

Oracle® ZFS Storage Appliance Customer Service Manual

For ZS5-x, ZS4-4, ZS3-x, 7x20 Controllers, and DEx-24, Sun Disk Shelves, Release OS8.7.0

ORACLE®

Part No: E78914-01
March 2017

Part No: E78914-01

Copyright © 2009, 2017, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Référence: E78914-01

Copyright © 2009, 2017, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer un risque de dommages corporels. Si vous utilisez ce logiciel ou ce matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour des applications dangereuses.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée de The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers, sauf mention contraire stipulée dans un contrat entre vous et Oracle. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> ou le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> si vous êtes malentendant.

Contents

Servicing the Oracle ZFS Storage Appliance	15
Getting Started with Servicing the Hardware	17
Introduction to the BUI	17
Introduction to the CLI	18
BUI Hardware View	18
▼ Viewing the Remaining SSD Endurance (BUI)	19
▼ Setting a Threshold Alert for SSD Endurance (BUI)	20
▼ Viewing CPU Details (BUI)	21
System Overview	21
System Chassis	22
Disk Shelves	23
Chassis Details View	24
Disk Shelf Details View	25
SSD Endurance	27
InfiniBand Host Controller Adapters	27
Service Processor	28
CLI Hardware View	28
Component Properties	29
▼ Viewing the Remaining SSD Endurance (CLI)	30
▼ Setting a Threshold Alert for SSD Endurance (CLI)	31
▼ Viewing CPU Details (CLI)	31
▼ Performing a Diagnostic Reboot (CLI)	32
Configuring the Management Port	32
Using Oracle ILOM to Diagnose Hardware Faults	33
▼ Identifying the Oracle ILOM Firmware Version (BUI)	34
▼ Identifying the Oracle ILOM Firmware Version (CLI)	34
▼ Logging in to Oracle ILOM Using a Local Serial Connection	34

▼ Logging in to Oracle ILOM Remotely Using a Web Interface	35
▼ Logging in to Oracle ILOM Remotely Using a Command Line Interface	35
▼ Viewing and Clearing CPU Faults from Oracle ILOM	36
Understanding Hardware Status	36
▼ Locating the Chassis Serial Number	37
▼ Locating a Failed Component (BUI)	37
▼ Locating a Failed Component (CLI)	37
Working with Problems	39
Active Problems Display	39
Repairing Active Problems	40
Related Features	40
Using Logs	41
Alert Log Entries	41
Fault Log Entries	42
System Log Entries	42
Audit Log Entries	43
Phone Home Log Entries	43
▼ Viewing Logs (BUI)	44
▼ Listing Logs (CLI)	44
▼ Viewing All Log Entries (CLI)	44
▼ Viewing Groups of Log Entries (CLI)	45
▼ Viewing Entry Details (CLI)	46
▼ Exporting Logs (BUI)	47
▼ Exporting Logs (CLI)	48
Using Workflows	49
Prerequisites for Servicing Hardware	51
Safety Information	51
Required Tools and Information	52
Powering Off the Controller	52
▼ Extending the Storage Controller from the Rack	53
▼ Removing the Top Cover	54
▼ Powering Off the Disk Shelf	55
Servicing the ZS5-4 Controller	57
▼ Replacing a ZS5-4 HDD	57

▼ Replacing a ZS5-4 Fan Module	59
▼ Replacing a ZS5-4 Power Supply	63
▼ Replacing a ZS5-4 DIMM	66
▼ Replacing a ZS5-4 Memory Riser Card	69
▼ Replacing or Adding a ZS5-4 PCIe Card	72
▼ Replacing a ZS5-4 System Battery	78
ZS5-4 Controller Hardware Overview	82
ZS5-4 Configuration Options	82
ZS5-4 Front and Rear Panel Components	83
ZS5-4 Specifications	87
ZS5-4 Internal Components	88
ZS5-4 CPUs and Memory Riser Card	96
ZS5-4 PCIe Card Configuration	100
ZS5-4 Attached Storage	101
Servicing the ZS5-2 Controller	103
▼ Replacing a ZS5-2 HDD	103
▼ Replacing a ZS5-2 Fan Module	106
▼ Replacing a ZS5-2 Power Supply	111
▼ Replacing a ZS5-2 DIMM	116
▼ Replacing or Adding a ZS5-2 PCIe Card	121
▼ Replacing a ZS5-2 System Battery	123
ZS5-2 Controller Hardware Overview	127
ZS5-2 Configuration Options	127
ZS5-2 Front and Rear Panel Components	128
ZS5-2 Specifications	131
ZS5-2 Internal Components	132
ZS5-2 PCIe Card Configuration	137
ZS5-2 Cable Management Arm	138
ZS5-2 Attached Storage	139
Servicing the ZS4-4 Controller	141
▼ Replacing a ZS4-4 HDD or SSD	141
▼ Replacing a ZS4-4 Power Supply	143
▼ Replacing a ZS4-4 Fan Module	146
▼ Identifying a Faulted ZS4-4 Memory Module	148
▼ Replacing a ZS4-4 DIMM	150

▼ Replacing a ZS4-4 Memory Riser Card	153
▼ Replacing a ZS4-4 PCIe Card	156
▼ Replacing a ZS4-4 Battery	160
ZS4-4 Controller Hardware Overview	163
ZS4-4 Base Configuration	163
ZS4-4 Front and Rear Panel Components	164
ZS4-4 PCIe I/O Cards	168
ZS4-4 Internal Components	172
ZS4-4 Attached Storage	178
Servicing the ZS3-4 Controller	179
▼ Replacing a ZS3-4 HDD or SSD	179
▼ Replacing a ZS3-4 Fan Module	181
▼ Replacing a ZS3-4 Power Supply	183
▼ Replacing a ZS3-4 DIMM	185
▼ Replacing a ZS3-4 PCIe Card or Riser	194
▼ Replacing a ZS3-4 Battery	198
ZS3-4 Controller Hardware Overview	201
ZS3-4 Base Configuration	201
ZS3-4 Specifications	201
ZS3-4 Front and Rear Panel Components	202
ZS3-4 Internal Components	206
ZS3-4 PCIe I/O Cards	212
ZS3-4 Attached Storage	216
Servicing the ZS3-2 Controller	219
▼ Installing Optional PCIe Slot Cards in Order	220
▼ Disconnecting the RJ-45 Cable from the Cluster Serial Port	221
▼ Replacing a ZS3-2 HDD or SSD	222
▼ Replacing a ZS3-2 Power Supply	225
▼ Replacing a ZS3-2 Fan Module	228
▼ Replacing a ZS3-2 DIMM	231
▼ Replacing a ZS3-2 PCIe Card or Riser	238
▼ Replacing a ZS3-2 Flash Drive	244
▼ Replacing a ZS3-2 Battery	246
ZS3-2 Controller Hardware Overview	248
ZS3-2 Base Configuration	248

ZS3-2 Front and Rear Panel Components	249
ZS3-2 Specifications	253
ZS3-2 Internal Components	254
ZS3-2 PCIe I/O Cards	261
ZS3-2 Optional Cable Management Arm	262
ZS3-2 Attached Storage	263
Servicing the 7x20 Controller	265
▼ Replacing a 7x20 HDD or SSD	265
▼ Replacing a 7x20 Fan Module	267
▼ Replacing a 7x20 Power Supply	272
▼ Replacing a 7x20 DIMM	276
▼ Replacing a 7x20 Controller PCIe Card or Riser	289
▼ Replacing a 7x20 Battery	300
7420 Controller Hardware Overview	306
7420 Configuration Options	306
7420 Front and Back Panel Components	307
7420 Specifications	310
7420 Internal Components	311
7420 PCIe I/O Cards	316
7420 Attached Storage	321
7320 Controller Hardware Overview	321
7320 Configuration Options	321
7320 Front and Back Panel Components	324
7320 Electrical Specifications	328
7320 Internal Components	328
7320 Attached Storage	333
7120 Controller Hardware Overview	333
7120 Base Configuration	334
7120 Chassis Overview	335
7120 Front Panel	335
7120 Rear Panel	336
7120 Physical Specifications	337
7120 Electrical Specifications	337
7120 Internal Boards	337
7120 CPU and Memory	338
7120 Power Distribution, Fan Module and Disk Components	339

7120 I/O Components	341
7120 NIC/HBA Options	342
7120 PCIe Options	343
7120 Connectors	344
7120 Cables	345
7120 Attached Storage	346
Servicing the Disk Shelves	347
Electrostatic Discharge Precautions	347
▼ Replacing a Faulty Disk Shelf Drive	348
▼ Replacing a Disk Shelf Power Supply	351
▼ Replacing a Disk Shelf I/O Module	357
▼ Replacing a Disk Shelf SIM Board	359
▼ Adding a New Disk Shelf	362
▼ Adding a Disk Shelf Cache Device	363
Connecting to Attached Storage	365
Disk Shelf Hardware Overview	365
Oracle Storage Drive Enclosure DE3-24	365
Oracle Storage Drive Enclosure DE2-24	373
Sun Disk Shelf	380
Disk Shelf Configurations	384
Maximum Disk Shelves per Controller Configuration	387
SAS-2 and SAS-3	388
Viewing System Disks Status	389
Working with Support Bundles	391
Support Bundle Options	391
▼ Generating and Uploading a Support Bundle (BUI)	392
▼ Generating and Uploading a Support Bundle (CLI)	392
▼ Viewing Support Bundle Details (BUI)	393
▼ Viewing Support Bundle Details (CLI)	394
▼ Uploading a Locally Generated Bundle (CLI)	394
▼ Retrying a Failed Support Bundle Upload (BUI)	395
▼ Retrying a Failed Support Bundle Upload (CLI)	395
▼ Canceling a Pending Support Bundle Operation (BUI)	396

▼ Canceling a Pending Operation (CLI)	396
▼ Deleting a Support Bundle (BUI)	396
▼ Deleting a Support Bundle (CLI)	396
Performing the Initial Setup	399
Restarting the Appliance	401
Performing a Factory Reset	403
Backing Up the Configuration	405
Configuration Backup Contents	405
Configuration Restore Impact	406
Security Considerations for Configuration Backups	408
▼ Creating a Configuration Backup (BUI)	408
▼ Creating a Configuration Backup (CLI)	409
▼ Restoring from a Saved Configuration (BUI)	409
▼ Restoring from a Saved Configuration (CLI)	409
▼ Deleting a Saved Configuration (BUI)	410
▼ Deleting a Saved Configuration (CLI)	410
▼ Exporting a Saved Configuration (BUI)	411
▼ Exporting a Saved Configuration (CLI)	411
▼ Importing a Saved Configuration (BUI)	411
▼ Importing a Saved Configuration (CLI)	412
▼ Viewing Appliance Configurations (BUI)	412
▼ Viewing Appliance Configurations (CLI)	412
Working with Software Notifications and Updates	415
▼ Setting Up Software Notification (BUI)	415
▼ Setting Up Software Notification (CLI)	416
▼ Checking for and Downloading Software Updates (BUI)	416
▼ Checking for and Downloading Software Updates (CLI)	417
▼ Removing System Updates (BUI)	418
▼ Removing System Updates (CLI)	418

Software Update States	419
Upgrading the Software	421
Preparing for a Software Upgrade	421
Upgrading Software on a Standalone Controller	422
▼ Upgrading Software on a Standalone Controller (BUI)	423
▼ Upgrading Software on a Standalone Controller (CLI)	425
Upgrading Software on Clustered Controllers	427
▼ Preparing to Upgrade Clustered Controllers (BUI)	428
▼ Preparing to Upgrade Clustered Controllers (CLI)	429
▼ Upgrading Clustered Controllers (BUI)	431
▼ Upgrading Clustered Controllers (CLI)	433
Updating the Platform Firmware	439
▼ Checking the Platform Firmware Versions (BUI)	440
▼ Checking the Platform Firmware Versions (CLI)	440
▼ Updating the Platform Firmware (BUI)	441
▼ Updating the Platform Firmware (CLI)	442
Working with Health Checks	445
▼ Checking System Health (BUI)	445
▼ Checking System Health (CLI)	447
Troubleshooting Update Health Check Failures	448
Resolving Health Check Alerts	448
▼ Resolving Health Check Alerts	449
Deferred Updates	451
▼ Applying Deferred Updates (BUI)	452
▼ Applying Deferred Updates (CLI)	452
Passthrough x Deferred Update	453
User Quotas Deferred Update	454
COMSTAR Deferred Update	455
Triple Parity RAID Deferred Update	455
Dedup Deferred Update	455
Replication Deferred Update	456
Received Properties Deferred Update	456

Slim ZIL Deferred Update	456
Snapshot Deletion Deferred Update	457
Resize Snapshots Deferred Update	457
Multi Replace Deferred Update	457
RAIDZ Mirror Deferred Update	458
Optional Child Dir Deferred Update	458
Multiple Initiator Groups per LUN Deferred Update	458
Large Block Sizes Deferred Update	459
Sequential Resilvering Deferred Update	459
NDMP-zfs Replica Backup Deferred Update	459
ACL Passthrough with Mode Preservation Deferred Update	460
LZ4 Data Compression Deferred Update	460
Fast Copy with Encryption Deferred Update	461
Data Deduplication v2 Deferred Update	461
Asynchronous Dataset Deletion Deferred Update	462
RAID Space Efficiency Improvements Deferred Update	462
Understanding Deferred Updates	462
Working with Firmware Updates	465
▼ Checking the Status of Firmware Updates (BUI)	467
▼ Checking the Status of Firmware Updates (CLI)	467
Rolling Back System Software	469
▼ Fail-safe Rollback	469
▼ Rolling Back the System Software (BUI)	470
▼ Rolling Back the System Software (CLI)	471
Removing an Update Package	473
▼ Removing an Update Package (BUI)	473
▼ Removing an Update Package (CLI)	473

Servicing the Oracle ZFS Storage Appliance

This document encompasses the following sections that describe maintenance procedures for the Oracle ZFS Storage Appliance.

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

- [“Getting Started with Servicing the Hardware” on page 17](#)
- [“Prerequisites for Servicing Hardware” on page 51](#)
- [“Servicing the ZS5-4 Controller” on page 57](#)
- [“Servicing the ZS5-2 Controller” on page 103](#)
- [“Servicing the ZS4-4 Controller” on page 141](#)
- [“Servicing the ZS3-4 Controller” on page 179](#)
- [“Servicing the ZS3-2 Controller” on page 219](#)
- [“Servicing the 7x20 Controller” on page 265](#)
- [“Servicing the Disk Shelves” on page 347](#)

To maintain the system, use the following sections.

- [“Viewing System Disks Status” on page 389](#)
- [“Working with Support Bundles” on page 391](#)
- [“Performing the Initial Setup” on page 399](#)
- [“Restarting the Appliance” on page 401](#)
- [“Performing a Factory Reset” on page 403](#)
- [“Backing Up the Configuration” on page 405](#)
- [“Working with Software Notifications and Updates” on page 415](#)
- [“Upgrading the Software” on page 421](#)
- [“Working with Health Checks” on page 445](#)
- [“Deferred Updates” on page 451](#)
- [“Working with Firmware Updates” on page 465](#)
- [“Rolling Back System Software” on page 469](#)
- [“Removing an Update Package” on page 473](#)

Getting Started with Servicing the Hardware

To get started with servicing the hardware, use the following sections:

- [“Introduction to the BUI” on page 17](#)
- [“Introduction to the CLI” on page 18](#)
- [“BUI Hardware View” on page 18](#)
- [“System Overview” on page 21](#)
- [“Chassis Details View” on page 24](#)
- [“Disk Shelf Details View” on page 25](#)
- [“CLI Hardware View” on page 28](#)
- [“Configuring the Management Port” on page 32](#)
- [“Using Oracle ILOM to Diagnose Hardware Faults” on page 33](#)
- [“Understanding Hardware Status” on page 36](#)
- [“Working with Problems” on page 39](#)
- [“Using Logs” on page 41](#)
- [“Using Workflows” on page 49](#)

Introduction to the BUI

The Browser User Interface (BUI) is the graphical tool for administration of the appliance. The BUI provides an intuitive environment for 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 OS8.7.0*](#). Direct your browser to the system using either the *IP address* or *host name* you assigned to the NET-0 port during initial configuration as follows: `https://ipaddress:215` or `https://hostname:215`. The login screen appears. The online help linked in the top right of the BUI is context-sensitive. For every top-level and second-level screen in the BUI, the associated help page appears when you click the Help button.

Introduction to the CLI

The Command Line Interface (CLI) is designed to mirror the capabilities of the BUI, while also providing a powerful scripting environment for performing repetitive tasks. The following sections describe details of the CLI. When navigating through the CLI, there are two principles to be aware of:

- Tab completion is used extensively: if you are not sure what to type in any given context, pressing the Tab key will provide you with possible options. Throughout the documentation, pressing Tab is presented as the word "tab" in bold italics.
- Help is always available: the help command provides context-specific help. Help on a particular topic is available by specifying the topic as an argument to help, for example help commands. Available topics are displayed by tab-completing the help command, or by typing help topics.

You can combine these two principles, as follows:

```
hostname:> helptab
builtins  commands  general  help      properties  script
```

BUI Hardware View

The BUI hardware view provides interactive illustrations that enable you to browse through the appliance and attached disk shelf components.

The buttons in the hardware view are described in the following table:

TABLE 1 Icons in the BUI Hardware View

Icon	Description	Icon	Description
	Show a more detailed view of this component		Toggle blinking of the locator LED for this component
	Leave this detailed view		Power off, reboot, or diagnostic reboot
	Click for more details		Offline disk
	Hardware component is ok (green)		Port active
	Hardware component is not present (grey)		Port inactive

Icon	Description	Icon	Description
	Hardware component is faulted (amber)		

▼ Viewing the Remaining SSD Endurance (BUI)

Use the following procedure to view the remaining SSD endurance.

1. **Go to Maintenance > 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. Go to Configuration > Alerts > Threshold Alerts.
2. Click the add icon .
3. In the Threshold field, select SSD: Endurance on any device and specify its parameters.



Add Threshold Alert [CANCEL] [APPLY]

Threshold

SSD: Endurance on any device falls below 95 percent

Timing

for at least 5 minutes only between 00:00 and 00:00 (UTC)
 only during Weekdays

Repost alert every 5 minutes while this condition persists.
 Also post alert when this condition clears for at least 5 minutes

Alert actions


Send email Send to [] Subject [] [TEST]

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

▼ Viewing CPU Details (BUI)

Use the following procedure to view CPU details.

