside in slots 0 and 1. Up to four read-optimized cache solid state drives (SSDs) fill slots 2 through 5, in order. Slots 6 and 7 are empty and must contain drive fillers.

Table 11-9 Sun ZFS Storage 7320 Controller Front Panel Drive Locations

Disk Drive Location	Disk Drive Location	Disk Drive Location	Disk Drive Location	Disk Drive Location
HDD1	HDD3	HDD5	(no drive)	(no drive)
HDD0	HDD2	HDD4	HDD6	HDD7

Sun ZFS Storage 7320 Back Panel Components - Following is an illustration of the Sun ZFS Storage 7320 storage controller rear panel. The Sun 375-3609 belongs in slot 2, cannot be installed in any other slots, and a second is not offered as an option.

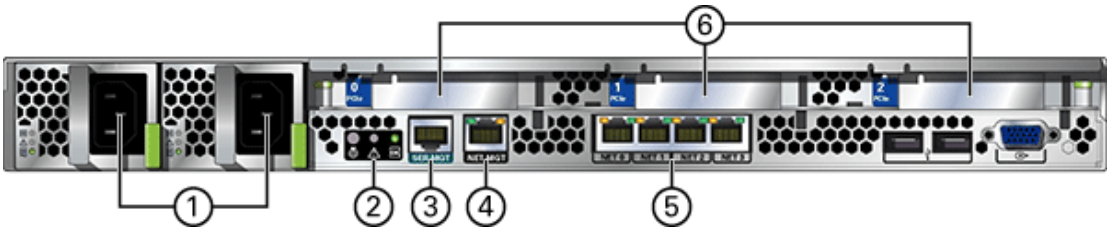

1 Power supplies	**4** Network management (NET MGT) 10/100 BASE-T Ethernet port
2 SC summary status LEDs	**5** Ethernet ports
3 Serial management port	**6** PCIe slots

Cluster Interface Card Compatibility - The Sun ZFS Storage 7320 controller contains a Version 2 cluster interface card. This card cannot be replaced with a Version 3 cluster interface card.

For information about how to connect cables to form a cluster, see Connecting Cluster Cables in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

Ethernet Ports

There are four RJ-45 Gigabit Ethernet connectors (NET0, NET1, NET2, NET3) located on the Sun ZFS Storage 7320 motherboard that operate at 10/100/1000 Mbit/sec. These network interfaces must be configured before use.

Network Management Port

The network management connector (NET MGT) is an RJ-45 10/100 BASE-T Ethernet port on the Sun ZFS Storage 7320 motherboard and provides an alternate terminal interface to the SP console.

Serial Management Port

The Sun ZFS Storage 7320 serial management connector (SERIAL MGT) is an RJ-45 connector and is a terminal connection to the SP console.

Sun ZFS Storage 7320 Electrical Specifications

The following list shows the electrical specifications for the Sun ZFS Storage 7320 controller.

> **Note:**
>
> The power dissipation numbers listed are the maximum rated power numbers for the power supply. The numbers are not a rating of the actual power consumption of the appliance.

Connectors

- Two C13 connectors which work on 110-220v outlets

Input

- Nominal frequencies: 50/60Hz
- Nominal voltage range: 100-120/200-240 VAC
- Maximum current AC RMS: 9.0 amps Max
- AC operating range: 90-264 VAC

Output

- 3.3 VDC STBY: 3.6A
- +12 VDC: 62.3A

Power Dissipation

- Max power consumption: 873 W
- Max heat output: 2977 BTU/hr
- Volt-Ampere rating: 891 VA @ 240 VAC, 0.98P.F.

Sun ZFS Storage 7320 Internal Components

The Sun ZFS Storage 7320 storage controller chassis has the following boards installed.

> **Note:**
>
> Field-replaceable units (FRUs) should only be replaced by trained Oracle service technicians.

- **PCIe Risers:** The storage controller contains three PCIe risers that are customer-replaceable units (CRUs) and are attached to the rear of the motherboard. Each riser supports one PCIe card.
- **Motherboard:** The motherboard is a FRU and includes CPU modules, slots for 18 DIMMs, memory control subsystems, and the service processor (SP) subsystem. The SP subsystem controls the host power and monitors host system events (power and environmental). The SP controller draws power from the host 3.3V standby supply rail, which is available whenever the system is receiving AC input power, even when the appliance is turned off.
- **Power Distribution Board:** The power distribution board is a FRU and distributes main 12V power from the power supplies to the rest of the storage controller. It is directly connected to the paddle board and to the motherboard through a bus bar and ribbon cable. It also supports a top cover interlock *kill* switch.
- **Paddle Board:** The paddle board is a FRU and serves as the interconnect between the power distribution board and the fan power boards, hard drive backplane, and I/O board.
- **Fan Board:** The fan boards are FRUs and carry power to the storage controller fan modules. In addition, they contain fan module status LEDs and transfer I2C data for the fan modules.
- **Disk Backplane:** The hard drive backplane is a FRU and includes the connectors for the hard disk drives, as well as the interconnect for the I/O board, Power and Locator buttons, and system/component status LEDs. The storage controller has an eight-disk backplane. Each drive has an LED indicator for Power/Activity, Fault, and OK-to-Remove (not supported).

The following list contains the replaceable system boards for the Sun ZFS Storage 7320 storage controller.

Table 11-10 Sun ZFS Storage 7320 Controller Replaceable System Boards

Part Number	Description	FRU/CRU
F541-2883-01	X8 PCIe Riser Card 1U	CRU
F541-2885-01	X16 PCIe Riser Card 1U	CRU
F541-4081-01	RoHS Motherboard and Tray	FRU
F511-1489-01	DB, Power Distribution Board	FRU
F511-1548-01	PCB, 8 Disk 1U Backplane	FRU
F541-4275-02	PCBA, Connector Board, 1U	FRU

Sun ZFS Storage 7320 Replaceable Components - The following list contains all of the replaceable power distribution, disk, and fan module components of Sun ZFS Storage 7320. Note that power supplies, disks, and fan modules are hot-pluggable on the storage controller.

Table 11-11 Sun ZFS Storage 7320 Controller Replaceable Components

Part Number	Description	FRU/CRU
F300-2233-02	RoHS 760W Power Supply	CRU
F541-2075-04	Buss Bar Power, 1U	FRU
F542-0184-01	DR, 3Gb SATA	CRU
F542-0330-01	2.5" Read Cache SSD	CRU
F541-276-01	ASSY,FAN Module	CRU
F541-4274-02	Fan Board (1U)	FRU

The following list contains the replaceable CPU and memory components of Sun ZFS Storage 7320.

Table 11-12 Sun ZFS Storage 7320 Controller Replaceable CPU and Memory Components

Part Number	Description	FRU/CRU
F371-4966-01	DIMM, 8GB, DDR3, 2RX4, 13	CRU
F371-4885-01	Intel E5620, 2.40G	FRU

Sun ZFS Storage 7320 CPU and Memory Components - The Sun ZFS Storage 7320 storage controller motherboard has 18 slots in two groups that hold industry-standard DDR3 DIMM memory cards. All sockets must be occupied by either a filler or a DDR3 DIMM.

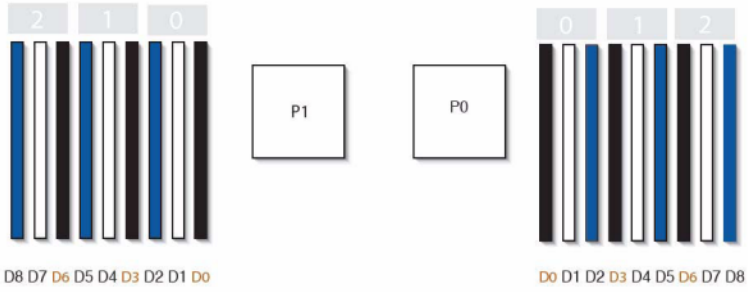

Sun ZFS Storage 7320 Power Distribution and Fan Module Components - The following figure and legend identify the Power Distribution/Fan Module components of the Sun ZFS Storage 7320 storage controller.

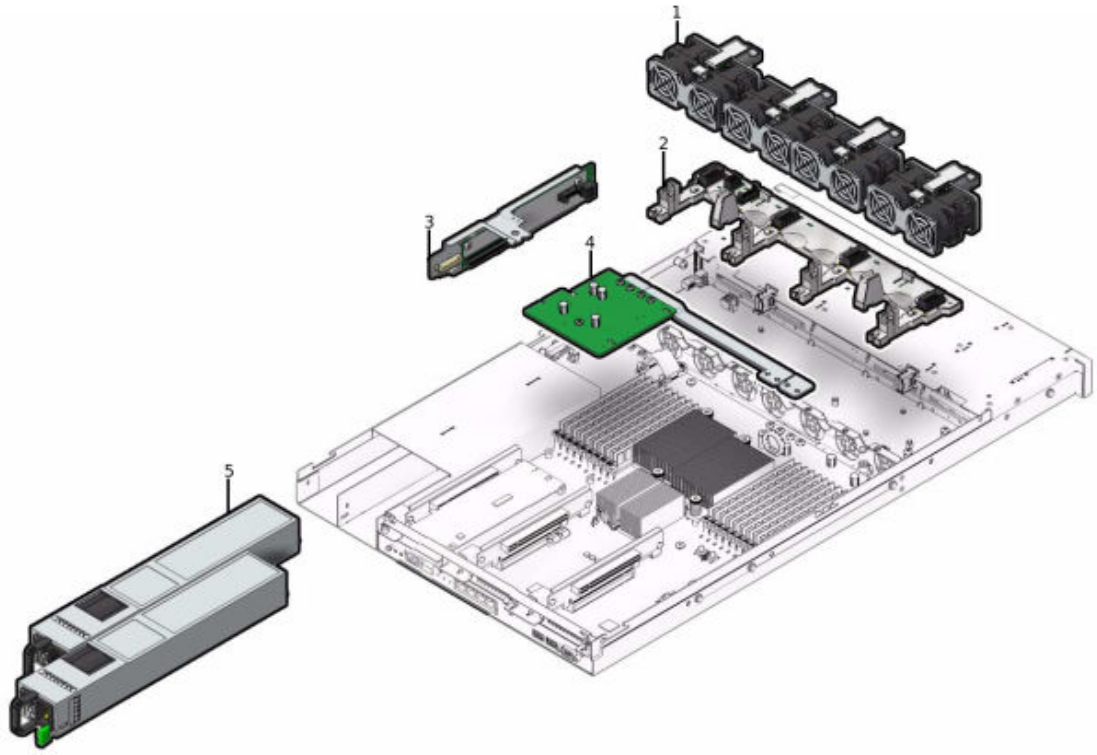

1 Fan Modules	**4** Power Distribution/Bus Bar Assembly
2 Fan Board	**5** Power Supplies
3 Paddle Board	

Sun ZFS Storage 7320 I/O Components - The following figure and legend identify the I/O components of the Sun ZFS Storage 7320 storage controller.

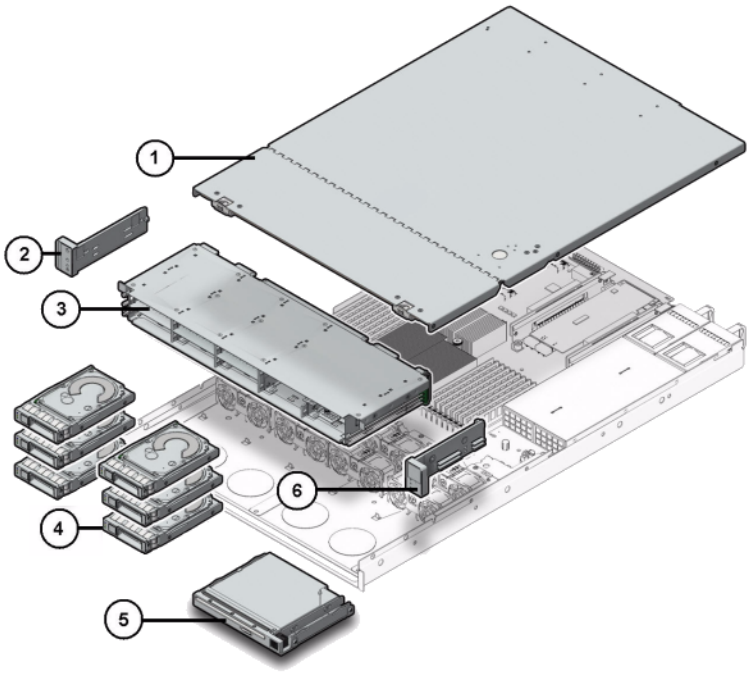

1 Top cover	**2** Left Control Panel Light Pipe Assembly
3 Drive Cage	**4** Solid State Drives
5 blank/USB Module	**6** Right Control Panel Light Pipe Assembly

Sun ZFS Storage 7320 PCIe Cards and Risers - Following is the complete list of replaceable PCIe cards for the Sun ZFS Storage 7320 system.

Table 11-13 Sun ZFS Storage 7320 Controller Replaceable PCIe Cards

Part Number	Description	FRU/CRU
F371-4325-01	8Gb FC HBA (PCIe)	CRU
F375-3609-02	PCA, SAS 6GBS 8 Port (PCIe)	CRU
F375-3606-03	Dual Port (x4) IB HCA (PCIe)	CRU
F375-3696-01	Dual Port CX2 4XQDR (PCIe)	CRU
F375-3617-01	2X10GbE SFP+, X8 (PCIe)	CRU
F375-3481-01	NIC Card Quad Port 1GigE Cu (PCIe)	CRU
F511-1496-04	Sun Fishworks Cluster Controller 200 (PCIe)	FRU

Sun ZFS Storage 7320 Cables - The following list contains the replaceable cables for the Sun ZFS Storage 7320 controller.

Table 11-14 Sun ZFS Storage 7320 Controller Replaceable Cables

Part Number	Description	FRU/CRU
F530-4228-01	Cable, Mini SAS	FRU (internal)

Table 11-14 (Cont.) Sun ZFS Storage 7320 Controller Replaceable Cables

Part Number	Description	FRU/CRU
F530-3927-01	FRU,CBL,PDB,MB,1U+2U,RIBBON	FRU (internal)
F530-4431-01	Cable, Fan data	FRU (internal)
F530-4417-01	FRU Cable, Fan paddle	FRU (internal)
F530-3880-01	Cable, Assembly, Ethernet, Shielded, RJ45-RJ45, 6m	CRU (external)
F530-3883-01	FRU,2M,4X Mini SAS Cable	CRU (external)

Sun ZFS Storage 7320 Attached Storage

The Sun ZFS Storage 7320 standalone and cluster controller configurations allow a single chain of one to six disk shelves. Any combination of disk-only and Logzilla-capable shelves may be combined within the chain in any order. The cabling configurations are unchanged. Half-populated shelf configurations are not supported. For more information, see Connecting to Attached Storage. See also Disk Shelf Hardware Overview for component specifications and diagrams.

Sun ZFS Storage 7120 Controller Hardware Overview

This section describes the internal and external components of the Sun ZFS Storage 7120 controller. See the following for details:

- Sun ZFS Storage 7120 Base Configuration
- Sun ZFS Storage 7120 Chassis Overview
- Sun ZFS Storage 7120 Front Panel
- Sun ZFS Storage 7120 Rear Panel
- Sun ZFS Storage 7120 Physical Specifications
- Sun ZFS Storage 7120 Electrical Specifications
- Sun ZFS Storage 7120 Internal Boards
- Sun ZFS Storage 7120 CPU and Memory
- Sun ZFS Storage 7120 Power Distribution, Fan Module and Disk Components
- Sun ZFS Storage 7120 I/O Components
- Sun ZFS Storage 7120 NIC/HBA Options
- Sun ZFS Storage 7120 PCIe Options
- Sun ZFS Storage 7120 Connectors
- Sun ZFS Storage 7120 Cables
- Sun ZFS Storage 7120 Attached Storage

Sun ZFS Storage 7120 Base Configuration

The following table shows the configuration options for a Sun ZFS Storage 7120 controller. All PCIe cards are low-profile, and must be fitted with low-profile mounting brackets.

This table describes base configurations for the Sun ZFS Storage 7120 controller with Aura Logzilla.

Table 11-15 Sun ZFS Storage 7120 Standalone Controller Base Configurations for Aura Logzilla

Mktg Part Number	Description	Mfg Part Number
TA7120-12TB	S7120, 1xCPU, 24GB, 12TB	597-0754-01
TA7120-24TB	S7120, 1xCPU, 24GB, 24TB	597-0755-01

The following table describes base configurations for the Sun ZFS Storage 7120 controller with Logzilla 3.5" SSD.

Table 11-16 Sun ZFS Storage 7120 Standalone Controller Base Configurations for Logzilla 3.5" SSD

Mktg Part Number	Description	Mfg Part Number
7101282	S7120, 1xCPU, 24GB, 3.3TB	7014523
7101284	S7120, 1xCPU, 24GB, 6.6TB	7014525

Sun ZFS Storage 7120 Chassis Overview

Sun ZFS Storage 7120 is an enterprise-class two-socket rackmount x64 system powered by the Intel Xeon processor. It packs high performance and room for growth with four PCIe slots and 18 DIMM slots into a compact 2U footprint.

Refer to the technical brief *Implementing Fibre Channel SAN Boot with Oracle's Sun ZFS Storage Appliance* for details on FC SAN boot solutions using Sun ZFS Storage 7120.

Sun ZFS Storage 7120 is a standalone controller that consists of an internal SAS-2 HBA providing disk shelf expansion, write flash acceleration, and 11 x 300GB 15K, 600GB 15K, 1TB 7.2K, 2TB 7.2K, or 3TB 7.2K hard drive storage. The SAS-2 storage fabric supports a greater number of targets, greater bandwidth, higher reliability, and bigger scale.

Sun ZFS Storage 7120 Front Panel

The following figure and legend show the Sun ZFS Storage 7120 front panel and the drive locations. The Logzilla 3.5" SSD belongs in slot 3 and is not supported in controllers configured with the internal Sun Aura flash HBA Logzilla.

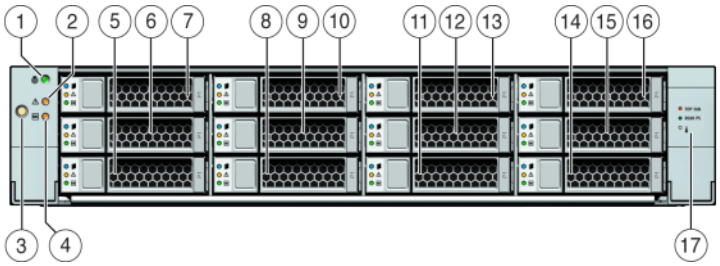

1 Locator LED/button (white)	**7** HDD 2	**13** HDD 8
2 Service Action Required LED (amber)	**8** HDD or SSD 3	**14** HDD 9
3 Power button	**9** HDD 4	**15** HDD 10
4 Power/OK LED (green)	**10** HDD 5	**16** HDD 11
5 HDD 0	**11** HDD 6	**17** Drive map
6 HDD 1	**12** HDD 7	

Sun ZFS Storage 7120 Rear Panel

The following figure and legend show the Sun ZFS Storage 7120 rear panel.

> **Note:**
>
> Optional Sun Dual Port 40Gb/sec 4x InfiniBand QDR HCAdapter PCIe cards (375-3606-01) may be located in slots 1, 2, or 3. 375-3606-01 HCA expansion cards are not supported in the 10Gb network configurations.

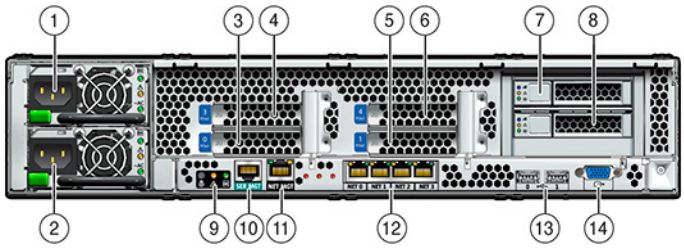

1 Power Supply Unit 1	**6** PCIe 4	**11** Network Management port
2 Power Supply Unit 0	**7** Boot HDD 1	**12** Gbit Ethernet ports NET 0, 1, 2, 3
3 PCIe 0	**8** Boot HDD 0	**13** USB 2.0 ports (0, 1)
4 PCIe 3	**9** Rear Panel System Status LEDs	**14** HD15 Video port
5 PCIe 1	**10** Serial Management port	

Sun ZFS Storage 7120 Physical Specifications

The Sun ZFS Storage 7120 2U chassis form factor dimensions are as follows:

Table 11-17 Sun ZFS Storage 7120 Controller Dimensions

Dimension	Measurement	Dimension	Measurement
Height	87.6 mm/3.45 in	Depth	765.25 mm/30.13 in
Width	436.8 mm/17.2 in	Weight	29.54 kg/65 lb

Sun ZFS Storage 7120 Electrical Specifications

The following list shows the electrical specifications for Sun ZFS Storage 7120. Note that the power dissipation numbers listed are the *maximum* rated power numbers for the power supply. The numbers are not a rating of the actual power consumption of the appliance.

Connectors

- Two C13 connectors which work on 110-220v outlets

Input

- Nominal frequencies: 50/60Hz
- Nominal voltage range: 100-120/200-240 VAC
- Maximum current AC RMS: 13.8A @ 100 VAC
- AC operating range: 90-264 VAC

Output

- 3.3 VDC STBY: 3.0A
- +12 VDC: 86.7A

Power Dissipation

- Max power consumption: 1235.3 W
- Max heat output: 4212 BTU/hr
- Volt-Ampere rating: 1261 VA @ 240 VAC, 0.98P.F.

Sun ZFS Storage 7120 Internal Boards

The Sun ZFS Storage 7120 chassis has the following boards installed.

> **Note:**
>
> Field-replaceable units (FRU) should only be replaced by trained Oracle service technicians.

- **PCIe Risers:** Each riser supports two PCIe cards that are customer-replaceable. There are two risers per system, each attached to the rear of the motherboard.
- **Motherboard:** The motherboard is a FRU and includes CPU modules, slots for 18 DIMMs, memory control subsystems, and the service processor (SP) subsystem. The SP subsystem controls the host power and monitors host system events (power and environmental). The SP controller draws power from the host 3.3V standby supply rail, which is available whenever the system is receiving AC input power, even when the system is turned off.
- **Power Distribution Board:** The power distribution board is a FRU and distributes main 12V power from the power supplies to the rest of the storage controller. It is directly connected to the connector break out board and to the motherboard through a bus bar and ribbon cable. It also supports a top cover interlock *kill* switch. The power supplies connect directly to the power distribution board.
- **Connector Break Out Board:** The connector break out board is FRU and serves as the interconnect between the power distribution board and the fan power boards, storage drive backplane, and I/O board. It also contains the top-cover interlock "kill" switch.
- **Fan Power Boards:** The two fan power boards are FRUs and carry power to the system fan modules. In addition, they contain fan module status LEDs and transfer I2C data for the fan modules.
- **Storage Drive Backplane:** The storage drive backplane is a FRU and includes the connectors for the storage drives, as well as the interconnect for the I/O board, power and locator buttons, and system/component status LEDs. The system has a 12-disk backplane. Each drive has an LED indicator for Power/Activity, Fault, and Locate.

Sun ZFS Storage 7120 CPU and Memory

The Sun ZFS Storage 7120 motherboard has 18 slots in two groups that hold industry-standard DDR3 DIMMs. The standard memory configuration is 48GB, 6x8GB DDR-1333 low voltage (LV) DIMMS.

Following are the replaceable CPU and memory components of the Sun ZFS Storage 7120 system.

Table 11-18 Sun ZFS Storage 7120 Controller Replaceable CPU and Memory Components

Part Number	Description	FRU/CRU
F371-4966-01	DIMM, 8GB, DDR3, 2RX4, 13	CRU
F371-4885-01	Intel E5620, 2.40G	FRU

All sockets must be occupied by either a filler or a DDR3 DIMM. All DDR3 DIMMs must be identical. DIMMs are pre-installed in P0 slots D1, D2, D4, D5, D7, and D8.

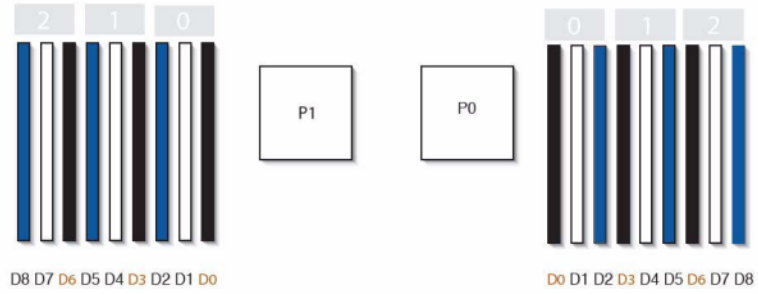

Sun ZFS Storage 7120 Power Distribution, Fan Module and Disk Components

The Sun ZFS Storage 7120 fan modules and LEDs are shown in the following illustration.

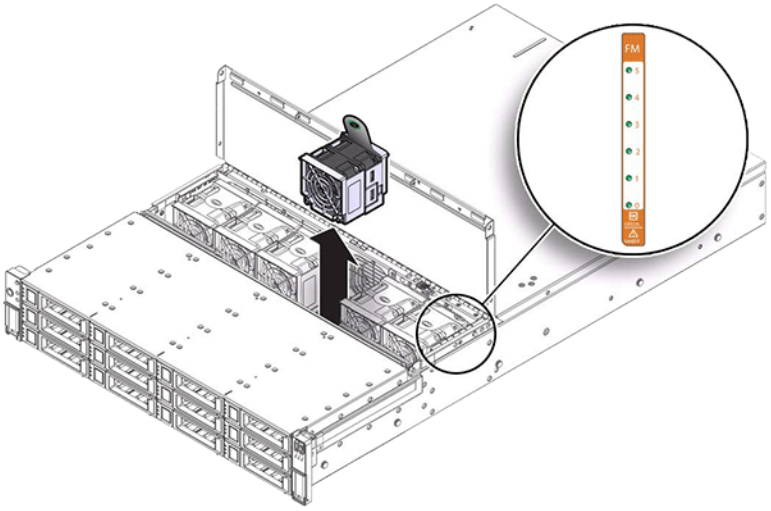

The following figure and legend show the power distribution board and associated components.

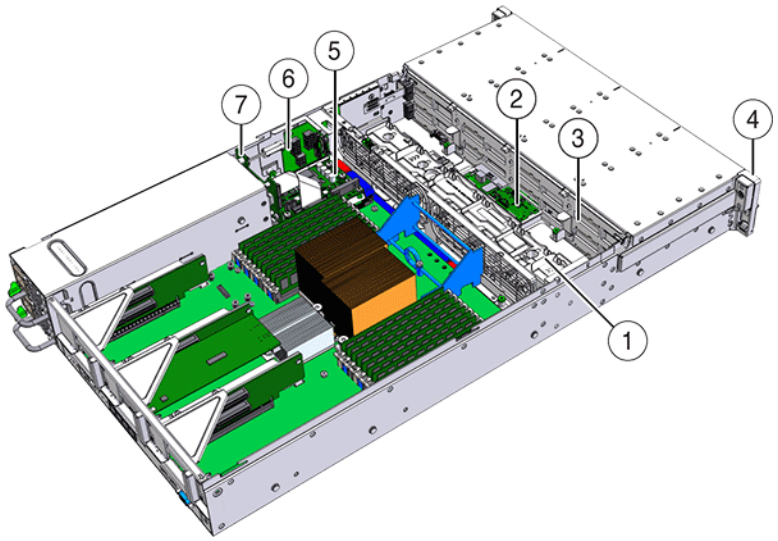

1 Fan Board	**5** Power Distribution Board
2 SAS Expander Board	**6** Connector Board
3 Disk Backplane	**7** Power Supply Backplane
4 Front Control Panel Light Pipe Assembly	

Sun ZFS Storage 7120 I/O Components

The following figure and legend show the I/O components of the Sun ZFS Storage 7120 system.

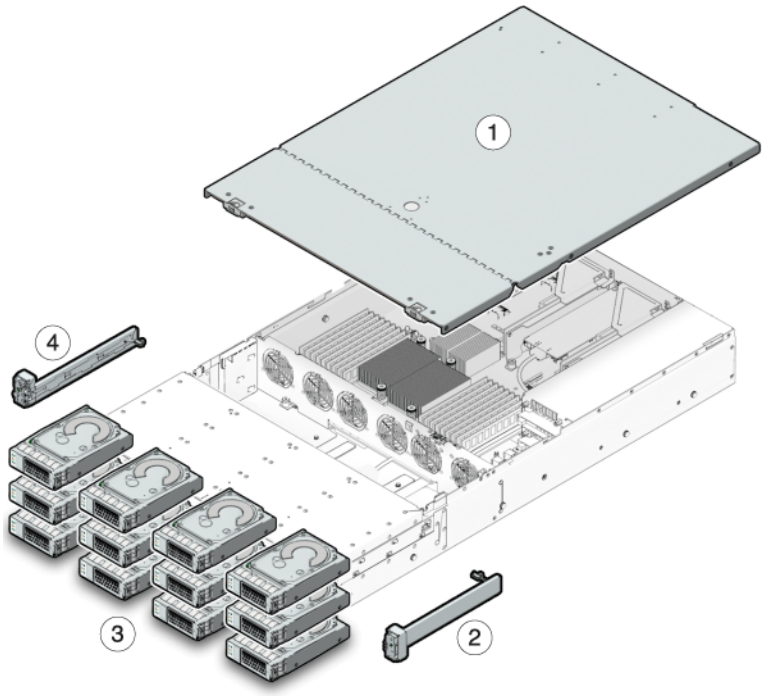

1 Top Cover	**3** Hard Disk Drives
2 Right Control Panel Light Pipe Assembly	**4** Left Control Panel Light Pipe Assembly

Sun ZFS Storage 7120 NIC/HBA Options

The following table describes NIC/HBA options for Sun ZFS Storage 7120.

