

Best Practices for Upgrading Oracle ZFS Storage

Oracle ZFS Storage Appliances and Oracle ZFS High-Availability Instances in OCI

December, 2023, Version 2.0

Copyright © 2023, Oracle and/or its affiliates

Public

Table of contents

Introduction	4
Purpose Statement	4
Disclaimer	4
About System Updates	5
System Update Version	5
Retrieving a System Update	5
Uploading a System Update	6
Remove Older System Updates	7
Planning an Upgrade	8
Data Services for Clustered Controllers	8
Maintenance Window	9
Network Interfaces	10
Data Interfaces	10
Management Interfaces	11
Disable Non-critical Services	11
Backup Configuration	12
Browser User Interface	12
Command Line Interface	12
Disk Events	13
Health Check	13
Browser User Interface	14
Command Line Interface	15
Deferred Updates	15
Upgrading	16
Standalone Controller	16
Upgrade the Controller	16
Clustered Controllers	18
Upgrade First Controller	18
Upgrade the Second Controller	20
Post Upgrade	23
Deferred Updates	23
Restart Data Services	24
Backup Configuration Again	24



Rollback	25
Standard Rollback	25
Fail-safe Rollback	26
Appendix A – Oracle ZFS Storage Appliance Standalone Planning Sheet	27
Appendix B – Oracle ZFS Storage Appliance Clustered Planning Sheet	28

Introduction



Purpose Statement

This document serves as a planning guide for upgrading Oracle ZFS Storage Appliance software. It contains step-by-step instructions for the upgrade process, as well as best practices for ensuring success during the entirety of the procedure. By reading the entire guide before performing an upgrade, you can devise a strategy that accurately reflects your ecosystem's requirements.

The Oracle ZFS Storage Appliance is a network-attached storage platform that provides high performance and high scalability. Its hybrid storage pool design invokes tiered caching between DRAM, SSD flash, and backend hard disks, providing low latency, high throughput, and high I/O access. Oracle ZFS Storage Appliance offers NFS, SMB, iSCSI, Object Storage, and Fibre Channel connectivity to allow storage and network administrators to design a deployment strategy that seamlessly integrates into an existing datacenter.

ZFS High-Availability Storage (ZFS-HA) in the Oracle Cloud Infrastructure (OCI) brings the software features and protocols of the hardware Oracle ZFS Storage Appliance to the cloud in a scalable form to match your OCI and on-premises clients' data needs. These virtual ZFS Storage Appliances run the same code as the hardware appliances and are upgraded through the same procedures detailed here.

Disclaimer

This document in any form, software, or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced, or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement, nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

About System Updates

A system update for Oracle ZFS Storage Appliance is a binary file that contains new management software as well as new hardware firmware for your storage controllers and disk shelves. Its purpose is to provide additional features, bug fixes, and security updates, allowing your storage environment to run at peak efficiency. It also adds features that may be specific to the ZFS-HA storage instances that run in OCI. Oracle releases system updates roughly once a month.

This section highlights the properties of a system update and how to download it onto your Oracle ZFS Storage Appliance or ZFS-HA instance.

System Update Version

The update feature is contained in the Software Updates section under Maintenance → System of the management interface. Each available system update is listed there with a version format that is contained in a numerical string as well as a simplified release name. To understand this format, look at the following browser-user interface example:

Software Updates			Check for updates on My Oracle Support
VERSION ▲	RELEASE DATE	RELEASE NAME	STATUS
2013.06.05.8.62,1-1.3	2023-10-5 14:41:46	OS8.8.62	Current system software
2013.06.05.8.63,1-1.1	2023-10-30 09:50:17	OS8.8.63	Uploaded at: 2023-12-5 16:12:05 ↻

The current running version on this system is **2013.06.05.8.62.1-1.3**, more commonly referred to as **OS8.8.62**. Here is a breakdown of the version string:

2013.06.05. Refers to the date the underlying software was synchronized with the base Oracle Solaris 11.4 operating system. Because Solaris went to a rolling release schedule with 11.4, this date is not likely to change. This can be considered the major release number and replaced with “OS8”. for the release name.

8. Refers to the minor version, now frozen at 8 due to the nature of the Solaris 11.4 rolling release numbering.

62, Refers to the “dot” version and is in step with the Solaris 11.4.x micro release numbering.

1-1.3 Refers to the build number. This is only relevant to Oracle Engineering.

Retrieving a System Update

Next to Software Updates, you can click on the link “Check for updates on My Oracle Support”, which will open the web page for My Oracle Support. From there you can search for the [Doc ID 2021771.1](#) which lists all the updates available for ZFS Storage.