1. **Go to Maintenance > 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.**

System Overview

The main hardware page lists the system chassis, a summary of its contents, and any attached disk shelves (on supported systems). This provides an overview of the hardware present on the system. The power icon , located in the upper left of the view, presents a dialog box to

either power off, reboot (power cycle), or reboot the appliance with diagnostics. Only select the diagnostic reboot option when instructed by Oracle Service personnel because it could take a long time to complete and could have adverse results if not performed properly. Do not perform a diagnostic reboot when system-affecting operations are occurring, such as upgrading firmware, executing commands, and configuring or unconfiguring storage. The diagnostic reboot option is not available when using the system chassis power icon.

System Chassis

The primary system chassis is shown on the top half of the view. At the top left, click the right-arrow icon  to get more detail about the chassis. The indicator notes if there are any faulted components within the chassis, and the name of the chassis. The chassis name is initially set to the appliance name during installation. To change the chassis name, use the entry field on the Configuration > Services > System Identity screen. For more information, see [“System Identity Configuration”](#) in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0*.

At the top right of the system chassis is the locator icon  to light the locator LED, and the power icon , which presents a dialog box to either power off or reboot (power cycle) the appliance.

A thumbnail of the controller is presented at left. Clicking on the thumbnail or the "Show Details" link takes you to a detailed view of the chassis, and is identical to clicking on the right-arrow icon  at the top left of the view.

The following information is presented in a summary view:

TABLE 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 is only valid for standalone systems. If there are no data disks present, "-" will be displayed.
Cache	Size and number of read-optimized cache devices in the system chassis. This is only valid for expandable systems that support additional disk shelves.

Property	Description
	If there are no cache devices installed in the system chassis, "-" will be displayed.
Log	Size and number of log disks in the system chassis. This is only valid for standalone systems. If there are no log devices installed in the system chassis, "-" will be displayed.
Total	Total size and count of all disks in the system.

Disk Shelves

A list of disk shelves, if supported, is displayed at the bottom of the Maintenance > Hardware view. The thumbnail to the left represents the front of the currently selected disk shelf. Clicking on the right-pointing arrow or double-clicking on a row within the list will provide complete details about the disk shelf. The state indicator will be orange if the chassis contains any faulted components.

TABLE 3 Disk Shelf Properties

Property	Description
Name	Name of the disk shelf, used in faults and alerts. This is initially set to the serial number of the disk shelf, but can be changed by clicking on the name within the list.
MFR/MODEL	The disk shelf manufacturer and the model number.
Type	Device type HDD or SSD.
RPM	Revolutions per minute; the speed of the disk drive.
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 within the shelf, then "-" is displayed. 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" on page 384 .
Log	Total size of all write-optimized cache devices within the drive shelf. If there are no log devices within the shelf, then "-" is displayed.
Paths	Total number of I/O paths to the disk shelf. The only supported configurations are those with multiple paths to all disks, so this should read "2" under normal operating circumstances. Clicking the information icon  will bring up a dialog with information about each path. This includes which HBAs are connected to the disk shelf, and the state of any paths. 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 it displays two paths to the chassis.
Locate 	Toggle the locator LED for this disk shelf. If the LED is currently on, then this indicator will flash.

Related Topics

- [“Disk Shelf Details View” on page 25](#)

Chassis Details View

To view the chassis details, click on the right-arrow icon  (or one of the alternative forms described above). This view includes some of the same controls in the upper left (state, name, locate, reset, power off), as well as listings of all the components in the chassis.

At the left is a set of images describing the chassis. If there are multiple views, then you can switch between them by clicking on the name of the view above the image.

For each view, faulted components will be highlighted in red. In addition, the currently selected component will be highlighted in the image. Clicking on a component within the image will select the corresponding component in the list to the right.

A tab is present for each component type in the following list. Each component type has a state icon which will be orange if there is a faulted component of the given type.

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

Clicking on a component type will display a list of all physical locations within the chassis where components may be present. Clicking on a component within the list will highlight it within the appropriate chassis image. Clicking on the information icon  while over a row or double-clicking a row will bring up a dialog with detailed information about the component. The information displayed in the list depends on the component type, but is a subset of the information available in the component detail. Disks and service processors support additional operations described below. Each component can report any or all of the following properties:

TABLE 4 Chassis Component Properties

Property	Description
Label	Human-readable identifier for this component within the chassis. This is typically, but not necessarily, equivalent to the label printed on the physical chassis.

Property	Description
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.
Active Problems	For a faulted component, links to active problems affecting the component.
Manufacturer	Component manufacturer.
Model	Component model.
Build	Manufacturing build identifier. This is used to identify a particular location or batch where the component was manufactured.
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.
Serial	Component serial number.
Revision	Firmware or hardware revision of the component.
Size	Total memory or storage, in bytes.
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.
Speed	Processor speed, in gigahertz.
Cores	Number of CPU cores.
GUID	Hardware global unique identifier.
Endurance	The lifetime percentage remaining on an SSD. Lifetime starts at 100% and decreases with disk usage.
Last Update	The date and time of the last endurance reading of the SSD.

Disk Shelf Details View

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




TABLE 5 Disk Device Detail Properties

Property	Description
ID	Unique identifier of the device.
MFR/MODEL	The device manufacturer and the model number.
RPM	Revolutions per minute; the speed of the disk drive.
Type	Disk type. Can be one of System, Data, Log, Cache, Meta, Spare, or Peer. Note - 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” on page 384 .
Locate 	Toggle the locator LED for this disk device. If the device LED is currently on, then this indicator will flash.
Online 	Offline the 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 it. 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

Property	Description
	<p>possible to offline a disk that would prevent access to data (i.e. offlining both halves of a mirror). If the device is an active hot spare, this 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.</p> <p>Note - Do not perform a pool configuration operation while a disk firmware upgrade is occurring. To check if an upgrade is in progress, navigate to Maintenance > System in the BUI, or <code>maintenance system updates</code> in the CLI.</p>

SSD Endurance

SSD endurance reports the remaining life expectancy of an SSD. Endurance properties are reported on the BUI Maintenance > Hardware chassis details page and in the CLI maintenance chassis disk context. You can set a threshold alert when an SSD exceeds a specified percentage. For example, set an alert to occur when one or more SSD devices exceeds a 95% threshold. For more information, see [“Threshold Alerts” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0.](#)

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.

InfiniBand Host Controller Adapters

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

TABLE 6 Additional Properties for InfiniBand Host Controller Adapters

Action	Description
State	<p>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.</p>
GUID	The hardware assigned port GUID.

Action	Description
Speed	The current port speed enabled: Single Data Rate (SDR), Dual Data Rate (DDR) or Quad Data Rate (QDR)

Service Processor

The service processor behaves differently from other component nodes. Instead of providing a list of components, it presents a set of network properties that can be configured from the storage appliance. The following properties control the behavior of the service processor network management port.

TABLE 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	Dotted decimal subnet, when using static IP configuration.
Default Gateway	IPv4 default gateway address.

Changing multiple values in conflicting ways (such as changing static IP assignments while in DHCP mode) has undefined behavior.

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 0839QCJ01A             ok        Sun Microsystems, Inc. Sun Storage 7320
cpu-000     CPU 0                       ok          AMD                Quad-Core AMD Op
```

```

cpu-001      CPU 1           ok      AMD           Quad-Core AMD Op
cpu-002      CPU 2           ok      AMD           Quad-Core AMD Op
cpu-003      CPU 3           ok      AMD           Quad-Core AMD Op
disk-000     HDD 0           ok      STEC          MACH8 IOPS
disk-001     HDD 1           ok      STEC          MACH8 IOPS
disk-002     HDD 2           absent  -             -
disk-003     HDD 3           absent  -             -
disk-004     HDD 4           absent  -             -
disk-005     HDD 5           absent  -             -
disk-006     HDD 6           ok      HITACHI       HTE5450SASUN500G
disk-007     HDD 7           ok      HITACHI       HTE5450SASUN500G
fan-000      FT 0            ok      unknown       ASY,FAN,BOARD,H2
fan-001      FT 0 FM 0       ok      Sun Microsystems, Inc. 541-2068
fan-002      FT 0 FM 1       ok      Sun Microsystems, Inc. 541-2068
fan-003      FT 0 FM 2       ok      Sun Microsystems, Inc. 541-2068
fan-004      FT 1            ok      unknown       ASY,FAN,BOARD,H2
fan-005      FT 1 FM 0       ok      Sun Microsystems, Inc. 541-2068
fan-006      FT 1 FM 1       ok      Sun Microsystems, Inc. 541-2068
fan-007      FT 1 FM 2       ok      Sun Microsystems, Inc. 541-2068
memory-000   DIMM 0/0        ok      HYNIX         4096MB DDR-II 66
memory-001   DIMM 0/1        ok      HYNIX         4096MB DDR-II 66
...

```

A 5th and 6th column for serial number ("SERIAL") and revolutions per minute ("RPM") have been truncated in the above example, as has the length of this list.

Component Properties

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

TABLE 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

CLI Property	BUI Property
offline (writeable)	(status indicator)
part	Part
present	(status indicator)
revision	Revision
serial	Serial
size	Size
speed	Speed
type	(combined with use)
use	Type

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

▼ 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 = 2014-3-3 22:04:14
```

```
locate = false
offline = false
```

As shown above, disk-015 is at 100% which indicate 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.

▼ 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 = AMD
    model = Quad-Core AMD Opteron(tm) Processor 8356
    part = 1002
    revision = 03
    cores = 4
    speed = 2.14G
```

▼ 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> diagreboot
```

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 using the Configuration > Cluster option in the BUI. 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.

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 <http://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](#)” on page 34
- “[Logging in to Oracle ILOM Remotely Using a Web Interface](#)” on page 35
- “[Logging in to Oracle ILOM Remotely Using a Command Line Interface](#)” on page 35
- “[Viewing and Clearing CPU Faults from Oracle ILOM](#)” on page 36

▼ 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 (defaults are root and changeme).**

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

▼ Logging in to Oracle ILOM Remotely Using a Web Interface

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

Before You Begin This procedure requires that you know the administrative account name and password (defaults are root and changeme), 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.

- 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

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

Before You Begin This procedure requires that you know the administrative account name and password (defaults are root and changeme), and that you know the IP address or host name of the controller Service Processor (SP).

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

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

Before You Begin Log in to the server as root using the ILOM CLI.

1. **To view server 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” on page 37](#)
- [“Locating a Failed Component \(BUI\)” on page 37](#)
- [“Locating a Failed Component \(CLI\)” on page 37](#)

▼ 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. **Go to the Maintenance > Hardware screen.**
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. **(Optional) 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 the maintenance hardware context:**

```
hostname:> maintenance hardware
```

2. List the appliance components:

```
hostname:maintenance hardware> list
```

	NAME	STATE	MODEL	SERIAL
chassis-000	hostname	ok	Sun Storage 7320	unknown
chassis-001	000000000C	faulted	J4410	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 disks:

```
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 above example.

5. Select the faulted disk and turn on the locator LED:

```

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, the appliance detects persistent hardware failures (*faults*) and software failures (*defects*, often included under faults) and reports them as active problems on this screen. If the phone home service is enabled, active problems are automatically reported to Oracle Support, where a support case may be opened depending on the service contract and the nature of the fault.

Active Problems Display

For each problem, the appliance reports what happened, when the problem was detected, the severity and type of the problem, and whether it has been phoned home. Below are some example faults as they would be displayed in the BUI:

TABLE 9 Example BUI Problem Displays

Date	Description	Type	Phoned Home
2009-09-16 13:56:36	SMART health-monitoring firmware reported that a disk failure is imminent.	Major Fault	Never
2009-09-05 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
2009-08-21 16:40:37	The ZFS pool has experienced currently unrecoverable I/O failures.	Major Error	Never
2009-07-16 22:03:22	A memory module is experiencing excessive correctable errors affecting large numbers of pages.	Major Fault	Never

This information can also be viewed in the CLI:

```
hostname:> maintenance problems show
Problems:
```

COMPONENT	DIAGNOSED	TYPE	DESCRIPTION
problem-000	2010-7-27 00:02:49	Major Fault	SMART health-monitoring firmware reported that a failure is imminent on disk 'HDD 17'.

Selecting any fault shows more information about the fault including the impact to the system, affected components, the system's automated response (if any), and the recommended action for the administrator (if any). In the CLI, only the "uuid", "diagnosed", "severity", "type", and "status" fields are considered stable. Other property values may change from release to release.

For hardware faults, you may be able to select the affected hardware component to locate it on the Maintenance > Hardware screen.

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. A repair typically involves replacing a physical component (for hardware faults) or reconfiguring and restarting the affected service (for software defects). After repairing a problem, it 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.

Related Features

- A persistent log of all faults and defects is available under Logs as the **Fault Log**. For more information, see [“Fault Log Entries” on page 42](#).
- Faults and defects are subcategories of Alerts. Filter rules can be configured to cause the 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 OS8.7.0](#).

Using Logs

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

- “Alert Log Entries” on page 41
- “Fault Log Entries” on page 42
- “System Log Entries” on page 42
- “Audit Log Entries” on page 43
- “Phone Home Log Entries” on page 43
- “Viewing Logs (BUI)” on page 44
- “Listing Logs (CLI)” on page 44
- “Viewing All Log Entries (CLI)” on page 44
- “Viewing Groups of Log Entries (CLI)” on page 45
- “Viewing Entry Details (CLI)” on page 46
- “Exporting Logs (BUI)” on page 47
- “Exporting Logs (CLI)” on page 48

Alert Log Entries

This is the appliance alert log, recording key events of interest during appliance operation.

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

TABLE 10 Example BUI Alert Log Entry Displays

Time	Event ID	Description	Type
2013-9-16 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
2013-9-16 13:01:51	 8fb8688c-08f2-c994-a6a5-ac6e755e53bb	A disk has been inserted into slot 'HDD 4' of chassis 'JBOD #1'.	Minor alert
2013-9-16 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.

The 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 OS8.7.0.](#)

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 11 Example BUI Fault Log Entry Displays

Date	Event ID	Description	Type
2013-9-5 17:42:35	9e46fc0b-b1a4-4e69-f10f-e7dbe80794fe	The device 'HDD 6' has failed or could not be opened.	Major Fault
2013-9-3 19:20:15	d37cb5cd-88a8-6408-e82d-c05576c52279	External sensors indicate that a fan is no longer operating correctly.	Minor Fault
2013-8-21 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” on page 39.](#)

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 12 Example BUI System Log Entry Displays

Time	Module	Priority	Description
2013-10-11 14:13:38	ntpddate	error	no server suitable for synchronization found

Time	Module	Priority	Description
2013-10-11 14:03:52	genunix	notice	SunOS Release 5.11 Version ak/generic@2013.10.10,1-0 64-bit
2013-10-11 14:02:04	genunix	notice	done
2013-10-11 14:02:01	genunix	notice	syncing file systems...
2013-10-11 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 (see [“Configuring Users” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0](#)), each audit entry should be noted with a reason.

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

TABLE 13 Example BUI Audit Log Entry Displays

Time	User	Host	Summary	Session Annotation
2013-10-12 05:20:24	root	deimos	Disabled ftp service	
2013-10-12 03:17:05	root	deimos	User logged in	
2013-10-11 22:38:56	root	deimos	Browser session timed out	
2013-10-11 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 OS8.7.0](#).

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

TABLE 14 Example BUI Phone Home Entry Display

Time	Description	Result
2013-10-12 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, navigate to the Maintenance > Logs screen.**
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         2013-10-16 02:44:04
audit        42        2013-10-16 18:19:53
fltlog       2         2013-10-16 02:44:04
scrk         0         -
system       100      2013-10-16 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  2013-1-9 15:13:50  root, 10.154.161.197, User logged in
entry-001  2013-1-9 15:17:42  root, 10.154.161.197, Released resources to
cluster peer
entry-002  2013-1-9 15:20:17  root, 10.154.161.197, Cluster takeover
requested
entry-003  2013-1-9 15:20:33  root, 10.154.161.197, Transferred ak:/net/igb1
to 2917b8aa-0b0a-4b74-f36b-ff0a8d150c3b
...
entry-2077 2013-12-17 05:24:43  osc_agent, 10.80.218.16, User logged in
entry-2078 2013-12-17 05:24:45  osc_agent, 10.80.218.16, User logged out
entry-2079 2013-12-17 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 2014-2-20 18:05:44  root, 10.159.134.211, Beginning system update
to ak-nas@2014.02.11,1-0
entry-2875 2014-2-20 18:08:18  root, 10.159.134.211, Rebooting system as part
of update to ak-nas@2014.02.11,1-0
entry-2876 2014-2-20 18:08:20  root, 10.159.134.211, User logged out
entry-2877 2014-2-20 18:16:06  root, 10.159.134.211, User logged in
entry-2878 2014-2-20 18:18:20  root, 10.159.134.211, Deleted update media
ak-nas@2011.04.24.8.0,1-1.43
entry-2879 2014-2-20 18:34:19  root, 10.159.134.211, Cluster takeover
requested
...
entry-2970 2014-2-27 17:40:37  root,
dhcp-amer-vpn-rmdc-anyconnect-10-159-100-157.vpn.oracle.
com,
User logged out of CLI
entry-2971 2014-2-27 17:48:23  root, 10.159.100.157, User logged in
entry-2972 2014-2-27 19:05:46  root, 10.159.100.157, Browser session timed out
entry-2973 2014-2-27 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> list
```

```
ENTRY      TIME                SUMMARY
entry-023  2013-1-16 15:06:36  d1dd862d-93ff-6e3d-a51c-fe5f81159a3e, System
software update cannot proceed: J4400 and J4500
disk shelves are not supported in this
...
entry-122  2013-2-17 20:34:06  65f79e6b-1a77-6041-9a92-9440dbba4c40, Finished
replicating project 'Test1' from appliance
'AdamZhang'., Minor Alert
```

```
hostname:maintenance logs alert> previous
```

```
ENTRY      TIME                SUMMARY
entry-000  2013-11-21 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  2013-1-16 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
```

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

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

```
ENTRY      TIME                SUMMARY
entry-721  2013-2-18 22:02:02  <system>, <system>, Request to create replication
package
entry-722  2013-2-18 22:02:08  <system>, <system>, Request to modify replication package
entry-723  2013-2-18 23:33:32  root, 10.135.69.243, User logged in
entry-724  2013-2-19 00:48:51  root, 10.135.69.243, Browser session timed out
entry-725  2013-2-19 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 = 2013-10-15 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 OS8.7.0.](#)

▼ 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, navigate to the **Maintenance > Logs** screen.
2. Click **Collect**.

The screenshot shows the Oracle ZFS Storage Appliance BUI interface. The top navigation bar includes 'Configuration', 'Maintenance' (highlighted), 'Shares', 'Status', and 'Analytics'. Below this, there are sub-tabs for 'Hardware', 'System', 'Problems', 'Logs', and 'Workflows'. The 'Logs' section is active, showing a table of alerts. The table has columns for 'Time', 'Event ID', 'Description', and 'Type'. A 'Collect' button is located in the top right corner of the alert list.