Table 11-19 Sun ZFS Storage 7120 Controller NIC/HBA Options

Mktg Part Number	Description	Mfg Part Number
SG-XPCIESAS-GEN2-Z	2-port External Sun Thebe SAS (x4) HBA, PCIe	594-5889-01
SG-XPCIE2FC-QF8-Z	2-port FC HBA, 8Gb, PCIe	594-5684-01
X4446A-Z	4-port PCI-E Quad GigE UTP	594-4024-01
X4237A-N	2-port 4X IB HCA PCIe	594-5862-02
X1109A-Z	2-port 10Gig SFP+ NIC, PCIe	594-6039-01

Sun ZFS Storage 7120 PCIe Options

The following table describes the supported PCIe configuration option summary for the Sun ZFS Storage 7120 controller.

Table 11-20 Sun ZFS Storage 7120 Controller PCIe Options

Slot	Type	Sun Part Number	Vendor Part Number	Description	Note
0	PCIe	540-7975-03	Sun Aura	Internal Flash HBA Logzilla	Base Configuration (OBSOLETE)
0	PCIe	375-3481-01	Intel EXPI9404PT	QP Copper NIC	Optional Recommended Front-end
0	PCIe	375-3617-01	Intel Niantic	DP Optical 10GE NIC	Optional Recommended Front-end
0	PCIe	371-4325-01	QLogic	8Gb DP FC HBA	Optional FC Target or Initiator (Backup)
0	PCIe	375-3606-01	Mellanox MHJH29-XTC	InfiniBand HCA	Optional Recommended Front-end
1	PCIe	375-3617-01	Intel Niantic	DP Optical 10GE NIC	Optional Recommended Front-end
1	PCIe	375-3606-01	Mellanox MHJH29-XTC	InfiniBand HCA	Optional Recommended Front-end
1	PCIe	375-3481-01	Intel EXPI9404PT	QP Copper NIC	Optional Recommended Front-end

Table 11-20 (Cont.) Sun ZFS Storage 7120 Controller PCIe Options

Slot	Type	Sun Part Number	Vendor Part Number	Description	Note
1	PCIe	371-4325-01	QLogic	8Gb DP FC HBA	Optional FC Target or Initiator (Backup)
3	PCIe	375-3665-01	Sun Thebe (INT)	Internal SAS HBA	Base Configuration
4	PCIe	375-3481-01	Intel EXPI9404PT	QP Copper NIC	Optional Recommended Front-end
4	PCIe	371-4325-01	QLogic	8Gb DP FC HBA	Optional FC Target or Initiator (Backup)
4	PCIe	375-3609-03	Sun Thebe (EXT)	8P 6Gb/s SAS HBA	Additional Optional Back-end

Sun ZFS Storage 7120 Connectors

The Sun ZFS Storage 7120 serial management connector (SER MGT) is an RJ-45 port and provides a terminal connection to the SP console.

TheSun ZFS Storage 7120 network management connector (NET MGT) is an RJ-45 port and provides an alternate terminal interface to the SP console.

There are four RJ-45 Gigabit Ethernet ports (NET0, NET1, NET2, NET3) located on the Sun ZFS Storage 7120 motherboard that operate at 10/100/1000 Mbit/sec. These network interfaces must be configured before use.

Sun ZFS Storage 7120 Cables

The following figure and legend show the Sun ZFS Storage 7120 controller internal cables.

> **Note:**
>
> The rear boot drives are not depicted in this illustration.

Sun ZFS Storage 7120 Controller Storage Drive Data Cable

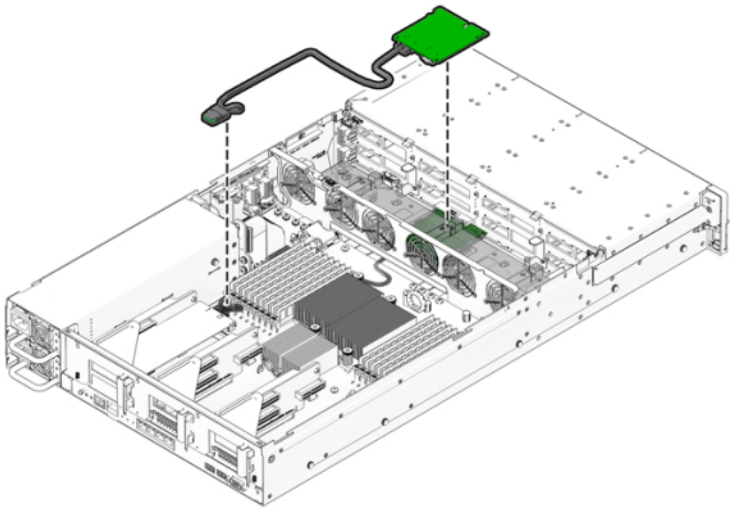

The Sun ZFS Storage 7120 controller storage drive data cable connects between the HBA PCIe card and the storage drive backplane.

Sun ZFS Storage 7120 Controller Ribbon Cable

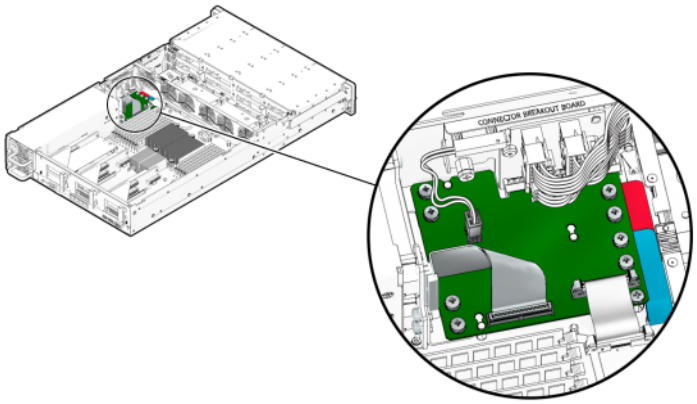

The Sun ZFS Storage 7120 controller ribbon cable connects between the power distribution board and the motherboard.

Sun ZFS Storage 7120 Attached Storage

The Sun ZFS Storage 7120 standalone configurations allow a single chain of 1 or 2 disk shelves. Write-optimized (Logzilla) SSDs are not supported in the expansion storage for Sun ZFS Storage 7120. The disk shelves must be fully populated with 24 HDDs. Half-populated shelf configurations are not supported. For more information, see Connecting to Attached Storage. See also Disk Shelf Hardware Overview for component specifications and diagrams.

12

Servicing the Disk Shelves

Before servicing a disk shelf, be sure to read sections Getting Started with Servicing the Appliance and Prerequisites for Servicing Hardware.

To replace disk shelf components, use the following tasks.

- Electrostatic Discharge Precautions
- Replacing a Faulty Disk Shelf Drive
- Manually Faulting a Disk Shelf Drive
- Replacing a Disk Shelf Power Supply
- Replacing a Disk Shelf I/O Module
- Replacing a Disk Shelf SIM Board
- Adding a New Disk Shelf
- Adding a Disk Shelf Cache Device
- Connecting to Attached Storage

To understand disk shelf components, see the following topics.

- Oracle Storage Drive Enclosure DE3-24
- Oracle Storage Drive Enclosure DE2-24
- Oracle Storage Drive Enclosure DE3-12C
- Sun Disk Shelf
- Disk Shelf Configurations
- Maximum Disk Shelves per Controller Configuration
- SAS-2 and SAS-3

Electrostatic Discharge Precautions

- Remove all plastic, vinyl, and foam material from the work area.
- Wear an antistatic wrist strap at all times when handling any customer-replaceable unit (CRU).
- Before handling any CRU, discharge any static electricity by touching a grounded surface.
- Do not remove a CRU from its antistatic protective bag until you are ready to install it.
- After removing a CRU from the chassis, immediately place it in an antistatic bag or antistatic packaging.
- Handle any card that is part of a CRU by its edges only and avoid touching the components or circuitry.
- Do not slide a CRU over any surface.

- Limit body movement (which builds up static electricity) during the removal and replacement of a CRU.

Replacing a Faulty Disk Shelf Drive

> **Caution:**
>
> Do not remove a component if you do not have an immediate replacement. The disk shelf must not be operated without all components in place, including filler panels. Refer to Disk Shelf Hardware Overview for component locations.

Use the following procedure to replace a faulty drive. The disk shelf drives are hot-swappable and may be replaced without removing power from the disk shelf. To avoid possible data loss when removing non-faulted drives, label each drive with the number of the slot from which it was removed and reinstall each drive into the same slot.

The replacement drive must be of the same type and rotational speed as the drive to be replaced. Typically, the replacement drive is the same capacity; however, Oracle might support and ship a specific larger capacity drive in certain cases. Contact Oracle Service with questions about allowable use.

> **Note:**
>
> Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** BUI menu, select **System**, or navigate to `maintenance system updates` in the CLI.

To add, not replace, a new cache device into an appropriate and empty drive slot, see Adding a Disk Shelf Cache Device.

1. Identify the faulted HDD.
   - If you are not physically at the system, from the **Maintenance** BUI menu, select **Hardware**, and open the detail view of the disk shelf that has the faulted drive. The disk shelf that has the faulted drive is marked with an amber status icon. On the detail view, the faulted drive has either an amber or a blue-on-amber (ready to remove) status icon. Note the ID of the faulted drive. Click the entry for that drive in the table to highlight the location of the drive in the disk shelf illustration. To view details about the faulted drive, click its information icon and then click **Active Problems**.
   - If you are physically at the system, the amber Service Action Required indicator on the HDD should be lit. Additionally, you can click the locator icon in the BUI to light the locator LED on the faulted drive.
2. Verify that the faulted drive is ready to be removed.

> **Caution:**
>
> Removing the drive before it is ready to be removed can result in loss of data.

Use any of the following methods to verify that the faulted drive is ready to be removed:

- In the BUI:
  - The status icon for the faulted drive in the detail view in the **Maintenance: Hardware** BUI menu has changed from amber to blue-on-amber.
  - The **Component Details** information window (click the information icon) for the faulted drive shows **Yes** as the value for the **Ready to remove** property.
- In the CLI (and RESTful API), the value of the `readytoremove` property for the faulted drive is `true`:

```
hostname:maintenance chassis-000 disk-001> get readytoremove
                readytoremove = true
```

3. Press the release button or latch to release the drive lever.

   For Oracle Storage Drive Enclosure DE3-12C, release the latch by pushing it toward the right with your finger or thumb.

4. Pull the drive lever fully open to unlock and partially eject the drive from the chassis.

   The following figure shows how to unlock and partially eject a horizontal disk drive.

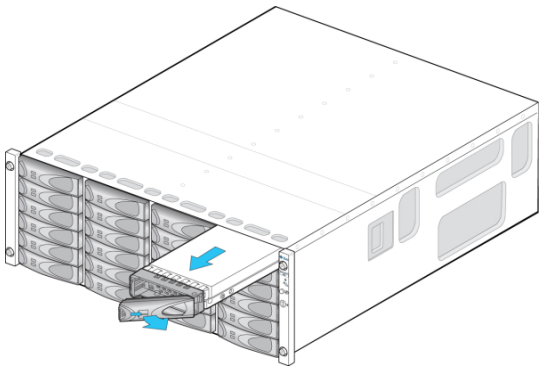

The following figure shows how to unlock and partially eject a vertical disk drive.

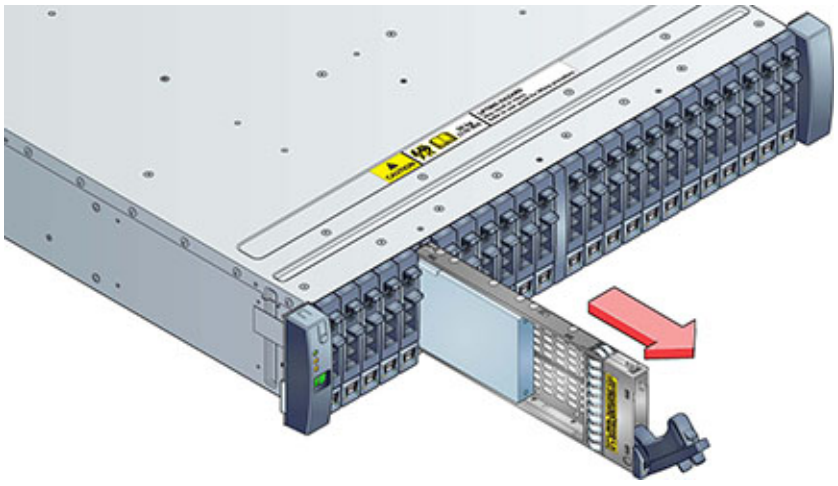

5. Grasp the middle of the drive body and pull it toward you to remove it from the chassis.
6. Wait at least 30 seconds.
7. Ensure the new drive lever is in the fully extended position.
8. While constantly pushing toward the pivot point of the lever, slide the drive fully into the chassis slot.

   The following figure shows how to slide a horizontal disk drive into the chassis slot.

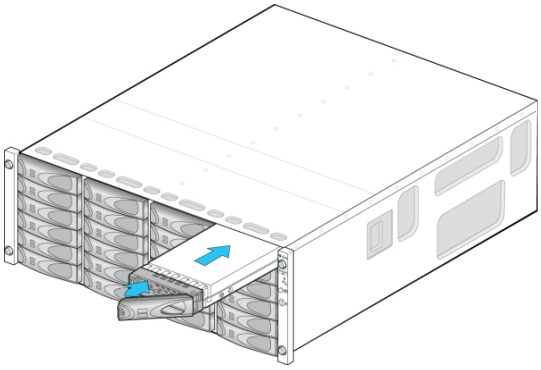

The following figure shows how to slide a vertical disk drive into the chassis slot.

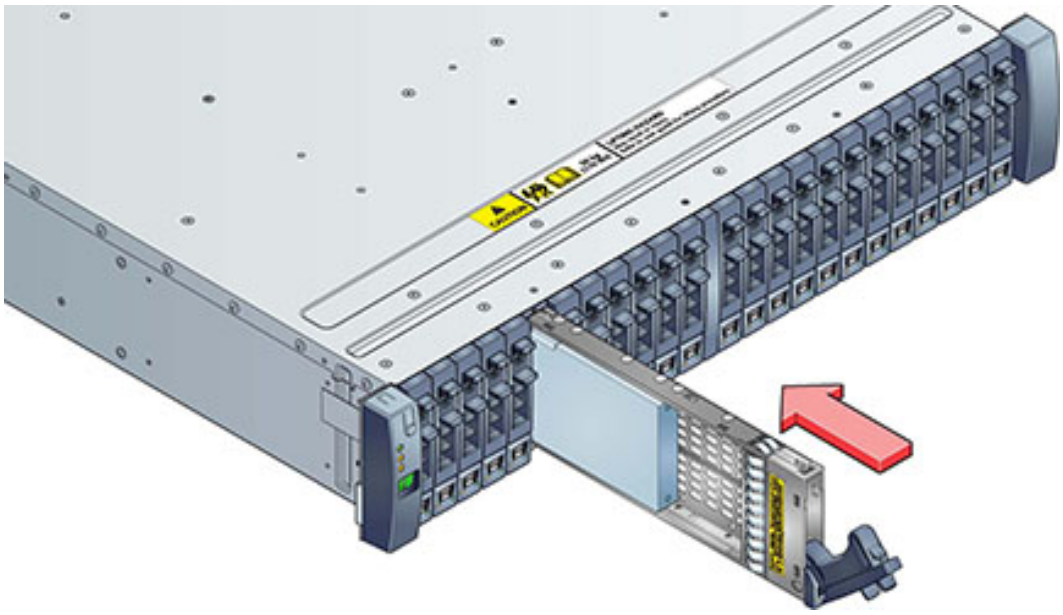

9. Press the drive lever closed until it locks in place. If a vertically oriented drive is higher than surrounding drives, push down on the top of the drive to properly seat it. The Activity LED will be steady green to indicate a ready state.
10. From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon at the beginning of the appropriate disk shelf row, and then click **Disk** to verify that the disk online icon is green for the newly installed disk.
11. If a fault does not self-clear, see Repairing Active Problems.

Manually Faulting a Disk Shelf Drive

For the user's role or exceptions, add the `disk-fault` BUI authorization or the `allow_disk-fault=true` CLI authorization under the hardware scope. For information about user configuration, including roles, authorizations, and exceptions, see Configuring Users in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Use the following procedure to manually fault a disk shelf drive, which is especially useful when a drive is causing a performance impact, but has not reached its fault threshold. This command-line-interface procedure, which is also available in the RESTful API but not in the BUI, allows a drive to be gracefully removed from the system. If the drive is in a pool with a configured spare drive, the spare drive will take over for the manually faulted drive.

To remove the manual fault, repair the associated drive problem or replace the faulted drive. The fault cannot be removed manually by setting the faulted parameter to `false`. To repair the faulted drive, go to the **Maintenance Problems** area of the software and follow procedures described in Repairing Active Problems. To replace a faulty drive, see Replacing a Faulty Disk Shelf Drive.

> **Note:**
>
> Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, from the **Maintenance** BUI menu, select **System**, or navigate to `maintenance system updates` in the CLI.

1. Identify the drive to manually fault by running, from the BUI, such analytics as "Disk: Disks: Broken down by percent utilization" and drilling down to the individual disks, and "Disks: I/O bytes: Broken down by disk".

   For details, refer to the Oracle ZFS Storage Appliance Analytics Guide, Release OS8.8.x.

2. In the CLI, go to `maintenance hardware` and enter `ls` to list the system hardware components.

```
hostname:> maintenance hardware
hostname:maintenance hardware> ls
hostname:maintenance hardware>
Chassis:
              NAME         STATE  MANUFACTURER  MODEL
chassis-000   hostname     ok     Oracle        Oracle ZFS Storage ZS7-2
...
chassis-001   1235FM401W   ok     Oracle        Oracle Storage DE2-24P
disk-000      HDD 0        ok     HGST          H7210A520SUN010T
disk-001      HDD 1        ok     HGST          H7210A520SUN010T
disk-002      HDD 2        ok     HITACHI       H109090SESUN900G
...
```

3. Select the chassis with the drive to be faulted and enter `ls` to list its properties and components.

```
hostname:maintenance hardware> select chassis-001
hostname:maintenance chassis-001> ls
Properties:
           name = 1227FMT00E
         faulted = false
   manufacturer = Oracle
          model = Oracle Storage DE2-24C
           part = 7046842
         serial = 1227FMT00E
       revision = 001E
           type = hdd
            rpm = 7200
           path = 2
         locate = false

Components:
           LABEL  STATE  MANUFACTURER  MODEL             SERIAL  RPM   TYPE
disk-000   HDD 0  ok     HITACHI       H7230AS60SUN3.0T  001...  7200  data
disk-001   HDD 1  ok     HITACHI       H7230AS60SUN3.0T  001...  7200  data
disk-002   HDD 2  ok     HITACHI       H7230AS60SUN3.0T  001...  7200  data
...
```

4. Enter `select disk`.

```
hostname:maintenance chassis-001> select disk
```

5. Enter the disk to be manually faulted, and enter `ls` to list its properties.

```
hostname:maintenance chassis-001 disk> select disk-002
hostname:maintenance chassis-001 disk-002> ls
Properties:
          label = HDD 2
        present = true
        faulted = false
   manufacturer = HITACHI
          model = H7230AS60SUN3.0T
         serial = 001218RRV9GD--------YHJRV9GD
       revision = A6C0
           size = 2.73T
           type = data
            rpm = 7200
         device = c0t5000CCA01A9A47ACd0
      pathcount = 2
      interface = SAS
         locate = false
        offline = false
```

6. Set the `faulted` property to `true`.

```
hostname:maintenance chassis-001 disk-002> set faulted=true
     faulted = true (uncommitted)
```

7. Enter `commit`.

```
hostname:maintenance chassis-001 disk-002> commit
```

Replacing a Disk Shelf Power Supply

> **Caution:**
>
> Do not remove a component if you do not have an immediate replacement. The disk shelf must not be operated without all components in place, including filler panels. Refer to Disk Shelf Hardware Overview for component locations.

Disk shelves are provided with redundant power supplies to prevent loss of service due to component failure. Each power supply is accompanied by one or more chassis cooling fans in one customer-replaceable unit (CRU). Power supplies are hot-swappable, meaning they can be replaced one at a time without removing power from the disk shelf.

> **Caution:**
>
> The modules can produce a high-energy hazard and should only be replaced by instructed individuals with authorized access to the equipment.

1. Separate indicator LEDs on the back panel represent the operational state of power supplies and fans individually; see the rear panel illustration for details. Failed components are indicated by amber LEDs as well as amber icons in the administrative BUI. From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon at the beginning of the appropriate disk shelf row, then select **PSU** or **Fan** to view details on the respective components. You can also click the locator icon to flash the chassis locator LED.

2. Locate the chassis and module containing the failed component.
3. Ensure the power supply on/off switch is in the "O" off position.

   Some power supplies for Oracle Storage Drive Enclosure DE3-24P and DE3-24C disk shelves have a disabled on/off switch with a do-not-use label (⃠) on the switch. If the disk shelf has this power supply, skip this step and continue with the next step.
4. Disconnect the power cord tie strap from the power cord, and unplug the power cord from the power supply.
5. Release the lever/ejection arms. Pull the module out of the chassis, being careful not to damage the connector pins in the back.

   Oracle Storage Drive Enclosure DE3-24P or DE3-24C:

   To release the lever, push the lever fully downward.

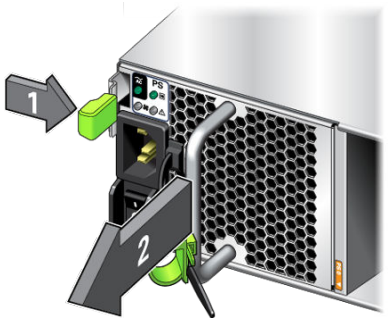

   Oracle Storage Drive Enclosure DE2-24P, DE2-24C or DE3-12C:

   To release the lever, grasp the latch and the opposite side of the module, and squeeze together to release the lever.

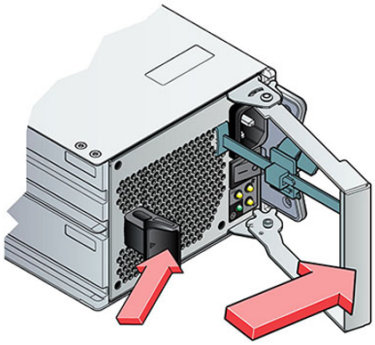

   Sun Disk Shelf:

   To release the arms, use your thumb and forefinger, unscrew both ejection arm captive screws until loose, and swing the ejection arms out until they are fully open.

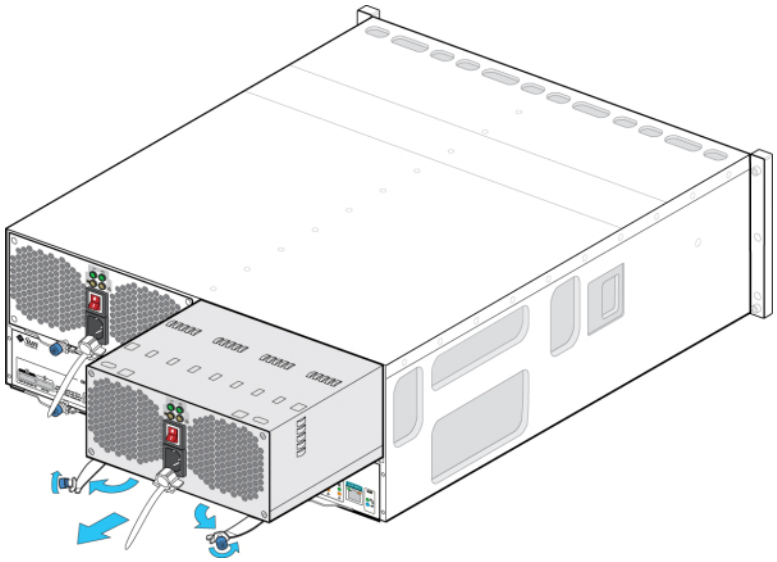

6. With the lever/ejection arms fully open (if applicable), slide the new module into the chassis slot until it contacts the chassis backplane, and the lever/ejection arms begin to engage. Close the lever/ejection arms. If an Oracle Storage Drive Enclosure DE3-24, the module clicks in place and the lever engages.

   Oracle Storage Drive Enclosure DE3-24P or DE3-24C:

   Push on the handle until you hear or feel a click.

   Oracle Storage Drive Enclosure DE2-24P, DE2-24C or DE3-12C:

   Push the lever fully closed until you hear or feel a click.

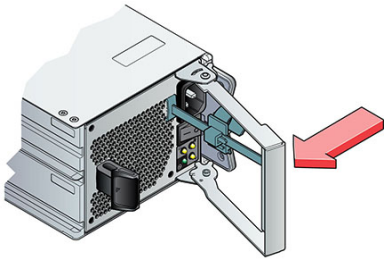

   Sun Disk Shelf:

Push the ejection arms fully closed and secure both captive screws to seat and secure the module in the chassis.

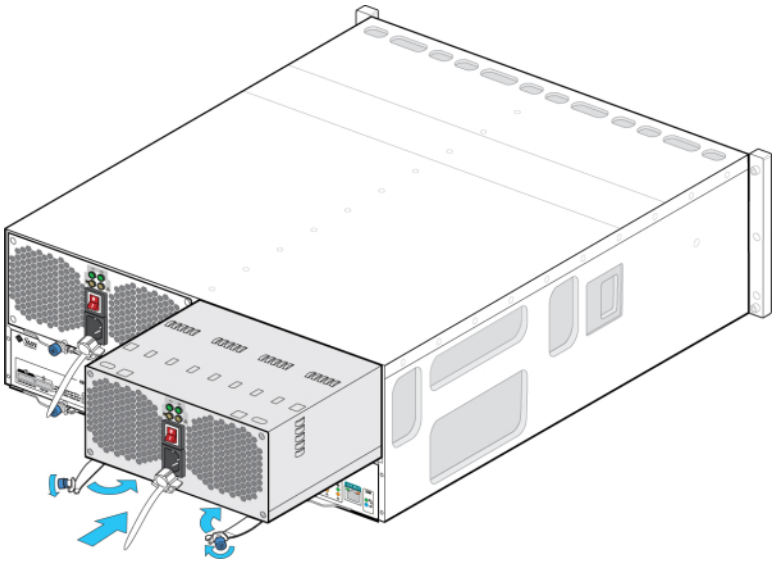

7. Ensure the power supply on/off switch is in the "O" off position. Plug the power cord into the new power supply and attach the power cord tie strap to the power cord.

   If the disk shelf has a power supply with a disabled on/off switch, do not place the switch in the off position.

8. Place the power supply on/off switch to the "I" on position. The Power/OK status LED should be a steady green, and all other indicators should be off.

   If the disk shelf has a power supply with a disabled on/off switch, do not place the switch in the on position.

   Oracle Storage Drive Enclosure DE3-24P or DE3-24C:

   Oracle Storage Drive Enclosure DE2-24P, DE2-24C or DE3-12C:

9. From the **Maintenance** BUI menu, select **Hardware**, and click the right-arrow icon at the beginning of the appropriate disk shelf row. As appropriate for the failure, click either **PSU** or **Fan** to verify that the online icon is green for the newly installed power supply with fan module.
10. If a fault does not self-clear, see Repairing Active Problems.

Replacing a Disk Shelf I/O Module

> **Caution:**
>
> Do not remove a component if you do not have an immediate replacement. The disk shelf must not be operated without all components in place, including filler panels. Refer to Disk Shelf Hardware Overview for component locations.

The following disk shelves have I/O Modules (IOMs):

- Oracle Storage Drive Enclosure DE3-24P
- Oracle Storage Drive Enclosure DE3-24C
- Oracle Storage Drive Enclosure DE2-24P
- Oracle Storage Drive Enclosure DE2-24C
- Oracle Storage Drive Enclosure DE3-12C

The I/O Modules (IOMs), which are similar to SIM boards, are hot-swappable so you can replace them without removing power to the system.

1. A faulted I/O Module is indicated by an amber LED. From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon at the beginning of the appropriate disk shelf row, and then click **Slot** to view details, or click the locator icon to turn on the locator LED.
2. Locate the failed I/O Module at the back of the disk shelf.
3. Label and disconnect the I/O Module interface cables.
4. Using your thumb and forefinger, squeeze the release button toward the lever hole to release the lever.

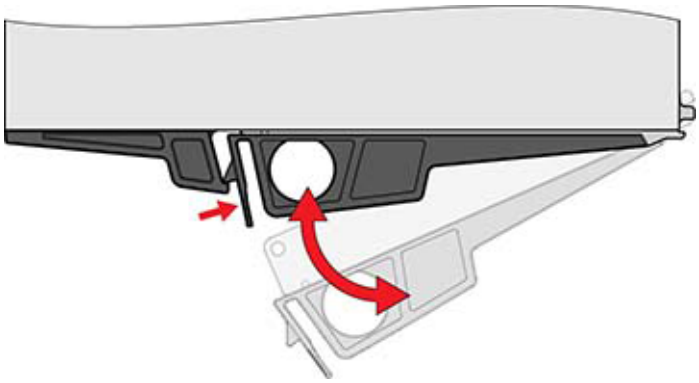

5. Grasp the lever and remove the I/O Module, being careful not to damage the connector pins in back.
6. With the lever of the new I/O Module in the open position, push on the front of the I/O Module to slide it into the disk shelf until the lever engages. Do not push on the lever, and be careful of the rear connector pins.
7. Push the lever fully closed until you hear or feel a click.
8. Reconnect the interface cables to their original locations.
9. Wait approximately 60 seconds for the I/O Module to complete its boot process, at which time the Power LED should be solid green and the Fault/Locate LED should be off. All four activity LEDs should be solid green for each SAS port in use.
10. From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon at the beginning of the appropriate disk shelf row, and then click **Slot** to verify that the I/O Module online icon is green for the newly installed I/O Module.
11. If a fault does not self-clear, see Repairing Active Problems.

Replacing a Disk Shelf SIM Board

> **Caution:**
>
> Do not remove a component if you do not have an immediate replacement. The disk shelf must not be operated without all components in place, including filler panels. Refer to Disk Shelf Hardware Overview for component locations.

The following disk shelves have SAS Interface Module (SIM) boards: Sun Disk Shelf.

The SIM boards, which are similar to I/O Modules, are hot-swappable so you can replace them without removing power to the system. The SIM boards are multi-pathed, so you can remove one of the SIM boards at any time, regardless of the state of the blue SIM OK indicator.