Software Updates			Check for updates on My Oracle Support
VERSION ▲	RELEASE DATE	RELEASE NAME	STATUS
2013.06.05.8.62,1-1.3	2023-10-5 14:41:46	OS8.8.62	Current system software
2013.06.05.8.63,1-1.1	2023-10-30 09:50:17	OS8.8.63	Uploaded at: 2023-12-5 16:12:05 ↻

System updates are delivered in a ZIP file which contains the release notes and the update’s PKG file. Unzip the downloaded and review the release notes in a browser window prior to upgrading to understand the impact on your

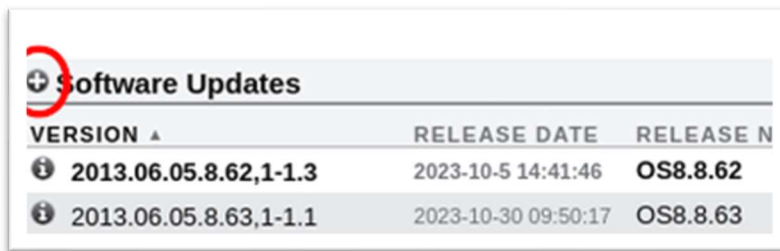
storage environment. It is possible that the uploaded system update is not supported on your Oracle ZFS Storage Appliance platform.

Uploading a System Update

Use the following browser user interface or command-line interface procedures to upload a system update onto the Oracle ZFS Storage controller. For clustered controllers, perform this on both controllers.

Browser User Interface

1. Navigate to Maintenance → System.
2. Click on the plus icon next to Software Updates.



VERSION ▲	RELEASE DATE	RELEASE N
2013.06.05.8.62,1-1.3	2023-10-5 14:41:46	OS8.8.62
2013.06.05.8.63,1-1.1	2023-10-30 09:50:17	OS8.8.63

3. Click Browse and locate the .pkg update file previously downloaded.



Command Line Interface

The command line interface requires a download from either an HTTP or FTP source. It supports username and password authentication if necessary.

1. zfs:> **maintenance system updates**
2. zfs:maintenance system updates download (uncommitted)> **set url=http://path/to/file.pkg**
3. zfs:maintenance system updates download (uncommitted)> **set user=** {if necessary}
4. zfs:maintenance system updates download (uncommitted)> **set password=** {if necessary}
5. zfs:maintenance system updates download (uncommitted)> **commit**

Remove Older System Updates


To avoid accruing too much space on the system disks, it is recommended that no more than two updates are retained at any given time. Follow these steps to remove older updates.

Browser User Interface

1. Navigate to Maintenance → System.
2. Select a Software Update to delete and click on its trash icon.

Software Updates		Check for updates on My Oracle Support	
VERSION ▲	RELEASE DATE	RELEASE NAME	STATUS
2013.06.05.8.62,1-1.3	2023-10-5 14:41:46	OS8.8.62	Current system software
2013.06.05.8.63,1-1.1	2023-10-30 09:50:17	OS8.8.63	Uploaded at: 2023-12-5 16:12:05

3. Click OK to confirm.



Command Line Interface

1. zfs:> **maintenance system updates**
2. zfs:maintenance system updates> **show**

Updates:

UPDATE	DATE	RELEASE NAME	STATUS
ak-nas@2013.06.05.8.61,1-1.2	2023-9-14 23:37:50	OS8.8.61	previous
ak-nas@2013.06.05.8.62,1-1.3	2023-10-5 18:41:46	OS8.8.62	current
ak-nas@2013.06.05.8.63,1-1.1	2023-10-30 13:50:17	OS8.8.63	waiting

3. zfs:maintenance system updates> **destroy ak-nas@2013.06.05.8.61,1-1.2**

This will destroy the update "ak-nas@2013.06.05.8.61,1-1.2". Are you sure? (Y/N) **Y**

Planning an Upgrade

Oracle ZFS Storage should be upgraded a minimum of once per year. However, at least two to four upgrades are strongly recommended. This helps to ensure your software and hardware firmware remain up to date, thus reducing the risk of unforeseen downtime.

Data Services for Clustered Controllers

While Oracle ZFS Storage features high availability using clustering, it cannot guarantee uptime for every protocol during the upgrade process. Some protocols, such as SMB or NDMP, are highly session-oriented and will not transfer over to the cluster peer in the event of a resource takeover. Others, such as NFS, feature timeout values that allow them to survive and migrate from one controller to the peer. Please read the following two sections carefully to understand which protocols will affect your storage availability during an upgrade.

Disruptive Protocols

The following table outlines the protocols that will experience an outage during an upgrade for a clustered environment. If your environment leverages these services, your upgrade plan must include scheduled downtime.

DATA SERVICES REQUIRING SCHEDULED DOWNTIME

Protocol	Action Plan
SMB 1.0	This includes workstations running Windows XP or earlier and servers running Windows Server 2003 or earlier. Network drives may need to be remounted after the upgrade process is complete.
SMB 2.0 / 2.1	This includes workstations running Windows Vista or later and servers running Windows Server 2008 or later. Network drives may need to be remounted after the upgrade process is complete.
FTP	All FTP sessions will be disconnected during the upgrade process.
TFTP	All TFTP sessions will be disconnected during the upgrade process.
NDMP	All NDMP backups will be disconnected during the upgrade process. It is recommended to complete all backups prior to upgrade and then disable the NDMP service to enforce a quiet period.