TIME	EVENT ID	DESCRIPTION	TYPE
2013-11-12 11:38:28	6a519404-e6a2-4ad4-bbf1-9f5b2239a745	Failed to upload system logs to: http://10.153.34.75:85, Error: access denied by host.	Minor Alert
2013-11-12 11:38:21	fe80e0fa-f05e-ed87-fd00-eb32f9ce8dd5	Collecting system logs for upload to: http://10.153.34.75:85.	Minor Alert
2013-11-12 11:37:40	a06c2c8e-5147-5227-a365-b1f614711bea	Failed to upload system logs to: http://10.153.34.75:85/shares/export/fs1/, Error: access denied by host.	Minor Alert
2013-11-12 11:37:33	da8bd68b-f2ba-ef6c-be23-8b51a6d5c909	Collecting system logs for upload to: http://10.153.34.75:85/shares/export/fs1/.	Minor Alert
2013-11-12 11:32:01	3cf8ae21-5c60-c928-c6e6-fe9348f62e4a	Successfully uploaded system logs to: http://10.153.34.75:85.	Minor Alert
2013-11-12 11:31:54	772a9827-d88e-4664-f6f2-d13075817534	Collecting system logs for upload to: http://10.153.34.75:85.	Minor Alert
2013-11-12 11:31:36	4f3f6008-4d94-c881-abc9-c02d0cb4a3bd	Successfully uploaded system logs to: http://10.153.34.75:86.	Minor Alert