1. A faulted SIM board is indicated by an amber LED. From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon at the beginning of the appropriate disk shelf row, and then click **Slot** to view details, or click the locator icon to turn on the locator LED.

2. Locate the failed SIM at the back of the disk tray.
3. Label and disconnect the tray interface cables.
4. Use two hands to disconnect the SAS cable. Grasp the metal body of the connector with one hand and firmly grasp and pull the tab gently toward the connector body with the other, then pull the connector body outward to extract it from the bulkhead. Do not twist or pull the tab in any direction other than parallel with the connector body or it may break. If the tab breaks, use a small sharp object (such as a fine-tipped screwdriver) to lift the metal spring at the top of the connector shell to unlatch it.
5. Loosen the two extraction arm captive screws using your thumb and forefinger. If the captive screws are too tight to loosen by hand, use a No.2 Phillips screwdriver to loosen each screw.
6. Pull each ejector tab outward, and push to the sides to release and partially eject the SIM from the chassis.

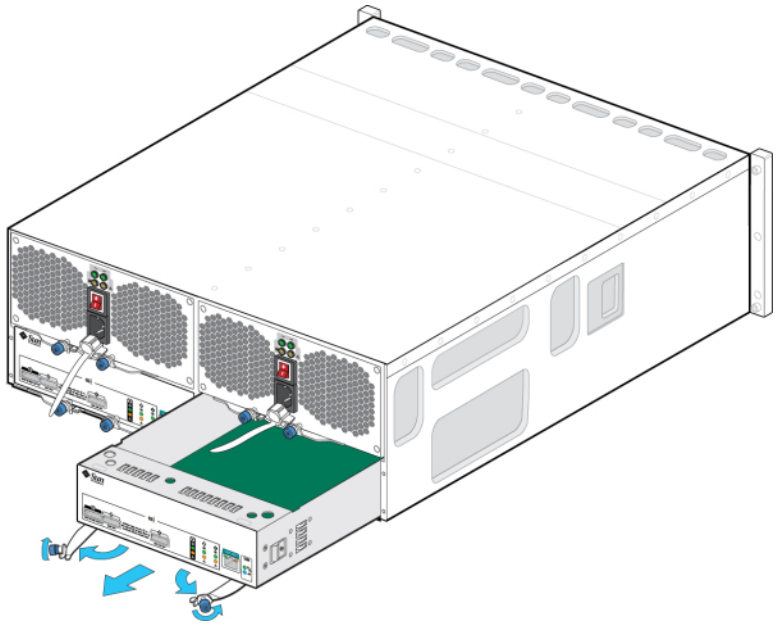

7. Grasp the middle of the SIM board, and slide it out of the slot.
8. With the ejector arms in the full open position, align the new SIM board with the open slot, and slide it into the tray until the ejector arms contact the tray connectors and begin to swing closed.

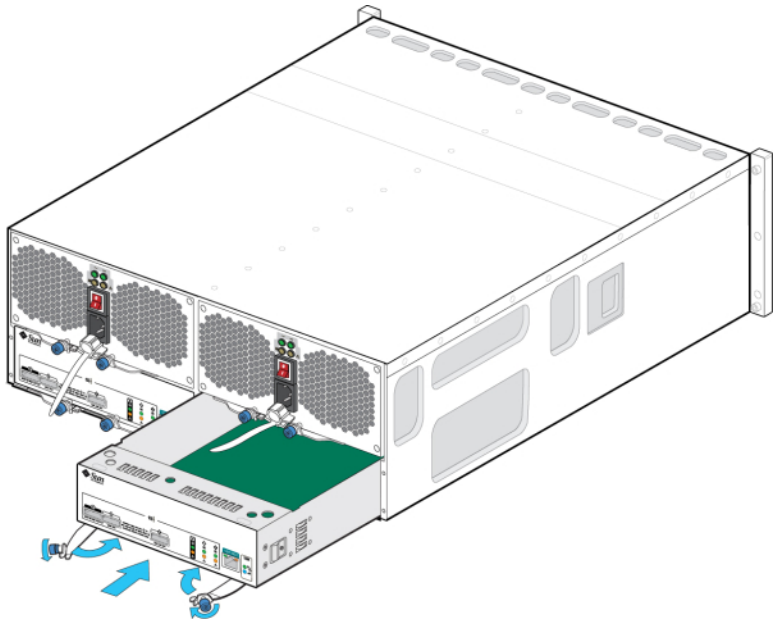

9. Swing both ejector arms in until they are flush with the SIM board panel to seat the board.
10. Tighten both captive screws to secure the board.
11. Reconnect the SAS interface cables to their original locations.
12. Wait approximately 60 seconds for the SIM board to complete its boot process, at which time the Power LED should be solid green and the SIM locate LED should be off.
13. From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon at the beginning of the appropriate disk shelf row, and then click **Slot** to verify that the SIM board online icon is green for the newly installed SIM board.
14. If a fault does not self-clear, see Repairing Active Problems.

Adding a New Disk Shelf

Caution:

Improper orientation of mini-SAS HD cables during installation can damage the HBA connector and cause the HBA to malfunction. These cables are used with 4X4 port SAS-2 HBAs and 4X4 port SAS-3 HBAs. Position the release tab downward for horizontally mounted HBAs or to the right for vertically mounted HBAs. Mini-SAS HD active optical cables (AOCs) should be oriented similarly to copper cables. .
For detailed information, see Cabling Disk Shelves Together in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x* and Cabling Controllers to Disk Shelves in a Base Cabinet in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

Use the following procedure to add a new disk shelf to your system. Disk shelves can be added to standalone or clustered storage controllers without powering down the appliance or loss of service to clients. If the new disk shelf requires that other components be moved within the rack, you must move those components first. For disk shelf guidelines as well as cabling diagrams, see Getting Started with Cabling in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.

1. Stabilize the appliance rack. If your rack includes an anti-tip foot, extend it from the bottom of the rack.
2. Install the new disk shelf into the rack, connect power cables, and supply power using the appropriate procedures in the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x. For disk shelf placement guidelines, see Cabinet and Cabling Guidelines in *Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x*.
3. Working with a single disk shelf cable at a time, connect the cables to the controller(s) and other disk shelves using redundant paths and following the diagrams in the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x. Always retain one active path while cabling the new disk shelf.
4. Verify that you have redundant paths to the disk shelf by examining the cabling.
5. To verify the disk shelf presence, from the **Maintenance** BUI menu, select **Hardware**, or navigate to CLI context `maintenance hardware` and issue the command `list`.
6. To add the new disk shelf to an existing or new pool, see Configuring Storage in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.
7. If the storage workload is mostly read operations, it is recommended that you perform a backup and restore operation to quickly balance the workload across all disk shelves. To back up and restore user data, you can use NDMP backup software, snapshots, and/or remote replication. Accordingly, see in the *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*: NDMP Configuration, Snapshots and Clones, or Remote Replication.

Adding a Disk Shelf Cache Device

Use the following procedure to install a new cache device, an SSD, into an available and appropriate drive slot. Cache devices can be added without removing power from the disk shelf. To replace a faulty drive, see Replacing a Faulty Disk Shelf Drive.

1. Remove the filler panel from the appropriate slot. For SSD locations and slot order, see the Disk Shelf Hardware Overview.
2. Ensure the new drive lever is in the fully extended position.

   For Oracle Storage Drive Enclosure DE3-12C, release the lever by pushing it toward the right with your finger or thumb.
3. While constantly pushing toward the pivot point of the lever, slide the drive fully into the chassis slot.

   The following figure shows how to slide a horizontal disk drive into the chassis slot.

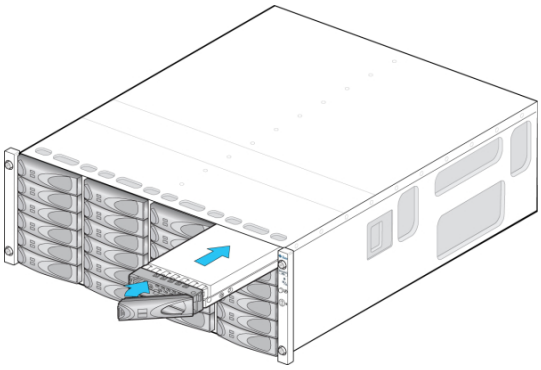

The following figure shows how to slide a vertical disk drive into the chassis slot.

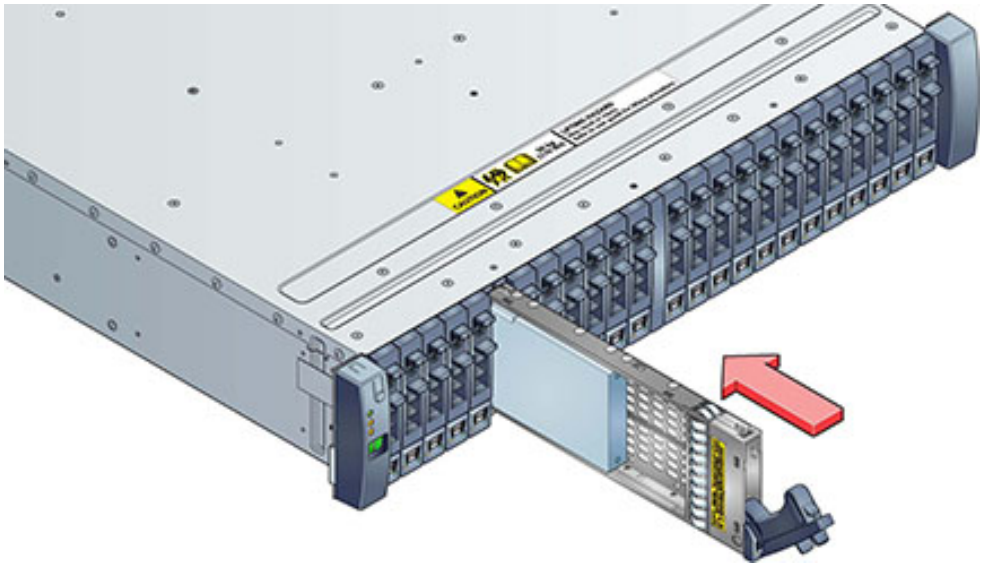

4. Press the drive lever closed until it locks in place. If a vertically oriented drive is higher than surrounding drives, push down on the top of the drive to properly seat it. The Activity LED will be steady green to indicate a ready state.
5. From the **Maintenance** BUI menu, select **Hardware**, click the right-arrow icon at the beginning of the appropriate disk shelf row, and then click **Disk**. Verify that the new disk is present and that its disk icon is green.
6. To add the new drive to a pool, from the **Configuration** BUI menu, select **Storage**.
7. From the **Available Pools** list, select the pool to which you are adding the drive.
8. Click the **Add** button to add the drive to the pool.
9. Select the drive you are adding to the pool, and click **Commit**.
10. Select the log profile (if applicable), and click **Commit**.

Connecting to Attached Storage

Disk shelves can be added to standalone or clustered storage controllers without powering down the appliance or loss of service to clients. After connecting expansion storage, verify that you have redundant paths to each disk shelf.

To connect one or more disk shelves, use the diagrams in the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.

When connecting SAS cables to two Oracle Storage Drive Enclosure DE3-12C disk shelves, cable tabs should be oriented downward toward the port's SAS logo for the top I/O Module (IOM A). For the bottom I/O Module (IOM B), cable tabs should be oriented upward toward the port's SAS logo.

> **Note:**
>
> Contact Oracle Service if you want to move a disk shelf between Oracle ZFS Storage Appliance systems. Certain limitations exist when importing storage pools, and can cause problems if not done properly by Oracle Service personnel.

Disk Shelf Hardware Overview

This section provides an overview of Oracle ZFS Storage Appliance disk shelves. For installation instructions, refer to Installing Disk Shelves in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*. For service instructions, refer to Servicing the Disk Shelves.

Oracle disk shelves are high-availability serial attached SCSI (SAS) devices that provide expanded storage. The main components are hot-swappable, including drives, I/O Modules (IOMs) or SAS Interface Module (SIM) boards for connecting to controllers and other disk shelves, and dual load-sharing power supply with fan modules. This provides a fault-tolerant environment with no single point of failure. Component status is indicated with lights on the disk shelf, and in the **Maintenance: Hardware** screen of the BUI.

To understand disk shelf components, see the following topics.

- Oracle Storage Drive Enclosure DE3-24
- Oracle Storage Drive Enclosure DE2-24
- Oracle Storage Drive Enclosure DE3-12C
- Sun Disk Shelf
- Disk Shelf Configurations
- Maximum Disk Shelves per Controller Configuration
- SAS-2 and SAS-3

Oracle Storage Drive Enclosure DE3-24

This section introduces the front and rear panel components of the Oracle Storage Drive Enclosure DE3-24P and Oracle Storage Drive Enclosure DE3-24C disk shelves.

Oracle Storage Drive Enclosure DE3-24P Disk Shelf

Oracle Storage Drive Enclosure DE3-24P is a 2U chassis that supports 24 2.5" SAS-3 drives. The high-performance HDDs provide reliable storage, and the SSDs provide accelerated write operations. This disk shelf features dual, redundant I/O Modules (IOMs), and dual power supply with fan modules.

Oracle Storage Drive Enclosure DE3-24P Drive Locations (Front View)

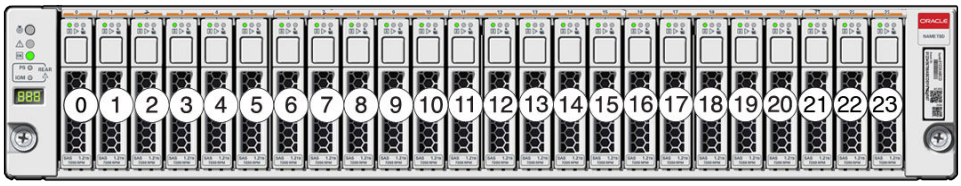

Up to four write-optimized log SSDs are supported per disk shelf. Log devices should be populated in order of slots 20, 21, 22, and 23.

External Read Cache Devices - Read cache devices, installed in Oracle Storage Drive Enclosure DE3-24 disk shelves, are external to the controller and referred to as external second level adaptive replacement cache (EL2ARC). When a cluster takeover occurs, read cache remains persistent and sharable between controllers, enabling read performance to be sustained. For slot configuration details and minimum software requirements, see Disk Shelf Configurations.

Oracle Storage Drive Enclosure DE3-24P Status Indicators (Front View)

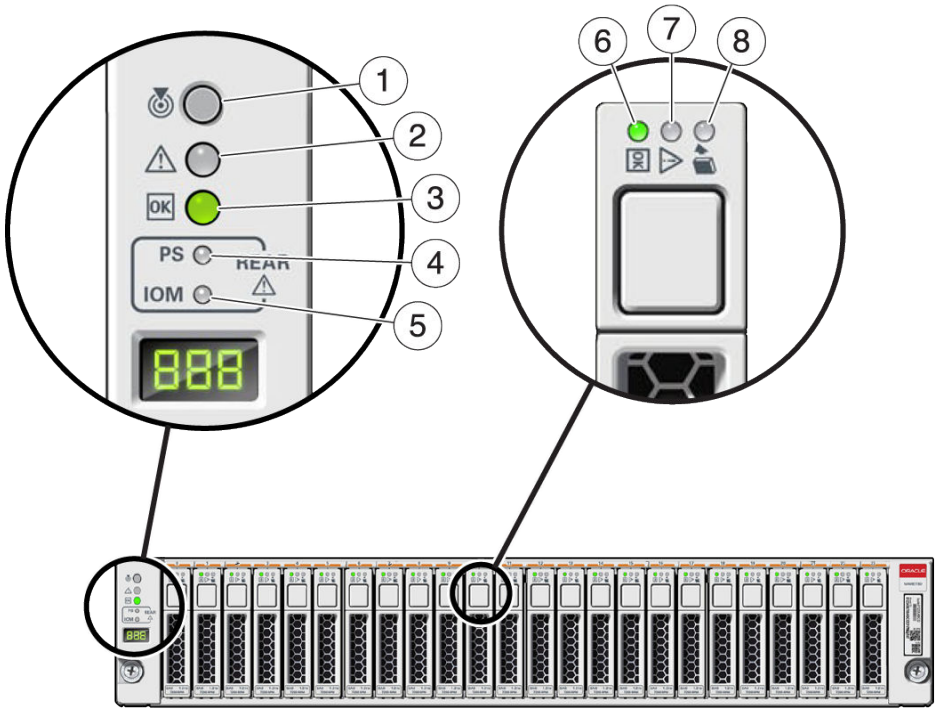

1 Locate indicator	**2** Service required indicator	**3** OK indicator
4 Power supply fault indicator	**5** SAS I/O Module fault indicator	**6** Power/Activity indicator
7 Drive service required indicator	**8** Ready-to-remove indicator	

Oracle Storage Drive Enclosure DE3-24P Rear Panel

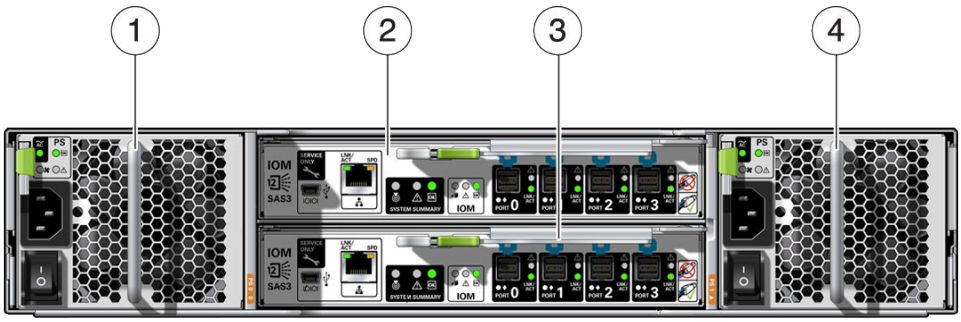

1 Power supply with fan module 0	**2** I/O Module 1
3 I/O Module 0	**4** Power supply with fan module 1

Oracle Storage Drive Enclosure DE3-24C Disk Shelf

Oracle Storage Drive Enclosure DE3-24C is a 4U chassis that supports 24 3.5" SAS-3 drives. The SSDs provide accelerated write operations, and the high-capacity HDDs provide reliable storage. This disk shelf features dual, redundant I/O Modules (IOMs), and dual power supply with fan modules.

Oracle Storage Drive Enclosure DE3-24C Drive Locations (Front View)

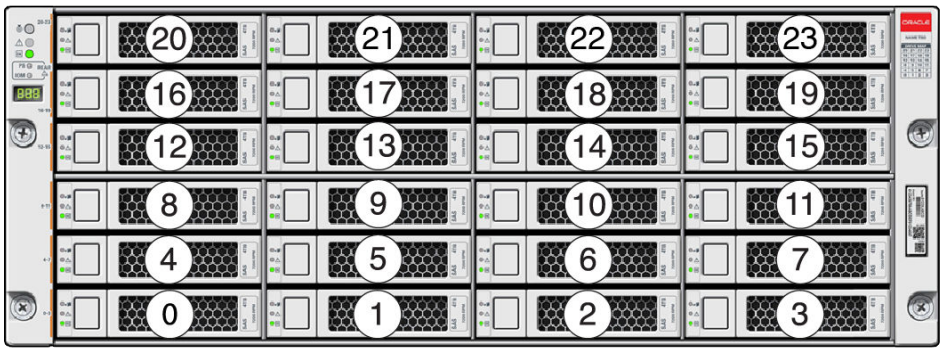

Up to four write-optimized log SSDs are supported per disk shelf. Log devices should be populated in order of slots 20, 21, 22, and 23.

External Read Cache Devices - Read cache devices, installed in Oracle Storage Drive Enclosure DE3-24 disk shelves, are external to the controller and referred to as external second level adaptive replacement cache (EL2ARC). When a cluster takeover occurs, read cache remains persistent and sharable between controllers, enabling read performance to be sustained. For slot configuration details and minimum software requirements, see Disk Shelf Configurations.

Oracle Storage Drive Enclosure DE3-24C Status Indicators (Front View)

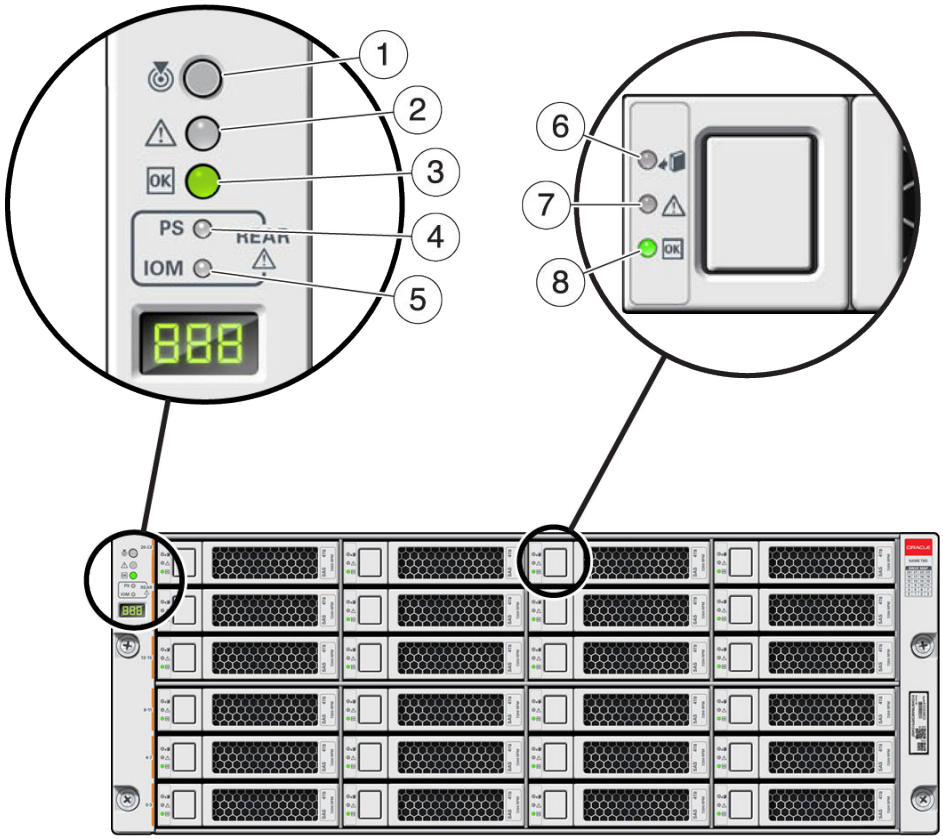

1 Locate indicator	**2** Service required indicator	**3** OK indicator
4 Power supply fault indicator	**5** SAS I/O Module fault indicator	**6** Ready-to-remove indicator
7 Drive service required indicator	**8** Power/Activity indicator	

Oracle Storage Drive Enclosure DE3-24C Rear Panel

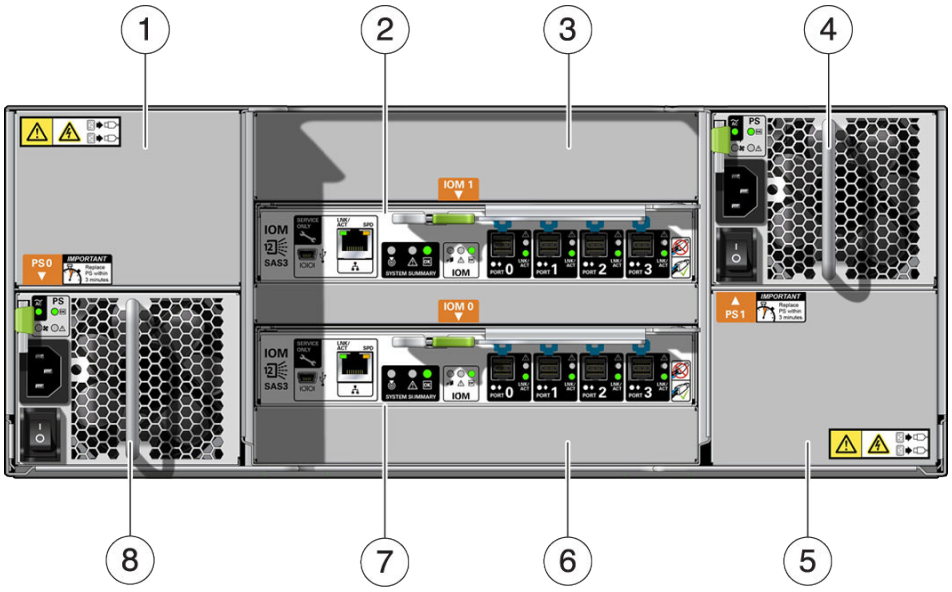

1 Power supply filler panel, slot 4	**2** I/O Module 1	**3** I/O Module filler panel
4 Power supply with fan module 2	**5** Power supply filler panel, slot 2	**6** I/O Module filler panel
7 I/O Module 0	**8** Power supply with fan module 1	

> **Note:**
>
> It is especially important that power supplies and their filler panels are in the correct slots.

Oracle Storage Drive Enclosure DE3-24 I/O Module Indicators

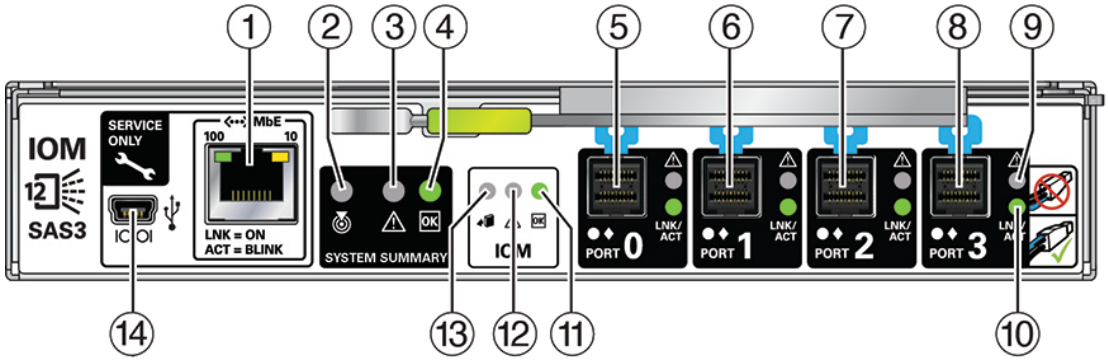

1 For Oracle service only	**2** Locate indicator	**3** Service required indicator

4 OK indicator	**5** SAS-3 Port 0	**6** SAS-3 Port 1
7 SAS-3 Port 2	**8** SAS-3 Port 3	**9** Host port service required indicator
10 Host port activity indicator	**11** Power/Activity indicator	**12** I/O Module service required indicator
13 Ready-to-remove indicator	**14** For Oracle service only	

The Oracle Storage Drive Enclosure DE3-24C and Oracle Storage Drive Enclosure DE3-24P disk shelves use the same I/O Module.

Oracle Storage Drive Enclosure DE3-24 Power Supply Indicators

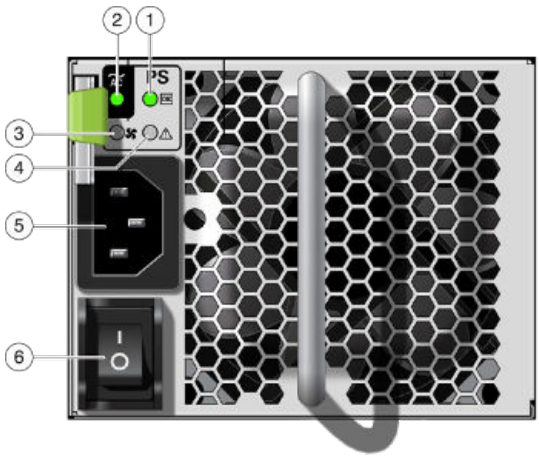

1 AC input OK indicator	**2** Power supply OK indicator	**3** Fan fail indicator
4 Service required indicator	**5** Universal power input connector	**6** Power on/off switch

The Oracle Storage Drive Enclosure DE3-24C and Oracle Storage Drive Enclosure DE3-24P disk shelves use the same power supply.

Oracle Storage Drive Enclosure DE2-24

This section introduces the front and rear panel components of the Oracle Storage Drive Enclosure DE2-24P and Oracle Storage Drive Enclosure DE2-24C disk shelves.

Oracle Storage Drive Enclosure DE2-24P Disk Shelf

Oracle Storage Drive Enclosure DE2-24P is a 2U chassis that supports 24 2.5" SAS-2 drives. The high-capacity HDDs provide reliable storage. The SSDs provide accelerated write operations, and for some controllers, read-optimized cache operations. This disk shelf features dual, redundant I/O Modules (IOMs), and dual power supply with fan modules.

Oracle Storage Drive Enclosure DE2-24P Drive Locations (Front View)

Up to four write-optimized log SSDs are supported per disk shelf. Log devices should be populated in order of slots 20, 21, 22, and 23.

External Read Cache Devices - Read cache devices, installed in Oracle Storage Drive Enclosure DE2-24P disk shelves, are external to the controller and referred to as external second level adaptive replacement cache (EL2ARC). When a cluster takeover occurs, read cache remains persistent and sharable between controllers, enabling read performance to be sustained. For slot configuration details and minimum software requirements, see Disk Shelf Configurations.