Non-disruptive Protocols

The following table outlines the protocols that do not require scheduled downtime during an upgrade for clustered controllers. These services will experience a brownout period while network and pool resources are transferred to the cluster peer. The timeout variables on these protocols ensure that services will recover after the takeover action.

DATA SERVICES NOT REQUIRING SCHEDULED DOWNTIME

Protocol	Action Plan
NFSv3 / NFSv4	Verify hard mounts are used for every device connected to a share. Soft mounts are inherently unreliable. This can be checked by running <code>nfsstat -m</code> on all attached devices.
Fibre Channel	Verify all attached devices have multiple active and standby paths. This can be verified with <code>multipath -ll</code> on Linux, <code>mpathadm show LU</code> on Solaris, and the Disk Management Utility on Windows.
iSCSI	Verify all attached devices have multiple active and standby paths. This can be verified with <code>multipath -ll</code> on Linux, <code>mpathadm show LU</code> on Solaris, and the Disk Management Utility on Windows.
iSER	Verify all attached devices have multiple active and standby paths. This can be verified with <code>multipath -ll</code> on Linux and <code>mpathadm show LU</code> on Solaris.
SRP	Verify all attached devices have multiple active and standby paths. This can be verified with <code>multipath -ll</code> on Linux.
Replication	No action is required provided the initial sync of a project or share is complete. If the initial sync has not been finished, it will need to be manually restarted after the upgrade process is complete.

Maintenance Window

Regardless of whether you use disruptive or non-disruptive protocols, it is recommended to schedule a maintenance window for the upgrading of your storage controllers. You should inform your users that storage will be either offline or functioning in a limited capacity for the duration of the upgrade. The minimum length of time should be set at one hour. This does not mean your storage will be offline for the entire hour. Instead, it helps to set expectations that storage performance and availability cannot be guaranteed during this period.

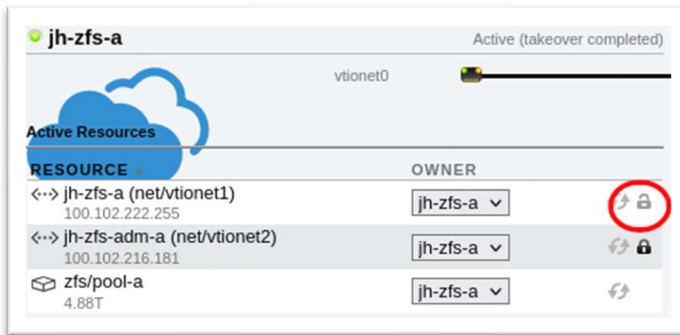
Network Interfaces

Data Interfaces

It is recommended that all shared data interfaces for clustered controllers be *open*, or *unlocked*, prior to upgrading. This ensures these interfaces migrate to the peer controller during a takeover or reboot. Failure to do so will result in downtime. Follow the procedures below on both controllers when in a clustered configuration.

Browser User Interface

1. Navigate to Configuration → Cluster.
2. Identify the interfaces dedicated for data and verify that their lock icons are grey.



3. If necessary, click APPLY for changes to take effect.



Command Line Interface

1. `zfs:> configuration cluster resources`
2. `zfs:configuration cluster resources> show`
Resources:

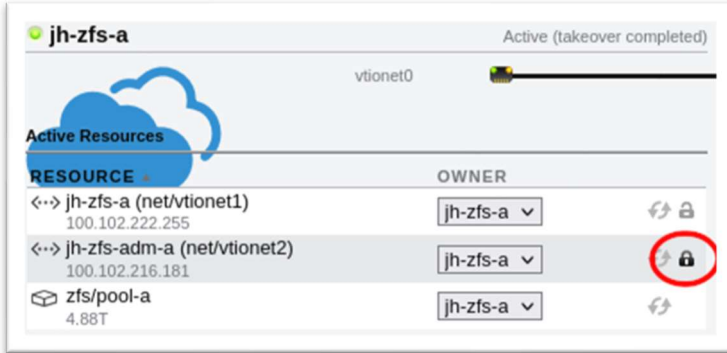
RESOURCE	OWNER	TYPE	LABEL	CHANGES	DETAILS
net/vt10net1	jh-zfs-a	singleton	jh-zfs-a	no	10.80.218.170
net/vt10net2	jh-zfs-a	private	jh-zfs-adm-a	no	10.80.216.46
zfs/pool-a	jh-zfs-a	singleton		no	4.88T
3. `zfs:configuration cluster resources> select net/vt10net2 set type=singleton`
type = singleton
4. `zfs:configuration cluster resources> commit`

Management Interfaces

It is recommended that all management interfaces for clustered controllers are set to *private*, or *locked*, prior to upgrading. This ensures the interfaces do not migrate to the peer controller during a takeover or reboot. Follow the procedures below on both controllers when in a clustered configuration.

Browser User Interface

1. Navigate to Configuration → Cluster.
2. Identify the interface dedicated for management and click on its lock icon. It should be black.