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	2013-2-18 23:02:29
audit	731	2013-2-20 16:13:04
fltlog	2	2013-2-3 06:29:02
scrk	0	-
system	44	2013-1-14 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: Wed Nov 13 03:50:15 UTC 2013
PLATFORM: Sun-Fire-X4240, CSN: 0000000000,
HOSTNAME: hpc-iwashi-01
SOURCE: ak-diagnosis, REV: 1.0
EVENT-ID: d83655cd-c03d-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: Wed Nov 13 03:50:15 UTC 2013
PLATFORM: Sun-Fire-X4240, CSN: 0000000000, HOSTNAME: hpc-iwashi-01
SOURCE: ak-diagnosis, REV: 1.0
EVENT-ID: d83655cd-c03d-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 flltlog>
```

Using Workflows

A workflow is a script that is uploaded to and managed by the Oracle ZFS Storage Appliance by itself. For information on scripting, see [“Working with CLI Scripting” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0](#). 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 OS8.7.0](#). Workflows may also be optionally executed as alert actions or at a designated time. As such, workflows allow for the 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 OS8.7.0](#).

Prerequisites for Servicing Hardware

This section describes the prerequisites for servicing hardware. See the followings for details:

- [“Safety Information” on page 51](#)
- [“Required Tools and Information” on page 52](#)
- [“Powering Off the Controller” on page 52](#)
- [“Extending the Storage Controller from the Rack” on page 53](#)
- [“Removing the Top Cover” on page 54](#)
- [“Powering Off the Disk Shelf” on page 55](#)

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 ZS3-2. 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
- 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.
- SSH in to the appliance and enter the `maintenance system poweroff` command.
- 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.

- SSH or use 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” on page 55](#)

▼ Extending the Storage 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” on page 52](#).

1. **Disconnect all cables from the back of the storage controller. Be sure the cables are labeled for proper connection later.**
2. **From the front of the storage controller, release the two slide release latches.**

3. **While squeezing the slide release latches, slowly pull the storage controller forward until the slide rails latch.**




▼ **Removing the Top Cover**

To access internal controller components for servicing, remove the top cover. Use the following procedure to remove the top cover.

1. **Fully loosen the two captive screws at the rear of the top cover (1).**

2. **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” on page 52](#).

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 "O" 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” on page 52](#)
- For details on NSPF profiles, see [“Storage Pool Concepts” in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0*](#)

Servicing the ZS5-4 Controller

Be sure to read sections [“Getting Started with Servicing the Hardware”](#) on page 17 and [“Prerequisites for Servicing Hardware”](#) on page 51.

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see [“Repairing Active Problems”](#) on page 40.

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

- [“Replacing a ZS5-4 HDD”](#) on page 57
- [“Replacing a ZS5-4 Fan Module”](#) on page 59
- [“Replacing a ZS5-4 Power Supply”](#) on page 63
- [“Replacing a ZS5-4 DIMM”](#) on page 66
- [“Replacing a ZS5-4 Memory Riser Card”](#) on page 69
- [“Replacing or Adding a ZS5-4 PCIe Card”](#) on page 72
- [“Replacing a ZS5-4 System Battery”](#) on page 78

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

- [“ZS5-4 Configuration Options”](#) on page 82
- [“ZS5-4 Front and Rear Panel Components”](#) on page 83
- [“ZS5-4 Specifications”](#) on page 87
- [“ZS5-4 Internal Components”](#) on page 88
- [“ZS5-4 CPUs and Memory Riser Card”](#) on page 96
- [“ZS5-4 PCIe Card Configuration”](#) on page 100
- [“ZS5-4 Attached Storage”](#) on page 101

▼ Replacing a 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 a ZS5-4 HDD. For information on disk compatibility, see the [Oracle Systems Handbook](#).

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, navigate to Maintenance > System in the BUI, or `maintenance system updates` in the CLI.


1. Identify the faulted HDD.

- If you are not physically at the system, go to the Maintenance > Hardware section of the BUI, click the right-arrow icon for the controller, and note the ID of the faulted drive, which has an amber icon. To highlight the location of the drive in the controller picture, click on the drive ID. 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 flash the locator indicator on the faulted drive.


2. Verify that the Ready to Remove indicator is lit on the faulted drive.

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, go to the **Hardware > Maintenance** screen, 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. Go to the **Maintenance > Hardware** screen of the BUI. 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” on page 40](#).

▼ Replacing a ZS5-4 Fan Module



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 a ZS5-4 fan module.

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

1. Identify the faulted fan module.


- If you are not physically at the system, go to the Maintenance > Hardware section of the BUI, 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 Storage Controller from the Rack” on page 53.](#)

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 server front panel are not lit.**

11. Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).

▼ Replacing a 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 a ZS5-4 power supply.

1. **Identify the faulted power supply.**
 - If you are not physically at the system, go to the Maintenance > Hardware section of the BUI, 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” on page 52](#).
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 server 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. Go to the Maintenance > Hardware screen of the BUI. 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” on page 40.](#)

▼ Replacing a 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.

Go to the Maintenance > Hardware section of the BUI, 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” on page 52.](#)

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 Storage Controller from the Rack” on page 53.](#)

5. Remove the top cover as described in [“Removing the Top Cover” on page 54.](#)


6. Remove the memory riser card in which the faulted DIMM resides. See [“Replacing a ZS5-4 Memory Riser Card” on page 69.](#)

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 a ZS5-4 Memory Riser Card”](#) on page 69.
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](#).
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. **Go to the Maintenance > Hardware screen of the BUI. 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”](#) on page 40.**

▼ Replacing a ZS5-4 Memory Riser Card

Use the following procedure to replace a ZS5-4 memory riser card.


1. **Shut down the controller using one of the power-off methods described in [“Powering Off the Controller”](#) on page 52.**
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 Storage Controller from the Rack”](#) on page 53.**
4. **Remove the top cover as described in [“Removing the Top Cover”](#) on page 54.**

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](#).
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. Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).

▼ Replacing or Adding a 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” on page 421](#). 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.
Go to the Maintenance > Hardware section of the BUI, 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 [“ZS5-4 PCIe Card Configuration” on page 100](#).

2. **Shut down the controller using one of the power-off methods described in [“Powering Off the Controller” on page 52](#).**

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 Storage Controller from the Rack” on page 53](#).**


5. **Remove the top cover as described in [“Removing the Top Cover” on page 54](#).**

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 server 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](#).
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. Go to the Maintenance > Hardware screen of the BUI. 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 [Oracle ZFS Storage Appliance Cabling Guide](#). 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](#) and connect the expansion storage as described in [Oracle ZFS Storage Appliance Cabling Guide](#).

▼ Replacing a 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 a ZS5-4 system battery.


1. **Shut down the controller using one of the power-off methods described in [“Powering Off the Controller”](#) on page 52.**
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 Storage Controller from the Rack”](#) on page 53.**
4. **Remove the top cover as described in [“Removing the Top Cover”](#) on page 54.**

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 a ZS5-4 PCIe Card”](#) on page 72.**
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 a ZS5-4 PCIe Card”](#) on page 72.
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](#).
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 OS8.7.0.](#)**

ZS5-4 Controller Hardware Overview

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

- [“ZS5-4 Configuration Options” on page 82](#)
- [“ZS5-4 Front and Rear Panel Components” on page 83](#)
- [“ZS5-4 Specifications” on page 87](#)
- [“ZS5-4 Internal Components” on page 88](#)
- [“ZS5-4 CPUs and Memory Riser Card” on page 96](#)
- [“ZS5-4 PCIe Card Configuration” on page 100](#)
- [“ZS5-4 Attached Storage” on page 101](#)

ZS5-4 Configuration Options

The 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


CPU	Memory	System Disks	PCIe Slots	HBA
				4-port (4x4) SAS-2 6 Gb/s external

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

ZS5-4 Front and Rear Panel Components

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

FIGURE 1 ZS5-4 Front Panel




<p>1 LEDs (left to right)</p> <ul style="list-style-type: none"> ■ Locator ■ Service action required ■ Power/OK 	<p>2 Power button</p>	<p>3 Alert indicators (left to right)</p> <ul style="list-style-type: none"> ■ Service Processor (SP) OK/Fault ■ Fan/ CPU/ Memory service action required ■ Power Supply (PS) service action required ■ Over-temperature warning
<p>4 Permanent HDD filler panels</p>	<p>5 System disk 1</p>	<p>6 System disk 0</p>

7 DVD filler panel	8 DB-15 VGA port	9 Two USB 2.0 ports
10 RFID tag		

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

FIGURE 2 ZS5-4 Rear Panel



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)	20 PSU 1 AC inlet	21 PSU 0 AC inlet
<ul style="list-style-type: none"> ■ Power/OK (green) ■ Attention (amber) ■ Locate (white) 		

Cluster Interface Card Compatibility - The 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](#).

Ethernet Ports

The ZS5-4 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 15 Gigabit Ethernet Port Status Indicators

Status Indicator Name	Location	Color	State and Meaning
Activity	Top left	Green	<ul style="list-style-type: none"> ■ ON – No link ■ OFF – No activity ■ FLASHING – Packet activity
Link speed	Top right	Bicolored: Amber/ Green	<ul style="list-style-type: none"> ■ OFF – 100BASE-T link ■ Amber ON – 1000BASE-T link ■ Green ON – 10GBASE-T link

Network Management Port

The 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 16 Network Management Port Status Indicators

Status Indicator Name	Location	Color	State and Meaning
Activity	Top left	Amber	<ul style="list-style-type: none"> ■ ON – Link up

Status Indicator Name	Location	Color	State and Meaning
Link speed	Top right	Green	<ul style="list-style-type: none">■ OFF – No link or down link■ FLASHING – Packet activity■ 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 server 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 ZS5-4, provides connectivity to external disk shelves. The HBA ports are numbered 3-0, top to bottom, as shown in the following figure.

FIGURE 3 ZS5-4 4x4 SAS HBA

ZS5-4 Specifications

ZS5-4 Physical Specifications: The 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

ZS5-4 Electrical Specifications: The 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

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).

ZS5-4 Internal Components


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

FIGURE 4 ZS5-4 Internal Components



Callout	Description	Callout	Description
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)		

FIGURE 5 ZS5-4 Internal Components (Exploded View)




Callout	Description	Callout	Description
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

Callout	Description	Callout	Description
6	CPU	14	Storage drive backplane board
7	Heatsink	15	Controller chassis
8	Cover		

ZS5-4 Storage Drive

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



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

ZS5-4 Fan Module

The six fan module slots are at the front of the 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.



Callout	Description	Callout	Description
1	Fan Module, FM 0	4	Fan Module, FM 3
2	Fan Module, FM 1	5	Fan Module, FM 4
3	Fan Module, FM 2	6	Fan Module, FM 5


ZS5-4 Power Supply

The two power supply slots are located on the right side of the controller (from the front of the server) 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.



Callout	Description
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.



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

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.



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

Memory Riser layout:


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



Call Out	Description
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.

ZS5-4 Storage Drive and USB

The storage drive slots are on the right side of the controller and are designated consecutively from bottom to top.




Callout	Description	Callout	Description
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		

ZS5-4 CPUs and Memory Riser Card

The four CPU sockets are located in the middle of the 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.



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

Memory Riser Card Components

The following illustration shows the components of the memory riser card.



Callout	Description	Callout	Description
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 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 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 above:



Memory Riser Card Population Rules


The memory riser population rules for the 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).

ZS5-4 PCIe Card Configuration

The eleven PCIe slots are located inside the 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.



Callout	Description	Callout	Description
1	PCIe 1	7	PCIe 7
2	PCIe 2	8	PCIe 8
3	PCIe 3	9	PCIe 9
4	PCIe 4	10	PCIe 10

Callout	Description	Callout	Description
5	PCIe 5	11	PCIe 11
6	PCIe 6		

PCIe Slot Length and Capabilities

The 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 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 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.

ZS5-4 Attached Storage

The 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 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” on page 384](#). The ZS5-4 does not support the Sun Disk Shelf.

Servicing the ZS5-2 Controller

Be sure to read sections [“Getting Started with Servicing the Hardware”](#) on page 17 and [“Prerequisites for Servicing Hardware”](#) on page 51.

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see [“Repairing Active Problems”](#) on page 40.

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

- [“Replacing a ZS5-2 HDD”](#) on page 103
- [“Replacing a ZS5-2 Fan Module”](#) on page 106
- [“Replacing a ZS5-2 Power Supply”](#) on page 111
- [“Replacing a ZS5-2 DIMM”](#) on page 116
- [“Replacing or Adding a ZS5-2 PCIe Card”](#) on page 121
- [“Replacing a ZS5-2 System Battery”](#) on page 123

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

- [“ZS5-2 Configuration Options”](#) on page 127
- [“ZS5-2 Front and Rear Panel Components”](#) on page 128
- [“ZS5-2 Specifications”](#) on page 131
- [“ZS5-2 Internal Components”](#) on page 132
- [“ZS5-2 PCIe Card Configuration”](#) on page 137
- [“ZS5-2 Cable Management Arm”](#) on page 138
- [“ZS5-2 Attached Storage”](#) on page 139

▼ Replacing a 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 a ZS5-2 HDD. For information on disk compatibility, see the [Oracle Systems Handbook](#).

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, navigate to Maintenance > System in the BUI, or `maintenance system updates` in the CLI.


1. Identify the faulted HDD.

- If you are not physically at the system, go to the Maintenance > Hardware section of the BUI, click the right-arrow icon for the controller, and note the ID of the faulted drive, which has an amber icon. To highlight the location of the drive in the controller picture, click on the drive ID. 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 flash the locator indicator on the faulted drive.


2. To remove a faulted drive, press the release button and pivot the release lever open to the right.



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




4. After at least 30 seconds, navigate to the Hardware > Maintenance screen, and click the right-arrow icon  on the system controller to verify that the software has detected that the drive is not present.
5. To insert the replacement drive, push the release button of the drive.
The release lever opens.
6. 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.

7. Slide the drive into the chassis until the release lever moves slightly inward.

8. **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.

9. **Go to the Maintenance > Hardware screen of the BUI. 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.**
10. **If a fault does not self-clear, see [“Repairing Active Problems” on page 40](#).**

▼ Replacing a 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 a ZS5-2 fan module.


1. **Identify the faulted fan module.**
 - If you are not physically at the system, go to the Maintenance > Hardware section of the BUI, 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” on page 52.](#)
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 Storage Controller from the Rack” on page 53.](#)
5. Remove the top cover as described in [“Removing the Top Cover” on page 54.](#)
6. Press the Fault Remind button on the motherboard to illuminate the fan module fault LED.




No.	Description
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. Go to the Maintenance > Hardware screen of the BUI. 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”](#) on page 40.

▼ Replacing a 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 a ZS5-2 power supply.

1. **Identify the faulted power supply.**
 - If you are not physically at the system, go to the Maintenance > Hardware section of the BUI, 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 faulty 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](#).
10. Attach the power cord to the power supply.
11. Verify that the green AC Present status indicator is lit.
12. Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).

▼ Replacing a 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.

Go to the Maintenance > Hardware section of the BUI, 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” on page 52.](#)

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 Storage Controller from the Rack” on page 53.](#)

5. Remove the top cover as described in [“Removing the Top Cover” on page 54.](#)

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 [ZS5-2 DIMM Configuration](#) 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](#).
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. Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).

▼ Replacing or Adding a 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 storage controller. If you are installing a newly released HBA, upgrade your system software before installing the HBA as described in [“Upgrading the Software” on page 421](#). All HBAs must be of the same type. Use the following procedure to replace or add a ZS5-2 PCIe card.


1. **Identify the PCIe card slot location.**
 - If replacing a faulted PCIe card, identify the PCIe card to be replaced.
Go to the Maintenance > Hardware section of the BUI, 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 [“ZS5-2 PCIe Card Configuration” on page 137](#).
2. **Shut down the controller using one of the power-off methods described in [“Powering Off the Controller” on page 52](#).**
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 Storage Controller from the Rack](#)” on page 53.**
6. **Remove the top cover as described in “[Removing the Top Cover](#)” on page 54.**
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](#).**
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. **Go to the Maintenance > Hardware screen of the BUI. 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 [Oracle ZFS Storage Appliance Cabling Guide](#).**
20. **If a fault does not self-clear, see [“Repairing Active Problems”](#) on page 40.**

▼ Replacing a 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 a ZS5-2 battery.

1. **Power off the storage controller as described in [“Powering Off the Controller”](#) on page 52.**
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 Storage Controller from the Rack](#)” on page 53.**
4. **Remove the top cover as described in “[Removing the Top Cover](#)” on page 54.**
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](#).**
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 OS8.7.0.](#)

ZS5-2 Controller Hardware Overview

This section describes the internal and external components of the Oracle ZFS Storage ZS5-2 controller.

- [“ZS5-2 Configuration Options” on page 127](#)
- [“ZS5-2 Front and Rear Panel Components” on page 128](#)
- [“ZS5-2 Specifications” on page 131](#)
- [“ZS5-2 Internal Components” on page 132](#)
- [“ZS5-2 PCIe Card Configuration” on page 137](#)
- [“ZS5-2 Cable Management Arm” on page 138](#)
- [“ZS5-2 Attached Storage” on page 139](#)

ZS5-2 Configuration Options

The 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

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

ZS5-2 Front and Rear Panel Components

ZS5-2 Front Panel Components - The ZS5-2 controller drive slots and front panel components are shown in the following figure.


FIGURE 6 ZS5-2 Front Panel



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)		

ZS5-2 Rear Panel Components - The ZS5-2 controller PCIe slots and rear panel components are shown in the following figure.

FIGURE 7 ZS5-2 Rear Panel



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 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](#).

Ethernet Ports

The ZS5-2 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 17 Gigabit Ethernet Port Status Indicators

Status Indicator Name	Location	Color	State and Meaning
Activity	Top left	Green	<ul style="list-style-type: none"> ■ ON – No link ■ OFF – No activity ■ FLASHING – Packet activity
Link speed	Top right	Bicolored: Amber/ Green	<ul style="list-style-type: none"> ■ OFF – 100BASE-T link ■ Amber ON – 1000BASE-T link ■ Green ON – 10GBASE-T link

Network Management Port

The 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 18 Network Management Port Status Indicators

Status Indicator Name	Location	Color	State and Meaning
Activity	Top left	Amber	<ul style="list-style-type: none"> ■ ON – Link up ■ OFF – No link or down link ■ FLASHING – Packet activity
Link speed	Top right	Green	<ul style="list-style-type: none"> ■ 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 server 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 ZS5-2, provides connectivity to external disk shelves. The HBA ports are numbered 3-0, top to bottom, as shown in the following figure.

FIGURE 8 ZS5-2 4x4 SAS HBA



ZS5-2 Specifications

ZS5-2 Physical Specifications: The 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

ZS5-2 Electrical Specifications: The 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
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

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.

ZS5-2 Internal Components

ZS5-2 Storage, Power, and Fan Modules - The ZS5-2 internal storage, power, and cooling components are described in the following figure and legend. The 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.

FIGURE 9 ZS5-2 Storage, Power, and Fan Modules




Figure Legend		
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)

ZS5-2 Memory, PCIe Cards, and System Battery - The 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.

FIGURE 10 ZS5-2 Memory, PCIe Cards, and Battery




Figure Legend		
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		

ZS5-2 DIMM Configuration - The 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.

FIGURE 11 ZS5-2 DIMM Configuration




Figure Legend

1 /SYS/MB/P0/D(0-5)	2 /SYS/MB/P0/D(6-11)	3 /SYS/MB/P1/D(0-5)
4 /SYS/MB/P1/D(6-11)	5 /System/Memory/DIMMs/ DIMM_(0-11)	6 /System/Memory/DIMMs/ DIMM_(12-23)


ZS5-2 PCIe Card Configuration

The 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 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.

ZS5-2 Cable Management Arm

The following figure identifies the components of the cable management arm (CMA). See the CMA installation instructions as described in [“Installing a ZS5-2 Cable Management Arm”](#) in [Oracle ZFS Storage Appliance Installation Guide](#).

FIGURE 12 ZS5-2 Controller Cable Management Arm**Figure Legend**

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	

ZS5-2 Attached Storage

The 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 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” on page 384](#). The ZS5-2 does not support the Sun Disk Shelf.

Servicing the ZS4-4 Controller

Be sure to read sections [“Getting Started with Servicing the Hardware”](#) on page 17 and [“Prerequisites for Servicing Hardware”](#) on page 51.

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see [“Repairing Active Problems”](#) on page 40.

To replace ZS4-4 hardware components, use the following tasks:

- [“Replacing a ZS4-4 HDD or SSD”](#) on page 141
- [“Replacing a ZS4-4 Power Supply”](#) on page 143
- [“Replacing a ZS4-4 Fan Module”](#) on page 146
- [“Identifying a Faulted ZS4-4 Memory Module”](#) on page 148
- [“Replacing a ZS4-4 DIMM”](#) on page 150
- [“Replacing a ZS4-4 Memory Riser Card”](#) on page 153
- [“Replacing a ZS4-4 PCIe Card”](#) on page 156
- [“Replacing a ZS4-4 Battery”](#) on page 160

To understand the ZS4-4 controller, use the following topics:

- [“ZS4-4 Base Configuration”](#) on page 163
- [“ZS4-4 Front and Rear Panel Components”](#) on page 164
- [“ZS4-4 PCIe I/O Cards”](#) on page 168
- [“ZS4-4 Internal Components”](#) on page 172
- [“ZS4-4 Attached Storage”](#) on page 178

▼ Replacing a 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 a ZS4-4 HDD or SSD. For information on disk compatibility, see the [Oracle Systems Handbook](#).

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, navigate to Maintenance > System in the BUI, or `maintenance system updates` in the CLI.

1. Identify the failed HDD or SSD.

If you are not physically at the system, go to the Maintenance > Hardware section of the BUI, click the right-arrow icon for the controller, and note the ID of the faulted drive, which has an amber icon. To highlight the location of the drive in the controller picture, click on the drive ID. 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 flash the locator indicator on the faulted drive.


2. Verify that the Ready to Remove indicator is lit on the faulted drive.

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, go to the Hardware > Maintenance screen, 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. Go to the Maintenance > Hardware screen of the BUI, 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”](#) on page 40.

▼ Replacing a ZS4-4 Power Supply


Storage 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 a 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” on page 52](#).**
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. Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).

▼ Replacing a 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 a ZS4-4 fan module.


1. **To flash the locator LED on the chassis to service, go to Maintenance > Hardware in the BUI 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, go to Maintenance > Hardware > 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 Storage Controller from the Rack”](#) on page 53.**
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. Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).

▼ Identifying a Faulted 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 ZS4-4 memory module.


1. To identify a general memory fault, go to the Maintenance > Hardware screen of the BUI, 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” on page 52](#).
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 Storage Controller from the Rack” on page 53](#).
5. Remove the top cover as described in [“Removing the Top Cover” on page 54](#).
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 a ZS4-4 Memory Riser Card”](#) on page 153 for removal and installation instructions.
- If a DIMM has failed, see [“Replacing a ZS4-4 DIMM”](#) on page 150 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 a ZS4-4 DIMM

Use the following procedure to replace a ZS4-4 DIMM.


1. **Shut down the controller using one of the power-off methods described in [“Powering Off the Controller” on page 52.](#)**
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 Storage Controller from the Rack” on page 53.](#)**


4. Remove the top cover as described in [“Removing the Top Cover” on page 54](#).
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 ZS4-4 Memory Module” on page 148](#).



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 a ZS4-4 Memory Riser Card” on page 153](#).
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](#).
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. Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).

▼ Replacing a ZS4-4 Memory Riser Card

Use the following procedure to replace a ZS4-4 memory riser card.


1. Shut down the controller using one of the power-off methods described in [“Powering Off the Controller” on page 52](#).
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 Storage Controller from the Rack” on page 53](#).
4. Remove the top cover as described in [“Removing the Top Cover” on page 54](#).
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](#).
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. Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).

▼ Replacing a ZS4-4 PCIe Card

Replacing 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. All HBAs must be of the same type. Use the following procedure to replace a ZS4-4 PCIe card.

1. Go to the Maintenance > Hardware screen of the BUI, 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” on page 52](#).
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 Storage Controller from the Rack](#)” on page 53.**
5. **Remove the top cover as described in “[Removing the Top Cover](#)” on page 54.**
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](#).

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. **Go to the Maintenance > Hardware screen of the BUI. 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”](#) on page 40.**
21. **If you replaced an HBA, connect the disk shelves to it as described in [“Connecting to Attached Storage”](#) on page 365. 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* and connect the expansion storage as described in [“Connecting to Attached Storage”](#) on page 365.**

▼ Replacing a ZS4-4 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 a ZS4-4 system battery.


1. **Shut down the controller using one of the power-off methods described in [“Powering Off the Controller”](#) on page 52.**
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 Storage Controller from the Rack” on page 53.](#)**
4. **Remove the top cover as described in [“Removing the Top Cover” on page 54.](#)**

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 a ZS4-4 PCIe Card” on page 156.](#)**
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 a 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](#).
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 OS8.7.0.](#)

ZS4-4 Controller Hardware Overview

To understand the ZS4-4 controller, use the following topics:

- [“ZS4-4 Base Configuration” on page 163](#)
- [“ZS4-4 Front and Rear Panel Components” on page 164](#)
- [“ZS4-4 PCIe I/O Cards” on page 168](#)
- [“ZS4-4 Internal Components” on page 172](#)
- [“ZS4-4 Attached Storage” on page 178](#)

ZS4-4 Base Configuration

The 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 19 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)

Refer to the Oracle ZFS Storage ZS4-4 data sheet at <http://www.oracle.com/goto/zs4-4> for the most recent component specification and physical, electrical, and environmental specifications.

For information about supported disks and capacities, refer to the [Oracle Systems Handbook](#).