Oracle Storage Drive Enclosure DE2-24P Status Indicators (Front View)

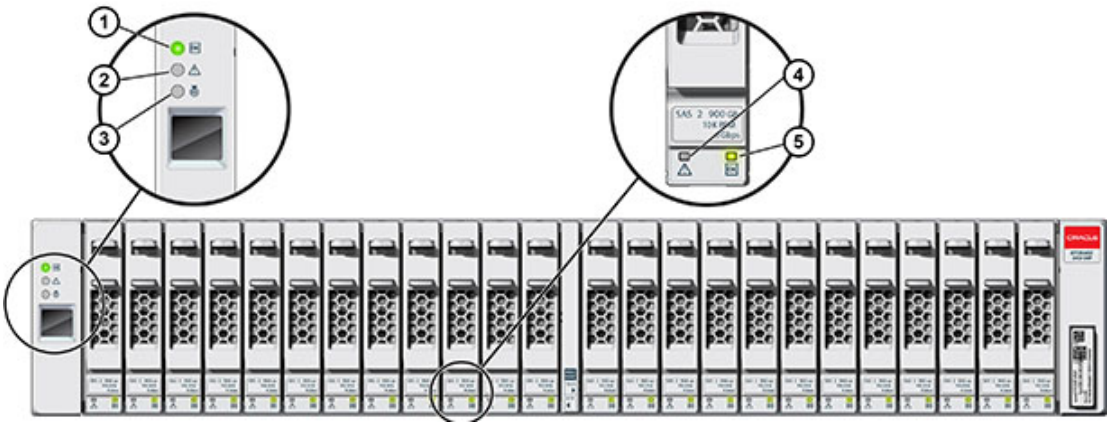

1 System power indicator	**2** Module fault indicator	**3** Locate indicator
4 Power/Activity indicator	**5** Drive fault indicator	

Oracle Storage Drive Enclosure DE2-24P Rear Panel

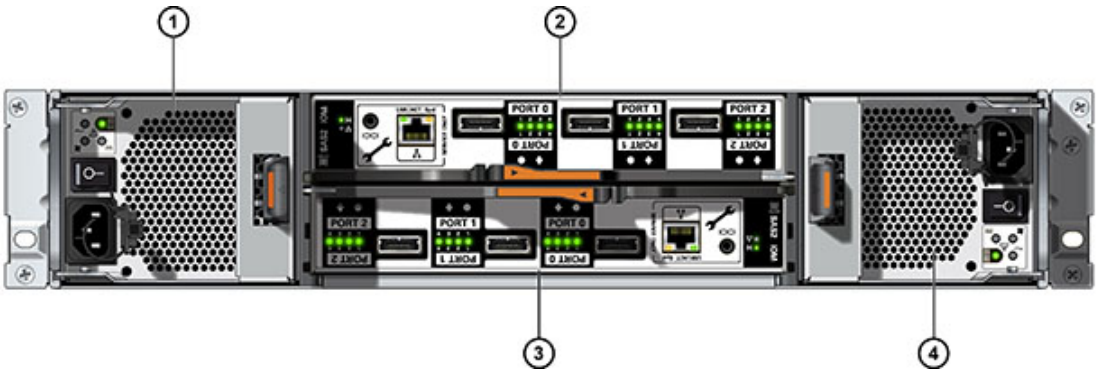

1 Power supply with fan module 0	**2** I/O Module 1
3 I/O Module 0	**4** Power supply with fan module 1

Oracle Storage Drive Enclosure DE2-24C Disk Shelf

Oracle Storage Drive Enclosure DE2-24C is a 4U chassis that supports 24 3.5" SAS-2 drives. The high-capacity HDDs provide reliable storage. The SSDs provide accelerated write operations, and for some controllers, read-optimized cache operations. This disk shelf features dual, redundant I/O Modules (IOMs), and dual power supply with fan modules.

Oracle Storage Drive Enclosure DE2-24C Drive Locations (Front View)

Up to four write-optimized log SSDs are supported per disk shelf. Log devices should be populated in order of slots 20, 21, 22, and 23.

External Read Cache Devices - Read cache devices, installed in Oracle Storage Drive Enclosure DE2-24C disk shelves, are external to the controller and referred to as external second level adaptive replacement cache (EL2ARC). When a cluster takeover occurs, read cache remains persistent and sharable between controllers, enabling read performance to be sustained. For slot configuration details and minimum software requirements, see Disk Shelf Configurations.

Oracle Storage Drive Enclosure DE2-24C Status Indicators (Front View)

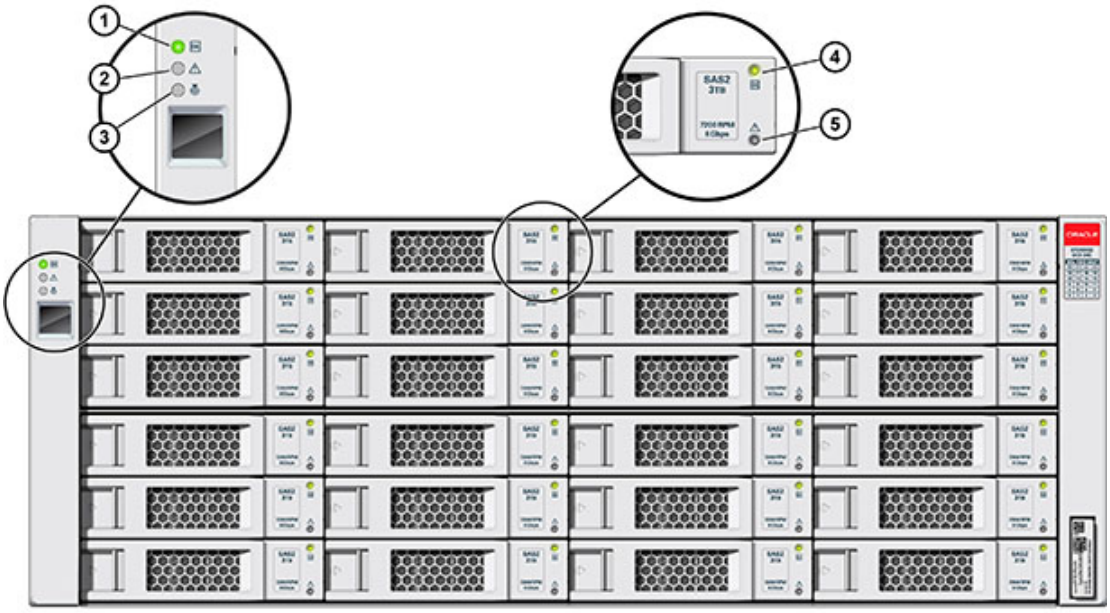

1 System power indicator	**2** Module fault indicator	**3** Locate indicator
4 Power/Activity indicator	**5** Drive fault indicator	

Oracle Storage Drive Enclosure DE2-24C Rear Panel

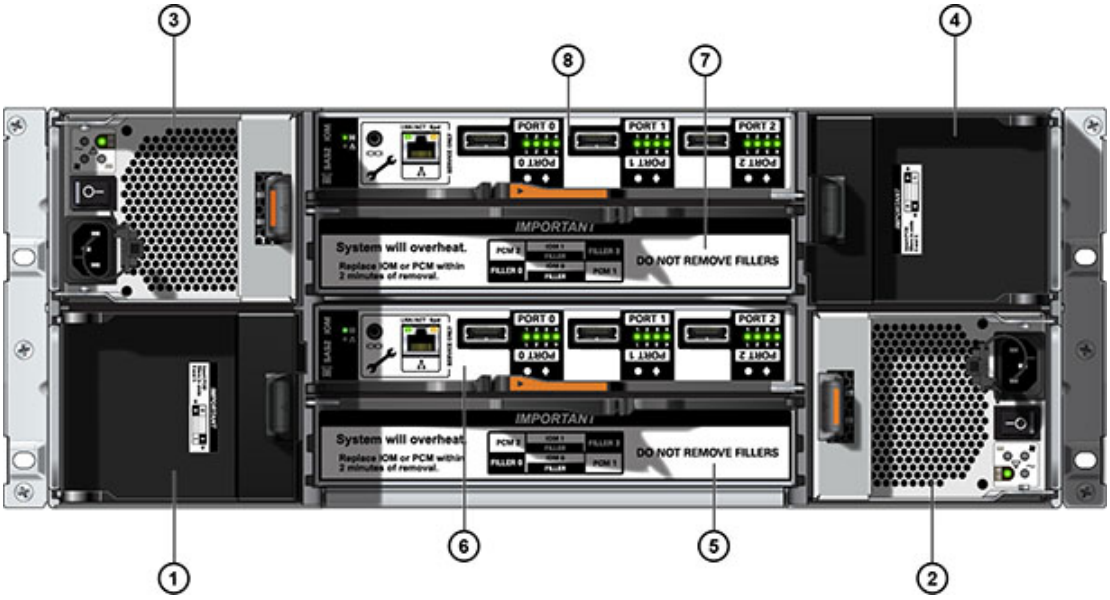

1 Power supply filler panel, slot 0	**2** Power supply with fan module 1	**3** Power supply with fan module 2

4 Power supply filler panel, slot 3	**5** I/O Module filler panel	**6** I/O Module 0
7 I/O Module filler panel	**8** I/O Module 1	

> **Note:**
>
> It is especially important that power supplies and their filler panels are in the correct slots.

Oracle Storage Drive Enclosure DE2-24 I/O Module Indicators

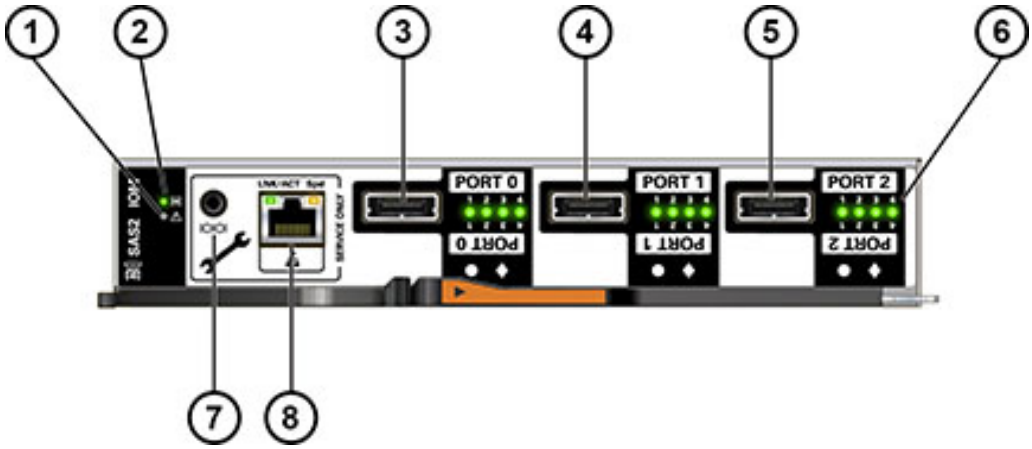

1 Fault/Locate indicator	**2** Power / OK indicator	**3** SAS-2 Port 0
4 SAS-2 Port 1	**5** SAS-2 Port 2	**6** Host port activity indicators
7 For Oracle service only	**8** For Oracle service only	

The Oracle Storage Drive Enclosure DE2-24C and Oracle Storage Drive Enclosure DE2-24P disk shelves use the same I/O Module.

Oracle Storage Drive Enclosure DE2-24 Power Supply Indicators

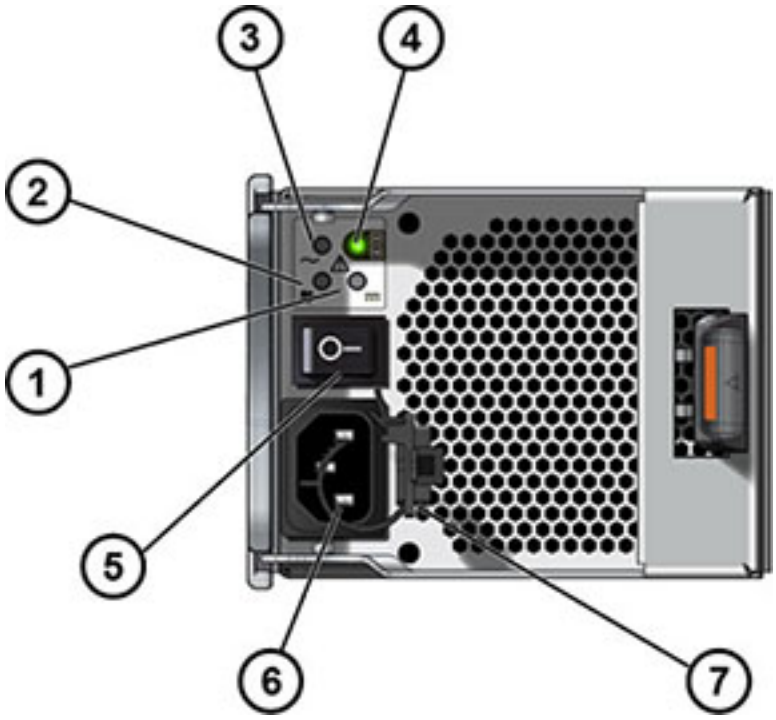

1 DC power fail indicator	**2** Fan fail indicator	**3** AC power fail indicator
4 Power supply status indicator	**5** Power on/off switch	**6** Universal power input connector
7 Power cord tie wrap		

The Oracle Storage Drive Enclosure DE2-24C and Oracle Storage Drive Enclosure DE2-24P disk shelves use the same power supply.

Oracle Storage Drive Enclosure DE3-12C

This section introduces the front and rear panel components of the Oracle Storage Drive Enclosure DE3-12C disk shelf.

Oracle Storage Drive Enclosure DE3-12C is a 2U chassis that supports 12 3.5" SAS-3 drives. The high-performance HDDs provide reliable storage, and the SSDs provide accelerated write operations. This disk shelf features dual, redundant I/O Modules (IOMs), and dual power supply with fan modules.

Oracle Storage Drive Enclosure DE3-12C Drive Locations (Front View)

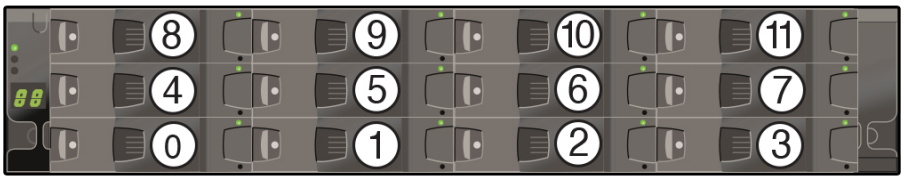

Up to two write-optimized log SSDs are supported per disk shelf. Log devices can be populated in slots 10 and 11.

External Read Cache Devices - Read cache devices, installed in Oracle Storage Drive Enclosure DE3-12C disk shelves, are external to the controller and referred to as external second level adaptive replacement cache (EL2ARC). When a cluster takeover occurs, read cache remains persistent and sharable between controllers, enabling read performance to be sustained. For slot configuration details and minimum software requirements, see Disk Shelf Configurations.

Oracle Storage Drive Enclosure DE3-12C Status Indicators (Front View)

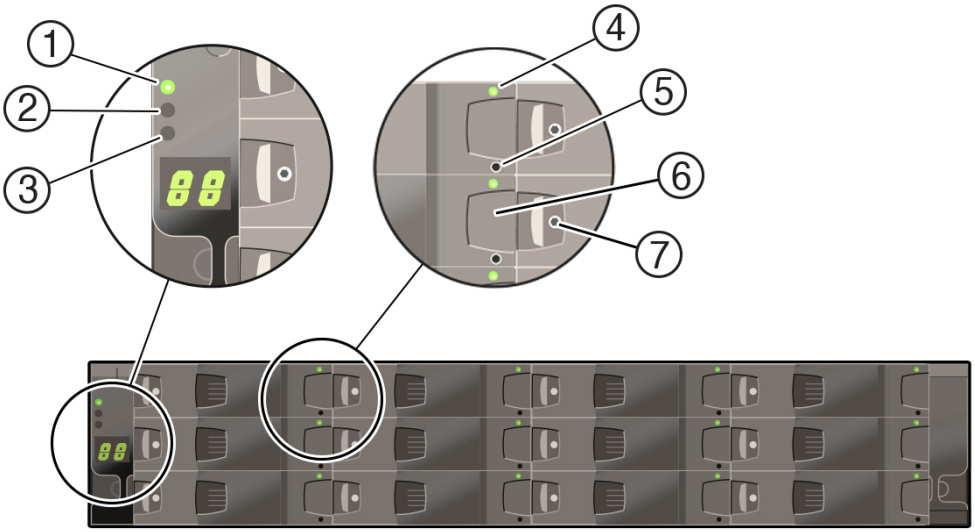

1 Power indicator	**2** Chassis service required indicator	**3** Locate indicator
4 Power/Activity indicator	**5** Drive service required indicator	**6** Serial number tag
7 Locking mechanism - Not used		

Oracle Storage Drive Enclosure DE3-12C Rear Panel

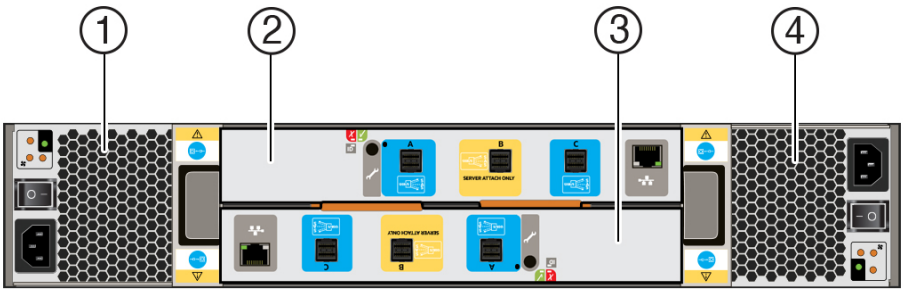

1 Power supply with fan module 0	**2** I/O Module A
3 I/O Module B	**4** Power supply with fan module 1

Oracle Storage Drive Enclosure DE3-12C I/O Module Indicators

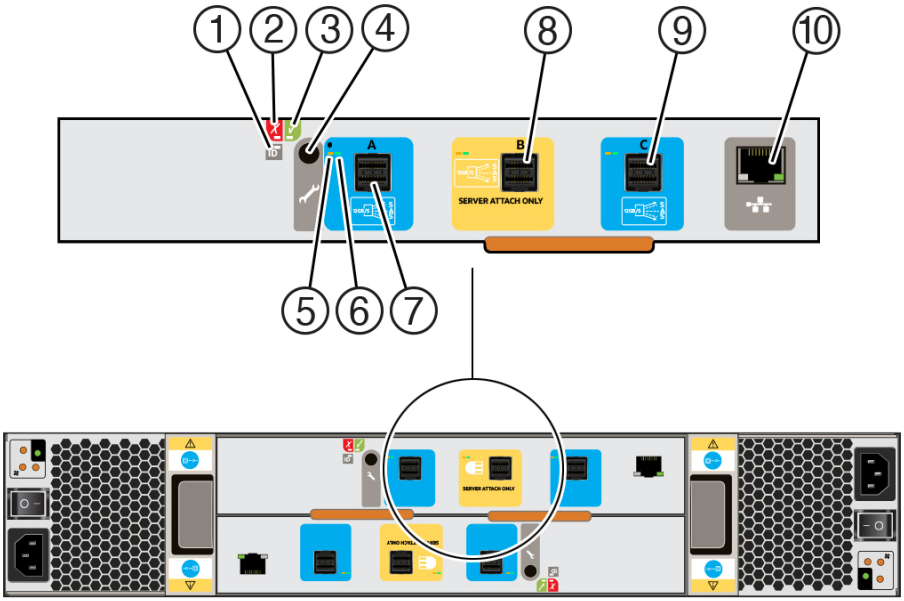

1 Identity indicator: Module is being identified if blinking	**2** I/O Module service required indicator	**3** Power/Activity indicator
4 For Oracle service only	**5** SAS-3 port fault indicator	**6** SAS-3 port activity indicator
7 SAS-3 Port A	**8** SAS-3 Port B - Not used	**9** SAS-3 Port C
10 Not used		

Oracle Storage Drive Enclosure DE3-12C Power Supply Indicators

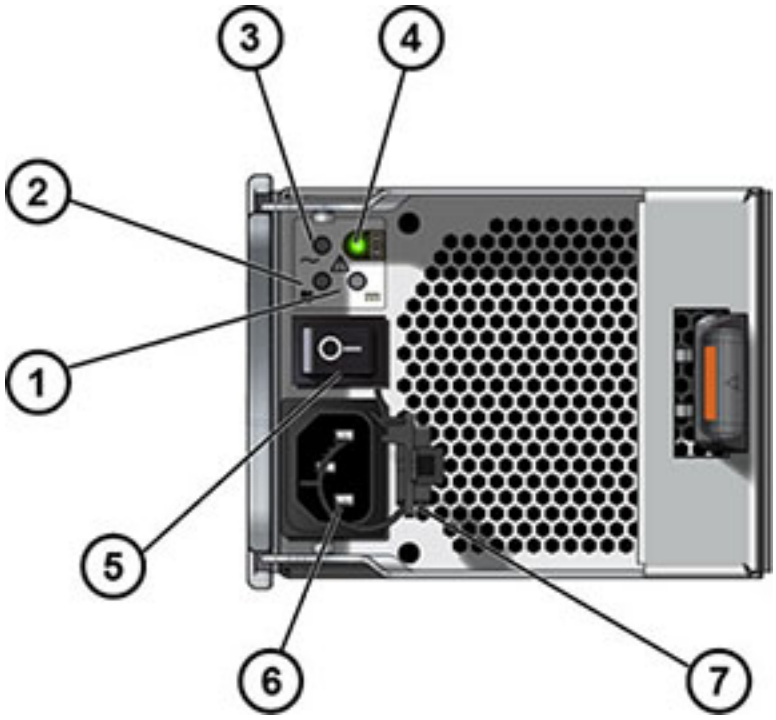

1 DC power fail indicator	**2** Fan fail indicator	**3** AC power fail indicator
4 Power supply status indicator	**5** Power on/off switch	**6** Universal power input connector
7 Power cord tie wrap		

Sun Disk Shelf

Sun Disk Shelf is a 4U chassis that supports 24 3.5" SAS-2 drives. The SSDs provide accelerated write operations, and the high-capacity HDDs provide reliable storage. This disk shelf features dual, redundant SAS Interface Module (SIM) boards, and dual power supply with fan modules.

Sun Disk Shelf Drive Locations (Front View)

Read cache devices cannot be installed in the Sun Disk Shelf.

Sun Disk Shelf Status Indicators (Front View)

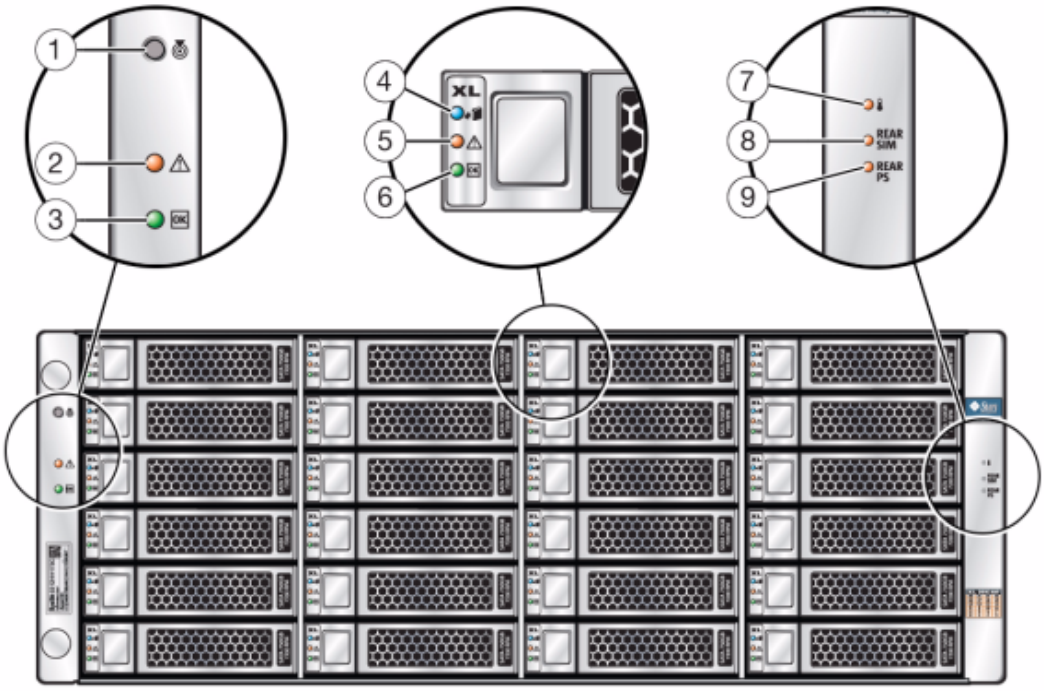

1 Locate button and indicator	**2** System fault indicator	**3** System power indicator
4 Disk ready to be removed indicator	**5** Disk fault indicator	**6** Disk activity indicator
7 Over temperature warning indicator	**8** SIM board fault indicator	**9** Power supply fault indicator

Sun Disk Shelf Rear Panel

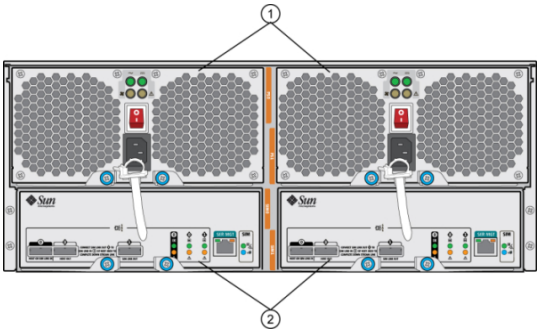

1 Power supply modules with built-in fans. Power supply 0 is on the left and power supply 1 is on the right.	**2** Removable SAS Interface Module (SIM) Boards. SIM 0 is on the left, and SIM 1 is on the right.

Sun Disk Shelf SIM Board Indicators

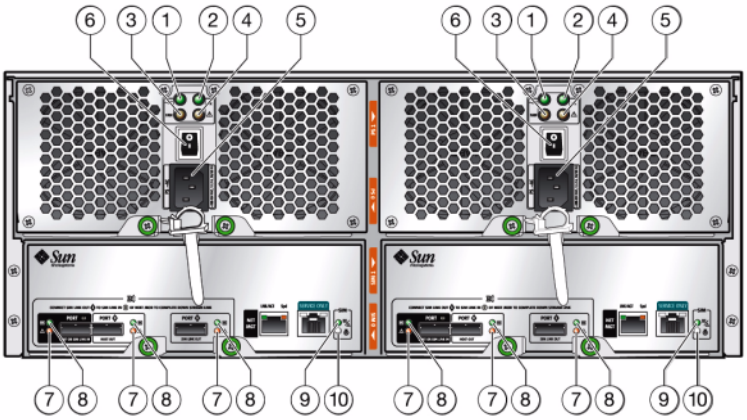

1 AC power indicator	**2** DC power indicator	**3** Fan fault indicator
4 Power supply fault indicator	**5** Universal power connector	**6** Power switch
7 Port fault indicator	**8** Port OK indicator	**9** SIM board OK indicator (green)/SIM board fault indicator (amber)
10 SIM locator indicator		

Sun Disk Shelf Power Supply and Fan Module Indicators

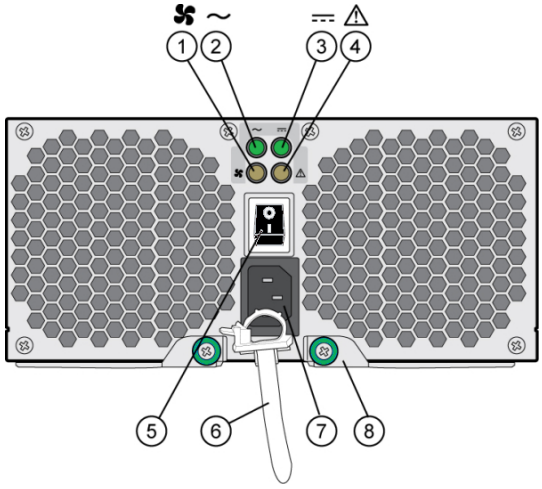

1 Cooling fan status indicator	**2** AC power status indicator	**3** DC power status indicator
4 Power supply status indicator	**5** Power on/off switch	**6** Power cord tie wrap
7 Universal power input connector	**8** Right ejection arm and captive screw latch	

Disk Shelf Configurations

The following tables describe device types supported for each disk shelf. For cabling information, see the Oracle ZFS Storage Appliance Cabling Guide, Release OS8.8.x.