3. Click APPLY for the change to take effect.



Command Line Interface

1. `zfs:> configuration cluster resources`
2. `zfs:configuration cluster resources> show`
Resources:

RESOURCE	OWNER	TYPE	LABEL	CHANGES	DETAILS
net/vt10net1	jh-zfs-a	singleton	jh-zfs-a	no	10.80.218.170
net/vt10net2	jh-zfs-a	private	jh-zfs-adm-a	no	10.80.216.46
zfs/pool-a	jh-zfs-a	singleton		no	4.88T
3. `zfs:configuration cluster resources> select net/vt10net2 set type=private`
type = private
4. `zfs:configuration cluster resources> commit`

Disable Non-critical Services


Disabling non-critical data services can shorten the upgrade time and ensures that the system has a minimal operation load during the update. Non-critical services may include NDMP, shadow migration, replication, or others.

Backup Configuration

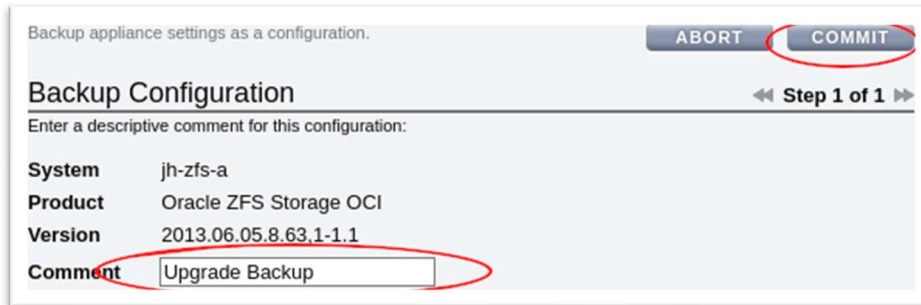
In the event of an unforeseen failure, it may be necessary to factory reset a storage controller. To minimize the downtime, it is recommended to maintain an up-to-date backup copy of the management configuration.

Browser User Interface

1. Navigate to Maintenance → System.
2. At the bottom, click on the BACKUP button.



3. Create a comment for the configuration and press COMMIT to create the backup configuration.



4. Download the configuration locally to your computer by clicking on the download icon.



Command Line Interface

1. `zfs:> maintenance system configs`
2. `zfs:maintenance system configs> backup`
3. `zfs:maintenance system configs conf_backup step0> set comment="Upgrade Backup"`
`comment = Upgrade Backup`
4. `zfs:maintenance system configs conf_backup step0> done`

Disk Events

To avoid unnecessary delays with the upgrade process, do not update your system whenever there are active disk resilvering events or scrub activities. Check if these activities are occurring and allow them to complete if they are in progress.

Browser User Interface

1. Navigate to Configuration → Storage.
2. Verify there are no disk resilvering events.
3. Verify there are no scrub activities.

Command Line Interface

```
1. zfs:> configuration storage
2. zfs:configuration storage> show
   Properties:
           pool = pool-a
           status = online
           errors = 0
           owner = jh-zfs-a
           scrub = scrub in progress for 0h8m, 27.6% done
```

Health Check

Oracle ZFS Storage Appliance has a health check feature that examines the state of your storage controller and disk shelves prior to upgrading. It is automatically run as part of the upgrade process but should also be run independently to check storage health prior to entering a maintenance window. This functionality examines the following potential issues:

- Controller is not supported for the desired software release
- Controller HBA is faulted
- Disk shelf is not supported for the desired software release
- Disk shelf IOM is missing, faulted, or offline
- Disk shelf IOM is missing one or more paths
- Drive inside disk shelf is missing one or more paths

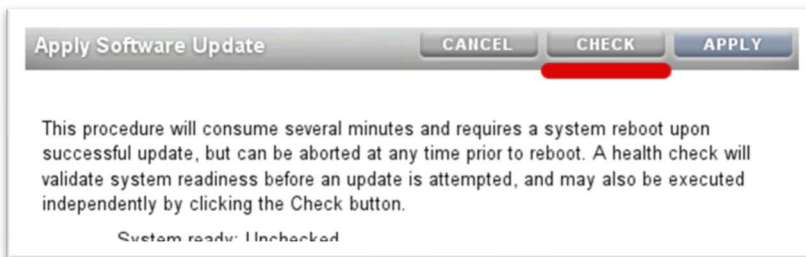
Use the following procedure to run a health check:

Browser User Interface

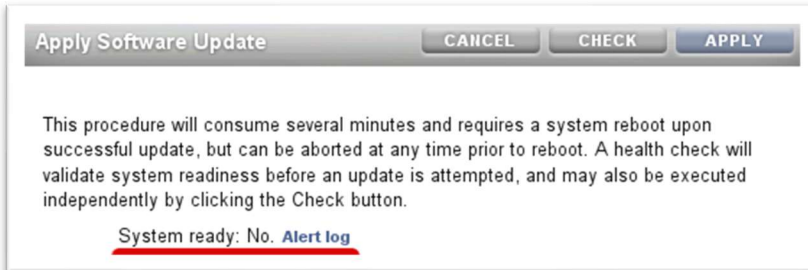
1. Navigate to Maintenance → System.
2. Click on the arrow icon next to the desired system update.