ZS4-4 Front and Rear Panel Components

Front Panel Components - The ZS4-4 controller drive slots and front panel components are shown in the following figure.

FIGURE 13 ZS4-4 Controller Front Panel




Figure Legend	Figure Legend
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)

Figure Legend	Figure Legend
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)

Rear Panel Components - The ZS4-4 rear panel is shown in the following figure. Base configuration PCIe cards are not depicted in this illustration.

FIGURE 14 ZS4-4 Controller Rear Panel




Figure Legend	Figure Legend
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 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](#).

Ethernet Ports


The ZS4-4 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 figure below. 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.

FIGURE 15 ZS4-4 Ethernet Ports



Network Management Port

The network management connector (NET MGT), shown in the figure below, is an RJ-45 10/100/1000 BASE-T Ethernet port and provides an alternate terminal interface to the service processor (SP) console.

FIGURE 16 ZS4-4 Network Management Port

Serial Management Port


The serial management connector (SER MGT), shown in the figure below, is an RJ-45 port and provides a terminal connection to the SP console.

FIGURE 17 ZS4-4 Serial Management Port

4x4 SAS HBA

The 4x4 SAS HBA provides connectivity to external disk shelves. The HBA ports are numbered 3-0, top to bottom, as shown in the following figure.

FIGURE 18 4x4 SAS HBA Port Numbers




ZS4-4 PCIe I/O Cards

The 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.

FIGURE 19 ZS4-4 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.

ZS4-4 PCIe Base and Optional Configurations - The following table describes the PCIe base and optional slot assignments for 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

Slot	Description	Max	Type	Note
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
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

Slot	Description	Max	Type	Note
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.

Legend for Interconnect Types and Options			
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

Legend for Interconnect Types and Options			
D	GbE NIC UTP 10GBase-T - Copper RJ-45 Unshielded Twisted Pair		

ZS4-4 Internal Components

The 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.

FIGURE 20 Internal Components (Exploded View)




Figure	Legend		
1	Power supplies (CRU)	8	Memory riser card (CRU)
2	Power supply backplane (FRU)	9	Motherboard (FRU)

Figure	Legend		
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

ZS4-4 System Drive - The 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” on page 384](#).

A filler panel must be installed in empty drive slots. The system drive LEDs are shown in the following figure.

FIGURE 21 System Drive LEDs



For information about supported disks and capacities, refer to the [Oracle Systems Handbook](#).

TABLE 20 System Drive LEDs

Figure Legend		
1 Blue (Ready to remove)	2 Service action required (amber)	3 OK/Activity (green)

ZS4-4 CPU and Memory - The 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.

FIGURE 22 ZS4-4 CPU and Memory




Figure	Legend		
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 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 a ZS4-4 DIMM” on page 150](#).

ZS4-4 Cooling Subsystem - The 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.

FIGURE 23 ZS4-4 Cooling Subsystem




Figure	Legend		
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)

Figure	Legend		
5	Temperature sensor TS_ZONE1 (U4507)	11	Temperature sensor TS_ZONE0_A (U4508)

ZS4-4 Attached Storage

The 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” on page 384](#).

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” on page 384](#).

Related Topics

- [“Connecting to Attached Storage” on page 365](#)
- [“Disk Shelf Hardware Overview” on page 365](#)
- [Oracle ZFS Storage Appliance Cabling Guide](#)

Servicing the ZS3-4 Controller

Be sure to read sections [“Getting Started with Servicing the Hardware” on page 17](#) and [“Prerequisites for Servicing Hardware” on page 51](#).

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see [“Repairing Active Problems” on page 40](#).

Oracle 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 controllers for disk shelf interoperability.

To service ZS3-4 hardware components, use the following tasks.

- [“Replacing a ZS3-4 HDD or SSD” on page 179](#)
- [“Replacing a ZS3-4 Fan Module” on page 181](#)
- [“Replacing a ZS3-4 Power Supply” on page 183](#)
- [“Replacing a ZS3-4 DIMM” on page 185](#)
- [“Replacing a ZS3-4 PCIe Card or Riser” on page 194](#)
- [“Replacing a ZS3-4 Battery” on page 198](#)

To understand the ZS3-4, see the following topics:

- [“ZS3-4 Base Configuration” on page 201](#)
- [“ZS3-4 Specifications” on page 201](#)
- [“ZS3-4 Front and Rear Panel Components” on page 202](#)
- [“ZS3-4 Internal Components” on page 206](#)
- [“ZS3-4 PCIe I/O Cards” on page 212](#)
- [“ZS3-4 Attached Storage” on page 216](#)


▼ Replacing a 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 a ZS3-4 HDD or SSD. For information on disk compatibility, see the [Oracle Systems Handbook](#).


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, navigate to Maintenance > System in the BUI, or `maintenance system updates` in the CLI.

1. Identify the failed HDD or SSD by going to the Maintenance > Hardware section of the BUI and clicking the drive information icon . If you are physically at the system, the amber Service Required indicator on the HDD or SSD should be illuminated.
2. If you are not physically at the system, turn on the locator indicator by clicking the locator icon .
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, navigate to the Hardware > Maintenance screen, 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. The Oracle ZFS Storage system software automatically detects and configures the new drive. The device appears in the BUI Maintenance > Hardware screen when you view details for the controller or drive shelf.
9. If a fault does not self-clear, see [“Repairing Active Problems”](#) on page 40.

▼ Replacing a 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 a ZS3-4 fan module.

1. To flash the locator LED on the chassis to service, go to Maintenance > Hardware in the BUI 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, go to Maintenance > Hardware > 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 Storage Controller from the Rack” on page 53](#).**
4. **Lift the latch at the top of the fan module to unlock the fan module, and pull the fan module out.**




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. **Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).**
11. **Push the release tabs on the side of each rail and slowly slide the storage controller into the rack.**

▼ Replacing a ZS3-4 Power Supply


Storage 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 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. The following figure shows the 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. Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).

▼ Replacing a 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, go to the Maintenance > Hardware screen of the BUI, 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” on page 52.](#)**
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 Storage Controller from the Rack” on page 53.](#)**
4. **Remove the top cover as described in [“Removing the Top Cover” on page 54.](#)**
5. **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 ZS3-4 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.



The following image shows the chassis.



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](#).**
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. Go to the Maintenance > Hardware screen of the BUI, 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” on page 40](#).

▼ Replacing a 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.

Go to the Maintenance > Hardware screen of the BUI and click the right-arrow icon  on the controller, and then click Slots 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” on page 52](#).
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 Storage Controller from the Rack” on page 53](#).
4. Remove the top cover as described in [“Removing the Top Cover” on page 54](#).
5. Locate the PCIe card position in the storage controller, see [“ZS3-4 PCIe I/O Cards” on page 212](#).
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](#).
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. Go to the Maintenance > Hardware screen of the BUI, 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” on page 40](#).
16. Install the disk shelf as described in [“Installation Prerequisites and Hardware Overviews” in Oracle ZFS Storage Appliance Installation Guide](#) and connect the expansion storage as described in [“Connecting to Attached Storage” on page 365](#).

▼ Replacing a ZS3-4 Battery

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.


1. Shut down the controller using one of the power-off methods described in [“Powering Off the Controller” on page 52](#).
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 Storage Controller from the Rack” on page 53](#).
4. Remove the top cover as described in [“Removing the Top Cover” on page 54](#).

5. Using a small, non-metallic screwdriver, press the latch and remove the battery from the motherboard. The battery is shown here.



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 OS8.7.0*.**

ZS3-4 Controller Hardware Overview

To understand the ZS3-4, see the following topics:

- [“ZS3-4 Base Configuration” on page 201](#)
- [“ZS3-4 Specifications” on page 201](#)
- [“ZS3-4 Front and Rear Panel Components” on page 202](#)
- [“ZS3-4 Internal Components” on page 206](#)
- [“ZS3-4 PCIe I/O Cards” on page 212](#)
- [“ZS3-4 Attached Storage” on page 216](#)

ZS3-4 Base Configuration

The 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 21 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

Refer to the product site at <https://www.oracle.com/storage/nas/index.html> for the most recent component specification.

ZS3-4 Specifications

ZS3-4 Physical Specifications - The 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 22 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

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.

ZS3-4 Front and Rear Panel Components

ZS3-4 Front Panel Components - The ZS3-4 controller drive slots and front panel components are shown in the following figure.

FIGURE 24 ZS3-4 Controller Front Panel




Figure Legend	Figure Legend
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)

ZS3-4 Rear Panel Components - The following graphic shows the rear panel. Base configuration HBAs are not depicted in this illustration.

FIGURE 25 ZS3-4 Controller Rear Panel




Figure Legend	Figure Legend
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 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](#).

Ethernet Ports

There are four RJ-45 Gigabit Ethernet connectors (NET0, NET1, NET2, NET3) located on the 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 and provides a LAN interface to the SP console.

FIGURE 26 ZS3-4 Controller Network Management Port



Serial Management Port

The serial management connector (SER MGT) is an RJ-45 connector and provides a terminal connection to the SP console.

FIGURE 27 ZS3-4 Controller Serial Management Port



ZS3-4 Internal Components

The components of the ZS3-4 controller are shown in the following figure and identified in the legend.

FIGURE 28 ZS3-4 Controller Components




Figure Legend	Figure 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	

ZS3-4 Internal Boards - The 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.

ZS3-4 System Drives - The 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 DE2-24 disk shelf slots.

Note - Read cache devices must be installed in either the controller or a 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” on page 384](#).

FIGURE 29 ZS3-4 Controller System Drive




Figure Legend

1 Ready to remove (blue)

2 Service action required (amber)

3 OK/Activity (green)


ZS3-4 4x4 SAS-2 HBA - The 4x4 SAS-2 HBA, installed in the ZS3-4, provides connectivity to external DE2 and Sun Disk Shelves. The HBA ports are numbered 3-0, top to bottom:

FIGURE 30 ZS3-4 Controller 4x4 SAS-2 HBA Port Numbers

See [“ZS3-4 PCIe I/O Cards” on page 212](#) for slot placement.


ZS3-4 CPU and Memory - The ZS3-4 controller supports 16GB DDR3 DIMMs installed on all eight risers, accommodating 1TB of memory.

FIGURE 31 ZS3-4 Controller DIMM 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*.


FIGURE 32 ZS3-4 Controller DIMM Placement



DIMM names in appliance logs and the Maintenance > Hardware view are displayed with the full name, such as /SYS/MB/P0/D7.

ZS3-4 Fan Modules - The Fan Modules and Fan Module LEDs of the controller are shown in the following figure. The following LEDs are lit when a fan module fault is detected:

FIGURE 33 ZS3-4 Controller Fan Modules



- Front and rear Service Action Required LEDs
- Fan Module Service Action Required (TOP) LED on the front of the server
- 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.

ZS3-4 PCIe I/O Cards

ZS3-4 PCIe Options - This table describes the PCIe base configuration and optional slot assignments for ZS3-4 standalone and cluster configurations.

TABLE 23 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 HBA

Slot	Mfg Part Number	Description	Max	Note
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)
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

Slot	Mfg Part Number	Description	Max	Note
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
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

Slot	Mfg Part Number	Description	Max	Note
9	7070195	4-port 1Gb Ethernet Copper UTP NIC	4	Optional recommended front-end
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)

ZS3-4 NIC/HBA Options - This table describes NIC/HBA PCIe card options for ZS3-4 standalone and cluster configurations.


TABLE 24 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

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.

FIGURE 34 ZS3-4 Controller PCIe Cards Slot Order



Install optional PCIe cards in the following order:

TABLE 25 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)

ZS3-4 Attached Storage

The 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 (DE2 family and legacy 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 controller or a DE2-24 disk shelf. Read cache devices cannot be installed in both controller and DE2 disk shelf slots at the same time. For software requirements and slot configuration rules, see [“Disk Shelf Configurations” on page 384](#).

Related Topics

- [“Connecting to Attached Storage” on page 365](#)
- [“Disk Shelf Hardware Overview” on page 365](#)

Servicing the ZS3-2 Controller

Be sure to read sections [“Getting Started with Servicing the Hardware” on page 17](#) and [“Prerequisites for Servicing Hardware” on page 51](#).

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see [“Repairing Active Problems” on page 40](#).

Oracle 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 controllers for disk shelf interoperability.

To service ZS3-2 hardware components, use the following tasks.

- [“Installing Optional PCIe Slot Cards in Order” on page 220](#)
- [“Disconnecting the RJ-45 Cable from the Cluster Serial Port” on page 221](#)
- [“Replacing a ZS3-2 HDD or SSD” on page 222](#)
- [“Replacing a ZS3-2 Power Supply” on page 225](#)
- [“Replacing a ZS3-2 Fan Module” on page 228](#)
- [“Replacing a ZS3-2 DIMM” on page 231](#)
- [“Replacing a ZS3-2 PCIe Card or Riser” on page 238](#)
- [“Replacing a ZS3-2 Flash Drive” on page 244](#)
- [“Replacing a ZS3-2 Battery” on page 246](#)

To understand the ZS3-2 hardware, use the following topics.

- [“ZS3-2 Base Configuration” on page 248](#)
- [“ZS3-2 Front and Rear Panel Components” on page 249](#)
- [“ZS3-2 Specifications” on page 253](#)
- [“ZS3-2 Internal Components” on page 254](#)
- [“ZS3-2 PCIe I/O Cards” on page 261](#)
- [“ZS3-2 Optional Cable Management Arm” on page 262](#)
- [“ZS3-2 Attached Storage” on page 263](#)

▼ Installing Optional PCIe Slot Cards in Order

Install 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 below.




▼ Replacing a 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 a ZS3-2 HDD or SSD. For information on disk compatibility, see the [Oracle Systems Handbook](#).


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, navigate to Maintenance > System in the BUI, or `maintenance system updates` in the CLI.

1. Identify the failed drive by going to the Maintenance > Hardware section of the BUI and clicking the drive information icon . If you are physically at the system, the amber Service Required indicator on the HDD or SSD should be illuminated.


2. If you are not physically at the system, turn on the locator indicator by clicking the locator icon .
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 it 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, navigate to the Hardware > Maintenance screen, 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. Go to the Maintenance > Hardware screen of the BUI. 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”](#) on page 40.

▼ Replacing a ZS3-2 Power Supply

Storage 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 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](#).

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. **Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).**

▼ Replacing a 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 a ZS3-2 fan module.

1. **To locate the chassis you want to service, click the associated locator icon  on the Maintenance > Hardware screen of the BUI or issue the `set /SYS/LOCATE status=on` command at the service processor (SP) prompt. The locate LED will flash on the controller chassis.**
2. **Identify the faulted fan module by locating the corresponding Service Required status indicator or by clicking the locator icon  on the Maintenance > Hardware screen of the BUI for the fan you want to replace.**
3. **Power off the storage controller as described in [“Powering Off the Controller” on page 52](#).**
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 Storage Controller from the Rack” on page 53.](#)**
6. **Remove the top cover as described in [“Removing the Top Cover” on page 54.](#)**
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](#).
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. Go to the Maintenance > Hardware screen of the BUI. Click the right-arrow icon  for the controller and then click Fan to verify that the online icon  is green for the newly fan module.
18. If a fault does not self-clear, see [“Repairing Active Problems” on page 40](#).

▼ Replacing a 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, go to the Maintenance > Hardware screen of the BUI, and click on the right-arrow icon  on the controller. Then click DIMMs 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” on page 52](#).

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 Storage Controller from the Rack” on page 53.](#)
5. Remove the top cover as described in [“Removing the Top Cover” on page 54.](#)
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](#).
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. Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).

▼ Replacing a 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, go to the Maintenance > Hardware screen of the BUI and click the right-arrow icon  on the controller, and then click Slots.
2. Power off the storage controller as described in [“Powering Off the Controller” on page 52](#).
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 Storage Controller from the Rack” on page 53](#).
5. Remove the top cover as described in [“Removing the Top Cover” on page 54](#).
6. Locate the PCIe card position in the storage controller, see [“ZS3-2 PCIe I/O Cards” on page 261](#).


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](#).
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. Go to the Maintenance > Hardware screen of the BUI, 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.
25. If a fault does not self-clear, see [“Repairing Active Problems” on page 40](#).


▼ Replacing a 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 a ZS3-2 flash drive.

1. Power off the storage controller as described in [“Powering Off the Controller” on page 52](#).
2. Extend the controller from the rack as described in [“Extending the Storage Controller from the Rack” on page 53](#).
3. Remove the top cover as described in [“Removing the Top Cover” on page 54](#).
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](#).
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 a 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 a ZS3-2 battery.

1. Power off the storage controller as described in [“Powering Off the Controller” on page 52.](#)
2. Extend the controller from the rack as described in [“Extending the Storage Controller from the Rack” on page 53.](#)
3. Remove the top cover as described in [“Removing the Top Cover” on page 54.](#)
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 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](#).
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 OS8.7.0*.

ZS3-2 Controller Hardware Overview

To understand the ZS3-2 hardware, use the following topics.

- [“ZS3-2 Base Configuration”](#) on page 248
- [“ZS3-2 Front and Rear Panel Components”](#) on page 249
- [“ZS3-2 Specifications”](#) on page 253
- [“ZS3-2 Internal Components”](#) on page 254
- [“ZS3-2 PCIe I/O Cards”](#) on page 261
- [“ZS3-2 Optional Cable Management Arm”](#) on page 262
- [“ZS3-2 Attached Storage”](#) on page 263

ZS3-2 Base Configuration

The 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 26 ZS3-2 Controller Features

CPU	Memory	Read Cache	Boot Drive	PCIe Slots	HBA
2x8-core, 2.1 GHz	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


Refer to the product site at <https://www.oracle.com/storage/nas/index.html> for the most recent component specification.

For information about supported disks and capacities, refer to the [Oracle Systems Handbook](#).

ZS3-2 Front and Rear Panel Components

ZS3-2 Front Panel Components - The ZS3-2 controller drive slots and front panel components are shown in the following figure.

FIGURE 35 ZS3-2 Controller Front Panel



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.

Figure Legend

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.

ZS3-2 Rear Panel Components - The ZS3-2 controller PCIe slots and rear panel components are shown in the following figure.

FIGURE 36 ZS3-2 Controller Rear Panel




Figure Legend

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 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 ZS3-2 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 Mbts/sec, 1000 Mbts/sec, or 10-Gbts/sec. Use these ports to connect the appliance to the network.

TABLE 27 ZS3-2 Ethernet Port LEDs Description

Connection Type	EEE Terminology	Speed LED Color	Transfer Rate
Fast Ethernet	100BASE-TX	Off	100 Mbts/sec
Gigabit Ethernet	1000BASE-T	Amber	1000 Mbts/sec
10 Gigabit Ethernet	10GBASE-T	Green	10000 Mbts/sec

The LEDs located above each NET port are Link/Activity (left) and Speed (right) indicators for each port.

FIGURE 37 ZS3-2 Controller Ethernet Ports



Network Management Port

The 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.

FIGURE 38 ZS3-2 Controller Network Management Port



Serial Management Port

The serial management connector (SER MGT) is an RJ-45 port and provides a terminal connection to the service processor (SP) console.

FIGURE 39 ZS3-2 Controller Serial Management Port




Cluster I/O Ports

Two cluster serial ports (0 and 1) and one Ethernet port provide communication between two controllers to form a cluster configuration. For information about how to connect cables to form a cluster, see [“ZS3-2 Controller Cluster I/O Ports” in Oracle ZFS Storage Appliance Cabling Guide](#).



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. To properly remove an RJ-45 cable from a cluster serial port, do the following:


FIGURE 40 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 ZS3-2, provides connectivity to external disk shelves. The HBA ports are numbered 0-3, as shown in the following figures.

FIGURE 41 ZS3-2 Controller 4x4 SAS HBA Port Numbers



For controllers with an additional 4x4 SAS HBA, see [“ZS3-2 PCIe I/O Cards” on page 261](#) for slot placement.

ZS3-2 Specifications

ZS3-2 Physical Specifications - The ZS3-2 controller 2U chassis dimensions are as follows:

TABLE 28 ZS3-2 Controller Physical Specifications

Dimension	Measurement	Dimension	Measurement
Height	3.44 in./87.4 mm	Depth	20.25 in./514 mm

Dimension	Measurement	Dimension	Measurement
Width	17.52 in./445 mm	Weight	41.23 lb/18.70 kg

ZS3-2 Electrical Specifications - The 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.

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).

ZS3-2 Internal Components

The 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.

ZS3-2 Motherboard, Memory, and PCIe Cards - The 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.

FIGURE 42 ZS3-2 Controller Motherboard, Memory, and PCIe Components




Figure Legend

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)

The ZS3-2 controller replaceable memory components and part numbers are listed below.

TABLE 30 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

ZS3-2 Storage, Power, and Fan Components - The 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.

FIGURE 44 ZS3-2 Controller Storage, Power, and Fan Components




Figure Legend

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)

Figure Legend

4 LED board (FRU)

8 PDB signal cable (FRU)

12 Power supplies

ZS3-2 Internal Cables - The 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.

FIGURE 45 ZS3-2 Controller Internal Cables




Figure Legend

1 PDB signal cable

3 Drive signal cable

5 Drive power cable

2 USB board cable

4 LED board cable

ZS3-2 PCIe I/O Cards

ZS3-2 PCIe I/O Options - The following table describes the supported PCIe configuration options for the 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 31 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
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

Slot	Mfg Part Number	Description	Max	Note
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

ZS3-2 NIC/HBA Options - The following table describes the NIC/HBA options for the ZS3-2 controller.

TABLE 32 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
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

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

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.

FIGURE 46 ZS3-2 Controller Optional Cable Management Arm**Figure Legend**

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	

ZS3-2 Attached Storage

The 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” on page 365](#). See also [“Disk Shelf Hardware Overview” on page 365](#) for component specifications and diagrams.

Servicing the 7x20 Controller

Be sure to read sections [“Getting Started with Servicing the Hardware” on page 17](#) and [“Prerequisites for Servicing Hardware” on page 51](#).

Most hardware faults clear after replacing the affected component. If a fault does not self-clear, see [“Repairing Active Problems” on page 40](#).

Oracle 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 controllers for disk shelf interoperability.

To service 7x20 hardware components, use the following tasks.

- [“Replacing a 7x20 HDD or SSD” on page 265](#)
- [“Replacing a 7x20 Fan Module” on page 267](#)
- [“Replacing a 7x20 Power Supply” on page 272](#)
- [“Replacing a 7x20 DIMM” on page 276](#)
- [“Replacing a 7x20 Controller PCIe Card or Riser” on page 289](#)
- [“Replacing a 7x20 Battery” on page 300](#)

To understand the 7x20 hardware, see the following topics:


- [“7420 Controller Hardware Overview” on page 306](#)
- [“7320 Controller Hardware Overview” on page 321](#)
- [“7120 Controller Hardware Overview” on page 333](#)

▼ Replacing a 7x20 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 a 7x20 HDD or SSD. For information on disk compatibility, see the [Oracle Systems Handbook](#).


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, navigate to Maintenance > System in the BUI, or `maintenance system updates` in the CLI.

1. Identify the failed HDD or SSD by going to the Maintenance > Hardware section of the BUI and clicking the drive information icon . If you are physically at the system, the amber Service Required indicator on the HDD or SSD should be illuminated.
2. If you are not physically at the system, turn on the locator indicator by clicking the locator icon .
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, navigate to the Hardware > Maintenance screen, 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. Go to the Maintenance > Hardware screen of the BUI. 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, see [“Repairing Active Problems” on page 40](#).


▼ Replacing a 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.

7120 or 7320: The fan modules and status indicators are hidden under a fan door in the 7120 and 7320 storage controllers. Components may differ slightly between the 7120 and 7320, however service procedures for each are identical. The illustration shows the 7320 controller.



Caution - The storage controller might overheat and shut down if the door is left open for more than 60 seconds.

7420: The following illustration shows the fan modules in the 7420 storage controller.




1. To flash the locator LED on the chassis to service, go to **Maintenance > Hardware** in the BUI 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, go to **Maintenance > Hardware > 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 Storage Controller from the Rack” on page 53](#).
4. **7120 or 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. **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 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. **Go to the Maintenance > Hardware screen of the BUI. 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”](#) on page 40.**

▼ Replacing a 7x20 Power Supply

Storage 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 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” on page 52.](#)**

7120 or 7320: Release the latch, then remove the power supply. Components may differ slightly between the 7120 and 7320, however service procedures for each are identical. The illustration shows the 7320 controller.




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 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. Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).

▼ Replacing a 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, go to the Maintenance > Hardware screen of the BUI, and click on the right-arrow icon  on the controller. Then click DIMMs to locate the faulted component, indicated by the warning icon .


1. **Power off the storage controller as described in “Powering Off the Controller” on page 52.**
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 Storage Controller from the Rack” on page 53.**
4. **7120 or 7320: Components may differ slightly between the 7120 and 7320, however service procedures for each are identical. The illustration shows the 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 server.




5. 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 7120 controller.




The following illustration shows the Fault remind button on the 7420 controller.




7. **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. 7120 or 7320: Components may differ slightly between the 7120 and 7320, however service procedures for each are identical. The illustration shows the 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. **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 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.**



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 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](#).**
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. Go to the Maintenance > Hardware screen of the BUI. 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” on page 40](#).

▼ Replacing a 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 7120 Sun Flash Accelerator F20 card is a FRU and *must* be replaced by an Oracle service representative.

Go to the Maintenance > Hardware screen of the BUI and click the right-arrow icon  on the controller, and then click Slots 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” on page 52](#).**
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 Storage Controller from the Rack” on page 53](#).**
4. **7120 or 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. **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 [“7320 Configuration Options”](#) on page 321 for the 7320, the [“7120 PCIe Options”](#) on page 343 for the 7120, or [“7420 PCIe I/O Cards”](#) on page 316 for the 7420.**
 7. **7120 or 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. 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. **7120 or 7320:** Components may differ slightly between the 7120 and 7320, however service procedures for each are identical. The illustration shows the 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.