Table 12-1 Oracle Storage Drive Enclosure DE3-24P/Oracle Storage Drive Enclosure DE2-24P Disk Shelf Configurations

Device Types	Data Slots	Cache/Log Slots	Notes/Examples
Data	0 - 23	none	Fully populated HDD data devices.
20 Data and 1-4 Log devices	0 - 19	20 - 23	Twenty HDD data devices and up to four log devices. Example slot configurations: Slot 20 (Log), Slot 21 (Log), Slot 22 (Log), Slot 23 (Log) Slot 20 (Log), Slot 21 (Log), Slot 22 (filler), Slot 23 (filler)
20 Data and 1-4 Cache or Meta devices	0 - 19	20 - 23	Software version OS8.6.0 or later is required for read cache devices. OS8.7.0 or later is required for meta devices, which must be a minimum of 3.2 TB each to support enhanced data deduplication. Up to four read cache or meta devices can be installed in consecutive slots. Two meta devices are recommended for four disk shelves. Example slot configuration: Slot 20 (Cache), Slot 21 (Cache), Slot 22 (filler), Slot 23 (filler)
20 Data and 1-4 Log, Cache, or Meta devices	0 - 19	20 - 23	Software version OS8.6.0 or later is required for read cache devices. OS8.7.0 or later is required for meta devices, which must be a minimum of 3.2 TB each to support enhanced data deduplication. Up to four read cache or meta devices can be installed in consecutive slots. Two meta devices are recommended for four disk shelves. Example slot configuration: Slot 20 (Cache), Slot 21 (Cache), Slot 22 (Log), Slot 23 (Log)

Table 12-2 Oracle Storage Drive Enclosure DE3-24P All-Flash Disk Shelf Configurations

Device Types	Data Slots	Cache/Log Slots	Notes/Examples
24 Data	0 - 23	none	Software version OS8.7.0 or later is required for all-flash pools.
20 Data and 1-4 Log devices	0 - 19	20 - 23	Twenty SSD data devices and up to four log devices installed in consecutive slots. Even with SSD data devices, log devices can boost performance by 15-20% for small block transactional workloads. Example slot configurations: Slot 20 (Log), Slot 21 (Log), Slot 22 (Log), Slot 23 (Log) Slot 20 (Log), Slot 21 (Log), Slot 22 (filler), Slot 23 (filler)

Table 12-3 Oracle Storage Drive Enclosure DE3-24C/Oracle Storage Drive Enclosure DE2-24C Disk Shelf Configurations

Device Types	Data Slots	Cache/Log Slots	Notes/Examples
Data	0 - 23	none	Fully populated HDD data devices.
20 Data and 1-4 Log devices	0 - 19	20 - 23	Twenty HDD data devices and up to four log devices installed in consecutive slots. Example slot configurations: Slot 20 (Log), Slot 21 (Log), Slot 22 (Log), Slot 23 (Log) Slot 20 (Log), Slot 21 (Log), Slot 22 (filler), Slot 23 (filler)
20 Data and 1-4 Cache or Meta devices	0 - 19	20 - 23	Software version OS8.6.0 or later is required for read cache devices. OS8.7.0 or later is required for meta devices, which must be a minimum of 3.2 TB each to support enhanced data deduplication. Up to four read cache or meta devices installed in consecutive slots. (Two meta devices recommended for four disk shelves.) Example slot configuration: Slot 20 (Cache), Slot 21 (Cache), Slot 22 (filler), Slot 23 (filler)
20 Data and 1-4 Log, Cache, or Meta devices	0 - 19	20 - 23	Software version OS8.6.0 or later is required for read cache devices. OS8.7.0 or later is required for meta devices, which must be a minimum of 3.2 TB each to support enhanced data deduplication. Twenty HDD data devices and up to four log, read cache, or meta devices installed in consecutive slots. (Two meta devices recommended for four disk shelves.) Example slot configuration: Slot 20 (Cache), Slot 21 (Cache), Slot 22 (Log), Slot 23 (Log)

Table 12-4 Oracle Storage Drive Enclosure DE3-12C Disk Shelf Configurations

Device Types	Data Slots	Cache/Log Slots	Notes/Examples
Data	0 - 11	none	Fully populated HDD data devices.
10 Data and 1-2 Log devices	0 - 9	10 - 11	Ten HDD data devices and up to two log devices. If no read cache/meta devices will be installed, then two log devices must be installed. Configurations do not include filler panels.
10 Data and 1-2 Cache or Meta devices	0 - 9	10 - 11	Ten HDD data devices and up to two read cache or meta devices. One log device and one read cache/meta device are allowed. Configurations do not include filler panels. Two meta devices are recommended for eight disk shelves. If using two DE3-12C disk shelves together, one meta device per disk shelf satisfies two meta devices for eight disk shelves. Example slot configurations: Slot 10 (Log), Slot 11 (Cache) Slot 10 (Cache), Slot 11 (Log) Slot 10 (Cache), Slot 11 (Cache)

Table 12-5 Sun Disk Shelf (DS2) Disk Shelf Configurations

Device Types	Data Slots	Cache/Log Slots	Notes/Examples
Data	0 - 23	none	Fully populated HDD data devices.
20 Data and 1-4 Log devices	0-23	20 - 23	20 HDD data devices and up to four SSD log devices installed in consecutive slots. **Read cache devices and meta devices cannot be installed in the Sun Disk Shelf.**

Maximum Disk Shelves per Controller Configuration

When determining how many disk shelves a controller configuration supports, it is important to remember that each disk shelf chain can support up to six disk shelves, and some controller configurations are limited to a specific number of total disk shelves. The following table details the number of supported disk shelf chains per HBA type.

Table 12-6 Supported Disk Shelf Chains per HBA

HBA Type	Number of Disk Shelf Chains per HBA
SAS-3 4x4 port	2
SAS-2 4x4 port	2
SAS-2 2x4 port	1

Controllers cannot simultaneously use 2x4 port SAS-2 HBAs and 4x4 port SAS-2 HBAs. As of software release OS8.7.0, 4x4 port SAS-3 HBAs and 4x4 port SAS-2 HBAs can be used together in the same system. Sun Disk Shelves are not supported in any system that includes a SAS-3 HBA, or a combination of SAS-3 and SAS-2 HBAs.

The following table shows the maximum number of HBAs per controller, and the maximum number of disk shelves supported per controller configuration. Both standalone and clustered controllers support the same maximum number of disk shelves.

Table 12-7 Maximum Disk Shelves per Controller Configuration

Controller	Max. Shelves	Max. 2x4 Port SAS-2 HBA	Max. 4x4 Port SAS-2 HBA	Max. 4x4 Port SAS-3 HBA	Guidelines
Oracle ZFS Storage ZS9-2 high-end model	48 (1152 drives, including filler panels)	NA	NA	4	Disk shelves can include 24x HDD/SSD, or combine 20x HDD/SSD plus 4x write-optimized and/or read-optimized flash accelerators (SSDs). Total number of chains is eight. Any combination of disk-only and write- or read-optimized shelves may be combined within the chain in any order. When employing Oracle Storage Drive Enclosure DE3-12C disk shelves, twoOracle Storage Drive Enclosure DE3-12C disk shelves are the equivalent of one Oracle Storage Drive Enclosure DEx-24 disk shelf. Thus, the maximum shelves is 96, which is still 1152 drives. The Oracle Storage Drive Enclosure DE3-12C disk shelf can include 12x HDD, or combine 10x HDD plus 2x write-optimized and/or read-optimized flash accelerators (SSDs). Total number of chains is still eight. When combining Oracle Storage Drive Enclosure DE3-12C and Oracle Storage Drive Enclosure DE3-24 disk shelves, the maximum allowed number of drives, including filler panels, is 1152 drives.
Oracle ZFS Storage ZS9-2 mid-range model	24 (576 drives, including filler panels)	NA	NA	2	Disk shelves can include 24x HDD/SSD, or combine 20x HDD/SSD plus 4x write-optimized and/or read-optimized flash accelerators (SSDs). Total number of chains is four. Any combination of disk-only and write- or read-optimized shelves may be combined within the chain in any order. When employing Oracle Storage Drive Enclosure DE3-12C disk shelves, two Oracle Storage Drive Enclosure DE3-12C disk shelves are the equivalent of one Oracle Storage Drive Enclosure DEx-24 disk shelf. Thus, the maximum shelves is 48, which is still 576 drives. The Oracle Storage Drive Enclosure DE3-12C disk shelf can include 12x HDD, or combine 10x HDD plus 2x write-optimized and/or read-optimized flash accelerators (SSDs). Total number of chains is still four. When combining Oracle Storage Drive Enclosure DE3-12C and Oracle Storage Drive Enclosure DE3-24 disk shelves, the maximum allowed number of drives, including filler panels, is 576 drives.

Table 12-7 (Cont.) Maximum Disk Shelves per Controller Configuration

Controller	Max. Shelves	Max. 2x4 Port SAS-2 HBA	Max. 4x4 Port SAS-2 HBA	Max. 4x4 Port SAS-3 HBA	Guidelines
Oracle ZFS Storage ZS7-2 high-end model	48	NA	NA	4	Disk shelves can include 24x HDD/SSD, or combine 20x HDD/SSD plus 4x write-optimized and/or read-optimized flash accelerators (SSDs). Total number of chains is eight. Any combination of disk-only and write- or read-optimized shelves may be combined within the chain in any order.
Oracle ZFS Storage ZS7-2 mid-range model	24	NA	NA	2	Disk shelves can include 24x HDD/SSD, or combine 20x HDD/SSD plus 4x write-optimized and/or read-optimized flash accelerators (SSDs). Total number of chains is four. Any combination of disk-only and write- or read-optimized shelves may be combined within the chain in any order.
Oracle ZFS Storage ZS5-4	48	NA	4	4	Disk shelves can include 24x HDD/SSD, or combine 20x HDD/SSD plus 4x write-optimized and/or read-optimized flash accelerators (SSDs). Total number of chains is eight. Any combination of disk-only and write- or read-optimized shelves may be combined within the chain in any order.
Oracle ZFS Storage ZS5-2	16	NA	2	2	Disk shelves can include 24x HDD/SSD, or combine 20x HDD/SSD plus 4x write-optimized and/or read-optimized flash accelerators (SSDs). Total number of chains is four. Any combination of disk-only and write- or read-optimized shelves may be combined within the chain in any order.
Oracle ZFS Storage ZS4-4	36	NA	4	4	Disk shelves can include 24x HDD/SSD, or combine 20x HDD/SSD plus 4x write-optimized and/or read-optimized flash accelerators (SSDs). Total number of chains is eight. Any combination of disk-only and write-optimized shelves can be combined within the chain in any order.
Oracle ZFS Storage ZS3-4	36	NA	4	NA	Disk shelves can include 24x HDD, or combine 20x HDD plus 4x write-optimized flash accelerators (SSDs). Total number of chains is eight. Any combination of disk-only and write-optimized shelves can be combined within the chain in any order. You can also connect mixed disk shelf types (DE2 and legacy Sun Disk Shelves) to the same controllers, but each chain must contain only the same disk shelf type. Directly connecting different disk shelf types is not supported.

Table 12-7 (Cont.) Maximum Disk Shelves per Controller Configuration

Controller	Max. Shelves	Max. 2x4 Port SAS-2 HBA	Max. 4x4 Port SAS-2 HBA	Max. 4x4 Port SAS-3 HBA	Guidelines
Oracle ZFS Storage ZS3-2	16	NA	2	2	Disk shelves can include 24x HDD, or combine 20x HDD plus 4x write-optimized flash accelerators (SSDs). Total number of chains is four. Any combination of disk-only and write-optimized shelves can be combined within the chain in any order. You can also connect mixed disk shelf types (DE2 and legacy Sun Disk Shelves) to the same controllers, but each chain must contain only the same disk shelf type. Directly connecting different disk shelf types is not supported.
Sun ZFS Storage 7420	36	6	6	NA	Disk shelves can include 24x HDD, or combine 20x HDD plus 4x write-optimized flash accelerators (SSDs). Total number of chains is six with 2X4 port SAS-2 HBAs, and 12 with 4X4 port SAS-2 HBAs. Any combination of disk-only and write-optimized shelves can be combined within the chain in any order. You can also connect mixed disk shelf types (DE2 and legacy Sun Disk Shelves) to the same controllers, but each chain must contain only the same disk shelf type. Directly connecting different disk shelf types is not supported.
Sun ZFS Storage 7320	6	1	1	NA	Any combination of disk-only and write-optimized shelves may be combined within the chain in any order. The cabling configurations are unchanged. You can also connect mixed disk shelf types (DE2 and legacy Sun Disk Shelves) to the same controllers, but each chain must contain only the same disk shelf type. Directly connecting different disk shelf types is not supported.
Sun ZFS Storage 7120	2	1	NA	NA	Write-optimized SSDs are not supported in the expansion storage for the 7120. The disk shelves must be fully populated with 24 HDDs. You can also connect mixed disk shelf types (DE2 and legacy Sun Disk Shelves) to the same controllers.

SAS-2 and SAS-3

With both the SAS-2 and SAS-3 fabrics, you are encouraged to apply entire shelves to pools, so you can gain the benefits of No Single Point of Failure (NSPF) configurations, and striping across the maximum possible number of devices.

SAS-2: The SAS-2 (Serial Attached SCSI 2.0) host bus adapter (HBA) is designed for Oracle ZFS Storage Appliance with a standard chip set to support a high-density of

target devices, up to 1024. You can daisy-chain SAS-2 disks to as many as 36 disk shelves for certain systems, for a total of 864 disks. The SAS-2 HBA has 2x4 (8 PHYs) or 4x4 port (16 PHYs) with each supporting 6 Gbit/s (SAS-2) devices. For more information about the SAS-2 HBA, see the hardware overview for your controller.

For supported disk shelf configurations, see Disk Shelf Configurations.

SAS-3: The SAS-3 (Serial Attached SCSI 3.0) host bus adapter (HBA) increases speed to 12 Gbit/s and is capable of supporting up to 4095 devices. You can daisy-chain SAS-3 disks to as many as 48 disk shelves for certain systems, for a total of 1152 disks. The SAS-3 HBA is a low-profile, 4x4 port (16 PHYs), with each port supporting 12 Gbit/s (SAS-3) devices. For more information about the SAS-3 HBA, see the hardware overview for your controller.

For supported disk shelf configurations, see Disk Shelf Configurations.

13
Viewing System Disks Status

The system disks section shows the status of the system disks and their current usage. The BUI displays this with a pie-chart, and the CLI as a text list. For example:

```
hostname:> maintenance system disks show
Properties:
                       profile = mirror
                          root = 5.10G
                           var = 374M
                        update = 1.86G
                         stash = 226M
                          dump = 36.0G
                         cores = 210M
                       unknown = 35.0G
                          free = 3.36T

Disks:

DISK        LABEL      STATE
disk-000    NVME 1     healthy
disk-001    NVME 0     healthy
```

> **Note:**
>
> The `DISK` column is not required by the BUI.

14

Working with Support Bundles

Oracle ZFS Storage Appliance can generate support bundles containing system configuration information and core files for use by Oracle support personnel for troubleshooting system problems. When the Phone Home service is enabled, support bundles will be generated automatically and securely uploaded in response to faults. You can also manually generate and upload bundles using the BUI or CLI.

To facilitate support bundle uploads, you must first:

- Create an Oracle Single Sign-On account. Go to My Oracle Support and click **Register** to create your account.
- Register the appliance with the Phone Home service. For information on Phone Home, see Phone Home Configuration in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

In a cluster configuration, the controller that originally diagnoses a fault uploads the support bundle to Oracle Service. For example, a fault occurring on a disk shelf that is shared by both controllers is reported on both controllers, but the fault is posted by the controller that initially diagnosed the fault.

Once generated, support bundles are authenticated and securely uploaded to Oracle Service. After a support bundle has been successfully uploaded, the support bundle and core files are automatically deleted from the appliance.

If you have not registered the appliance with Phone Home before a support bundle is uploaded, you will be given the option to continue collect the bundle for local analysis or to cancel the upload.

Support Bundle Options

For each support bundle currently being generated or uploaded or which has failed to upload, the following options may be available:

Table 14-1 Support Bundle Options

Icon	Description
	Cancel the current operation. If the bundle is being generated, it will be deleted. If the bundle is being uploaded, the upload will be cancelled and the appliance will not retry it later.
	Download the support bundle. This icon appears after the bundle is built and the upload is cancelled or fails.
	Try again to upload the bundle to support.
	Cancel any pending operation and delete the support bundle.

Generating and Uploading a Support Bundle (BUI)

Be sure the Phone Home service is configured for Oracle ZFS Storage Appliance. If Phone Home is not configured, the following message is displayed:

```
You must register for Phone Home before uploading
support bundles to Oracle Support. Continue to collect the
bundle for local analysis, or Cancel to register your
account and activate Phone Home.
```

Use the following procedure to generate and upload a support bundle.

1. From the **Maintenance** BUI menu, select **System**.
2. Click the add icon ⊕ next to **Support Bundles**.
3. Enter the Service Request (SR) number that was provided to you by Oracle Support. The SR number format is 3-nnnnnnnnnn, where n is a digit between 0 and 9.

   (Optional) Uncheck **Send to Oracle Support** to create a local bundle for analysis. This bundle can be uploaded to Oracle Support later if an SR is provided.

Generating and Uploading a Support Bundle (CLI)

Use the following procedure to generate and upload a support bundle.

1. Go to `maintenance system bundles`.

   ```
   hostname:> maintenance system bundles
   ```

2. Enter `bundle`, space, and the Service Request (SR) number.

   ```
   hostname:maintenance system bundles> bundle 3-9999999999
   A support bundle is being created and sent to Oracle. You will receive an
   alert
   when the bundle has finished uploading. Please save the following filename,
   as
   Oracle support personnel will need it in order to access the bundle:
      /upload/issue/3-9999999999/3-9999999999_ak.edd08324-f091-cd19-b4c9-
   bb508d6ecc59.tar.gz
      /upload/issue/3-9999999999/3-9999999999_ak.fee4b98a-e3f7-e600-877d-
   b6d7933255f3.tar.gz
   ```

   > **Note:**
   >
   > If the SR number is not currently available, the system can build a local bundle for later upload. See the next step.

3. To build a local bundle, enter the `bundle` command:

   ```
   hostname:maintenance system bundles> bundle
   The support data you requested is being built in 2 files. Use 'send <srn>'
   with
   each bundle to associate the bundle with a Service Request Number and send it
   to Oracle Support. Alternatively, you may download the bundles via the
   appliance BUI.
   ```

```
ak.ce9e680f-15ef-635d-bb77-9888205da52f.tar.gz
ak.32d3494a-e1ae-4c80-ce59-da107e0eaccf.tar.gz
```

Related Topics

- Viewing Support Bundle Details (CLI)
- Retrying a Failed Support Bundle Upload (CLI)
- Canceling a Pending Operation (CLI)
- Deleting a Support Bundle (CLI)

Viewing Support Bundle Details (BUI)

Use the following procedure to view support bundle details.

1. From the **Maintenance** menu, select **System**.
2. Hover over the appropriate support bundle row, and click its information icon .
3. Click **OK**.

Viewing Support Bundle Details (CLI)

Use the following procedure to view support bundle details.

1. Go to `maintenance system bundles` and enter `list`.

```
hostname:maintenance system bundles> list

BUNDLE                                  SRNUMBER          STATUS        PROGRESS
edd08324-f091-cd19-b4c9-bb508d6ecc59    3-9999999999      Building      6  %
fee4b98a-e3f7-e600-877d-b6d7933255f3    3-9999999999      Building      66 %
```

2. Enter `select`, space, and the UUID.

```
hostname:maintenance system bundles> select edd08324-f091-cd19-b4c9-
bb508d6ecc59
```

3. Enter `show`.

   In this example, the read-only properties for support bundle `edd08324-f091-cd19-b4c9-bb508d6ecc59` indicate that the appliance has uploaded 6% of the file.

```
hostname:maintenance system bundles edd08324-f091-cd19-b4c9-bb508d6ecc59> show
Properties:
                    filename = /upload/uuid/edd08324-f091-cd19-b4c9-
bb508d6ecc59/ak.
       edd08324-f091-cd19-b4c9-bb508d6ecc59.tar.gz
                      status = building
                        date = 2022-1-9 17:42:09
                        type = User initiated
                step_progress = 6
```

Related Topics

- Retrying a Failed Support Bundle Upload (CLI)
- Canceling a Pending Operation (CLI)
- Deleting a Support Bundle (CLI)

Uploading a Locally Generated Bundle (CLI)

Use the following procedure to upload a locally generated support bundle.

1. To select a support bundle, enter `select`, space, and the bundle's UUID.

   ```
   hostname:maintenance system bundles> select fee4b98a-e3f7-e600-877d-b6d7933255f3
   ```

2. To upload the selected bundle, use the `send` command, space, and the SR number:

   ```
   hostname:maintenance system bundles fee4b98a-e3f7-e600-877d-b6d7933255f3> send 3-9999999999
   ```

> **Note:**
>
> An SR number is required when running `send` on a locally generated bundle or an error is thrown.

Retrying a Failed Support Bundle Upload (BUI)

Use the following procedure to retry a failed support bundle upload.

1. From the **Maintenance** menu, select **System**.
2. Hover over the appropriate support bundle row, and click its send icon .

   The **Send support bundle** dialog box is displayed.

Retrying a Failed Support Bundle Upload (CLI)

Use the following procedure to select and retry a failed upload using the CLI.

1. Go to `maintenance system bundles`.

   ```
   hostname:> maintenance system bundles
   ```

2. Enter `select`, space, and the Service Request (SR) number:

   ```
   hostname:maintenance system bundles> select fee4b98a-e3f7-e600-877d-b6d7933255f3
   ```

3. To retry a failed upload, use the `send` command:

   ```
   hostname:maintenance system bundles fee4b98a-e3f7-e600-877d-b6d7933255f3> send
   ```

> **Note:**
>
> An SR number is required when sending a locally generated bundle. If an SR number is not entered, an error occurs. See Uploading a Locally Generated Bundle (CLI).

4. To retry a bundle upload with a new SR number, include the SR number after the `send` command:

```
hostname:maintenance system bundles fee4b98a-e3f7-e600-877d-b6d7933255f3> send
3-9999999999
```

Canceling a Pending Support Bundle Operation (BUI)

Use the following procedure to cancel a pending support bundle operation.

1. From the **Maintenance** menu, select **System**.
2. Hover over the appropriate support bundle row, and click its cancel icon .

Canceling a Pending Operation (CLI)

Use the following procedure to cancel a pending operation using the CLI.

1. To select a support bundle to cancel, enter `select`, space, and the SR number:

```
hostname:maintenance system bundles> select fee4b98a-e3f7-e600-877d-
b6d7933255f3
```

2. To cancel the pending operation, use the `cancel` command:

```
hostname:maintenance system bundles 5ff532a2-2377-e72d-b0fe-f2efc2aa8aaf> cancel
```

Deleting a Support Bundle (BUI)

Use the following procedure to delete a support bundle.

1. From the **Maintenance** menu, select **System**.
2. Hover over the appropriate support bundle row, and click its trash icon .

Deleting a Support Bundle (CLI)

Use the following procedure to delete a support bundle using the CLI.

1. To select a support bundle to delete, enter `select`, space, and the SR number:

```
hostname:maintenance system bundles> select fee4b98a-e3f7-e600-877d-
b6d7933255f3
```

2. To delete the support bundle, use the `destroy` command:

```
hostname:maintenance system bundles 5ff532a2-2377-e72d-b0fe-f2efc2aa8aaf>
destroy
```

15

Performing the Initial Setup

You can repeat the initial setup process after you have already configured Oracle ZFS Storage Appliance for the first time, as described in Configuring the Appliance for the First Time in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*. This does not change any of the current settings unless explicitly requested. User data on the storage pool (including projects and shares) is not affected.

To perform an initial setup:

- **BUI:** From the **Maintenance** menu, select **System**, and click the **Initial Setup** button.
- **CLI:** Go to `maintenance system`, and enter command `setup`.

16

Restarting the Appliance

> **Caution:**
>
> Only issue the `restart` command as instructed by Oracle Service personnel.

This function is only available via the CLI and is not the same as a reboot (power cycle) via the BUI. The `restart` command is a software-only operation that restarts the management controller, which could impact some client services, like replication. During execution, both the CLI and BUI are not available; wait for Oracle ZFS Storage Appliance to return to normal operation.

To restart Oracle ZFS Storage Appliance, go to `maintenance system`, and enter the `restart` command:

```
hostname:maintenance system> restart
```

17

Performing a Factory Reset

A factory reset operation resets Oracle ZFS Storage Appliance in the following ways and reboots the appliance.

- **Software version:** The factory reset operation sets the appliance to the current version. The current version is shown in the following ways:
  - BUI: The version that is shown as the current version in the **Software Updates** section in **Maintenance: System** prior to the factory reset.
  - CLI: The version that is marked `current` when you use the following command from the `maintenance system updates` context prior to the factory reset:

    ```
    hostname:maintenance system updates> show
    ```
- **Appliance configuration:** The factory reset operation resets the Oracle ZFS Storage Appliance configuration back to factory settings for the current version described in “Software version” earlier. All configuration changes are discarded and initial configuration is performed again.
- **Deferred updates:** The factory reset operation applies all deferred updates. Deferred updates are shown in the following ways:
  - BUI: Updates that are shown in the **Deferred Updates** section in **Maintenance: System** prior to the factory reset.
  - CLI: Updates that are marked `waiting` when you use the `show` command prior to the factory reset.

User data on the storage pool, including projects and shares, is not affected; however, the pool must be imported during the configuration process.

Loaded software images remain on the system and can still be used to upgrade.

Factory Resetting a Standalone Controller

Use one of the following methods to perform a factory reset of a standalone Oracle ZFS Storage Appliance controller. During the storage pool configuration step, import the existing storage pool.

- BUI: From the **Maintenance** menu, select **System**, and click **Factory Reset**.
- CLI: Go to `maintenance system` and enter the command `factoryreset`.
- SSH or Serial Console: While rebooting, select `Factory Reset` from the boot menu.

Use one of the following methods to perform initial configuration:

- Performing Initial Configuration (BUI) in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*
- Performing Initial Configuration (CLI) in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*

Use one of the following methods to import the pool:

- Importing an Existing Storage Pool (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*
- Importing an Existing Storage Pool (CLI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*

Factory Resetting Clustered Controllers

Because factory reset of an Oracle ZFS Storage Appliance controller while in a clustered configuration is not supported, use the following procedure to unconfigure the clustered controllers, factory reset each standalone controller, re-cluster the controllers as part of initial configuration, and import the existing storage pool. Both controllers must be factory reset before they can be clustered together again.

1. Use one of the following methods to perform a factory reset on controller A. When controller A starts to reboot, immediately turn off its power.
   - BUI: From the **Maintenance** menu, select **System**, and click **Factory Reset**.
   - CLI: Go to `maintenance system` and enter the command `factoryreset`.
   - SSH or Serial Console: While rebooting, select `Factory Reset` from the boot menu.
2. On controller B, use one of the following methods to unconfigure the cluster.
   - BUI: From the **Configuration** menu, select **Cluster**, and click **UNCONFIG**.
   - CLI: Go to `configuration cluster` and enter the command `unconfig`.
3. On controller B, use one of the following methods to perform a factory reset.
   - BUI: From the **Maintenance** menu, select **System**, and click **Factory Reset**.
   - CLI: Go to `maintenance system` and enter the command `factoryreset`.

   After completion, the `*` prompt is displayed.
4. Turn on the power for controller A.

   After completion, the `*` prompt is displayed.
5. Use one of the following methods to perform initial configuration on controller A and to make it into a clustered configuration with controller B. During the storage pool configuration step, import the existing storage pool.
   - BUI: Performing Initial Configuration (BUI) in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*, and Importing an Existing Storage Pool (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.
   - CLI: Performing Initial Configuration (CLI) in *Oracle ZFS Storage Appliance Installation Guide, Release OS8.8.x*, and Importing an Existing Storage Pool (CLI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

18

Backing Up the Configuration

The configuration backup function enables an administrator to:

- **Back up** the Oracle ZFS Storage Appliance configuration, consisting of system metadata only (such as the network configuration, local users and roles, service settings, and other appliance metadata).
- **Restore** a previously saved configuration from a backup.
- **Export** a saved configuration, as a plain file, so that it may be stored on an external server, or included in a backup of a share on the appliance itself.
- **Import** a saved configuration that was previously exported from this system or another system, making it available for a restore operation.

To back up the configuration and manage the backups, use these tasks:

- Creating a Configuration Backup - BUI, CLI
- Restoring from a Saved Configuration - BUI, CLI
- Deleting a Saved Configuration - BUI, CLI
- Exporting a Saved Configuration - BUI, CLI
- Importing a Saved Configuration - BUI, CLI
- Viewing Appliance Configurations - BUI, CLI

To understand configuration backups, use these topics:

- Configuration Backup Contents
- Configuration Restore Impact
- Security Considerations for Configuration Backups

Configuration Backup Contents

A configuration backup *does* include:

- Metadata associated with the system as a whole, such as settings for NTP, NIS, LDAP, and other services.
- Network device, datalink, and interface configuration.
- User accounts, roles and privileges, preferences, and encrypted passwords for local users (not directory users).
- Alerts and thresholds and their associated rules.
- iSCSI targets and initiators.
- iSCSI service properties.
- SRP targets and initiators.

- Encryption configuration: All contents of the LOCAL keystore, including the keys (stored in an encrypted form). Configuration details, including the Registration PIN for OKM and the names only of the keys.

> **Note:**
>
> In a cluster configuration, configuration backups are visible only on the node on which they were created or imported.

A configuration backup *does not* include:

- User data (shares and LUNs). Your user data must be backed up separately, using NDMP backup software, snapshots, and/or remote replication.
- User passwords for directory users. These remain stored solely in your separate network directory service, such as LDAP or Active Directory, and are not stored in the backup or restored.
- Metadata directly associated with user data, such as snapshot schedules, user quotas, compression settings, and other attributes of shares and LUNs.
- Analytics and logs. Events can be redirected to external SNMP trap receivers or e-mail destinations using Alerts rules.
- System software. The system software is automatically backed up as part of the System Update capability.
- Replication targets.

Configuration Restore Impact

The restore operation takes a selected configuration backup, and modifies all of the corresponding system settings to reflect those in the backup, including removing aspects of the configuration that were not present at the time of the backup. Administrators should adhere to the following guidelines when planning a restore operation:

- **Cross-Platform Restore:** A configuration backed up from one Oracle ZFS Storage Appliance platform cannot be restored to a different Oracle ZFS Storage Appliance platform. For example, an Oracle ZFS Storage ZS9-2 configuration backup cannot be restored to an Oracle ZFS Storage ZS7-2. Attempting to restore a configuration to a different platform results in several failures that requires a factory reset to return to a normal state.
- **Scheduled downtime:** The restore process takes several minutes to complete and impacts service to clients, as the active networking configuration and data protocols are reconfigured. Therefore, a configuration restore should only be used on a development system, or during a scheduled downtime.
- **Service interruption:** Clients accessing data on the system through a data protocol, such as NFS, experience service interruption as the network is reconfigured and the service is restarted. If the selected backup copy was taken when a service was disabled by the administrator, that setting is restored, and therefore client sessions are terminated for that protocol.
- **Session interruption:** If restore is initiated from a web browser, that web browser session is also disconnected during the restore process as the network is

reconfigured. If the restored configuration does not include the same routing and network address settings used by the current browser connection, or if the browser is connected to a network address managed by DHCP, the browser session is interrupted during the restore. The restore process completes in the background, but you need to reload or point the browser at a new, restored network address to continue. For this reason, it may be desirable to initiate a complex configuration restore from the service processor serial console using the CLI.

- **Un-cluster, restore, and re-cluster:** Configuration backups may be initiated for appliances that are joined in a cluster, but a configuration restore may *not* be used while systems are actively clustered. The clustering process means that settings are being synchronized between cluster peers, and each peer appliance also is maintaining private settings. For this reason, you must first un-cluster the two systems. Because of the possibility of data loss to unconfigure clustering, contact Oracle support. Then, restore the configuration backup on a selected head, and then re-cluster the two systems, at which point the other system automatically synchronizes itself with the restored configuration.
- **Authorizations required:** Configuration backups contain sensitive information, including user and role authorization assignments. To reduce the risk of a single, non-root administrator using the configuration backup tasks to elevate their own authorizations, each of the four tasks (backup, export, import, restore) has a separate authorization. The `root` user has all possible privileges and authorizations. For security-related information, see Security Considerations for Configuration Backups. For user authorization details, see the Appliance scope in User Authorizations in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.
- **Verify setting for new features:** It is permitted to restore a configuration that was saved before applying a system update to a new version of the appliance software. In some cases, services and properties that were present at the time of the backup may have different effects, and new services and properties may exist in the newer software that did not exist at the time of the backup. Similar to the system update process, the configuration restore process makes every effort to transfer applicable settings, and apply reasonable defaults to those properties that did not exist at the time of the backup. When restoring across software versions, administrators should manually verify settings for new features following the restore.
- **Password maintenance:** The `root` or authorized user's password is *not* changed or reverted to the password at the time of the backup if it was different. The current password is maintained on the system across the restore operation. For more details about passwords, see Security Considerations for Configuration Backups.