Software Updates		Check for updates on My Oracle Support	
VERSION	RELEASE DATE	RELEASE NAME	STATUS
2013.06.05.8.62,1-1.3	2023-10-5 14:41:46	OS8.8.62	Current system software
2013.06.05.8.63,1-1.1	2023-10-30 09:50:17	OS8.8.63	Uploaded at: 2023-12-5 16:12:05

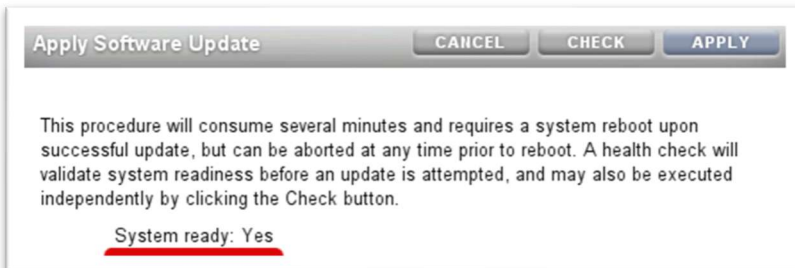
3. Click CHECK to proceed with the health check. Do *not* click APPLY because that will begin the upgrade process.



4. When the health check encounters an issue preventing upgrade, it will report that the system is not ready and post an event to the alert log. To correct the problem, click on the Alert log link.



5. If the health check does not encounter any issues, it will report that the system is ready.



Command Line Interface

1. `zfs:> maintenance system updates`
2. `zfs:maintenance system updates> show`

Updates:

UPDATE	DATE	RELEASE NAME	STATUS
ak-nas@2013.06.05.8.61,1-1.2	2023-9-14 23:37:50	OS8.8.61	previous
ak-nas@2013.06.05.8.62,1-1.3	2023-10-5 18:41:46	OS8.8.62	current
ak-nas@2013.06.05.8.63,1-1.1	2023-10-30 13:50:17	OS8.8.63	waiting

3. `zfs:maintenance system updates> select ak-nas@2013.06.05.8.63,1-1.1`
4. `zfs:maintenance system updates ak-nas@2013.06.05.8.63,1-1.1> 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?

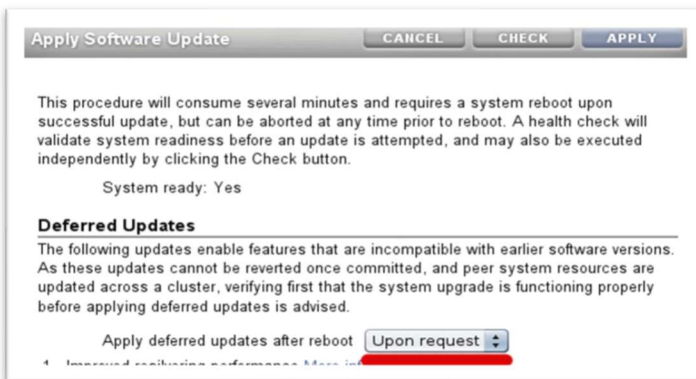
Are you sure? (Y/N) **Y**
Healthcheck running ...

Healthcheck completed. Conditions were reported which would cause an attempted update to this media to abort; see the alert log for details.

Deferred Updates

The majority of new features and bug fixes introduced in a system update are backwards-compatible and allow you to roll back to previous updates without issue. These are applied automatically by the storage management software. On occasion, a deferred update will be included as part of the update package. This is optional functionality that must be manually enabled by a storage administrator. It is necessary to enable a deferred update on only one cluster controller. The change is automatically propagated to the peer controller. If multiple deferred updates are available, you cannot apply individual deferred updates; you must apply them all. It is recommended that you read the release notes prior to accepting deferred updates so you understand their effects on your system.

Browser User Interface



When installing a system update that contains a deferred update, a new option appears on the confirmation screen. You may choose either “Upon request” or “Automatically”. It is recommended to choose “Upon request” and apply the deferred updates during the post-update process.

Command Line Interface

1. `zfs:> maintenance system updates`
2. `zfs:maintenance system updates> show`
 Updates:

UPDATE	DATE	RELEASE NAME	STATUS
ak-nas@2013.06.05.8.61,1-1.2	2023-9-14 23:37:50	OS8.8.61	previous
ak-nas@2013.06.05.8.62,1-1.3	2023-10-5 18:41:46	OS8.8.62	current
ak-nas@2013.06.05.8.63,1-1.1	2023-10-30 13:50:17	OS8.8.63	waiting
3. `zfs:maintenance system updates> select ak-nas@2013.06.05.8.63,1-1.1`
4. `zfs:maintenance system updates ak-nas@2013.06.05.8.63,1-1.1> set update_deferred=onrequest`

Upgrading

Before upgrading, activate a console on each controller to view the entire progress of the procedure. This is accessed either directly with a serial connection or by using the service processor, or Oracle Integrated Lights Out Manager (Oracle ILOM), of each storage controller. Use `ssh` to login.