- 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](#).
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. Go to the Maintenance > Hardware screen of the BUI, 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” on page 40](#).
19. Install the disk shelf as described in [“Installing Disk Shelves” in Oracle ZFS Storage Appliance Installation Guide](#) and connect the expansion storage as described in [“Connecting to Attached Storage” on page 365](#).

▼ Replacing a 7x20 Battery

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 a 7x20 battery.


1. Power off the storage controller as described in [“Powering Off the Controller” on page 52](#).
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 Storage Controller from the Rack” on page 53](#).
4. 7120 or 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. **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 7420 battery is shown here.




The following figure shows the 7120 battery.



7. **Press the replacement battery into the motherboard with the positive side (+) facing upward.**
8. **7120 or 7320: Components may differ slightly between the 7120 and 7320, however, service procedures for each are identical. The illustration shows the 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.**

- 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](#).
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 OS8.7.0*.

7420 Controller Hardware Overview

This section describes the internal and external components of the Oracle ZFS Storage 7420 controller. See the following for details:

- [“7420 Configuration Options”](#) on page 306
- [“7420 Front and Back Panel Components”](#) on page 307
- [“7420 Specifications”](#) on page 310
- [“7420 Internal Components”](#) on page 311
- [“7420 PCIe I/O Cards”](#) on page 316
- [“7420 Attached Storage”](#) on page 321

7420 Configuration Options

The Sun ZFS Storage 7420 Appliance consists of a standalone storage controller, or two storage controllers in a high-availability cluster configuration, and one to 36 disk shelves. The following table describes the 7420 configuration options:

Note - 7420 M2 controllers cannot be clustered with 7420 controllers.

TABLE 33 7420 Controller Configuration Options

Product ID	Mktg Part Number	CPU	DIMMs	Read Cache	Boot Drive	Mfg Part Number
7420 M2	7107089	4x8-core, 2.0 GHz	8GB, 16GB	1.6TB SAS	900GB SAS	7075466
7420 M2	7107090	4x10-core, 2.40GHz	8GB, 16GB	1.6TB SAS	900GB SAS	7075470

Product ID	Mktg Part Number	CPU	DIMMs	Read Cache	Boot Drive	Mfg Part Number
7420	7100566	4x8-core, 2.0 GHz	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 [Figure 52, “7420 Controller System Drive,”](#) on page 313.

Refer to the product site at <https://www.oracle.com/storage/nas/index.html> for the most recent component specification.

7420 Front and Back Panel Components

7420 Front Panel Components - The 7420 controller drive slots and front panel components are shown in the following figure.

FIGURE 47 7420 Controller Front Panel




Figure Legend	Figure Legend	Figure Legend
1 Locator LED and button (white)	7 Power Supply (PS) Service Required LED	13 SSD 2 (optional)
2 Service Required LED (amber)	8 Over Temperature Warning LED	14 SSD 3 (optional)
3 Power/OK LED (green)	9 USB 2.0 Connectors	15 SSD 4 (optional)
4 Power button	10 DB-15 video connector	16 SSD 5 (optional)
5 Service Processor (SP) OK LED (green)	11 Boot drive 0 (mirrored)	
6 Fan/CPU/Memory Service Required LED	12 Boot drive 1 (mirrored)	

7420 Back Panel Components - The following graphic shows the rear panel of the controller. Base configuration HBAs are not depicted in this illustration.

FIGURE 48 7420 Controller Rear Panel




Figure Legend	Figure Legend
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

Figure Legend	Figure Legend
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 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](#).

Ethernet Ports

There are four RJ-45 Gigabit Ethernet connectors (NET0, NET1, NET2, NET3) located on the 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 and provides a LAN interface to the SP console.

FIGURE 49 7420 Controller Network Management Port



Serial Management Port

The serial management connector (SER MGT) is an RJ-45 connector and provides a terminal connection to the SP console.

FIGURE 50 7420 Controller Serial Management Port

7420 Specifications

7420 Physical Specifications - The 3U chassis form factor dimensions are as follows:

TABLE 34 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

7420 Electrical Specifications - The following list shows the electrical specifications for 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.

7420 Internal Components

The controller components are shown in the following figure and identified in the table.

FIGURE 51 7420 Controller Components




Figure Legend	Figure 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	


7420 Internal Boards - The 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.

7420 System Drives - The 7420 M2 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 7420 M2 clustered configuration, the number of SSDs installed in each controller can vary.

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 7420 clustered configuration, the number of SSDs installed in both controllers must match.

For information about supported disks and capacities, refer to the [Oracle Systems Handbook](#).

FIGURE 52 7420 Controller System Drive**Figure Legend**

1 Ready to remove (blue)

2 Service action required (amber)

3 OK/Activity (green)


The 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.


7420 CPU and Memory - Previous 7420 controller configurations support two or four (1.86 GHz 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.

FIGURE 53 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*.

FIGURE 54 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.

7420 Fan Modules - The Fan Modules and Fan Module LEDs of the storage controller are shown in the following figure. The following LEDs are lit when a fan module fault is detected:

FIGURE 55 7420 Controller Fan Module LEDs

- Front and rear Service Action Required LEDs
- Fan Module Service Action Required (TOP) LED on the front of the server
- 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.

7420 PCIe I/O Cards

7420 PCIe I/O Options - The 7420 controller has ten PCIe slots. This table describes base and optional PCIe card slots for standalone and cluster configurations.

Note - The 7420 M2 reserves slot 0 for a SAS-2 internal HBA, and slots 1 and 8 for two SAS-2 external HBAs.

TABLE 35 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

Slot	Mfg Part Number	Description	Max	Note
				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
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

Slot	Mfg Part Number	Description	Max	Note
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)
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)

Slot	Mfg Part Number	Description	Max	Note
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
9	7054739	4-port copper NIC	6	Optional recommended front-end

7420 NIC/HBA Options - This table describes NIC/HBA PCIe card options for standalone and cluster configurations.

TABLE 36 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

Mktg Part Number	Description	Mfg Part Number
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

7420 PCIe Slot Order - When installing additional PCIe cards to your base 7420 configuration, cards must be added in a specific order.

Note - The 7420 M2 reserves slot 0 for a SAS-2 internal HBA, and slots 1 and 8 for two SAS-2 external HBAs (see [“7420 PCIe I/O Cards”](#) on page 316).

FIGURE 56 7420 Controller PCIe Cards and Slot Order



Install optional PCIe cards in the following order:

TABLE 37 7420 Controller Optional PCIe Cards and Slot Order

PCIe Card	Slot Order
1. SAS-2 external HBA	Slot 2, 7, 3, 6

PCIe Card	Slot Order
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)

7420 Attached Storage

Primary storage for 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” on page 365](#). See also [“Disk Shelf Hardware Overview” on page 365](#) for component specifications and diagrams.

7320 Controller Hardware Overview

This section describes the internal and external components of the Oracle ZFS Storage 7320 controller. See the following for details:

- [“7320 Configuration Options” on page 321](#)
- [“7320 Front and Back Panel Components” on page 324](#)
- [“7320 Electrical Specifications” on page 328](#)
- [“7320 Internal Components” on page 328](#)
- [“7320 Attached Storage” on page 333](#)

7320 Configuration Options

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 supports one to six disk shelves.

The 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 7320 controller included 24GB (base), 48GB, or 72GB memory options. The clustered configuration simply uses two servers and a cluster card in each server for a heartbeat connection between the servers.

All user-accessible storage is provided by one to six disk shelves that are external to the server (s). RAID functions are managed by software. Solid state SAS-1 drives (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 four of the 24 drives in a disk shelf. The remaining 20 drives are available for storage.

Refer to the product site at <https://www.oracle.com/storage/nas/index.html> for the most recent component specification.

The 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. For information about supported disks and capacities, refer to the [Oracle Systems Handbook](#).

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 7320.

TABLE 38 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 39 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

Slot	Type	Part Number	Vendor Part	Description	Note
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

7320 Cluster Configurations - The 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 7320 or two 7420 controllers to a cluster.

The following options are available for clustered storage controllers.


Note - When you cluster a 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 40 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
0	PCIe	375-3606-03	MHJH29-XTC	InfiniBand HCA	Optional Recommended Front-end
0	PCIe	375-3481-01	Intel EXP19404PT	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

7320 Front and Back Panel Components

7320 Front Panel Components - The following figure and legend identify the front panel LEDs.

FIGURE 57 7320 Controller Front Panel LEDs**Figure Legend**

- | | |
|---|------------------------------|
| 1 | Locate Button/LED |
| 2 | Service Required LED (amber) |
| 3 | Power/OK LED (green) |
| 4 | Power Button |
| 5 | Rear Power Supply |
| 6 | System Overtemperature LED |
| 7 | Top Fan |

The following figure and legend identify the 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.

FIGURE 58 7320 Controller Front Panel Drive Locations



TABLE 41 7320 Controller Front Panel Drive Locations

Disk Drive Locations				
HDD1	HDD3	HDD5		
HDD0	HDD2	HDD4	HDD6	HDD7

7320 Back Panel Components - Following is an illustration of the 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.

FIGURE 59 7320 Controller Rear Panel




Figure Legend

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 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](#).

Ethernet Ports

There are four RJ-45 Gigabit Ethernet connectors (NET0, NET1, NET2, NET3) located on the 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 motherboard and provides an alternate terminal interface to the SP console.

FIGURE 60 7320 Controller Network Management Port



Serial Management Port

The serial management connector (SERIAL MGT) is an RJ-45 connector and is a terminal connection to the SP console.

FIGURE 61 7320 Controller Serial Management Port



7320 Electrical Specifications

The following list shows the electrical specifications for the 7320.

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.

7320 Internal Components

The 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 7320 storage controller.

TABLE 42 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

7320 Replaceable Components - The following list contains all of the replaceable power distribution, disk, and fan module components of the 7320. Note that power supplies, disks, and fan modules are hot-pluggable on the storage controller.

TABLE 43 7320 Controller Replaceable Components

Part Number	Description	FRU/CRU
F300-2233-02	RoHS 760W Power Supply	CRU

Part Number	Description	FRU/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 the 7320.

TABLE 44 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

7320 CPU and Memory Components - The 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.

FIGURE 62 7320 Controller CPU and Memory Components



7320 Power Distribution and Fan Module Components - The following figure and legend identify the Power Distribution/Fan Module components of the storage controller.

FIGURE 63 7320 Controller Power Distribution and Fan Module Components**Figure Legend**

1 Fan Modules

2 Fan Board

3 Paddle Board

Figure Legend

4 Power Distribution/Bus Bar Assembly

5 Power Supplies

7320 I/O Components - The following figure and legend identify the I/O components of the storage controller.

FIGURE 64 7320 Controller I/O Components




Figure Legend	Figure Legend
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

7320 PCIe Cards and Risers - Following is the complete list of replaceable PCIe cards for the 7320 system.

TABLE 45 7320 Controller Replaceable PCIe Cards

Part Number	Description	FRU/CRU
F371-4325-01	8Gb FC HBA (PCIe)	CRU

Part Number	Description	FRU/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

7320 Cables - The following list contains the replaceable cables for the 7320 storage controller.

TABLE 46 7320 Controller Replaceable Cables

Part Number	Description	FRU/CRU
F530-4228-01	Cable, Mini SAS	FRU (internal)
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)

7320 Attached Storage

The 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” on page 365](#). See also [“Disk Shelf Hardware Overview” on page 365](#) for component specifications and diagrams.

7120 Controller Hardware Overview

This section describes the internal and external components of the Oracle ZFS Storage 7120 controller. See the following for details:

- [“7120 Chassis Overview” on page 335](#)
- [“7120 Base Configuration” on page 334](#)
- [“7120 Front Panel” on page 335](#)
- [“7120 Rear Panel” on page 336](#)
- [“7120 Physical Specifications” on page 337](#)
- [“7120 Electrical Specifications” on page 337](#)
- [“7120 Internal Boards” on page 337](#)
- [“7120 CPU and Memory” on page 338](#)
- [“7120 Power Distribution, Fan Module and Disk Components” on page 339](#)
- [“7120 I/O Components” on page 341](#)
- [“7120 NIC/HBA Options” on page 342](#)
- [“7120 PCIe Options” on page 343](#)
- [“7120 Connectors” on page 344](#)
- [“7120 Cables” on page 345](#)
- [“7120 Attached Storage” on page 346](#)

7120 Base Configuration

The following table shows the configuration options for a 7120 controller. All PCIe cards are low-profile, and must be fitted with low-profile mounting brackets.

This table describes base configurations for the 7120 with Aura Logzilla.

TABLE 47 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 7120 with Logzilla 3.5" SSD.

TABLE 48 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

7120 Chassis Overview

The 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 product site at <https://www.oracle.com/storage/nas/index.html> for the most recent component specification.

Refer to the white paper *Implementing Fibre Channel SAN Boot with Oracle's Sun ZFS Storage Appliance* at <http://www.oracle.com/technetwork/articles/servers-storage-admin/fbsanboot-365291.html> for details on FC SAN boot solutions using the Sun ZFS Storage 7120.

The 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.

7120 Front Panel

The following figure and legend show the 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.

FIGURE 65 7120 Controller Front Panel




Figure Legend

1 Locator LED/button (white)

7 HDD 2

13 HDD 8

Figure Legend		
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	

7120 Rear Panel

The following figure and legend show the 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.

FIGURE 66 7120 Controller Rear Panel




Figure Legend		
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	

7120 Physical Specifications

The 2U chassis form factor dimensions are as follows:

TABLE 49 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

7120 Electrical Specifications

The following list shows the electrical specifications for the 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.

7120 Internal Boards

The 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.

7120 CPU and Memory

The 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 7120 system.

TABLE 50 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.


FIGURE 67 7120 Controller CPU and Memory Components




7120 Power Distribution, Fan Module and Disk Components

The fan modules and LEDs are shown in the following illustration.

FIGURE 68 7120 Controller Fan Modules and LEDs



The following figure and legend show the power distribution and associated components.

FIGURE 69 7120 Controller Power Distribution Board and Associated Components**Figure Legend**

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	

7120 I/O Components

The following figure and legend show the I/O components of the 7120 system.

FIGURE 70 7120 Controller I/O Components




Figure Legend

1 Top Cover	3 Hard Disk Drives
2 Right Control Panel Light Pipe Assembly	4 Left Control Panel Light Pipe Assembly

7120 NIC/HBA Options

The following table describes NIC/HBA options for the 7120.

TABLE 51 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

Mktg Part Number	Description	Mfg Part Number
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

7120 PCIe Options

The following table describes the supported PCIe configuration option summary for the 7120.

TABLE 52 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
1	PCIe	371-4325-01	QLogic	8Gb DP FC HBA	Optional FC Target or Initiator (Backup)

Slot	Type	Sun Part Number	Vendor Part Number	Description	Note
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

7120 Connectors

The serial management connector (SER MGT) is an RJ-45 port and provides a terminal connection to the SP console.

FIGURE 71 7120 Controller Serial Management Port



The network management connector (NET MGT) is an RJ-45 port and provides an alternate terminal interface to the SP console.

FIGURE 72 7120 Controller Network Management Port




There are four RJ-45 Gigabit Ethernet ports (NET0, NET1, NET2, NET3) located on the motherboard that operate at 10/100/1000 Mbit/sec. These network interfaces must be configured before use.

7120 Cables


The following figure and legend show the storage controller internal cables.

Note - The rear boot drives are not depicted in this illustration.

FIGURE 73 7120 Controller Storage Drive Data Cable



Cable	Connection
1 Storage Drive Data Cable	Connection between the HBA PCI-Express Card and the storage drive backplane.

FIGURE 74 7120 Controller Ribbon Cable

Cable	Connection
2 Ribbon cable	Connection is between the power distribution board and the motherboard.

7120 Attached Storage

The 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 the 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” on page 365](#). See also [“Disk Shelf Hardware Overview” on page 365](#) for component specifications and diagrams.

Servicing the Disk Shelves

Before servicing a disk shelf, be sure to read sections [“Getting Started with Servicing the Hardware” on page 17](#) and [“Prerequisites for Servicing Hardware” on page 51](#).

To replace disk shelf components, use the following tasks.

- [“Replacing a Faulty Disk Shelf Drive” on page 348](#)
- [“Replacing a Disk Shelf Power Supply” on page 351](#)
- [“Replacing a Disk Shelf I/O Module” on page 357](#)
- [“Replacing a Disk Shelf SIM Board” on page 359](#)
- [“Adding a New Disk Shelf” on page 362](#)
- [“Adding a Disk Shelf Cache Device” on page 363](#)

To understand disk shelf components, see the following topics.

- [“Oracle Storage Drive Enclosure DE3-24” on page 365](#)
- [“Oracle Storage Drive Enclosure DE2-24” on page 373](#)
- [“Sun Disk Shelf” on page 380](#)
- [“Disk Shelf Configurations” on page 384](#)
- [“Maximum Disk Shelves per Controller Configuration” on page 387](#)
- [“SAS-2 and SAS-3” on page 388](#)

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” on page 365](#) 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. The replacement drive must be of the same capacity and type as the drive to be replaced. 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.


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

1. **Faulted drives are indicated by an amber LED. Go to the Maintenance > Hardware section of the BUI, click the right-arrow icon  at the beginning of the appropriate disk shelf row, and click the information icon  for the appropriate drive to view details, or click the locator icon  to turn on the locator LED. To add, not replace, a new cache device into an appropriate and empty drive slot, see [“Adding a Disk Shelf Cache Device” on page 363](#).**
2. **Locate the failed disk drive with an amber LED at the front of the chassis.**
3. **Press the release button or latch to release the drive lever.**

4. Pull the drive lever fully open to unlock and partially eject the drive from the chassis.




- OR -




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.



- OR -



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. Go to the Maintenance > Hardware section of the BUI, 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” on page 40](#).

▼ 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” on page 365](#) 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. Go to the Maintenance > Hardware section of the BUI, 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.
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.


Oracle Storage Drive Enclosure DE3-24P or DE3-24C:

Push the lever fully downward.




Oracle Storage Drive Enclosure DE2-24P or DE2-24C:

Grasp the latch and the opposite side of the module, and squeeze together to release the lever.



Sun Disk Shelf:

Using your thumb and forefinger, unscrew both ejection arm captive screws until loose and swing the ejection arms out until they are fully open.



6. **Pull the module out of the chassis, being careful not to damage the connector pins in the back.**
7. **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. If an Oracle Storage Drive Enclosure DE3-24, the module clicks in place and the lever engages; then go to step 9.**

Oracle Storage Drive Enclosure DE3-24P or DE3-24C:


Push on the handle until you hear or feel a click.



8. Close the lever/ejection arms.


Oracle Storage Drive Enclosure DE2-24P or DE2-24C:

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.



9. **Ensure the power supply on/off switch is in the "O" off position.**
10. **Plug the power cord into the new power supply and attach the power cord tie strap to the power cord.**
11. **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.**

Note - After the disk shelf powers on, which could take several minutes, ensure that the system power indicator on the front of the disk shelf is lit a steady green, and that each power supply with fan module status indicator on the rear is lit a steady green. If not, power off the disk shelf and reseat the faulted module. If the module status indicator is not a steady green after module reinsertion, contact Oracle Customer Service.

Oracle Storage Drive Enclosure DE3-24P or DE3-24C:



Oracle Storage Drive Enclosure DE2-24P or DE2-24C:



12. Go to the Maintenance > Hardware section of the BUI, 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.
13. If a fault does not self-clear, see [“Repairing Active Problems” on page 40](#).

▼ 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” on page 365](#) 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

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. Go to the Maintenance > Hardware section of the BUI, 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, slide the I/O Module into the disk shelf, being careful of the 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. Go to the Maintenance > Hardware section of the BUI, 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” on page 40](#).

▼ 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” on page 365](#) 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. Go to the Maintenance > Hardware section of the BUI, 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. Go to the Maintenance > Hardware section of the BUI, 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” on page 40](#).

▼ Adding a New Disk Shelf



Caution - Improper orientation of disk shelf 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 blue release tab downward for horizontally mounted HBAs or to the right for vertically mounted HBAs. See [My Oracle Support](#) Doc ID 1643673.1.

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*.


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 Installation Guide](#). For disk shelf placement guidelines, see “[Cabinet and Cabling Guidelines](#)” in *Oracle ZFS Storage Appliance Cabling Guide*.**
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](#). 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. **Verify the disk shelf presence by navigating to Maintenance > Hardware in the BUI, or by going to the CLI context `maintenance hardware` and issuing 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 OS8.7.0*.**
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 “[NDMP Configuration](#)” in *Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0*, “[Snapshots and Clones](#)” in *Oracle ZFS Storage Appliance***

Administration Guide, Release OS8.7.0, or “Remote Replication” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0.


▼ 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” on page 348](#).

1. **Remove the filler panel from the appropriate slot. For SSD locations and slot order, see the [“Disk Shelf Hardware Overview” on page 365](#).**
2. **Ensure the new drive lever is in the fully extended position.**
3. **While constantly pushing toward the pivot point of the lever, slide the drive fully into the chassis slot.**



- OR -



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. Go to the Maintenance > Hardware section of the BUI, 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, go to Configuration > Storage.
7. From the Available Pools list, select the pool you are adding the drive to.
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](#).

Note - Contact Oracle Service if you want to move a disk shelf between Oracle ZFS Storage Appliances. 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](#). For service instructions, refer to [“Servicing the Disk Shelves” on page 347](#). For information on disk compatibility, refer to [Oracle Systems Handbook](#).

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” on page 365](#)
- [“Oracle Storage Drive Enclosure DE2-24” on page 373](#)
- [“Sun Disk Shelf” on page 380](#)
- [“Disk Shelf Configurations” on page 384](#)
- [“Maximum Disk Shelves per Controller Configuration” on page 387](#)
- [“SAS-2 and SAS-3” on page 388](#)

Oracle Storage Drive Enclosure DE3-24

This section introduces the front and rear panel components of the DE3-24P and DE3-24C disk shelves.

DE3-24P Disk Shelf

The 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.

FIGURE 75 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 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” on page 384](#).

FIGURE 76 DE3-24P Status Indicators (Front View)




Figure Legend

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	

FIGURE 77 DE3-24P Rear Panel





Figure Legend

1 Power supply with fan module 0	2 I/O Module 1
3 I/O Module 0	4 Power supply with fan module 1

DE3-24C Disk Shelf

The 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.

FIGURE 78 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 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” on page 384](#).

FIGURE 79 DE3-24C Status Indicators (Front View)




Figure Legend

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	

FIGURE 80 DE3-24C Rear Panel




Figure Legend

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.

FIGURE 81 DE3 I/O Module Indicators





Figure Legend

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 activity indicator
10 Power/Activity indicator	11 Drive service required indicator	12 Ready-to-remove indicator
13 For Oracle service only		

The DE3-24C and DE3-24P disk shelves use the same I/O Module.

FIGURE 82 DE3 Power Supply Indicators**Figure Legend**

1 Power supply status indicator	2 AC power fail indicator	3 Fan fail indicator
4 DC power fail indicator	5 Universal power input connector	6 Power on/off switch

The DE3-24C and 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 DE2-24P and DE2-24C disk shelves.

DE2-24 Disk Shelf

The 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.

FIGURE 83 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 DE2-24 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” on page 384](#).

FIGURE 84 DE2-24P Status Indicators (Front View)




Figure Legend

1 System power indicator

2 Module fault indicator

3 Locate indicator

Figure Legend

4 Power/Activity indicator	5 Drive fault indicator
----------------------------	-------------------------

FIGURE 85 DE2-24P Rear Panel





Figure Legend

1 Power supply with fan module 0	2 I/O Module 1
3 I/O Module 0	4 Power supply with fan module 1

DE2-24C Disk Shelf

The 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.

FIGURE 86 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 DE2-24 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” on page 384](#).

FIGURE 87 DE2-24C Status Indicators (Front View)




Figure Legend

1 System power indicator	2 Module fault indicator	3 Locate indicator
4 Power/Activity indicator	5 Drive fault indicator	

FIGURE 88 DE2-24C Rear Panel




Figure Legend

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.

FIGURE 89 DE2 I/O Module Indicators




Figure Legend

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 DE2-24C and DE2-24P disk shelves use the same I/O Module.

FIGURE 90 DE2 Power Supply Indicators




Figure Legend

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 DE2-24C and DE2-24P disk shelves use the same power supply.

Sun Disk Shelf

The 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.

FIGURE 91 Sun Disk Shelf Drive Locations (Front View)



Read cache devices cannot be installed in the Sun Disk Shelf.

FIGURE 92 Sun Disk Shelf Status Indicators (Front View)




Figure Legend

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

FIGURE 93 Sun Disk Shelf Rear Panel




Figure Legend

- | | |
|--|---|
| <p>1 Power supply modules with built-in fans. Power supply 0 is on the left and power supply 1 is on the right.</p> | <p>2 Removable SAS Interface Module (SIM) Boards. SIM 0 is on the left, and SIM 1 is on the right.</p> |
|--|---|

FIGURE 94 Sun Disk Shelf SIM Board Indicators




Figure Legend		
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		

FIGURE 95 Sun Disk Shelf Power Supply and Fan Module Indicators




Figure Legend		
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 table describes device types supported for each disk shelf. For cabling information, see [Oracle ZFS Storage Appliance Cabling Guide](#). For information on disk compatibility, see the [Oracle Systems Handbook](#).

TABLE 53 Disk Shelf Configurations

Shelf	Device Types	Data Slots	Cache/Log Slots	Notes/Examples
DE3-24P/DE2-24P	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. 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. 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)
DE3-24P All-Flash	24 Data 20 Data and 1-4 Log devices	0 - 23 0 - 19	none 20 - 23	Software version OS8.7.0 or later is required for all-flash pools. Fully populated SSD data devices or data devices and optional log devices.

Shelf	Device Types	Data Slots	Cache/Log Slots	Notes/Examples
DE3-24C/DE2-24C	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. 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. 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)
Sun Disk Shelf (DS2)	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.

Shelf	Device Types	Data Slots	Cache/Log Slots	Notes/Examples
				Note - 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 54 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. For information on disk compatibility, see the [Oracle Systems Handbook](#).

TABLE 55 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
ZS5-4	48	N/A	4	4
ZS5-2	16	N/A	2	2
ZS4-4	36	N/A	4	4
ZS3-4	36	N/A	4	N/A
ZS3-2	16	N/A	2	2

Controller	Max. Shelves	Max. 2X4 Port SAS-2 HBA	Max. 4X4 Port SAS-2 HBA	Max. 4x4 Port SAS-3 HBA
7420	36	6	6	N/A
7320	6	1	1	N/A
7120	2	1	N/A	N/A

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 the 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” on page 384](#).

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” on page 384](#).

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 = 1.14G
var = 52.4M
update = 2.52M
stash = 14.8M
dump = 16.0G
cores = 18K
unknown = 39.0G
free = 401G
```