Security Considerations for Configuration Backups

A configuration backup contains information that is normally only accessible to the local administrative `root` user on Oracle ZFS Storage Appliance or a local administrative user with authorization for creating, restoring, importing or exporting a saved configuration. Because a non-root user with one or more of the aforementioned discrete authorizations can read the backup file's contents, and because the contents might contain sensitive information, carefully plan how to grant these non-root user authorizations.

Also, note that non-root users with the create authorization or the restore authorization have the same full privileges as the `root` user; this is not true for the configuration export and import authorizations. Therefore, especially exercise caution when granting the create and restore authorizations to non-root users. To limit security concerns when non-root users perform configuration backup or restore operations, use the following guidelines:

- **Configuration Backup:** Grant the create authorization (`configBackup`) to one user, and grant the export authorization (`configExport`) to another user.
- **Configuration Restore:** Grant the import authorization (`configImport`) to one user, and grant the restore authorization (`configRestore`) to another user.

For user authorization details, see the Appliance scope in User Authorizations in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Any configuration backup that is exported to another system or into a filesystem share must apply security restrictions to the backup file to ensure that unauthorized users cannot read the backup file.

Local user passwords are stored in the backup file in encrypted (hashed) format, not as clear text. However, on the system, access to these password hashes is restricted, as they could be used as input to dictionary attacks. Therefore, administrators must carefully protect configuration backups that are exported, either by restricting file access to the backup, or by applying an additional layer of encryption to the entire backup file, or both.

Directory user passwords are *not* stored in the appliance, and therefore are not stored in the configuration backup. If you have deployed a directory service such as LDAP or AD for administrative user access, there are no copies of directory service password hashes for directory users stored in the configuration backup. Only the user name, user ID, preferences, and authorization settings for directory users are stored in the backup and then restored.

After a configuration restore operation is performed by a local administrative `root` user or a local administrative user with the configuration restore authorization, the user's password is *not* modified to the password at the time of the backup if it was different. The password is left as-is, unmodified, by the restore process to ensure that the retained password is for the user who executed the restore operation (and, thus, logged in with that password). If the user's intent was to also change the password at the time of the configuration restore, that step must be executed manually after the restore operation, using the normal password change procedure.

Creating a Configuration Backup (BUI)

Use the following procedure to create a configuration backup.

1. From the **Maintenance** menu, select **System**.
2. Click the **Backup** button above the list of saved configurations, and follow the instructions. You are prompted to enter a descriptive comment for the backup.
3. Click **APPLY** to create the backup configuration.
4. Click the download icon to download the configuration locally.

Creating a Configuration Backup (CLI)

Use the following procedure to create a configuration backup.

1. To back up a configuration, use the `backup` command. You are prompted to enter a descriptive comment for the backup. Enter `set comment="`*comment*`"`.

```
hostname:maintenance system configs> backup
Backup Configuration. Enter a descriptive comment for this configuration, and
```

```
click Commit to backup current appliance settings:
hostname:maintenance system configs conf_backup step0> set comment="pre-upgrade"
                    comment = pre-upgrade
```

2. To execute the backup operation, enter the `done` command:

```
hostname:maintenance system configs conf_backup step0> done
hostname:maintenance system configs>
```

Restoring from a Saved Configuration (BUI)

Before restoring a saved Oracle ZFS Storage Appliance configuration, review the guidelines in Configuration Restore Impact.

Use the following procedure to restore from a saved configuration.

> **Caution:**
>
> Backing up a configuration on one appliance platform, then restoring that configuration on a different appliance platform is unsupported. For more information, see Configuration Restore Impact.

1. From the **Maintenance** menu, select **System**.
2. Click the rollback icon on any saved configuration to begin the process of reverting the system to that saved configuration.
3. Confirm that it is okay to proceed after reviewing Configuration Restore Impact.

Restoring from a Saved Configuration (CLI)

Before restoring a saved Oracle ZFS Storage Appliance configuration, review the guidelines in Configuration Restore Impact.

Use the following procedure to restore from a saved configuration.

> **Caution:**
>
> Backing up a configuration on one appliance platform, then restoring that configuration on a different appliance platform is unsupported. For more information, see Configuration Restore Impact.

1. To revert the system to a saved configuration, use the command `restore`. You are prompted to enter the universal unique identifier (UUID) for the backup. Enter `set uuid=`*UUID*.

```
hostname:maintenance system configs> restore
Restore. Select the configuration to restore:
hostname:maintenance system configs conf_restore step0>
    set uuid=36756f96-b204-4911-8ed5-fefaf89cad6a
                        uuid = 36756f96-b204-4911-8ed5-fefaf89cad6a
```

2. To execute the restore operation, enter the command `done`:

```
hostname:maintenance system configs conf_restore step0> done
```

> **Note:**
>
> Storage pools are not automatically unconfigured when you execute the `restore` command.

Deleting a Saved Configuration (BUI)

Use the following procedure to delete a saved configuration.

1. From the **Maintenance** menu, select **System**.
2. For the appropriate configuration, click the trash icon .

Deleting a Saved Configuration (CLI)

Use the following procedure to delete a saved configuration.

- To delete a saved configuration, enter the command `destroy`, space, and the universal unique identifier (UUID) for the backup. Confirm your action:

```
hostname:maintenance system configs> destroy cb2f005f-cf2b-608f-90db-fc7a0503db2a
Are you sure you want to delete the saved configuration "new"? y
hostname:maintenance system configs>
```

Exporting a Saved Configuration (BUI)

Use the following procedure to export a saved configuration. You can also use the export function to export the file to a share on Oracle ZFS Storage Appliance, itself, that has the HTTP or FTP protocol enabled for writing. You can enter a username and password for authentication to the remote controller if one is required.

1. From the **Maintenance** menu, select **System**.
2. Mouse over the configuration list entry you want to export, and click the download icon .
3. The browser prompts you to save the file locally. The file is a compressed archive whose contents are versioned and may vary over time. You should not attempt to unpack or modify the content of the archive, and doing so makes it impossible to import the archive back to the appliance successfully.

Exporting a Saved Configuration (CLI)

Use the following procedure to export a saved configuration. You can also use the export function to export the file to a share on the appliance itself, that has the HTTP or FTP protocol enabled for writing. You can enter a username and password for authentication to the remote controller if one is required.

- To export a saved configuration, enter the `export` command, space, and the universal unique identifier (UUID) for the backup:

```
hostname:maintenance system configs> export cb2f005f-cf2b-608f-90db-
fc7a0503db2a
```

Importing a Saved Configuration (BUI)

Use the following procedure to import a saved configuration. You can also use the import function to import a configuration stored in a share on the appliance itself, that has the HTTP or FTP protocol enabled for reading. You can enter a username and password for authentication to the remote controller if one is required.

1. From the **Maintenance** menu, select **System**.
2. Click the add icon at the top of the saved configurations list, and then use your web browser's file selection dialog box to locate the previously exported configuration.

> **Note:**
>
> You should upload the single, compressed archive file previously saved using the export function.

Importing a Saved Configuration (CLI)

Use the following procedure to import a saved configuration. You can also use the import function to import a configuration stored in a share on the appliance, itself, that has the HTTP or FTP protocol enabled for reading. You can enter a username and password for authentication to the remote controller if one is required.

- To import a saved configuration, enter the `import` command, space, and the universal unique identifier (UUID) for the backup:

```
hostname:maintenance system configs> import cb2f005f-cf2b-608f-90db-
fc7a0503db2a
```

Viewing Appliance Configurations (BUI)

Use the following procedure to view appliance configurations.

- From the **Maintenance** menu, select **System**. The appliance configurations are listed at the bottom of the page.

Viewing Appliance Configurations (CLI)

Use the following procedure to view appliance configurations.

- To view appliance configurations, enter the `list` command:

```
hostname:maintenance system configs> list
CONFIG                               DATE                 SYSTEM    VERSION
bfa614d7-1db5-655b-cba5-bd0bb0a1efc4 2021-04-21 17:14:28  hostname
2013.06.05.4.0,1-1.6
cb2f005f-cf2b-608f-90db-fc7a0503db2a 2022-03-28 17:56:53  hostname
2013.06.05.7.0,1-1.23
```

19

Working with Software Updates

This section contains information for viewing the currently loaded software versions, removing unnecessary software updates, and downloading a software update. In addition, descriptions of the software update states are provided.

It is recommended that you retain no more than two software updates, which avoids error message "Insufficient Space" when downloading an update. Therefore, if there are “previous” updates, it is recommended that you delete them before downloading a new software update.

To manage software updates, use the following tasks:

- Viewing Loaded Software Versions - BUI, CLI
- Removing System Updates - BUI, CLI
- Downloading Software Updates
- Software Update States

After completing this section, continue to Upgrading the Software to upgrade your system with the downloaded software.

Viewing Loaded Software Versions (BUI)

Use the following procedure to determine which software versions are loaded onto Oracle ZFS Storage Appliance. The current software update is labeled “Current,” and older loaded updates are labeled “Previous.”

1. From the **Maintenance** menu, select **System**.
2. Under **Software Updates**, verify your current software version.
3. Delete non-current updates by following the procedure Removing System Updates (BUI).

Related Topics

- Removing System Updates (BUI)
- Software Update States

Viewing Loaded Software Versions (CLI)

Use the following procedure to determine which software versions are loaded onto Oracle ZFS Storage Appliance. The current software update is labeled `current`, and older loaded updates are labeled `previous`.

1. Go to `maintenance system updates`.
2. Enter `show` and verify your current software version.
3. Delete non-current updates by following the procedure Removing System Updates (CLI).

Related Topics

- Removing System Updates (CLI)

- Software Update States

Removing System Updates (BUI)

Use the following procedure to remove extraneous system updates. To avoid accruing too much space on the system disks, maintain no more than two updates at any given time.

1. From the **Maintenance** menu, select **System**.
2. Select a software update to delete, and click its trash icon .
3. Click **OK** to confirm your action.

Related Topics

- Software Update States

Removing System Updates (CLI)

Use the following procedure to remove extraneous system updates. To avoid accruing too much space on the system disks, maintain no more than two updates at any given time.

1. Go to `maintenance system updates`.

```
hostname:> maintenance system updates
```

2. Enter `show`.

```
hostname:maintenance system updates> show
Updates:

UPDATE                                  RELEASE DATE           STATUS
ak-nas@2013.06.05.1,1-1.2               2021-12-6 23:37:50     previous
ak-nas@2013.06.05.2.0,1-1.9             2022-5-28 15:20:06     current
```

3. Enter `destroy`, space, and the update that you want to remove.

```
hostname:maintenance system updates> destroy ak-nas@2013.06.05.1,1-1.2
```

4. Enter `Y` to confirm your action.

```
This will destroy the update "ak-nas@2013.06.05.1,1-1.2". Are you sure?
(Y/N) Y
```

Related Topics

- Software Update States

Downloading Software Updates

There are two methods for downloading a software update: 1) Download the software update using a link provided in My Oracle Support document Oracle ZFS Storage Appliance: Software Updates (2021771.1); or 2) Use the following procedure. Both methods entail signing in to your My Oracle Support account, and both provide release notes that should be read before downloading a software update.

1. Go to My Oracle Support and sign in.

2. Click tab **Patches & Updates**.
3. Click tab **Product or Family (Advanced)**.
4. In field **Product is**, start typing the following product name until it appears in the pull-down list, and then select it from the list: `Sun ZFS Storage Appliance Kit (AK) Software`
5. Click **Search**.
6. Click the patch name for the appropriate software release.
7. To review the release notes, click **Read Me**.
8. To download the software, click the link for the file name, and save the file to a local filesystem accessible from your desktop.
9. To upgrade your system with the downloaded software, see Upgrading the Software.

Related Topics

- Software Update States
- Upgrading the Software

Software Update States

The following table describes software update states as indicated in the BUI. CLI states are similar in name to the BUI states.

Table 19-1 BUI Software Update States

State	Description
Current	Currently installed software version
Previous	Previous software version available for rollback
Unpacking	Update unpacking after download
Waiting	Update available for upgrade from "Current" software version
Pending	Update will run soon
Healthcheck	Actively running manual health check
Running	Update is installing
Deleting	Update is being deleted.
Unavailable	Update unavailable because intermediate upgrade needed first
Damaged	Update present, but incomplete or invalid

20

Upgrading the Software

This section describes how to upgrade the system software for Oracle ZFS Storage Appliance. You can upgrade software using either the BUI or CLI.

Use the following topics and tasks to upgrade the software:

- Preparing for a Software Upgrade
- Upgrading Software on a Standalone Controller
- Upgrading Software on Clustered Controllers

> **Note:**
>
> During software upgrade that is initiated by the user, problem notifications are automatically suspended. For detailed information about what is suspended, see Suspending and Resuming Problem Notification. Problem notification is automatically resumed when the upgrade is complete.

Related Topics

- Working with Storage Firmware Updates

Preparing for a Software Upgrade

Before you upgrade the software, perform the following actions for standalone or clustered Oracle ZFS Storage Appliance controllers.

During the update process, some protocols may experience an outage. See "Related Topics" for more information.

> **Note:**
>
> Some selected SSL/TLS protocol versions and/or ciphers are removed after a software upgrade if they are no longer supported. If all previously configured protocol versions and/or ciphers are removed after an upgrade, you are prompted to change the settings or confirm the action and use default settings.

1. If not already completed, verify your current software version.
   - In the BUI, from the **Maintenance** menu, select **System**.
   - In the CLI, go to `maintenance system updates` and enter `show`.
2. If not already completed, remove extraneous system updates. See Removing System Updates - BUI, CLI.

3. Check the most recent release notes for additional preconditions that should be observed for the software release. If skipping some software releases, also review the release notes from all applicable previous releases. See My Oracle Support document Oracle ZFS Storage Appliance: Software Updates (2021771.1).
4. Disable non-critical data services. These services may include replication, NDMP, shadow migration, or others. Disabling these services can shorten the upgrade time, and ensures that the system has a minimal operation load during the update. See Disabling a Service - BUI, CLI.
5. Create a backup copy of the management configuration to minimize downtime in the event of an unforeseen failure. See Creating a Configuration Backup - BUI, CLI.
6. Ensure that any resilvering and scrub operations have completed.
   - In the BUI, from the **Configuration** menu, select **Storage**, and check the **STATUS** column next to each pool.
   - In the CLI, go to `configuration storage`, enter `set pool=` and the name of the pool you want to check, and then enter `show`.

     The `scrub` property indicates whether scrub or resilver operations are active or completed.
7. Ensure that there are no active problems.
   - In the BUI, from the **Maintenance** menu, select **Problems**.
   - In the CLI, go to `maintenance problems` and enter `show`.

   For more information, see Viewing Active Problems.
8. Perform a health check. See Checking System Health - BUI, CLI.

   A health check is automatically run as part of the update process, but should also be run independently to check storage health prior to entering a maintenance window.
9. Schedule a maintenance window of at least one hour to allow for disruptions in storage performance and availability during update.

Next Steps

Take one of the following next steps:

- Upgrading Software on a Standalone Controller
- Upgrading Software on Clustered Controllers

Related Topics

- Technical brief "Best Practices for Upgrading Oracle ZFS Storage Appliance," located at Oracle ZFS Storage Technical Papers and Solution Briefs.

Upgrading Software on a Standalone Controller

This section describes how to upgrade the system software for Oracle ZFS Storage Appliance with a standalone controller.

Use one of the following tasks to upgrade software on standalone controllers using either the BUI or CLI:

- Upgrading Software on a Standalone Controller (BUI)

- Upgrading Software on a Standalone Controller (CLI)

Upgrading Software on a Standalone Controller (BUI)

Complete Preparing for a Software Upgrade.

Use the following procedure to upgrade software on a standalone Oracle ZFS Storage Appliance controller.

1. On your local system, download the software update as described in Downloading Software Updates.

   The update package is a single compressed file named with the version number, such as `ak-nas-2013-06-05-0-0.0.pkg.zip`. You can rename the file if needed; the version number is recorded within the image. The updates can vary in size, but are typically several hundred megabytes in size.
2. From the **Maintenance** menu, select **System**.
3. Click the add icon ⊕ next to **Software Updates**.
4. Click **Browse** and locate the `pkg.zip` update file previously downloaded.

A progress bar indicates the progress of the media being uploaded, unpacked, and verified.

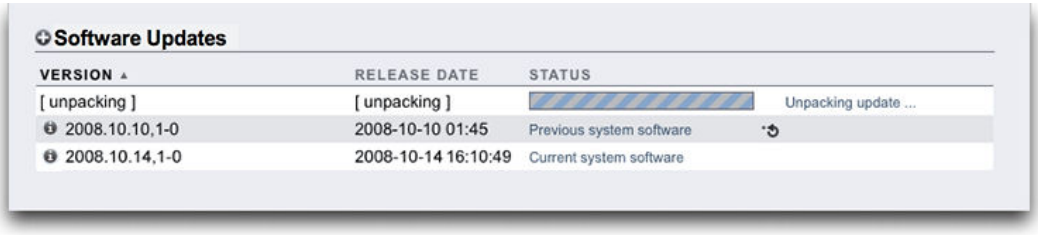

Note that on some older browsers, the progress bar may not be updated continuously during the upload; if you see a "watch" cursor, wait a minute. In the worst case, the upload will proceed all the way to completion, and you may not see the progress bar. After the media is done uploading, it is automatically unpacked and verified.

When the verification completes, the software update is displayed:

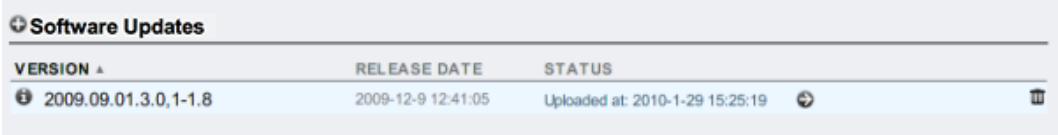

5. Click the information icon to view the Release Notes for the software update, and the installation date.

   The **Install Date** reflects the most recent installation. If the software update has not been installed, the **Install Date** is when the software update was downloaded to the appliance.

6. To begin the update, click the apply icon .

   During this process, an update health check is performed to verify the appliance is ready to upgrade. As the update progresses, you see the most recent message in the status field of the update. To cancel the update at any time (and without an adverse effect), click the cancel icon .

   > **Note:**
   >
   > If an update problem is marked as repaired during an upgrade from software release OS8.8.0 to a later release, the upgrade completion notice is not sent.

7. Enable any data services that were disabled before the upgrade, as described in Enabling a Service (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Upgrading Software on a Standalone Controller (CLI)

Complete Preparing for a Software Upgrade.

Use the following procedure to upgrade software on a standalone Oracle ZFS Storage Appliance controller.

1. On your local system, download the software update as described in Downloading Software Updates.

   The update package is a single compressed file named with the version number, such as `ak-nas-2013-06-05-0-0.0.pkg.zip`. You can rename the file if needed; the version number is recorded within the image. The updates can vary in size, but are typically several hundred megabytes in size.

2. To download the software update onto the appliance using the CLI, go to `maintenance system updates` and enter `download`:

```
hostname:maintenance system updates> download
hostname:maintenance system updates download (uncommitted)> get
                        url = (unset)
                       user = (unset)
                   password = (unset)
```

3. Set the `url` property to a valid URL for the download.

   This may be either local to your network or over the internet. The URL can be either HTTP (beginning with `http://`) or FTP (beginning with `ftp://`). If user authentication is required, it can be a part of the URL (for example, `ftp://myusername:mypasswd@myserver/export/mydirectory`), or you can leave the username and password out of the URL and instead set the user and password properties, and enter `commit`.

   After the file has been transferred, it is automatically unpacked and verified.

```
hostname:maintenance system updates download (uncommitted)> set url=ftp://
mydirectory/update.pkg.zip

                          url = ftp://mydirectory/update.pkg.zip
hostname:maintenance system updates download (uncommitted)> set user=bmc
                         user = bmc
hostname:maintenance system updates download (uncommitted)> set password
Enter password:
                     password = (set)
hostname:maintenance system updates download (uncommitted)> commit
Transferred 157M of 484M (32.3%) ...
Transferred 484M of 484M (100%) ... done
Unpacking ... done
```

4. Enter `list` to view the status.

```
hostname:maintenance system updates> list
UPDATE                                  RELEASE DATE              STATUS
ak-nas@2013.06.05.4.0,1-1.7             2015-5-1 20:13            waiting
...
```

5. From this context, you can set any properties specific to the update, including applying deferred updates.

   You can also apply deferred updates after updating the software. When installing a system update that contains a deferred update, you have the option to apply deferred updates upon request or automatically. It is recommended to select upon request and apply the deferred updates during the post-update process.

   For more information on the set of properties available for the particular update, run the `help properties` command. User-controllable properties begin with the `update_` prefix:

```
hostname:maintenance system updates ak-nas@2013.06.05.4.0,1-1.6> help properties
Properties that are valid in this context:

   version              => Update media version
   release_date         => Update release date
   install_date         => Update install date
   status               => Update media status
   update_zfs_upgrade   => Apply incompatible storage pool update

hostname:maintenance system updates ak-nas@2013.06.05.4.0,1-1.6> get
                      version = 2013.06.05.4.0,1-1.6
                 release_date = 2015-04-21 02:38:15:45
                 install_date = 2015-04-23 09:56:03:44
                       status = waiting
            update_zfs_upgrade = deferred
hostname:maintenance system updates ak-nas@2013.06.05.4.0,1-1.6> set
update_zfs_upgrade=onreboot
            update_zfs_upgrade = onreboot
hostname:maintenance system updates ak-nas@2013.06.05.4.0,1-1.6>
```

6. After you set any properties, enter the `upgrade` command and then enter `y` to start the update.

```
hostname:maintenance system updates> select ak-nas@2013.06.05.4.0,1-1.7-nd
hostname:maintenance system updates ak-nas@2013.06.05.4.0,1-1.7-nd> upgrade
The selected software update requires a system reboot in order to take effect.
The system will automatically reboot at the end of the update process. The
update will take several minutes. At any time during this process, you can
cancel the update with [Control-C].
```

```
Are you sure? (Y/N) y
Updating from ... ak/nas@2013.06.05.4.0,1-1.6
Backing up smf(5) ... done.
Loading media metadata ... done.
Selecting alternate product ... SUNW,iwashi
Installing Sun Storage 7120 2013.06.05.4.0,1-1.7
pkg://sun.com/ak/SUNW,iwashi@2013.06.05.4.0,1-1.7:20091014T084500Z
Creating system/boot/ak-nas@2013.06.05.4.0,1-1.7 ... done.
Creating system/root/ak-nas@2013.06.05.4.0,1-1.7 ... done.
...
```

As the update proceeds, the latest message is printed. You can cancel the update at any time by pressing **Control-C** and confirming your action:

```
Updating from ... ak/nas@2013.06.05.4.0,1-1.6
Backing up smf(5) ... done.
Loading media metadata ... ^C
This will cancel the current update. Are you sure? (Y/N) y
error: interrupted by user
hostname:maintenance system updates ak-nas@2013.06.05.4.0,1-1.7-nd>
```

> **Note:**
>
> If an update problem is marked as repaired during an upgrade from software release OS8.8.0 to a later release, the upgrade completion notice is not sent.

7. Enable any data services that were disabled before the upgrade, as described in Enabling a Service (CLI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Upgrading Software on Clustered Controllers

This section describes how to upgrade the system software for Oracle ZFS Storage Appliance with clustered controllers. Cluster update is also referred to as a rolling update. Controllers are upgraded one at a time, and each controller is upgraded while the cluster peer provides service to clients. Clients will encounter brief periods of service interruption while cluster resources transfer between the cluster peers.

> **Note:**
>
> As of software release OS8.7.22 or later, configuration changes are supported on a controller that is running down-level software and has active resources while its cluster peer is being upgraded.

Use the following tasks to upgrade software on clustered controllers using either the BUI or CLI:

- Preparing to Upgrade Clustered Controllers - BUI, CLI
- Upgrading Clustered Controllers - BUI, CLI

Preparing to Upgrade Clustered Controllers (BUI)

Complete Preparing for a Software Upgrade.

> **Note:**
>
> For the purpose of this procedure, the first controller to be upgraded is referred to as controller A and its peer is controller B. If one of the controllers is in a Stripped state (it has no active resources), upgrade that controller first to avoid availability delays. If both controllers in a cluster have active resources, choose either controller to upgrade first.

Use the following procedure to prepare to upgrade clustered Oracle ZFS Storage Appliance controllers.

1. On your local system, download the software update as described in Downloading Software Updates.

   The update package is a single compressed file named with the version number, such as `ak-nas-2013-06-05-0-0.0.pkg.zip`. You can rename the file if needed; the version number is recorded within the image. The updates can vary in size, but are typically several hundred megabytes in size.
2. On controller A, from the **Configuration** menu, select **Cluster**, and check the active resources and status of each peer.

   For information about the cluster states, see Cluster Takeover and Failback in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.
3. If the management interface is unlocked , click the lock icon next to the management interface.
4. From the **Maintenance** menu, select **System**, and click the add icon next to **Software Updates**.
5. In the **Add Software Update** dialog box, browse and select the update you want to install.

   A progress bar is displayed showing the progress of the upload.
6. Repeat steps 2-5 on controller B.

Next Steps

- Upgrading Clustered Controllers (BUI)

Preparing to Upgrade Clustered Controllers (CLI)

Complete Preparing for a Software Upgrade.

> **Note:**
>
> For the purpose of this procedure, the first controller to be upgraded is referred to as controller A and its peer is controller B. If one of the controllers is in a Stripped state (it has no active resources), upgrade that controller first to avoid availability delays. If both controllers in a cluster have active resources, choose either controller to upgrade first.

Use the following procedure to prepare to upgrade clustered Oracle ZFS Storage Appliance controllers.

1. On your local system, download the software update as described in Downloading Software Updates.

   The update package is a single compressed file named with the version number, such as `ak-nas-2013-06-05-0-0.0.pkg.zip`. You can rename the file if needed; the version number is recorded within the image. The updates can vary in size, but are typically several hundred megabytes in size.

2. Perform the following steps to verify that each controller has a dedicated management interface listed as a private resource:

   The following examples show the cluster in Active/Active mode, with example output if the management interface is listed as a private resource.

   a. On controller A, go to `configuration cluster resources show`.

   ```
   controller-a:> configuration cluster resources show
   Resources:

   RESOURCE        OWNER           TYPE       LABEL         CHANGES  DETAILS
   net/igb0        controller-a     private   management... no
   10.80.2xx.112
   ```

   b. On controller B, go to `configuration cluster resources show`.

   ```
   controller-b:> configuration cluster resources show
   Resources:

   RESOURCE        OWNER           TYPE       LABEL         CHANGES  DETAILS
   net/igb1        controller-b    private    management... no
   10.80.2xx.114
   ```

3. If the management interface for controller A does not show type `private`, lock the management interface, as shown in the following example:

   ```
   controller-a:> configuration cluster resources select net/igb0
       configuration cluster resources (uncommitted)> set owner=controller-a
       configuration cluster resources (uncommitted)> set type=private
       configuration cluster resources (uncommitted)> commit
   ```

   Note that each controller must be the current owner of the management interface.

4. On controller B, lock the management interface on controller B, as in step 3.

5. Download the update to both controllers.

   a. On controller A, go to `maintenance system updates download` and enter the following example commands:

```
controller-a:maintenance system updates download (uncommitted)> set url=
   ftp://mydirectory/update.pkg.zip
                          url = ftp://mydirectory/update.pkg.zip
 controller-a:maintenance system updates download (uncommitted)> set user=bmc
                         user = bmc
 controller-a:maintenance system updates download (uncommitted)> set password
 Enter password:
                     password = (set)
 controller-a:maintenance system updates download (uncommitted)> commit

 Transferred 157M of 484M (32.3%) ...
```

b. On controller B, go to `maintenance system updates download` and enter the following example commands:

```
controller-b:maintenance system updates download (uncommitted)> set url=
   ftp://mydirectory/update.pkg.zip
                          url = ftp://mydirectory/update.pkg.zip
 controller-b:maintenance system updates download (uncommitted)> set user=bmc
                          user = bmc
  controller-b:maintenance system updates download (uncommitted)> set password
  Enter password:
                       password = (set)
  controller-b:maintenance system updates download (uncommitted)> commit
  Transferred 157M of 484M (52.1%) ...
```

6. Wait for the package to fully download and unpack before proceeding.

Next Steps

- Upgrading Clustered Controllers (CLI)

Upgrading Clustered Controllers (BUI)

- Complete Preparing for a Software Upgrade.
- Complete Preparing to Upgrade Clustered Controllers (BUI).

> **Note:**
>
> For the purpose of this procedure, the first controller to be upgraded is referred to as controller A and its peer is controller B. If one of the controllers is in a Stripped state (it has no active resources), upgrade that controller first to avoid availability delays. If both controllers in a cluster have active resources, choose either controller to upgrade first.

1. Log in to controller A, and failover its resources to controller B by performing a reboot:
   a. Click on the power icon in the upper left section of the navigation bar.
   b. Select **Reboot** from the drop-down menu, and click **OK** to confirm.

      Do *not* select **Diagnostic Reboot**. Wait for the reboot to complete before proceeding.
2. On controller A, from the **Maintenance** menu, select **System**, and click the arrow icon next to the name of the update you want to install.
3. Click **CHECK** to perform health checks.

   For information about health checks, see Working with Health Checks.

4. Click **APPLY** to begin the update process.

> **Note:**
>
> If an update problem is marked as repaired during an upgrade from software release OS8.8.0 to a later release, the upgrade completion notice is not sent.

5. Wait for controller A to fully reboot, and log back in to controller A.
6. From the **Configuration** menu, select **Cluster**, and verify that controller A is in the **Ready (waiting for failback)** state.