1. `$ ssh -l root <zfs-ilom-ip-address>`
2. `-> start /SP/console`
 Are you sure you want to start /SP/console (y/n)? **y**

Standalone Controller

A standalone Oracle ZFS Storage Appliance is a controller that is not clustered and has no adjacent cluster peer. Follow these directions only if you are using a standalone controller.

Note that upgrading a standalone controller will cause an interruption of all services and announcing a maintenance window is highly recommended.

Upgrade the Controller

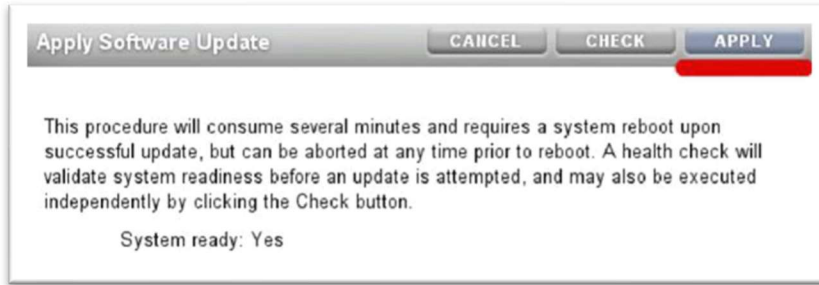
Browser User Interface

1. Navigate to Maintenance → System.
2. Click on the arrow icon next to the desired system update.

Software Updates		Check for updates on My Oracle Support	
VERSION ▲	RELEASE DATE	RELEASE NAME	STATUS
2013.06.05.8.62,1-1.3	2023-10-5 14:41:46	OS8.8.62	Current system software
2013.06.05.8.63,1-1.1	2023-10-30 09:50:17	OS8.8.63	Uploaded at: 2023-12-5 16:12:05 

3. (Optional) Run Health Check on the first controller. Refer to the section [Health Check](#).

4. Select Apply in the left side of the navigation bar while still within **Maintenance** → **System**.



The upgrade including a reboot will be executed. The progress can be observed until the reboot starts.



Command Line Interface

1. `zfs:> maintenance system updates`
2. `zfs:maintenance system updates> show`
 Updates:

UPDATE	DATE	RELEASE NAME	STATUS
ak-nas@2013.06.05.8.61,1-1.2	2023-9-14 23:37:50	OS8.8.61	previous
ak-nas@2013.06.05.8.62,1-1.3	2023-10-5 18:41:46	OS8.8.62	current
ak-nas@2013.06.05.8.63,1-1.1	2023-10-30 13:50:17	OS8.8.63	waiting
3. `zfs:maintenance system updates> select ak-nas@2013.06.05.8.63,1-1.1`
`zfs:> maintenance system updates ak-nas@2013.06.05.8.63,1-1.1> 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**

When using the CLI to manage the upgrade of the ZFS Storage Appliance, then the upgrade can also be observed on the console until the reboot will be initiated.

```
Are you sure? (Y/N)
Updating from ... ak-nas@2013.06.05.8.62,1-1.3
Loading media ... done.
Selecting alternate product ... SUNW,ankimo
```

Monitor Firmware Updates

The following hardware components can be updated with new firmware:

- Controller SAS HBA
- Disk shelf IOM
- Data disk
- Read cache device
- Write flash accelerator

Each update event will be held in either a Pending, In Progress, or Failed state. Contact Oracle Support if a Failed state is reported. These firmware updates can be monitored using the browser user interface or the command-line interface.

Browser User Interface

Navigate to Maintenance → System and click on Firmware Updates.

Command Line Interface

```
zfs:> maintenance system updates firmware show
```

Run Health Check

Refer to the section [Health Check](#).

Clustered Controllers

A clustered Oracle ZFS Storage Appliance has two storage controllers that ensure high availability during the upgrade process. Do not use the following procedures if you have a standalone controller.

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.

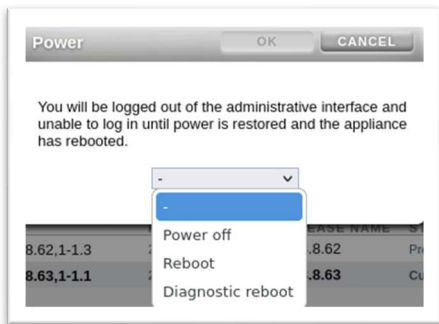
Upgrade First Controller

Browser User Interface

1. Log into controller A and failover its resources to controller B by performing a reboot by clicking on the power icon in the upper left section of the navigation bar.



2. Select Reboot from the drop-down menu and click OK to confirm.




Do not select Diagnostic Reboot. Wait for the reboot to complete before proceeding, then log back in to controller A.

3. Navigate to Maintenance → System.

- Click on the arrow icon next to the desired system update.