Disks:

DISK	LABEL	STATE
disk-000	HDD 7	healthy
disk-001	HDD 6	healthy

Note - The "disk" column is not required by the Graphical User Interface (GUI).

Working with Support Bundles

The 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 OS8.7.0.](#)

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 56 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)

Use the following procedure to generate and upload a support bundle.

Before You Begin Be sure the [Phone Home](#) service is configured for the appliance. If Phone Home is not configured, the following message is displayed:

Note - 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.

1. **Go to Maintenance > 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 and the 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 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 \(CLI\)” on page 394](#)
- [“Retrying a Failed Support Bundle Upload \(CLI\)” on page 395](#)
- [“Canceling a Pending Operation \(CLI\)” on page 396](#)
- [“Deleting a Support Bundle \(CLI\)” on page 396](#)

▼ Viewing Support Bundle Details (BUI)

Use the following procedure to view support bundle details.

1. **Go to Maintenance > 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 and the UUID.

```
hostname:maintenance system bundles> select edd08324-f091-cd19-b4c9-bb508d6ecc59
```

3. Enter show.

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 = 2014-1-9 17:42:09
      type = User initiated
      step_progress = 6

```

Related Topics

- [“Retrying a Failed Support Bundle Upload \(CLI\)” on page 395](#)
- [“Canceling a Pending Operation \(CLI\)” on page 396](#)
- [“Deleting a Support Bundle \(CLI\)” on page 396](#)

▼ 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 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 with 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 else an error will be thrown.

▼ Retrying a Failed Support Bundle Upload (BUI)

Use the following procedure to retry a failed support bundle upload.

1. **Go to Maintenance > System.**
2. **Hover over the appropriate support bundle row and click its send icon .**
The 'Send support bundle' dialog box appears.

▼ 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 and the 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 \(CLI\)” on page 394](#).

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. **Go to Maintenance > 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, use the following CLI command:**

```
hostname:maintenance system bundles> select fee4b98a-e3f7-e600-877d-b6d7933255f3
```

2. **To cancel a pending operation, use the following CLI 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. **Go to Maintenance > 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, use the following CLI command:**

```
hostname:maintenance system bundles> select fee4b98a-e3f7-e600-877d-b6d7933255f3
```

2. **To delete a support bundle, use the following CLI command:**

```
hostname:maintenance system bundles 5ff532a2-2377-e72d-b0fe-f2efc2aa8aaf> destroy
```


Performing the Initial Setup

You can repeat the initial setup process after you have already configured the appliance for the first time, as described in [“Configuring the Appliance for the First Time” in *Oracle ZFS Storage Appliance Installation Guide*](#). 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:** Click the Initial Setup button on the Maintenance > System screen.
- **CLI:** Enter the maintenance system context, then enter the setup command.

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 server, which could impact some client services, like replication. During execution, both the CLI and BUI are not available; wait for the appliance to return to normal operation.

To restart the appliance, use the following CLI command:

```
hostname:maintenance system> restart
```


Performing a Factory Reset

Factory reset resets the appliance configuration back to factory settings of the current software version, and reboots the appliance. All configuration changes are lost, and the appliance needs to go through initial configuration again, as when it was first installed. User data on the storage pool (including projects and shares) is not affected - however the pool needs to be imported as part of the initial setup process.

Note - Factory reset of a standalone controller while configured into a cluster is not supported. The controller must be unclustered first.

To perform a factory reset:

- **BUI:** Click the Factory Reset button on the Maintenance > System screen.
- **CLI:** Enter the `maintenance system` context, then issue the `factoryreset` command.
- **SSH or Serial Console:** When rebooting, select the Factory Reset entry from the boot menu.

Backing Up the Configuration

The configuration backup function enables an administrator to:

- **Back up** the 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” on page 405](#)
- [“Configuration Restore Impact” on page 406](#)
- [“Security Considerations for Configuration Backups” on page 408](#)

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:

- **Cross-Platform Restore:** A configuration backed up from one appliance platform cannot be restored to a different appliance platform. For example, a ZS4-4 configuration backup cannot be restored to a ZS3-2 appliance. 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.
- **Root privileges required** - Configuration backups include all system metadata, and therefore require all possible privileges and authorizations to create or apply. Therefore, unlike other delegated administrative options, *only* the root user is authorized to perform a configuration backup or restore.
- **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 password is *not* changed or reverted to the password at the time of the backup if it was different. The current root password is maintained on the system across the restore. For more details about passwords, refer to the summary of Security Considerations.

Security Considerations for Configuration Backups

A configuration backup contains information that is normally only accessible to the root administrative user on the appliance. Therefore, 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.

Following a configuration restore, the local **root** administrative user password is *not* modified to the root password at the time of the backup. The root password is left as-is, unmodified, by the restore process, to ensure that the password used by the administrator who is executing the restore process (and thus has logged in, using that password) is retained. If the administrator's intent was to also change the root password at the time of configuration restore, that step must be executed manually following the restore, using the normal administrative password change procedure.

▼ **Creating a Configuration Backup (BUI)**

Use the following procedure to create a configuration backup.

1. **Go to Maintenance > 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 CLI command. You are prompted to enter a descriptive comment for the backup.**

```
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, use the done command:**

```
hostname:maintenance system configs conf_backup step0> done
hostname:maintenance system configs>
```

▼ Restoring from a Saved Configuration (BUI)

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” on page 406](#).

Before You Begin Before restoring a saved appliance configuration, review the guidelines in [“Configuration Restore Impact” on page 406](#).

1. **Go to Maintenance > 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” on page 406](#).**

▼ Restoring from a Saved Configuration (CLI)

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” on page 406](#).

Before You Begin Before restoring a saved appliance configuration, review the guidelines in [“Configuration Restore Impact” on page 406](#).

1. **To revert the system to a saved configuration, use the following CLI command. You are prompted to enter the universal unique identifier for the backup.**

```
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, use the following CLI command:**

```
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. **Go to Maintenance > 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, use the following CLI command and specify the universal unique identifier for the backup:**

```
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 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 server if one is required.

1. **Go to Maintenance > System.**
2. **Mouse over the configuration list entry you wish to export and click the download icon .**
3. **Your 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 renders it unable to be imported 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 server if one is required.

- **To export a saved configuration, use the following CLI command and specify the universal unique identifier 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 server if one is required.

1. **Go to Maintenance > System.**
2. **Click the add icon  at the top of the saved configurations list and then use your web browser's file selection dialog 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 server if one is required.

- **To import a saved configuration, use the following CLI command and specify the universal unique identifier 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.

- **Go to Maintenance > 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, use the following CLI command:**

```
hostname:maintenance system configs> list
CONFIG          DATE          SYSTEM  VERSION
bfa614d7-1db5-655b-cba5-bd0bb0a1efc4 2009-8-5 17:14:28 hostname 2009.08.04,1-0
cb2f005f-cf2b-608f-90db-fc7a0503db2a 2009-8-24 17:56:53 hostname 2009.08.18,1-0
```


Working with Software Notifications and Updates

You can set up periodic checks for software updates, or you can check for updates at any time. When an updated software package is available, you are directed to download the latest package from My Oracle Support (MOS). It is recommended that you remove older system updates to avoid accruing unnecessary space on the system disks.

To use the update notification feature, the Phone Home service must be enabled, as described in [“Phone Home Configuration” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0](#).

To manage software updates, use the following tasks:

- Setting Up Software Notification - [BUI](#), [CLI](#)
- Checking for and Downloading Software Updates - [BUI](#), [CLI](#)
- Removing System Updates - [BUI](#), [CLI](#)

To understand the software update states, use the following section:

- [“Software Update States” on page 419](#)

▼ Setting Up Software Notification (BUI)

Use the following procedure to set up software notifications using the BUI.

Before You Begin Be sure the Phone Home service is enabled.

1. **Go to Maintenance > System > Software Updates.**
2. **Click the Check box, and select the frequency: daily, weekly, or monthly. When a software update is available on MOS, you will receive an alert notification similar to the following:**

An update is available on My Oracle Support version 2013.1.0.1.

3. **Go to Maintenance > System > Software Updates, and click the download icon  at the far right.**
4. **Click the download link shown in the popup window, and save the zip file.**

▼ **Setting Up Software Notification (CLI)**

Use the following procedure to set up software notifications using the CLI.

Before You Begin Be sure the Phone Home service is enabled.

1. **At the command line, enter `configuration services scrk`. The following example checks for new software versions every 30 days:**

```
hostname:configuration services scrk> ls
...
updatecheck_on = false
time_updatecheck = 7d
...
hostname:configuration services scrk> set updatecheck_on=true
updatecheck_on = true (uncommitted)
hostname:configuration services scrk> set time_updatecheck=30d
time_updatecheck = 30d (uncommitted)
```

2. **When you are notified of an update, go to MOS and download the package.**

▼ **Checking for and Downloading Software Updates (BUI)**

Use the following procedure to check for software updates using the BUI.

Before You Begin Be sure the Phone Home service is enabled.

1. **Go to Maintenance > System > Software Updates.**
2. **Notice the date of the last software check, displayed next to the Check now link.**
3. **Click the Check now link.**
4. **When an update is available on MOS, you see the following under STATUS:**
Update available for download

5. Click the download icon , at the far right.
6. Click the download link shown in the popup window, and save the zip file.

Related Topics

- [“Software Update States” on page 419](#)

▼ Checking for and Downloading Software Updates (CLI)

Use the following procedure to check for software updates using the CLI.

Before You Begin Be sure the Phone Home service is enabled.

- **At the command line, enter `show` in the `maintenance system updates` context:**

```
hostname:maintenance system updates> show
Updates:
UPDATE                RELEASE DATE          STATUS
2013.1.1.5            2014-2-18 08:00      downloadable
ak-nas@2014.01.15,1-0 2014-1-15 14:38:53   previous
ak-nas@2014.02.01,1-0 2014-2-1 19:38:55     previous
ak-nas@2014.02.08,1-0 2014-2-8 08:59:04    current
hostname:maintenance system updates>

hostname:maintenance system updates> select status=downloadable
hostname:maintenance system 2013.1.1.5> show
Properties:

        version = 2013.1.1.5
        release_date = 2014-2-18 08:00
        install_date = 2014-2-23 15:34
        status = downloadable
        url = https://updates.oracle.com/Orion/Services/download
/p18269573_20131_Generic.zip?aru=17312483&patch_file=p18269573_20131_Generic.zip

        checkdate = 2014-3-4
hostname:maintenance system 2013.1.1.5>
```

Related Topics

- [“Software Update States” on page 419](#)

▼ 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 three updates at any given time.

1. **Go to Maintenance > System.**
2. **Select a software update to delete and click its trash icon .**
3. **Click OK to confirm your action.**

▼ 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 three 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	2013-12-6 23:37:50	previous
ak-nas@2013.06.05.2.0,1-1.9	2014-5-28 15:20:06	current
ak-nas@2013.06.05.3.0,1-1.14	2014-12-19 14:31:49	waiting

3. **Enter destroy 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
```

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 57 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.

Upgrading the Software

This section describes how to upgrade the system software for an 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” on page 421](#)
- [“Upgrading Software on a Standalone Controller” on page 422](#)
- [“Upgrading Software on Clustered Controllers” on page 427](#)

Preparing for a Software Upgrade

Before you upgrade the software, perform the following actions for standalone or clustered controllers.

During the update process, some protocols may experience an outage. See related topics for more information.

1. Verify your current software version.
 - In the BUI, go to Maintenance > System.
 - In the CLI, go to `maintenance system updates` and enter `show`.
2. 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, go to Configuration > 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, go to Maintenance > Problems.
 - In the CLI, go to `maintenance problems show`.

For more information, see [“Active Problems Display” on page 39](#).

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

- [“Upgrading Software on a Standalone Controller” on page 422](#)
- or
- [“Upgrading Software on Clustered Controllers” on page 427](#)

Related Topics

- White paper "Best Practices for Upgrading Oracle ZFS Storage Appliance," located at <http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html>.

Upgrading Software on a Standalone Controller

This section describes how to upgrade the system software for an 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 \(BUI\)” on page 423](#)

- [“Upgrading Software on a Standalone Controller \(CLI\)” on page 425](#)

▼ Upgrading Software on a Standalone Controller (BUI)

Use the following procedure to upgrade software on a standalone controller.

Before You Begin Complete [“Preparing for a Software Upgrade” on page 421](#).


1. **On your local system, download the software update as described in [“Checking for and Downloading Software Updates \(BUI\)” on page 416](#).**

The update package is a single compressed file named with the version number, such as `ak-nas-2013-06-05-0-0.0.pkg.gz`. 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. **Go to Maintenance > System.**
3. **Click the add icon  next to Software Updates.**
4. **Click Browse and locate the pkg.gz 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 .

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 OS8.7.0.](#)

▼ Upgrading Software on a Standalone Controller (CLI)

Use the following procedure to upgrade software on a standalone controller.

Before You Begin Complete “[Preparing for a Software Upgrade](#)” on page 421.

1. **On your local system, download the software update as described in “[Checking for and Downloading Software Updates \(CLI\)](#)” on page 417.**

The update package is a single compressed file named with the version number, such as `ak-nas-2013-06-05-0-0.0.pkg.gz`. 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.