   For information about the cluster states, see Cluster Takeover and Failback in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.
7. To monitor firmware updates, from the **Maintenance** menu, select **System**, and check the update counter.
8. Log in to controller B, and from the **Configuration** menu, select **Cluster**, and verify that controller B is in the **Active (takeover completed)** state.

   For information about the cluster states, see Cluster Takeover and Failback in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.
9. From the **Configuration** menu, select **Cluster**, and click **FAILBACK** to change the cluster to an Active/Active configuration.

> **Note:**
>
> This is not necessary if you want an Active/Passive configuration.

10. From the **Maintenance** menu, select **System**, and click the arrow icon next to the name of the update that you want to install.
11. Click **CHECK** to perform health checks.

    For information about health checks, see Working with Health Checks.
12. Click **APPLY** to begin the update process.

> **Note:**
>
> If an update problem is marked as repaired during an upgrade from software release OS8.8.0 to a later release, the upgrade completion notice is not sent.

13. Wait for controller B to fully reboot, and then log back in to controller B.
14. From the **Configuration** menu, select **Cluster**, and verify that controller B is in the **Ready (waiting for failback)** state.

    For information about the cluster states, see Cluster Takeover and Failback in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.
15. To monitor firmware updates, from the **Maintenance** menu, select **System**, and check the update counter.

16. Log in to controller A, and from the **Configuration** menu, select **Cluster**, and verify that controller A is in the **Active (takeover completed)** state.

    For information about the cluster states, see Cluster Takeover and Failback in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

17. To verify that all firmware updates are complete, from the **Maintenance** menu, select **System**, and check the update counter.

    > **Note:**
    >
    > Do not begin the next step until all firmware updates are complete.

18. From the **Configuration** menu, select **Cluster**, and click **FAILBACK** to change the cluster to an Active/Active configuration.

    > **Note:**
    >
    > This is not necessary if you want an Active/Passive configuration.

    Both controllers are now upgraded.

19. From the **Maintenance** menu, select **Hardware**, and verify that all disks are online.

    All lights should be green.

20. Verify there are no controller and disk shelf component errors.

    All lights should be green. An amber light indicates a component error.

21. If any components have errors, check for pool errors: From the **Configuration** menu, select **Storage**, and check the **STATUS** and **ERRORS** columns for each pool.

    Pools should be online and have no errors.

22. Log in to controller B, and repeat steps 18-20 for controller B.

23. Enable any data services that were disabled before the upgrade, as described in Enabling a Service (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Upgrading Clustered Controllers (CLI)

- Complete Preparing for a Software Upgrade.
- Complete Preparing to Upgrade Clustered Controllers (CLI).

> **Note:**
>
> For the purpose of this procedure, the first controller to be upgraded is referred to as controller A and its peer is controller B. If one of the controllers is in a Stripped state (it has no active resources), upgrade that controller first to avoid availability delays. If both controllers in a cluster have active resources, choose either controller to upgrade first.

1. Log in to controller A, and failover its resources to controller B by performing a reboot:
   a. Click on the power icon in the upper left section of the navigation bar.
   b. Select **Reboot** from the drop-down menu, and click **OK** to confirm.

      Do *not* select **Diagnostic Reboot**. Wait for the reboot to complete before proceeding.
2. On controller A, go to `maintenance system updates`.
3. Enter `select` and the name of the update you want to install.

   ```
   controller-a:maintenance system updates> select ak-
   nas@2013.06.05.6.3,1-2.1.7.1
   ```
4. Enter `check`, and then enter `y` to perform health checks.

   ```
   controller-a:maintenance system updates ak-nas@2013.06.05.6.3,1-2.1.7.1>
   check
   You have requested to run checks associated with waiting upgrade media. This
   will execute the same set of checks as will be performed as part of any
   upgrade
   attempt to this media, and will highlight conditions that would prevent
   successful upgrade. No actual upgrade will be attempted, and the checks
   performed are of static system state and non-invasive. Do you wish to
   continue?

   Are you sure? (Y/N) y
   ```
5. Enter `upgrade`, and then enter `y` to begin the update process.

   > **Note:**
   >
   > You can cancel the update at any time in the CLI by pressing **Control-C**. If an update problem is marked as repaired during an upgrade from software release OS8.8.0 to a later release, the upgrade completion notice is not sent.

   As part of the update, controller A automatically reboots.

   ```
   controller-a:maintenance system updates ak-nas@2013.06.05.6.3,1-2.1.7.1>
   upgrade
   This procedure will consume several minutes and requires a system reboot upon
   successful update, but can be aborted with [Control-C] at any time prior to
   reboot. A health check will validate system readiness before an update is
   attempted, and may also be executed independently using the check command.

   Are you sure? (Y/N) y
   ```
6. After controller A has fully rebooted, log back in to controller A, and enter `configuration cluster show` to verify that controller A is in the `Ready (waiting for failback)` state.

   In the following example, the state of controller A is `AKCS_STRIPPED`, with a description of `Ready (waiting for failback)`.

   ```
   controller-a:> configuration cluster show
       state = AKCS_STRIPPED
       description = Ready (waiting for failback)
   ```

```
    ...
    peer_state = AKCS_OWNER
    peer_description = Active (takeover completed)
```

For information about the cluster states, see Cluster Takeover and Failback in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

7. Enter `maintenance system updates firmware show` to monitor firmware updates.

```
controller-a:> maintenance system updates firmware show
Pending
No Pending Updates

Failed
No Failed Updates

In Progress
No Updates in Progress
```

8. Log in to controller B, and enter `configuration cluster show` to verify that controller B is in the `Active (takeover completed)` state.

   For information about the cluster states, see Cluster Takeover and Failback in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

   In the following example, the state of controller B is `AKCS_OWNER`, with a description of `Active (takeover completed)`.

```
controller-b:> configuration cluster show
    state = AKCS_OWNER
    description = Active (takeover completed)
    ...
    peer_state = AKCS_STRIPPED
    peer_description = Ready (waiting for failback)
```

9. Enter `configuration cluster failback`, and then enter `y` to change the cluster to an Active/Active configuration.

> **Note:**
>
> This is not necessary if you want an Active/Passive configuration.

```
controller-b:> configuration cluster failback
Continuing will immediately fail back the resources assigned to the cluster
peer. This may result in clients experiencing a slight delay in service.

Are you sure? (Y/N) y
```

10. Enter `maintenance system updates`, and enter `show` to view the available updates.

```
controller-b:> maintenance system updates
hostname:maintenance system updates> show
Updates:

UPDATE                                 RELEASE DATE           STATUS
ak-nas@2013.06.05.6.2,1-1.1            2016-7-26  15:20:06    current
ak-nas@2013.06.05.6.3,1-2.1.7.1        2016-8-24  14:10:02    waiting
ak-nas@2013.06.05.6.1,1-1.1            2016-6-24  23:37:50    previous
```

11. Enter `select` and the name of the update you want to install.

```
controller-b:maintenance system updates> select ak-
nas@2013.06.05.6.3,1-2.1.7.1
```

12. Enter `check`, and then enter `y` to perform health checks.

```
controller-b:maintenance system updates ak-nas@2013.06.05.6.3,1-2.1.7.1>
check
You have requested to run checks associated with waiting upgrade media. This
will execute the same set of checks as will be performed as part of any
upgrade
attempt to this media, and will highlight conditions that would prevent
successful upgrade. No actual upgrade will be attempted, and the checks
performed are of static system state and non-invasive. Do you wish to
continue?

Are you sure? (Y/N) y
```

13. Enter `upgrade`, and then enter `y` to begin the update process.

> **Note:**
>
> You can cancel the update at any time in the CLI by pressing **Control-C**. If an update problem is marked as repaired during an upgrade from software release OS8.8.0 to a later release, the upgrade completion notice is not sent.

As part of the update, controller B automatically reboots.

```
controller-b:maintenance system updates ak-nas@2013.06.05.6.3,1-2.1.7.1>
upgrade
This procedure will consume several minutes and requires a system reboot upon
successful update, but can be aborted with [Control-C] at any time prior to
reboot. A health check will validate system readiness before an update is
attempted, and may also be executed independently using the check command.

Are you sure? (Y/N) y
```

14. After controller B has fully rebooted, log back in to controller B, and enter `configuration cluster show` to verify that controller B is in the `Ready (waiting for failback)` state.

    For information about the cluster states, see Cluster Takeover and Failback in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

    In the following example, the state of controller B is `AKCS_STRIPPED`, with a description of `Ready (waiting for failback)`.

```
controller-b:> configuration cluster show
    state = AKCS_STRIPPED
    description = Ready (waiting for failback)
    ...
    peer_state = AKCS_OWNER
    peer_description = Active (takeover completed)
```

15. Enter `maintenance system updates firmware show` to monitor firmware updates.

```
controller-b:> maintenance system updates firmware show
Pending
No Pending Updates

Failed
No Failed Updates

In Progress
No Updates in Progress
```

16. Log in to controller A, and enter `configuration cluster show` to verify that controller A is in the `Active (takeover completed)` state.

    In the following example, the state of controller A is `AKCS_OWNER`, with a description of `Active (takeover completed)`.

    ```
    controller-a:> configuration cluster show
        state = AKCS_OWNER
        description = Active (takeover completed)
        ...
        peer_state = AKCS_STRIPPED
        peer_description = Ready (waiting for failback)
    ```

    For information about the cluster states, see Cluster Takeover and Failback in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

17. To verify that all firmware updates are complete, enter `maintenance system updates firmware show` to monitor firmware updates.

    ```
    controller-a:> maintenance system updates firmware show
    Pending
    No Pending Updates

    Failed
    No Failed Updates

    In Progress
    No Updates in Progress
    ```

    > **Note:**
    >
    > Do not begin the next step until all firmware updates are complete.

18. Enter `configuration cluster failback`, and then enter `y` to change the cluster to an Active/Active configuration.

    > **Note:**
    >
    > This is not necessary if you want an Active/Passive configuration.

    ```
    controller-a:> configuration cluster failback
    Continuing will immediately fail back the resources assigned to the cluster
    peer. This may result in clients experiencing a slight delay in service.

    Are you sure? (Y/N) y
    ```

Both controllers are now upgraded.

19. Enter `maintenance hardware show` to verify that all disks are online.

```
controller-a:> maintenance hardware show
Chassis:
         NAME    STATE  MANUFACTURER  MODEL
SERIAL                RPM
...
disk-000 HDD 0   ok     SEAGATE       ST330057SSUN300G  001117G1G929
6SJ1G929 15000
disk-001 HDD 1   ok     SEAGATE       ST330057SSUN300G  001117G1GA8Y
6SJ1GA8Y 15000
disk-002 HDD 2   ok     SEAGATE       ST330057SSUN300G  001117G1KDPZ
6SJ1KDPZ 15000
...
```

20. Enter `configuration storage show` to check the status of each pool.

    Pools should be online and have no errors.

```
controller-a:> configuration storage show
Pools:

   POOL     OWNER          DATA PROFILE     LOG PROFILE     STATUS     ERRORS
-> pool25   controller-a    -                -               online     0
   pool26   controller-b    -                -               online     0
```

21. Log in to controller B, and repeat steps 18-19 for controller B.

22. Enable any data services that were disabled before the upgrade, as described in Enabling a Service (CLI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

21

Working with Storage Firmware Updates

Software updates can contain automatic firmware updates for Oracle ZFS Storage Appliance hardware. Following a software upgrade, any storage hardware for which the update includes newer versions of storage firmware is automatically upgraded. Before the upgrade window, it is recommended to run a scrub as described in Scrubbing a Storage Pool – Manual (BUI) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

> **Note:**
>
> If the storage firmware update occurs after a software upgrade, the scrub has already been performed.

Storage hardware, each with distinct characteristics, include disks, disk shelves, and certain internal SAS devices. These devices are upgraded in the background. When a firmware update is occurring, the progress is displayed in the left panel of the **Maintenance: System** BUI view, or in the `maintenance system updates` CLI context. These firmware updates are almost always hardware related, though it may briefly show some number of outstanding updates when applying certain deferred updates to components other than hardware.

During the storage firmware update process, hardware connections and status may be intermittent. Related hardware alerts are suppressed during this process. However, if you are viewing the **Maintenance: System** screen or the **Configuration: Storage** screen, you may see the effects of these updates in the form of missing or offline devices. If a device remains offline or missing for an extended period of time (several minutes or more) even after refreshing the hardware view, this may be an indication of a problem with the device. Check the **Maintenance: Problems** view for any relevant faults that may have been identified.

In some cases, the controllers in the disk shelves may stay offline during a storage firmware update. If this occurs, no other controllers are updated until this condition is fixed. If a disk shelf is listed as having only a single path for an extended period of time, check the physical disk shelf to determine whether the green link lights on the back of the SIM or I/O Module are active. If they are not, remove and re-insert the SIM or I/O Module to re-establish the connection. Verify that all disk shelves are reachable by two paths.

Related Topics

- Storage Firmware Updates for Clustered Controllers
- Outstanding Storage Firmware Updates
- Checking the Status of Storage Firmware Updates (BUI)
- Checking the Status of Storage Firmware Updates (CLI)

Storage Firmware Updates for Clustered Controllers

For Oracle ZFS Storage Appliance clustered controllers, updates that are pending are shown only on the primary controller. Therefore, firmware updates displayed on the peer controller do not include firmware updates for the primary controller.

Hardware updates are always applied in a completely safe manner. This means that the system may be in a state where hardware updates cannot be applied. This is particularly important in the context of clustered configurations. During takeover and failback operations, any in-progress firmware updates are completed, and any pending firmware updates are suspended until the takeover or failback has completed. Then the state restrictions are re-evaluated in the context of the new cluster state and, if possible, firmware updates resume.

> **Caution:**
>
> Do not perform takeover or failback operations while firmware updates are in progress.

The rolling update procedure for clustered controllers meets all of the following best practices and addresses the per-device-class restrictions. It should always be followed when performing updates in a clustered environment. In both clustered and standalone environments, these criteria are also re-evaluated upon any reboot or diagnostic system software restart, which may cause previously suspended or incomplete firmware updates to resume. For more information, see Upgrading Software on Clustered Controllers.

- Components internal to the storage controller (such as HBAs and network devices) other than disks and certain SAS devices are generally upgraded automatically during a reboot. These updates are not visible and will have completed by the time the management interfaces become available.
- Updating disk or flash device firmware causes the device to be taken offline during the process. If there is insufficient redundancy in the containing storage pool to allow this operation, the firmware update will not complete and may appear "stalled." However, if the storage pools are in an exported state, the disks will update as expected. Disks and flash devices that are part of a storage pool which is currently in use by the cluster peer, if any, are not upgraded.

  > **Note:**
  >
  > When updating the firmware on a system with striped pools, make sure both controllers are running the same version before attempting a pool unconfigure.

- Updating disk shelf firmware requires that both back-end storage paths be active to all disks within all disk shelves, and for storage in all shelves to be configured. For clusters with at least one active pool on each controller, these restrictions mean that disk shelf firmware updates can be performed only by a controller that is in the owner state. For more information, see Cluster Takeover and Failback in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Outstanding Storage Firmware Updates

When there are outstanding updates, an informational or warning icon appears next to the number of updates remaining. Clicking the icon opens the **Firmware Updates** dialog box, which lists the current remaining updates. Each update shows the current version of the component, the time of the last attempted update, and the reason why the last attempt did not succeed.

Outstanding updates can be in one of three states: Pending, In Progress, and Failed. An update begins in the Pending state and is periodically retried, then moves into the In Progress state. If there is a failure due to a transient condition, the update is moved back to the Pending state, and otherwise to the Failed state.

A problem exists in the following cases:

- Any updates are in the Failed state.
- Updates remain in the Pending state.
- Updates switch between the Pending and In Progress states for an extended period of time (more than half an hour), without the number of remaining updates decreasing.

The following condition does not indicate a problem: When upgrading multiple chassis, the number of remaining updates decrease, indicating that the update is progressing. Some of the chassis will appear with the Pending status, indicating only one path exists for one or more disks. When a chassis expander is reset during the update, causing some disks to temporarily have only one path, it can delay updates to other chassis until it is safe to do so again non-disruptively.

> **Note:**
>
> The **Firmware Updates** dialog box does not automatically refresh. To get an updated view, close and reopen the dialog box.

Checking the Status of Storage Firmware Updates (BUI)

1. From the **Maintenance** menu, select **System**.
2. In the left panel under **System Disks**, look for the firmware update progress.

   For clustered controllers, the status only shows the updates that are pending on the primary controller. Therefore, firmware updates displayed on the peer controller do not include firmware updates for the primary controller.

Checking the Status of Storage Firmware Updates (CLI)

1. Go to `maintenance system updates`, and enter `firmware`.

   ```
   hostname:maintenance system updates> firmware
   ```

2. Enter `get` to see the number of pending, failed, and in-progress firmware updates.

   ```
   hostname:maintenance system updates firmware> get
                upgrades_pending = 1
   ```

```
             upgrades_failed = 1
          upgrades_in_progress = 0
```

3. Enter `show` to view firmware update details.

```
hostname:maintenance system updates firmware> show
Properties:
             upgrades_pending = 1
              upgrades_failed = 1
          upgrades_in_progress = 0

Status:
Pending
Component                              Current Version  Status
Chassis 0945QCQ01A (SAS Expander 00)  3232             some disks only have 1
path(s) (2 paths required)

Failed
Component                              Current Version  Status
Chassis 0945QCQ00A (SAS Expander 00)  3232             some disks only have 1
path(s) (2 paths required)

In Progress
No Updates in Progress
```

For clustered controllers, the status only shows the updates that are pending on the primary controller. Therefore, firmware updates displayed on the peer controller do not include firmware updates for the primary controller.

22

Working with Platform Firmware Updates

This section describes how to check and update the Oracle ILOM and/or BIOS firmware for Oracle ZFS Storage Appliance. This process is known as a platform firmware update, and it is available with software version OS8.7.0 or later. BUI and CLI procedures are provided in this section for checking the current firmware versions and for updating the platform firmware.

> **Note:**
>
> During platform firmware update, problem notifications are automatically suspended. For detailed information about what is suspended, see Suspending and Resuming Problem Notification. Problem notification is automatically resumed when the update is complete.

The automatic platform firmware update feature is supported on the following platforms:

- Oracle ZFS Storage ZS9-2
- Oracle ZFS Storage ZS7-2
- Oracle ZFS Storage ZS5-4
- Oracle ZFS Storage ZS5-2
- Oracle ZFS Storage ZS4-4
- Oracle ZFS Storage ZS3-2

For automatic platform firmware updates, the appliance must have a minimum Oracle ILOM version, as shown in the following table:

Platform	Minimum Required Oracle ILOM Version
Oracle ZFS Storage ZS9-2	Contact Oracle Support
Oracle ZFS Storage ZS7-2	4.0.4.21 r127001
Oracle ZFS Storage ZS5-4	3.2.6.20 r109303
Oracle ZFS Storage ZS5-2	3.2.7.24 r113383
Oracle ZFS Storage ZS4-4	3.2.4.18.a r95773
Oracle ZFS Storage ZS3-2	3.2.4.58 r101476

To view the Oracle ILOM version in the BUI, click the Sun/Oracle logo in the upper left corner of the BUI masthead. In the **System Info** window, the Oracle ILOM version is listed for the "Service Processor."

To view the Oracle ILOM version in the CLI, enter `configuration version get sp_version`.

```
hostname:> configuration version get sp_version
                   sp_version = 3.2.7.24 r113383
```

If the Oracle ILOM version does not meet the minimum required level, contact Oracle Service. Otherwise, use the following tasks to update the platform firmware.

- Checking the Platform Firmware Versions - BUI, CLI
- Updating the Platform Firmware - BUI, CLI

Checking the Platform Firmware Versions (BUI)

1. Fro the **Maintenance** menu, select **Problems**.
2. Under **Active Problems**, look for a major alert stating the chassis is running an older (downrev) platform firmware version.

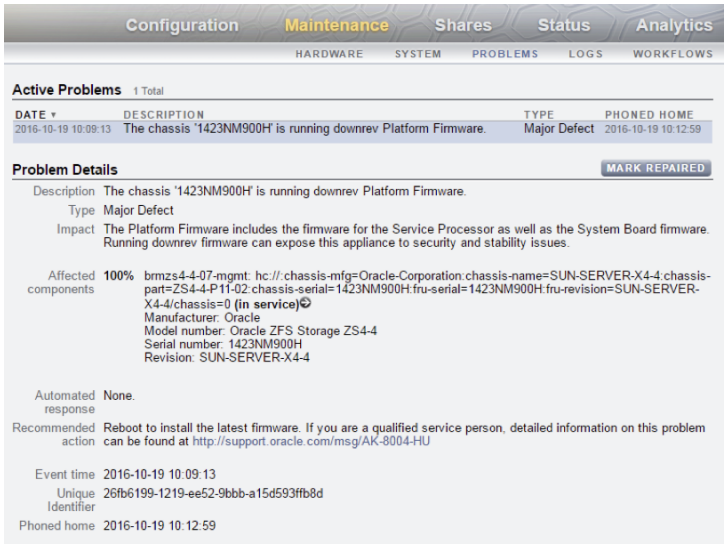

3. If this alert is present, proceed to Updating the Platform Firmware (BUI).

   If the alert is not present, the platform firmware is up to date, and no further action is needed.

Checking the Platform Firmware Versions (CLI)

1. Go to `maintenance system updates platform`.
2. Enter the `get` command:

```
hostname:maintenance system updates platform> get
                  update_needed = true
              power_down_needed = false
```

   If `updated_needed` is `true`, the Oracle ILOM firmware is ready to be updated.

If `power_down_needed` is also `true`, the BIOS firmware is also ready for an update, and it will require a power cycle during the platform firmware update.

> **Note:**
>
> If `update_needed` is `unknown`, the appliance might need a software upgrade to version OS8.7.0 or later, or automatic platform firmware upgrades are not supported on your platform.

3. If `update_needed` is `true`, proceed to Updating the Platform Firmware (CLI).

   If `update_needed` is `false`, the platform firmware is up to date, and no further action is needed.

Updating the Platform Firmware (BUI)

Before you update the Oracle ILOM/host firmware, review these prerequisites:

- Ensure you are using software version OS8.7.0 or later. If not, see Upgrading the Software.
- View the platform firmware versions to determine if this task is necessary. See Checking the Platform Firmware Versions (BUI).

Use this task to update platform firmware during an appliance reboot.

By default, if the platform firmware is ready to update, it will be automatically applied as part of performing a graceful reboot.

1. Select the power off/reboot appliance icon located in the upper left corner of the BUI masthead.

2. In the **Power** dialog box, select **Update Platform Firmware and reboot** from the drop-down list.

   A power cycle might be necessary.

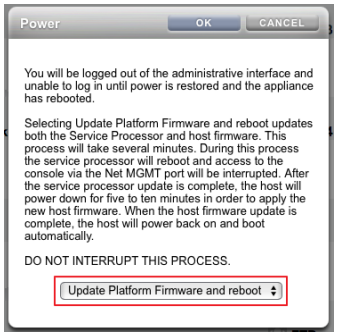

3. Click **OK** to start the update process.

   Wait for the service processor or appliance to reboot.

> **Caution:**
>
> Do not attempt to power on the appliance during the platform update and reboot process. Doing so might interrupt the process, and the system's state could be adversely affected.

Updating the Platform Firmware (CLI)

Before you update the Oracle ILOM/host firmware, review these prerequisites:

- Ensure you are using software version OS8.7.0 or later. If not, see Upgrading the Software.
- View the platform firmware versions to determine if this task is necessary. See Checking the Platform Firmware Versions (CLI).

Use this task to update platform firmware during an appliance reboot.

By default, if the platform firmware is ready to update, it will be automatically applied as part of performing a graceful reboot.

1. Enter `maintenance system reboot`.

```
hostname:> maintenance system reboot
Upgrading both the Service Processor and host firmware and rebooting.

This process will take several minutes. During this process the service
processor will reboot and access to the console via the Net MGMT port will be
interrupted. Data services to this appliance will not be affected. After the
service processor upgrade is complete the host will power down for five to
ten
minutes in order to apply the new host firmware. When the host firmware
upgrade is complete the host will power back on and boot automatically.

DO NOT INTERRUPT THIS PROCESS.

Are you sure? (Y/N)
```

2. Type `Y` to start the update process.

   Wait for the service processor or appliance to reboot.

> **Caution:**
>
> Do not attempt to power on the appliance during the platform update and reboot process. Doing so might interrupt the process, and the system's state could be adversely affected.

23
Working with Health Checks

System-level health checks are provided to help ensure that no pathologies interfere with the software update. If a problem is encountered, it is noted in the Alert Log and the update process is aborted. System software updates do not proceed until all problems have been corrected.

You can manually run the same health checks in advance of any planned update. This allows you to check the state of the system prior to scheduling an update maintenance window so you can correct any problems that could interfere with the update process. Any problem report that is issued by a manual health check is identical to that issued by the health checks integrated in the update process. As with the integrated health checks, you are presented with a link to the Alert Log, as described in Alert Log Entries, when problems are found. If no problems are found, the `System Ready` state transitions to `Yes` to indicate that the system is ready for software updates.

> **Note:**
>
> Running an update health check does not replace meeting required preconditions. Precondition checks, described in Preparing for a Software Upgrade, must also be executed and problems resolved prior to upgrading the system software.

Checking System Health (BUI)

Use the following procedure to check system health using the BUI.

1. After you select and start an update, update health checks may be issued from the **Apply Software Update** dialog box in the BUI by clicking **CHECK**. The system remains in the `Unchecked` state until the **CHECK** button is clicked.

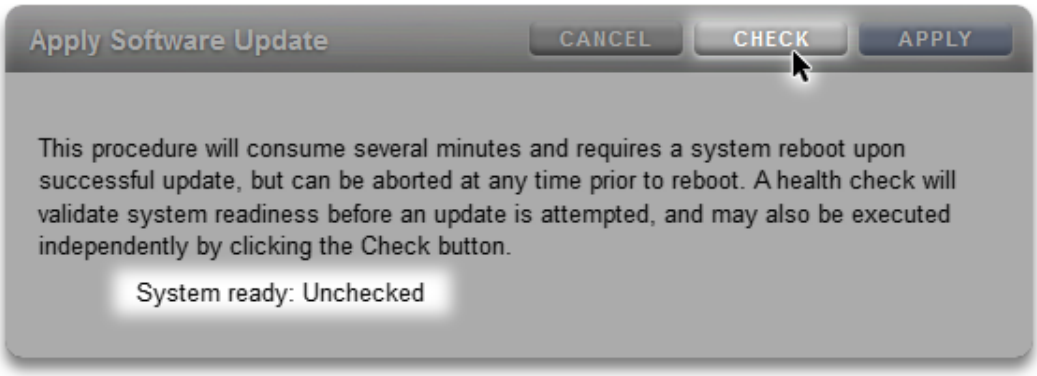

2. During the health check operation, an indicator shows its progress.

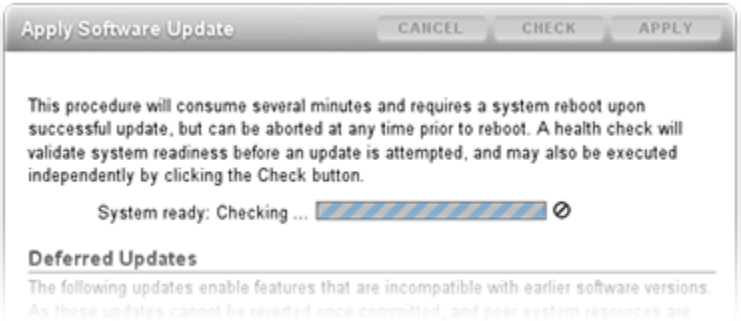

3. After completion, the `System Ready` state changes to `Yes` or `No` with a link to the Alert Log.

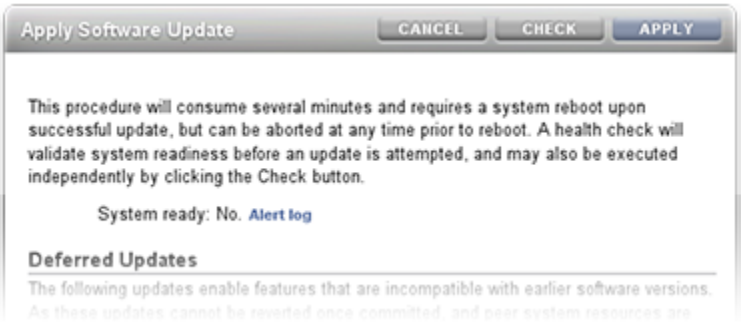

Checking System Health (CLI)

Use the following procedure to check system health using the CLI.

1. Go to `maintenance system updates.`
2. To select the update media, enter `select`, space, and the update media name:

```
hostname:maintenance system updates> select ak-nas@2013.06.05.0.0,1-1.6
```

3. To upgrade the software and check system health, use the `upgrade` command:

```
hostname:maintenance system updates:ak-nas@2013.06.05.0.0,1-1.6> upgrade
This procedure will consume several minutes and requires a system reboot upon
successful update, but can be aborted with [Control-C] at any time prior to
reboot. A health check will validate system readiness before an update is
attempted, and may also be executed independently using the check command.
Are you sure? (Y/N)
 Healthcheck running ... /
 Healthcheck completed.  There are no issues at this time which
 would cause an upgrade to this media to be aborted.
```

Troubleshooting Update Health Check Failures

Prior to the actual software update, health checks are performed automatically when an update is started. If an update health check fails, it can cause an update to abort (see following example). Update health checks only validate issues that can impact updates.