Software Updates		Check for updates on My Oracle Support	
VERSION ▲	RELEASE DATE	RELEASE NAME	STATUS
2013.06.05.8.62,1-1.3	2023-10-5 14:41:46	OS8.8.62	Current system software
2013.06.05.8.63,1-1.1	2023-10-30 09:50:17	OS8.8.63	Uploaded at: 2023-12-5 16:12:05 

- (Optional) Run Health Check on the first controller. Refer to the section [Health Check](#).
- Click APPLY to begin the upgrade process.



- Wait for controller A to fully reboot, and log back in.
- Navigate to Configuration → Cluster and verify that the first controller is in the **"Ready (waiting for failback)"** state.

Command Line Interface

- `zfs:> maintenance system updates`
- `zfs:maintenance system updates> show`


```
Updates:
UPDATE                                DATE                                RELEASE NAME    STATUS
ak-nas@2013.06.05.8.61,1-1.2          2023-9-14 23:37:50             OS8.8.61       previous
ak-nas@2013.06.05.8.62,1-1.3          2023-10-5 18:41:46             OS8.8.62       current
ak-nas@2013.06.05.8.63,1-1.1          2023-10-30 13:50:17             OS8.8.63       waiting
```
- `zfs:maintenance system updates> select ak-nas@2013.06.05.8.63,1-1.1`
- (Optional) Run Health Check on controller A. Refer to the section [Health Check](#).
- `zfs:> maintenance system updates ak-nas@2013.06.05.8.63,1-1.1> 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**

When using the CLI to manage the upgrade of the ZFS Storage Appliance, then the upgrade can also be observed on the console until the reboot will be initiated.

```
Are you sure? (Y/N)
Updating from ... ak-nas@2013.06.05.8.62,1-1.3
Loading media ... done.
Selecting alternate product ... SUNW,ankimo
```

- `zfs:> configuration cluster show`


```
state = AKCS_STRIPPED
description = Ready (waiting for failback)
```

Issue Failback on the Second Controller

If the controllers were in an Active / Active configuration before updating, perform a failback operation to return them to that state. This is not necessary if you have an Active / Passive configuration.

Browser User Interface

1. Navigate to Configuration → Cluster to verify that the second controller is in the “**Active (takeover completed)**” state.
2. Click FAILBACK.



Command Line Interface

1. `zfs:> configuration cluster show`
`state = AKCS_OWNER`
`description = Active (takeover completed)`
2. `zfs:> 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`

Upgrade the Second Controller


Browser User Interface

1. Navigate to Maintenance → System on controller B.
2. Click on the arrow icon next to the desired system update.

Software Updates			Check for updates on My Oracle Support
VERSION ▲	RELEASE DATE	RELEASE NAME	STATUS
2013.06.05.8.62,1-1.3	2023-10-5 14:41:46	OS8.8.62	Current system software
2013.06.05.8.63,1-1.1	2023-10-30 09:50:17	OS8.8.63	Uploaded at: 2023-12-5 16:12:05 

3. (Optional) Run Health Check on the first controller. Refer to the section [Health Check](#).

- Click APPLY to begin the upgrade process.



- Wait for controller B to fully reboot, and log back in.

Command-Line Interface

- `zfs:> maintenance system updates`
- `zfs:maintenance system updates> show`
 Updates:

UPDATE	DATE	RELEASE NAME	STATUS
ak-nas@2013.06.05.8.61,1-1.2	2023-9-14 23:37:50	OS8.8.61	previous
ak-nas@2013.06.05.8.62,1-1.3	2023-10-5 18:41:46	OS8.8.62	current
ak-nas@2013.06.05.8.63,1-1.1	2023-10-30 13:50:17	OS8.8.63	waiting
- `zfs:maintenance system updates> select ak-nas@2013.06.05.8.63,1-1.1`
- (Optional) Run Health Check on controller A. Refer to the section [Health Check](#).
- `zfs:> maintenance system updates ak-nas@2013.06.05.8.63,1-1.1> 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**

When using the CLI to manage the upgrade of the ZFS Storage Appliance, then the upgrade can also be observed on the console until the reboot will be initiated.

```
Are you sure? (Y/N)
Updating from ... ak-nas@2013.06.05.8.62,1-1.3
Loading media ... done.
Selecting alternate product ... SUNW,ankimo
```

- `zfs:> configuration cluster show`
 state = AKCS_STRIPPED
 description = Ready (waiting for failback)

Monitor Firmware Updates

An update package may contain firmware updates for on-board hardware components:

- Disk shelf IOM
- Data disk
- Read cache device
- Write flash accelerator
- Integrated Lights Out Module (ILOM)

Firmware updates will only occur after both controllers have been updated but when the cluster is in an Active/Stripped mode. At this point the Active controller will push the firmware updates to both heads.

Each update event will be held in either a Pending, In Progress, or Failed state. Contact Oracle Support if a Failed state is reported. These firmware updates can be monitored using the browser user interface or the command-line interface.

Browser User Interface

Navigate to Maintenance → System and click Firmware Updates to check the update counter.