After the file has been transferred, it is automatically unpacked and verified.

```
hostname:maintenance system updates download (uncommitted)> set url=ftp://mydirectory/
update.pkg.gz
      url = ftp://mydirectory/update.pkg.gz
hostname:maintenance system updates download (uncommitted)> set user=bmc
      user = bmc
hostname:maintenance system updates download (uncommitted)> set password
Enter password:
      password = *****
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>
```

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 OS8.7.0.](#)**

Upgrading Software on Clustered Controllers

This section describes how to upgrade the system software for an 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.

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)

Use the following procedure to prepare to upgrade clustered controllers.

Before You Begin Complete [“Preparing for a Software Upgrade”](#) on page 421.

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. **On your local system, download the software update as described in [“Checking for and Downloading Software Updates \(BUI\)”](#) on page 416.**

The update package is a single compressed file named with the version number, such as `ak-nas-2013-06-05-0-0.0.pkg.gz`. 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, go to Configuration > 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 OS8.7.0*.

3. **If the management interface is unlocked , click the lock icon  next to the management interface.**

4. **Go to Maintenance > System and click the add icon  next to Software Updates.**

5. **In the Add Software Update 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 \(BUI\)” on page 431](#)

▼ Preparing to Upgrade Clustered Controllers (CLI)

Use the following procedure to prepare to upgrade clustered controllers.

Before You Begin Complete [“Preparing for a Software Upgrade” on page 421](#).

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. **On your local system, download the software update as described in [“Checking for and Downloading Software Updates \(CLI\)” on page 417](#).**

The update package is a single compressed file named with the version number, such as `ak-nas-2013-06-05-0-0.0.pkg.gz`. 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.gz
  url = ftp://mydirectory/update.pkg.gz
controller-a:maintenance system updates download (uncommitted)> set user=bmc
  user = bmc
controller-a:maintenance system updates download (uncommitted)> set password
Enter password:
  password = *****
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.gz
  url = ftp://mydirectory/update.pkg.gz
controller-b:maintenance system updates download (uncommitted)> set user=bmc
  user = bmc
controller-b:maintenance system updates download (uncommitted)> set password
Enter password:
  password = *****
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 \(CLI\)” on page 433](#)

▼ Upgrading Clustered Controllers (BUI)

Before You Begin

- Complete [“Preparing for a Software Upgrade” on page 421](#).
- Complete [“Preparing to Upgrade Clustered Controllers \(BUI\)” on page 428](#).

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, go to Maintenance > System, and click the arrow icon next to the name of the update you want to install.**
2. **(Optional) Click CHECK to perform health checks.**
For information about health checks, see [“Working with Health Checks” on page 445](#).
3. **Click APPLY to begin the update process.**
4. **Wait for controller A to fully reboot, and log back in to controller A.**
5. **Go to Configuration > 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 OS8.7.0](#).
6. **(Optional) To monitor firmware updates, go to Maintenance > System and check the update counter.**
7. **Log in to controller B and go to Configuration > Cluster 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 OS8.7.0](#).

8. **Go to Configuration > 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.

9. **Go to Maintenance > System, and click the arrow icon  next to the name of the update you want to install.**
10. **(Optional) Click CHECK to perform health checks.**
For information about health checks, see [“Working with Health Checks”](#) on page 445.
11. **Click APPLY to begin the update process.**
12. **Wait for controller B to fully reboot, and then log back in to controller B.**
13. **Go to Configuration > Cluster 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 OS8.7.0*.
14. **(Optional) To monitor firmware updates, go to Maintenance > System and check the update counter.**
15. **Log in to controller A and go to Configuration > Cluster to 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 OS8.7.0*.
16. **To verify that all firmware updates are complete, go to Maintenance > System and check the update counter.**

Note - Do not begin the next step until all firmware updates are complete.

17. **Go to Configuration > 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.

18. **Go to Maintenance > Hardware to verify that all disks are online.**

All lights should be green.

19. **Verify there are no controller and disk shelf component errors.**
All lights should be green. An amber light indicates a component error.
20. **If any components have errors, check for pool errors by going to Configuration > Storage, and check the STATUS and ERRORS columns for each pool.**
Pools should be online and have no errors.
21. **Log in to controller B and repeat steps 18-20 for controller B.**
22. **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 OS8.7.0*.**

▼ Upgrading Clustered Controllers (CLI)

- Before You Begin**
- Complete “[Preparing for a Software Upgrade](#)” on page 421.
 - Complete “[Preparing to Upgrade Clustered Controllers \(CLI\)](#)” on page 429.

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, go to maintenance system updates, and enter show to view the available updates.**

```
controller-a:> 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

2. **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
```

3. (Optional) 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
```

4. 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.

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
```

5. 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 OS8.7.0.](#)

6. (Optional) 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

7. **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 OS8.7.0*.

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)
```

8. **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**

9. **Go to 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

10. **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
```

11. (Optional) 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
```

12. 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.

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
```

13. 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 OS8.7.0.](#)

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)
```

14. (Optional) 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

```

15. **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 OS8.7.0.](#)

16. **To verify that all firmware updates are complete, go to 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.

17. **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.

18. Go to 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
...
```

19. Go to 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
```

20. Log in to controller B and repeat steps 18-19 for controller B.

21. 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 OS8.7.0*.

Updating the Platform Firmware

This section describes how to automatically update the Oracle ILOM and/or BIOS firmware for an Oracle ZFS Storage Appliance. This process is known as a platform firmware update, and 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.

The automatic platform firmware update feature is supported on the following platforms:

- ZS5-4
- ZS5-2
- ZS4-4
- 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
ZS5-4	3.2.6.20 r109303
ZS5-2	3.2.6.34 r111631
ZS4-4	3.2.4.18.a r95773
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 = 3.2.7.32.a r112581
```

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. Go to Maintenance > Problems.
2. Under Active Problems, look for a major alert stating the chassis is running an older (downrev) platform firmware version.

The screenshot shows the Oracle ILOM BUI interface. The top navigation bar includes 'Configuration', 'Maintenance', 'Shares', 'Status', and 'Analytics'. Below this, there are sub-tabs for 'HARDWARE', 'SYSTEM', 'PROBLEMS', 'LOGS', and 'WORKFLOWS'. The 'Active Problems' section shows 1 total problem. The table below lists the problem details:

DATE	DESCRIPTION	TYPE	PHONED HOME
2016-10-19 10:09:13	The chassis '1423NM900H' is running downrev Platform Firmware.	Major Defect	2016-10-19 10:12:59

Below the table, the 'Problem Details' section provides further information:

- Description:** The chassis '1423NM900H' is running downrev Platform Firmware.
- Type:** Major Defect
- Impact:** The Platform Firmware includes the firmware for the Service Processor as well as the System Board firmware. Running downrev firmware can expose this appliance to security and stability issues.
- Affected components:** 100% brmzs4-4-07-mgmt: hc://chassis-mfg=Oracle-Corporation;chassis-name=SUN-SERVER-X4-4;chassis-part=ZS4-4-P11-02;chassis-serial=1423NM900H;fru-serial=1423NM900H;fru-revision=SUN-SERVER-X4-4;chassis=0 (in service)
- Manufacturer:** Oracle
- Model number:** Oracle ZFS Storage ZS4-4
- Serial number:** 1423NM900H
- Revision:** SUN-SERVER-X4-4
- Automated response:** None.
- Recommended action:** Reboot to install the latest firmware. If you are a qualified service person, detailed information on this problem can be found at <http://support.oracle.com/msg/AK-8004-HJ>
- Event time:** 2016-10-19 10:09:13
- Unique Identifier:** 26fb6199-1219-ee52-9bbb-a15d593ffb8d
- Phoned home:** 2016-10-19 10:12:59

3. If this alert is present, proceed to “Updating the Platform Firmware (BUI)” on page 441.
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. Enter maintenance system updates platform.

2. Enter get:

```
hostname:maintenance system updates platform> get
      update_needed = true
      power_down_needed = false
```

If `update_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 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)” on page 442.

If `update_needed` is `false`, the platform firmware is up to date and no further action is needed.

▼ Updating the Platform Firmware (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.

Before You Begin 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” on page 421](#).
- View the platform firmware versions to determine if this task is necessary. See [“Checking the Platform Firmware Versions \(BUI\)” on page 440](#).

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)

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.

Before You Begin 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” on page 421](#).
- View the platform firmware versions to determine if this task is necessary. See [“Checking the Platform Firmware Versions \(CLI\)” on page 440](#).

1. In the CLI, 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.

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” on page 41](#), 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” on page 421](#), 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. To select the update media, use the following CLI command:

```
hostname:maintenance system updates> select ak-nas@2013.06.05.0.0,1-1.6
```


2. To upgrade the software and check system health, use the following CLI 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 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.

FIGURE 96 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 an 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 58 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
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
PAN1 "Slot <slot> has a Revision B3 SAS HBA; Revision C0 (or later) required."	A revision B3 SAS HBA is present.	4
PAN2 "J4400 and J4500 disk shelves are not supported in this release."	Disk shelf/shelves are present that are not supported for this software release.	
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 unlit, check all connections and replace cables as needed.**
- 2. Identify affected chassis, then disconnect and remove faulted SIM or IOM. After two minutes, re-seat SIM or IOM and wait for steady Power LED before reconnecting cables.**
- 3. Identify affected chassis, and remove faulted disk. After 30 seconds, re-seat disk and wait for steady or flashing LED.**
- 4. Contact Oracle Support for component service or replacement.**

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” on page 462](#).

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 OS8.7.0](#).

To apply deferred updates, use one of the following tasks:

- Applying Deferred Updates [BUI](#), [CLI](#)

For a description of a deferred update, select a deferred update from the following table.

TABLE 59 Deferred Updates

Feature	Version Introduced
“Passthrough x Deferred Update” on page 453	2009.Q2.0.0
“User Quotas Deferred Update” on page 454	2009.Q3.0.0
“COMSTAR Deferred Update” on page 455	2009.Q3.0.0
“Triple Parity RAID Deferred Update” on page 455	2009.Q3.0.0
“Dedup Deferred Update” on page 455	2010.Q1.0.0
“Replication Deferred Update” on page 456	2010.Q1.0.0
“Received Properties Deferred Update” on page 456	2010.Q1.0.0
“Slim ZIL Deferred Update” on page 456	2010.Q3.1.0
“Snapshot Deletion Deferred Update” on page 457	2010.Q3.1.0
“Resize Snapshots Deferred Update” on page 457	2010.Q3.1.0
“Multi Replace Deferred Update” on page 457	2010.Q3.1.0
“RAIDZ Mirror Deferred Update” on page 458	2011.1.0.0
“Optional Child Dir Deferred Update” on page 458	2011.1.0.0

Feature	Version Introduced
“Multiple Initiator Groups per LUN Deferred Update” on page 458	2011.1.8.0
“Large Block Sizes Deferred Update” on page 459	2013.1.1.0
“Sequential Resilvering Deferred Update” on page 459	2013.1.2.0
“NDMP-zfs Replica Backup Deferred Update” on page 459	2013.1.4.0
“ACL Passthrough with Mode Preservation Deferred Update” on page 460	OS8.6.0
“LZ4 Data Compression Deferred Update” on page 460	OS8.7.0
“Fast Copy with Encryption Deferred Update” on page 461	OS8.7.0
“Data Deduplication v2 Deferred Update” on page 461	OS8.7.0
“Asynchronous Dataset Deletion Deferred Update” on page 462	OS8.7.0
“RAID Space Efficiency Improvements Deferred Update” on page 462	OS8.7.0

▼ Applying Deferred Updates (BUI)

Use the following procedure to apply deferred updates using the BUI.

- 1. Go to Maintenance > 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” on page 462](#)

▼ Applying Deferred Updates (CLI)

Use the following procedure to apply deferred updates using the CLI.

- **To see if there are any available deferred updates, run the `show` command. If deferred updates are available, you can use the `apply` command:**

```
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

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)

```
clownfish:maintenance system updates> apply
```

Related Topics

- [“Understanding Deferred Updates” on page 462](#)

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 introduces 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” on page 462](#)

User Quotas Deferred Update

With the 2009.Q3 software release, the system now supports user and group quotas on a per-share basis. In order 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 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 OS8.7.0.](#)

Related Topics

- Applying Deferred Updates [BUI, CLI](#)
- [“Understanding Deferred Updates” on page 462](#)

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](#)” on page 462

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.

In order 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](#)” on page 462

Dedup Deferred Update

This update provides the ability to use data deduplication. In order 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” on page 462](#)

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 resync. The system also sends incremental replication updates for actions configured under earlier releases.

Related Topics

- Applying Deferred Updates [BUI](#), [CLI](#)
- [“Understanding Deferred Updates” on page 462](#)

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” on page 462](#)

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” on page 462](#)

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” on page 462](#)

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” on page 462](#)

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” on page 462](#)

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” on page 462](#)

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” on page 462](#)

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” on page 462](#)

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” on page 462](#)

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” on page 462](#)

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](#)” on page 462

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](#)” on page 462

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 OS8.7.0](#).

Related Topics

- Applying Deferred Updates [BUI, CLI](#)
- ["Understanding Deferred Updates" on page 462](#)

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 \(2021771.1\)](#) (<https://support.oracle.com/epmos/faces/DocumentDisplay?id=2021771.1>).

Related Topics

- [Oracle ZFS Storage Appliance Plug-in Downloads](http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html) (<http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html>)
- Applying Deferred Updates [BUI, CLI](#)
- ["Understanding Deferred Updates" on page 462](#)

Data Deduplication v2 Deferred Update

This update provides a new data structure and algorithm for storing the deduplication metadata used by ZFS. The appliance is now able to store specific types of metadata on a designated set of devices (e.g., SSD) to optimize use cases like deduplication. In order to use data deduplication after the update, metadata devices will have to first be configured in a storage pool.

In order to use this mode, the storage pool must be upgraded. When a pool is upgraded, the existing entries will be migrated from the ZAP to the new data structure. This will happen in the background and will be throttled to have minimal impact on the system.

Related Topics

- Applying Deferred Updates [BUI, CLI](#)
- [“Understanding Deferred Updates” on page 462](#)

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” on page 462](#)

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” on page 462](#)

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 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](#)

Working with Firmware Updates

Following the application of a software update, any hardware for which the update includes newer versions of firmware is upgraded. Before the upgrade window, it is recommended to run a scrub as described in [“Scrubbing a Storage Pool \(BUI\)” in Oracle ZFS Storage Appliance Administration Guide, Release OS8.7.0](#).

There are several types of devices for which firmware updates may be made available; each has distinct characteristics. Disks, storage enclosures, and certain internal SAS devices are upgraded in the background. When this is occurring, the firmware update 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.

For clustered controllers, the status only shows the updates that are pending on the local controller. For example, firmware updates displayed on the peer controller do not include firmware updates for the primary controller.

As of 2010Q3.4, when there are outstanding updates, an informational or warning icon appears next to the number of updates remaining. Clicking the icon brings up the Firmware Updates dialog, which lists the current remaining updates. For each update we also show the current version of the component, the time of the last attempted update, as well as the reason why the last attempt did not succeed.

We consider any outstanding updates to be in one of 3 states: Pending, In Progress and Failed. An update begins in the Pending state, and is periodically retried, at which time it moves into the In Progress state. If we fail to upgrade, due to a transient condition, the update is moved back to the Pending state, and otherwise to the Failed state.

In general, there is only an indication of a problem if:

- There are updates in the Failed state.
- Updates remain in the Pending state (or in limbo 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:

-
- There are multiple chassis being upgraded, we are making progress (the number of remaining updates decreases), and some of the chassis transiently appear pending with a status indicating that some disk has only one path. This is also expected, since when we upgrade a chassis, we may reset one of its expanders. Resetting an expander causes some disks to temporarily have only one path, and as a result, updates to other chassis are held back until it is safe to do so again non-disruptively.

Note that currently the Firmware Updates dialog does not automatically refresh, so you would have to close it and re-open it to get an updated view.

Applying hardware updates is always done 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 update is completed; pending firmware updates are suspended until the takeover or failback has completed, at which time the restrictions described below are reevaluated in the context of the new cluster state and, if possible, firmware updates resume.



Caution - Unless absolutely necessary, takeover and failback operations should not be performed while firmware updates are in progress.

The rolling update procedure documented later meets all of these best practices and addresses the per-device-class restrictions described later. It should always be followed when performing updates in a clustered environment. In both clustered and standalone environments, these criteria are also reevaluated upon any reboot or diagnostic system software restart, which may cause previously suspended or incomplete firmware updates to resume.

- Components internal to the storage controller (such as HBAs and network devices) other than disks and certain SAS devices are generally upgraded automatically during boot; these updates are not visible and will have completed by the time the management interfaces become available.
- Upgrading disk or flash device firmware requires that the device 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.

- Upgrading the firmware in a disk shelf requires that both back-end storage paths be active to all disks within all enclosures, and for storage to be configured on all shelves to be upgraded. For clusters with at least one active pool on each controller, these restrictions

mean that disk shelf firmware update can be performed only by a controller that is in the "owner" state.

During the firmware update process, hardware may appear to be removed and inserted, or offlined and onlined. While alerts attributed to these actions are suppressed, if you are viewing the Maintenance > Hardware screen or the Configuration > Storage screen, you may see the effects of these updates in the UI in the form of missing or offline devices. This is not a cause for concern; however, 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. Additionally, in some cases, the controllers in the disk shelves may remain offline during firmware update. If this occurs, no other controllers are upgraded until this condition is fixed. If an enclosure is listed as only having a single path for an extended period of time, check the physical enclosure to determine whether the green link lights on the back of the SIM or IOM are active. If not, remove and re-insert the SIM or IOM to re-establish the connection. Verify that all enclosures are reachable by two paths.

▼ Checking the Status of Firmware Updates (BUI)

- **In the BUI, navigate to Maintenance > System view.**

The firmware update progress is displayed in the left panel of this view.

For clustered controllers, the status only shows the updates that are pending on the local controller. For example, firmware updates displayed on the peer controller do not include firmware updates for the primary controller.

▼ Checking the Status of 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 local controller. For example, firmware updates displayed on the peer controller do not include firmware updates for the primary controller.

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 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 fail-safe rollback procedure.

▼ 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, i.e., has failed to even 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 above. The confirmation step looks like this:

```
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:**
 - 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.

- **To roll back, locate a previous image and click on the rollback icon . You are asked to confirm that you wish to execute a rollback, and then the system reboots and 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” on page 469](#).**

▼ Rolling Back the System Software (CLI)

Use the following procedure to roll back the system software using the CLI.

- **To roll back to a previous software version, select the update that corresponds to that version and execute the rollback command. You will be asked to confirm to execute a rollback, and then the system will reboot and execute the rollback. Unlike fail-safe rollback, you will not be asked for further confirmation when the system reboots.**

Removing an Update Package

This section describes how to remove an 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 on 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, use the `destroy` command, specifying 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>
```