Example BUI and CLI Update Health Checks Failures

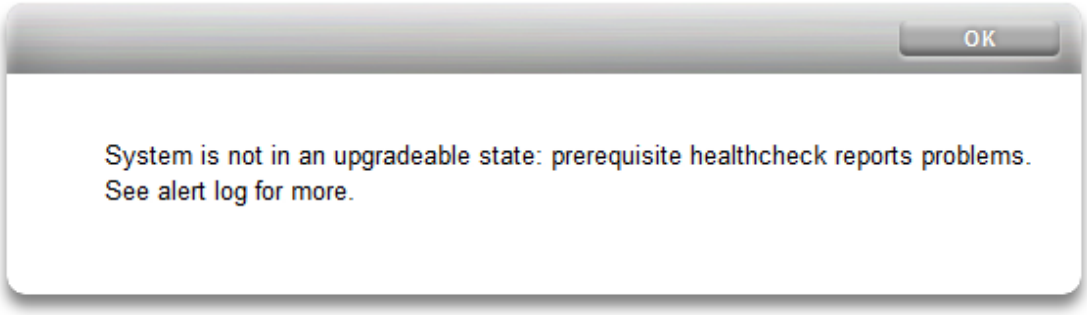


```
hostname:maintenance system updates ak-nas@2013.06.05.0.0,1-1.6> upgrade
This procedure will consume several minutes and requires a system reboot upon
successful update, but can be aborted with [Control-C] at any time prior to
reboot. A health check will validate system readiness before an update is
attempted, and may also be executed independently using the check command.
Are you sure? (Y/N)
error: System is not in an upgradeable state: prerequisite healthcheck reports
problems.
See alert log for more.
```

Resolving Health Check Alerts

After a software update health check failure, you can review the Alert Log and take action to resolve each failure based on the message in the log. The following table lists the update health check failures that can block an update, and describes the associated Alert Log message and recommended order of steps you can take to resolve the issue. For component faults, follow the instructions for removal and installation found in the maintenance procedures for your controller.

Table 23-1 Resolving Health Check Alerts

ID and Alert Log Message	Failure	Resolution Steps
B1 "System software update cannot proceed: Slot <label> in disk shelf <name> is reported as absent."	SIM or IOM cannot be detected.	1, 2, 4
B2 "System software update cannot proceed: Slot <label> in disk shelf <name> is faulted."	SIM or IOM is faulted.	1, 2, 4
C1 "System software update cannot proceed: Some slots of disk shelf <name> have no firmware revision information."	SIM or IOM is missing firmware revision information.	1, 4
C2 "System software update cannot proceed: The slots of disk shelf <name> have non-uniform part numbers."	SIMs or IOMs report different part numbers.	2, 4
C5 "System software update cannot proceed: The slots of disk shelf <name> have mixed firmware revisions <rev1> and <rev2>."	SIMs or IOMs report different firmware revisions.	4
E1 "System software update cannot proceed: Disk shelf <name> has <just one path *or* zero paths>."	Disk shelf does not have two paths.	1, 2, 4

Table 23-1 (Cont.) Resolving Health Check Alerts

ID and Alert Log Message	Failure	Resolution Steps
E2 "System software update cannot proceed: Disk shelf <name> path <pathname> is <state>."	Disk shelf path is not online.	1, 2, 4
E3 "System software update cannot proceed: Pool <data *or* log> disk <label> in disk shelf <name> has <just one path *or* zero paths>."	Disk or log device that is configured in a pool does not have two paths.	3, 4
E5 "System software update cannot proceed: Disk shelf <name> has disk <label> with the same iport being used for multiple paths."	Disk path redundancy at the SAS iport level.	5
V1 "<product> is not supported in this release."	Controller (product) is not supported for this software release.	

Resolving Health Check Alerts

Use the following procedure to resolve health check alerts.

1. If a SAS port LED is not lit, check all connections and replace cables as needed.
2. Identify the affected chassis, then disconnect and remove the faulted SIM or IOM. After two minutes, re-seat the SIM or IOM, and wait for the steady Power LED before reconnecting cables.
3. Identify the affected chassis, and remove the faulted disk. After 30 seconds, re-seat the disk, and wait for the steady or flashing LED.
4. Contact Oracle Support for component service or replacement.
5. Identify the affected chassis and disk. Ensure that each HBA cable on a controller is attached to a unique I/O Module (IOM) on the disk shelf, and that IOM0 of a disk shelf is never cabled to IOM1 of the same disk shelf. Proceeding before resolving this SAS cabling problem might lead to an outage, especially if the upgrade involves new disk shelf firmware.

24
Deferred Updates

Each software upgrade might contain new firmware or updates to external resources. If an update is not backward compatible, it is presented as a "deferred update" and is not automatically applied. To understand the implications of applying deferred updates, see Understanding Deferred Updates.

> **Note:**
>
> Replication failures can occur if the target appliance does not support a feature used by the project or share being replicated. See the "Incompatible Target" failure in section Replication Failures in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

To apply deferred updates, use one of the following tasks:

- Applying Deferred Updates (BUI)
- Applying Deferred Updates (CLI)

For a description of a deferred update, select a deferred update from the following table.

Table 24-1 Deferred Updates

Feature	Version Introduced
Passthrough x Deferred Update	2009.Q2.0.0
User Quotas Deferred Update	2009.Q3.0.0
COMSTAR Deferred Update	2009.Q3.0.0
Triple Parity RAID Deferred Update	2009.Q3.0.0
Dedup Deferred Update	2010.Q1.0.0
Replication Deferred Update	2010.Q1.0.0
Received Properties Deferred Update	2010.Q1.0.0
Slim ZIL Deferred Update	2010.Q3.1.0
Snapshot Deletion Deferred Update	2010.Q3.1.0
Recursize Snapshots Deferred Update	2010.Q3.1.0
Multi Replace Deferred Update	2010.Q3.1.0
RAIDZ Mirror Deferred Update	2011.1.0.0
Optional Child Dir Deferred Update	2011.1.0.0
Multiple Initiator Groups per LUN Deferred Update	2011.1.8.0
Large Block Sizes Deferred Update	2013.1.1.0
Sequential Resilvering Deferred Update	2013.1.2.0

Table 24-1 (Cont.) Deferred Updates

Feature	Version Introduced
NDMP-zfs Replica Backup Deferred Update	2013.1.4.0
ACL Passthrough with Mode Preservation Deferred Update	OS8.6.0
LZ4 Data Compression Deferred Update	OS8.7.0
Fast Copy with Encryption Deferred Update	OS8.7.0
Data Deduplication v2 Deferred Update	OS8.7.0
Asynchronous Dataset Deletion Deferred Update	OS8.7.0
RAID Space Efficiency Improvements Deferred Update	OS8.7.0
Support for Windows Compatible Automated Snapshot Names	OS8.7.3
Enable Pool Encryption Deferred Update	OS8.8.0
Deadlist Performance Deferred Update	OS8.8.6
Compact File Metadata for Encryption with Enhanced Replication Support Deferred Update	OS8.8.6
Reserved Pool Version Deferred Update	OS8.8.6
Improve LUN Attach/Detach Performance Deferred Update	OS8.8.22
Support for Snapshot Retention Deferred Update	OS8.8.39
File Retention Deferred Update	OS8.8.45
Unicode Versioning Deferred Update	OS8.8.51
Raw Crypto Replication Deferred Update	OS8.8.57
File Retention on Expiry Deferred Update	OS8.8.63

Applying Deferred Updates (BUI)

Use the following procedure to apply all available deferred updates using the BUI. You must apply all deferred updates or none; you cannot select individual deferred updates.

1. From the **Maintenance** menu, select **System**.

   The **Deferred Updates** section lists the available deferred updates, a description, and what effects the update will have on the system.

2. Click the **APPLY** button to apply all available deferred updates.

   Deferred updates will apply to both nodes in a cluster, and the cluster peer must be up and available to apply any deferred updates.

Related Topics

- Understanding Deferred Updates

Applying Deferred Updates (CLI)

Use the following procedure to apply all available deferred updates using the CLI. You must apply all deferred updates or none; you cannot select individual deferred updates.

1. Go to `maintenance system updates`.
2. To determine if there are any available deferred updates, enter `show`:

```
hostname:maintenance system updates> show
Updates:

UPDATE                                  DATE                    STATUS
ak-nas@2011.04.24.3.0,1-2.19.11.2       2012-6-24 17:14:19      current
ak-nas@2011.04.24.4.2,1-1.28            2012-11-5 03:11:34      waiting
ak-nas@2013.06.05.0.0,1-1.2             2013-6-19 12:58:18      unavailable
```

3. If deferred updates are available, you can use the `apply` command to apply all available deferred updates. Waiting media is listed as unavailable if the currently active software is not sufficiently recent to be able to upgrade to that waiting media. In this example, media has been downloaded for both 2011.1.4.2 and 2013.1 but the 2013.1 update is unavailable because you first need to upgrade to at least 2011.1.4.2.

```
Deferred updates:

The following incompatible updates are available. Applying these updates will
enable new software features as described below, but will prevent older
versions of the software from accessing the underlying resources. You should
apply deferred updates once you have verified that the current software update
is functioning and a rollback is not required. Applying deferred updates in a
cluster will also update any resources on the cluster peer.

1. Support for the "passthrough-x" aclinherit property for shares.
clownfish:maintenance system updates> apply
Applying deferred updates will prevent rolling back to previous versions of
software.

Are you sure? (Y/N)
hostname:maintenance system updates> Y
```

Related Topics

- Understanding Deferred Updates

Passthrough x Deferred Update

For filesystems, ACLs are inherited according to the `aclinherit` property on the filesystem (or inherited from the project). Previous versions of software allowed four options for this setting: "discard", "noallow", "restricted", and "passthrough". The 2009.Q2.0.0 release introduced a new option, "passthrough-x", with slightly different semantics as described in the product documentation:

Same as "passthrough", except that the owner, group, and everyone ACL entries inherit the execute permission only if the file creation mode also requests the execute bit.

The "passthrough" mode is typically used to cause all "data" files to be created with an identical mode in a directory tree. An administrator sets up ACL inheritance so that all files are created with a mode, such as 0664 or 0666. This all works as expected for data files, but

you might want to optionally include the execute bit from the file creation mode into the inherited ACL. One example is an output file that is generated from tools, such as "cc" or "gcc". If the inherited ACL does not include the execute bit, then the output executable from the compiler is not executable until you use `chmod(1)` to change the permissions for the file.

To use this mode, the storage pool must be upgraded. If you choose not to upgrade the pool and attempt to use this property, you will get an error indicating that the storage pool needs to be upgraded first. There is no other implication of applying this update, and it can be ignored if there is no need to use this new setting. Applying this update is equivalent to upgrading the on-disk ZFS pool to ZFS Pool Version Summary.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

User Quotas Deferred Update

With the 2009.Q3 software release, the system supports user and group quotas on a per-share basis. To make use of this feature, a deferred update must be applied to upgrade all shares in the system to support this feature. Applying this deferred update also allows the current usage (user or group) to be queried on a per-filesystem or per-project basis. To quote the product documentation:

Quotas can be set on a user or group at the filesystem level. These enforce physical data usage based on the POSIX or Windows identity of the owner or group of the file or directory. There are some significant differences between user and group quotas and filesystem and project data quotas.

This update is applied in the background, and it takes time proportional to the number of shares and amount of data on the system. Until this deferred update is finished, attempts to apply user quotas produce an error indicating that the update is still in progress.

Before attempting to use user or group quotas, read Space Management for Shares in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

COMSTAR Deferred Update

The COMSTAR framework relies on a ZFS pool update for complete support of persistent group reservations (PGRs). Before this update has been applied, the number of reservations stored with each LUN is severely limited, and may even be zero. Applying this update is equivalent to upgrading the on-disk ZFS pool to version 16.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Triple Parity RAID Deferred Update

This update provides the ability to use the triple-parity RAID storage profile, RAID-Z with three parity sectors per stripe. Triple-parity offers increased protection against drive failures and additional overall availability.

To use this mode, the storage pool must be upgraded. If you choose not to upgrade the pool and attempt to use this property, you will get an error indicating that the storage pool needs to be upgraded first. There is no other implication of applying this update, and it can be ignored if there is no need to use this new setting. Applying this update is equivalent to upgrading the on-disk ZFS pool to version 17.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Dedup Deferred Update

This update provides the ability to use data deduplication. To use this mode, the storage pool must be upgraded. If you choose not to upgrade the pool and attempt to use this property, you will get an error indicating that the storage pool needs to be upgraded first. There is no other implication of applying this update, and it can be ignored if there is no need to use this new setting. Applying this update is equivalent to upgrading the on-disk ZFS pool to version 21.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Replication Deferred Update

The 2010.Q1 release stores replication configuration differently than 2009.Q3 and earlier releases. This update migrates existing target, action, and replica configuration created under an earlier release to the new form used by 2010.Q1 and later.

After upgrading to 2010.Q1, but before this update is applied, incoming replication updates for existing replicas fail. Replicas received under earlier releases are not manageable via the BUI or CLI, although they occupy space in the storage pool. Additionally, the system does not send replication updates for actions configured on earlier releases.

After applying this update, incoming replication updates for replicas originally received on earlier releases continue normally and without a full resynchronization. The system also sends incremental replication updates for actions configured under earlier releases.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Received Properties Deferred Update

The 2010.Q1 feature that enables administrators to customize properties on replicated shares relies on a ZFS pool update. This update provides support of persistent local changes to received properties. Before this update has been applied, the system does not allow administrators to change properties on replicated shares. Applying this update is equivalent to upgrading the on-disk ZFS pool to version 22.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Slim ZIL Deferred Update

This deferred update changes the layout of ZFS intent log blocks to improve synchronous write performance. Before this update has been applied, log records continue to be written in the old format and performance may be reduced. Applying this update is equivalent to upgrading the on-disk ZFS pool to version 23.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Snapshot Deletion Deferred Update

This deferred update increases snapshot deletion parallelism and reduces the size of transaction groups associated with snapshot deletion to improve systemic responsiveness. Before this update has been applied, new snapshot data is stored in the old format and deleted using the old algorithm. Note that any snapshots created before this update is applied are also deleted using the old algorithm. Applying this update is equivalent to upgrading the on-disk ZFS pool to version 26.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Recursize Snapshots Deferred Update

This deferred update allows recursive snapshots to be taken without suspending the ZFS intent log, which greatly improves snapshot creation performance, especially on heavily loaded controllers. Before this update has been applied, the system is able to create snapshots, but does so using the old, much slower, algorithm. Applying this update is equivalent to upgrading the on-disk ZFS pool to version 27.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Multi Replace Deferred Update

This deferred update allows importing a pool with a missing log device and corrects the behavior of the system when a device that is being resilvered is itself removed or replaced. These fixes rely on a ZFS pool upgrade provided by this update. Before this update has been applied, the system is unable to import pools with missing log devices and does not correctly handle replacement of resilvering devices (see CR 6782540). Applying this update is equivalent to upgrading the on-disk ZFS pool to version 28.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

RAIDZ Mirror Deferred Update

This deferred update improves both latency and throughput on several important workloads. Applying this update is equivalent to upgrading the on-disk ZFS pool to version 29.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Optional Child Dir Deferred Update

This deferred update improves list retrieval performance and replication deletion performance by improving dataset rename speed. Before this update has been applied, the system is able to retrieve lists and delete replications, but does so using the old, much slower, recursive rename code. Applying this update is equivalent to upgrading the on-disk ZFS pool to version 31.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Multiple Initiator Groups per LUN Deferred Update

This deferred update allows a LUN to be associated with more than one initiator group at a time. If replicating LUNs with this update applied, also apply the update to the replication target system(s).

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Large Block Sizes Deferred Update

This feature provides improved performance when blocksize/recordsize larger than 128K are used. The supported block sizes are: 256k, 512k, and 1M. If large block sizes are used for projects or shares, replication of those projects or shares to systems that do not support large block sizes fail.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Sequential Resilvering Deferred Update

The previous resilvering algorithm repairs blocks from oldest to newest, which can degrade into a lot of small random I/O. The new resilvering algorithm uses a two-step process to sort and resilver blocks in LBA order.

The amount of improvement depends on how pool data is laid out. For example, sequentially written data on a mirrored pool shows no improvement, but randomly written data or sequentially written data on RAID-Z improves significantly - typically reducing time by 25 to 50 percent.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

NDMP-zfs Replica Backup Deferred Update

With this deferred update, NDMP backup type zfs can be performed directly within target appliance replication packages. It is not necessary to first clone a replica dataset (project or share) in order to back it up. You specify the project or share in the Data Management Application (DMA), and the appliance automatically backs up the latest system-generated replica snapshot for that project or share. Replicas of user-generated snapshots can also be backed up.

After applying this update on a target appliance, source appliances running software versions earlier than 2011.1.0.x will not be able to send replication updates to the target.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

ACL Passthrough with Mode Preservation Deferred Update

With this deferred update, there is improved interoperability between ZFS, NFS, and SMB sharing. Previous versions of software allowed five options for this setting: `discard`, `noallow`, `restricted`, `passthrough`, and `passthrough-x`. Release OS8.6.0 introduces the option `passthrough-mode-preserve`.

This option allows inheritable ACL entries to be inherited, while preserving the creation mode specified by the application. This preserves the inheritance properties so SMB creates ACLs that interoperate with shares accessed over NFS and SMB simultaneously.

If this deferred update is not applied and you use the `passthrough-mode-preserve` option in the CLI, an error appears indicates that the deferred update needs to be applied. In the BUI, this option will be grayed out and cannot be selected. A project or share with this option set can only be replicated to a target with the same deferred update applied.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

LZ4 Data Compression Deferred Update

The LZ4 compression feature is a data compression algorithm that compresses data using adaptive Lempel-Ziv coding. The LZ4 compression algorithm typically consumes less CPU than GZIP-2, but compresses better than LZJB, depending on the data that is compressed.

When LZ4 compression is set for projects and shares, replication to a target will fail if the target appliance does not support LZ4 compression. The LZ4 compression deferred update must be applied to both source and target appliances.

For more information, see "Data compression" described in Inherited Properties in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Fast Copy with Encryption Deferred Update

The fast copy with encryption feature is used for certain VMware appliance plug-ins. The plug-ins use the fast file cloning feature to offload the VM cloning operations to the appliance. This deferred update frees up NFS network bandwidth and improves VM cloning performance.

Install this deferred update to support these appliance plug-ins:

- Oracle ZFS Storage Appliance Plug-in for VMware Storage APIs for Array Integration
- Oracle ZFS Storage Appliance Provider for VMware vSphere APIs for Storage Awareness (VASA)

For information about plug-in availability, see the release notes in the My Oracle Support document Oracle ZFS Storage Appliance: Software Updates (Doc ID 2021771.1).

Related Topics

- Oracle ZFS Storage Appliance Plug-in Downloads
- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Data Deduplication v2 Deferred Update

This update provides a new data structure and algorithm for storing deduplication metadata. Oracle ZFS Storage Appliance is now able to store specific types of metadata on a designated set of meta devices to optimize use cases like deduplication.

Designate SSD meta cache devices that are at least 3.2 TB each for a new or existing storage pool that will use deduplication, so the meta devices can hold deduplication metadata tables. Shares and projects within the storage pool that are already configured for deduplication will be automatically and non-disruptively migrated in the background.

To preserve full replication compatibility if a pre-OS8.7.0 version of deduplication was in use, do not accept OS8.7.0 deferred updates, including this Data Deduplication v2 deferred update, on either sources or targets until software release OS8.7.0 or later has been installed on all. After both the sources and targets have been upgraded at least OS8.7.0, all OS8.7.0 deferred updates can be accepted.

When replicating a deduplicated source share to an OS8.7.0 target pool that does not meet deduplication requirements, the deduplication property of the replica share on the target is disabled. If the target pool is later upgraded by, for example, adding a meta cache device to the pool, the replication target then allows the administrator to set the deduplication property for projects and shares within existing replication packages. Any new replication packages created in the upgraded target pool will retain the deduplication settings of the source projects and shares.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Asynchronous Dataset Deletion Deferred Update

The Asynchronous Dataset Deletion feature improves performance of pool exports during a failback operation, if datasets are in the process of being destroyed.

Without this feature, a failback could take significantly longer to complete if a dataset destroy is in progress.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

RAID Space Efficiency Improvements Deferred Update

This deferred update significantly reduces the storage overhead for pools using the single-parity RAID data profile, thus increasing storage capacity.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Support for Windows Compatible Automated Snapshot Names

This deferred update will change the time stamp in automated snapshot names to be compatible with Microsoft Windows clients, by removing the ":" (colon) character. Applying this deferred update will change all existing automated snapshots to use the new time-stamp format. This change will not be propagated into replication packages until the packages have been reversed or severed. At that point, the automated snapshots will be renamed using the new time-stamp format.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Enable Pool Encryption Deferred Update

This deferred update provides the ability to create pools that are encrypted. Any project or share (filesystem or LUN) created in such a pool is encrypted by default.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Deadlist Performance Deferred Update

This deferred update delays creation of deadlists (ZFS metadata) for snapshots until they are actually needed, thus improving performance when snapshots are created and destroyed.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Compact File Metadata for Encryption with Enhanced Replication Support Deferred Update

> **Caution:**
>
> Accepting this deferred update could interrupt replication relations for encrypted data.

This deferred update enhances storage efficiency of encrypted data and also enables replication to transform data from unencrypted to encrypted form. Until the deferred update is also accepted on the replication target, existing replication relations that include encrypted data will fail after the deferred update is accepted on the replication source. The deferred update should be accepted after careful consideration of all existing replication relations that

include encrypted data. When possible, the deferred update should be accepted on replication targets first.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Reserved Pool Version Deferred Update

This deferred update is reserved for a change that was not implemented. There is no harm in accepting this deferred update.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Improve LUN Attach/Detach Performance Deferred Update

This deferred update improves LUN attach and detach speeds, and enhances Hyper-V Persistent Group Reservation performance.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Support for Snapshot Retention Deferred Update

This deferred update allows you to add a retention policy with a retention hold to a manual or automatic snapshot of a filesystem, LUN, or project to prevent snapshot deletion. Setting and releasing these retention holds requires specific role authorizations.

A retention hold for a manual snapshot, which is especially beneficial for legal holds, has no specified duration, and the snapshot cannot be deleted until the off option is manually set by an authorized user.

A retention hold for an automatic snapshot, or scheduled snapshot, is especially advantageous for meeting data governance needs, such as for internal retention and regulatory compliance requirements. A snapshot cannot be manually deleted until the "retention hold" value has been exceeded.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

File Retention Deferred Update

Before enabling file retention, see File Retention Management in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

File retention is advantageous for meeting data governance needs, especially for records retention compliance, and legal holds on files. The file retention feature provides the ability to mark individual files in a share to be retained until a future date, each retained file has its own expiration timestamp. A retained file can never be modified, even after its expiration, and those modifications include its name, attributes, and ACL. After the retention timestamp has passed, however, the file can be deleted. The retention policy is set at share creation and defaults to the same retention policy and periods as the parent project. For more information on the advantages of file retention, as well as implementation details, see File Retention Management in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

File retention can be set to off (default), privileged, or mandatory. After setting privileged or mandatory file retention, you define the retention periods: minimum, maximum, default, and optional grace period.

> **Caution:**
>
> Mandatory file retention affects the filesystem, project, and storage pool. Even the root user cannot shorten the retention period nor delete a retained file, its filesystem, its project, nor its storage pool before the retention timestamp has passed for all retained files. Carefully plan mandatory usage so that storage resources, especially pools and their associated drives, are not consumed for longer than necessary or overfilled.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Unicode Versioning Deferred Update

This deferred update provides support for the Unicode version 14.0.0 character set, with modified Turkish character support for the "dotless i" character. The previously supported version, Unicode 5.0.0, is available, supported, and is the default Unicode character set. Unicode 14 is selectable per file system.

The Unicode 14 character set has subtle changes for case insensitivity and normalization. It is not necessary to upgrade to Unicode 14 to use recently added characters; it is only necessary if you want its case-comparison and normalization changes. Most applications are not affected by upgrading to Unicode 14.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Raw Crypto Replication Deferred Update

Raw crypto replication improves the security and efficiency of replicating encrypted data by avoiding both decrypting the data on the source appliance and reencrypting it on the target appliance. Security is enhanced by sending the data encrypted instead of unencrypted.

The raw crypto option is enabled by default for new replication actions if the raw crypto deferred update has been applied on the source appliance. Similarly, the deferred update must be applied on target appliances to receive a replication with raw crypto enabled.

For implementation details, see Raw Crypto Replication in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

File Retention on Expiry Deferred Update

This deferred update adds on to the capability of the file retention feature, which requires the File Retention Deferred Update. Before enabling file retention, see File Retention Management in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Use the file retention on expiry policy to automatically delete files after expiration or to place a hold on all retained files. The default setting is `off`, which does not apply any behavior after a file's retention expires, and the file remains on the system. User authorization `retentionOnexpiry` is required to set property `retention.policy.onexpiry`.

For more information, see File Retention on Expiry Policy in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.8.x*.

Related Topics

- Applying Deferred Updates - BUI, CLI
- Understanding Deferred Updates

Understanding Deferred Updates

Each update may come with new firmware or updates to external resources. In general, these updates are backward compatible and applied automatically without user intervention. There are exceptions, however, for non-reversible updates. These updates involve updating a resource external to the system software in a way that is incompatible with older software releases. Once the update is applied, rolling back to previous versions will result in undefined behavior. For these updates, you will always be given an explicit option of applying them automatically during upgrade or applying them after the fact. They are therefore referred to as "deferred updates."

When applying an update to a version with incompatible version changes, you are given an option to apply these version changes as part of the update. For each version change, the benefits of applying the change are presented to you. The default is to not apply them, requiring you to return to the **Updates** view and apply them once the system has rebooted after the update is applied. This allows you to verify that the rest of the software is functional, and a rollback is not required before applying the update.

If you elect to not apply deferred updates during an update, you can return to the **Updates** view at any point to apply the update. If deferred updates are available for the current software version, they appear as a list below the current set of available

updates, with an **Apply** button to apply the updates. Deferred updates in a cluster take effect on both storage controllers simultaneously, and can only be applied while both Oracle ZFS Storage Appliance controllers are operational. Because deferred updates are listed only for resources present on the local storage controller, in a cluster it may be the case that deferred updates are available only for resources active on the peer controller. In a cluster, it is therefore necessary to check both storage controllers to determine the availability of deferred updates.

Related Topics

- Applying Deferred Updates - BUI, CLI

25
Rolling Back System Software

The rollback procedure reverts all of the system software and all of the metadata settings of the system back to their state just prior to applying an update. This feature is implemented by taking a snapshot of various aspects of the system before the new update is applied, and rolling back this snapshot to implement the rollback. The implications of rollback are as follows:

- Any Oracle ZFS Storage Appliance configuration changes are reverted and lost. For example, assume you are running version V, and then you upgrade to V+1, and then you change your DNS server. If you execute a rollback, then your DNS server setting modification is effectively undone and removed from the system permanently.
- Conversely, any changes made to *user data* are *not* reverted: if you upgrade from V to V+1, and clients then create directories or modify shares in any way, those changes still exist after the rollback (as you would expect).
- If the appliance is running version V, and has previous rollback targets V-1 and V-2, and you revert all the way to version V-2 (thereby "skipping" V-1), then you not only are removing the system software settings and system software for V, but also for V-1. That is, after a rollback to V-2, it is as if updates V-1 and V never happened. However, the software upload images for V-1 and V are saved on the system, and you can apply them again after the rollback, if you wish, by re-executing the update.

If after applying an update, the system is back up and running, you can use either the BUI or the CLI to initiate a rollback to one of two previously applied updates. If the system is not able to run at all after an update, then use the procedure Fail-safe Rollback.

Fail-safe Rollback

Administrators can execute a fail-safe rollback of the system software from the serial console by selecting one of the other boot menu entries, if present. Although rollback can also be requested from the BUI or CLI, rollback is offered from the boot menu because it is possible that rollback will be needed in scenarios where the new system software has completely failed, that is, has failed to boot.

1. To rollback from the console, access the serial console as usual, and during boot, before the ten-second timeout, use the arrow key to move the menu selection down to one of the earlier entries.

```
   GNU GRUB  version 0.97  (639K lower / 2620352K upper memory)
+-------------------------------------------------------------------------+
| Sun ZFS Storage 7120 2013.06.05.0.0,1-1.6                               |
| Sun ZFS Storage 7120 2011.04.24.4.2,1-1.28                                  |
|                                                                         |
+-------------------------------------------------------------------------+
     Use the ^ and v keys to select which entry is highlighted.
     Press enter to boot the selected OS, 'e' to edit the
     commands before booting, or 'c' for a command-line.
```

   After the rollback boot menu entry is selected, the system boots the old kernel software, but the rollback must be manually confirmed on the console to commit the rollback, which

effectively removes all changes to the system that have happened since, as described earlier. The confirmation step is similar to the following example:

```
SunOS Release 5.11 Version ak/generic@2011.04.24.4.2,1-1.28 64-bit
Copyright (c) 1983, 2010, Oracle and/or its affiliates. All rights reserved.

System rollback in progress.
Rollback to: ak/nas@2011.04.24.4.2,1-1.28

Proceed with system rollback [y,n,?]
```

2. When prompted to proceed with system rollback, enter one of the following commands:
   a. Enter `y` to proceed with the rollback.

      The system reboots using the prior snapshot.
   b. Enter `n` to cancel the rollback and immediately reboot the system.

      This action allows the administrator to select the current software version or an older snapshot boot image.

Rolling Back the System Software (BUI)

Use the following procedure to roll back the system software using the BUI.

1. To roll back, locate a previous image, and click on the rollback icon .
2. You are asked to confirm rollback request, the system reboots, and it executes the rollback.

   Unlike fail-safe rollback, you are not asked for further confirmation when the system reboots. For information on the fail-safe rollback, see Fail-safe Rollback.

Rolling Back the System Software (CLI)

Use the following procedure to roll back the system software using the CLI.

1. To roll back to a previous software version, select the update that corresponds to that version and execute the `rollback` command.
2. You are asked to confirm rollback request, the system reboots, and it executes the rollback.

   Unlike fail-safe rollback, you are not asked for further confirmation when the system reboots. For information on the fail-safe rollback, see Fail-safe Rollback.

26

Removing an Update Package

This section describes how to remove an Oracle ZFS Storage Appliance update package.

Removing an Update Package (BUI)

Use the following steps to remove an update package using the BUI.

- To remove an update package, highlight the corresponding row, and click the trash icon  .

Removing an Update Package (CLI)

Use the following steps to remove an update package using the CLI.

- To remove an update package, enter `destroy`, space, and specify the update to be removed:

```
hostname:maintenance system updates> destroy ak-nas@2009.10.14,1-0-nd
This will destroy the update "ak-nas@2009.10.14,1-0-nd". Are you sure? (Y/N) y
hostname:maintenance system updates>
```