Command Line Interface

```
zfs:> maintenance system updates firmware show
```

Check both controllers to verify that any firmware updates that were run have completed.


Note: it is important to wait until all firmware updates have been completed before continuing to the next step.

Issue Failback on the First Controller

If the controllers were in an Active / Active configuration before updating, perform a failback operation to return them to that state. This is not necessary if you have an Active / Passive configuration.

Browser User Interface

1. On controller A, navigate to Configuration → Cluster to verify that the second controller is in the “**Active (takeover completed)**” state.
2. Verify all firmware updates are complete. Refer to section “[Monitor Firmware Updates on First Controller.](#)”
Note: Do not begin the next step until all firmware updates are complete.
3. Click FAILBACK.



Command Line Interface

1. `zfs:> configuration cluster show`

```
state = AKCS_OWNER
description = Active (takeover completed)
```
2. `zfs:> 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.

From the Maintenance menu on each controller, select Hardware and verify that all disks are online - all lights should be green.

Verify there are no controller and disk shelf component errors - all lights should be green. An amber light indicates a component error.

If any components have errors, check for pool errors: From the Configuration menu, select Storage, and check the STATUS and ERRORS columns for each pool.

Pools should be online and have no errors.

Post Upgrade

Your environment needs to be returned to its original state after the upgrade process has completed. This section outlines the steps needed, which should be included in the maintenance window.

Deferred Updates

If “Upon request” was chosen during the initial system update sequence, deferred updates can be applied after upgrade.

Browser User Interface

1. Navigate to Maintenance → System.
2. Click APPLY next to Deferred Updates.

Deferred Updates APPLY

The following updates enable features that are incompatible with earlier software versions. As these updates cannot be reverted once committed, and peer system resources are updated across a cluster, verifying first that the system upgrade is functioning properly before applying deferred updates is advised.

1. Improved resilvering performance [More info ...](#)

The appliance is currently configured as part of a cluster. The cluster peer may have shared resources for which deferred updates are available. After all updates are completed, check both cluster peers for any deferred updates.

Command Line Interface

1. zfs:> **maintenance system updates**
2. zfs:maintenance system updates> **show**

```

Updates:
UPDATE                                DATE                                RELEASE NAME    STATUS
ak-nas@2013.06.05.8.61,1-1.2         2023-9-14 23:37:50             OS8.8.61       previous
ak-nas@2013.06.05.8.62,1-1.3         2023-10-5 18:41:46             OS8.8.62       previous
ak-nas@2013.06.05.8.63,1-1.1         2023-10-30 13:50:17             OS8.8.63       current

Deferred updates:

The following updates enable features that are incompatible with earlier
software versions. As these updates cannot be reverted once committed, and peer
system resources are updated across a cluster, verifying first that the system
upgrade is functioning properly before applying deferred updates is advised.

1. Support for associating multiple initiator groups with a LUN. Applying this
   update may disrupt replication. See the online help before applying this
   update.

2. Better I/O throughput

```
3. zfs:maintenance system updates> **apply**

```

Applying deferred updates will prevent rolling back to previous versions of
software.

Are you sure? (Y/N) Y

```

Restart Data Services

Regardless of whether you have exclusively disruptive or non-disruptive protocols in your environment, you should check each attached device for storage connectivity at the conclusion of an upgrade. It may be necessary to remount network shares and restart data services on these hosts.

Re-enable any non-critical data services previously disabled in the upgrade process. Non-critical services may include NDMP, shadow migration, replication, or others.

Backup Configuration Again

Because system configuration backups are OS version specific, it is recommended that another backup be taken after the upgrade process. Follow the process in [Backup Configuration](#) to create and save this backup.

Rollback

A rollback is a procedure that brings your storage controller back to a previous system update. During an upgrade, a snapshot is taken that retains the original system settings. In the event of a rollback, this snapshot is invoked and will override the currently running management properties. This means that a rollback will not necessarily retain all updates made to the storage controller. However, this does not affect any changes made to the storage pools themselves. Your data will remain intact during upgrades and rollbacks.

There are two types of rollbacks: standard and fail-safe. Use the fail-safe rollback if an update failed and the standard rollback is not successful.

Standard Rollback

Use the following directions to execute a typical rollback.

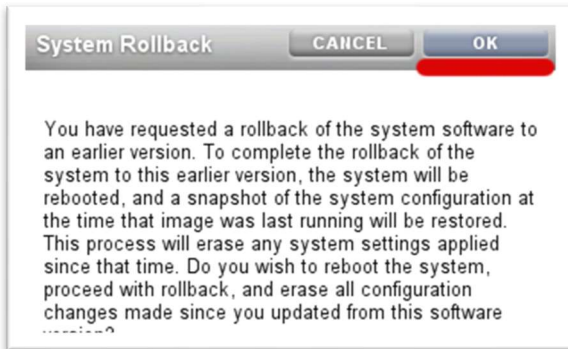
Browser User Interface

1. Navigate to Maintenance → System.
2. Click on the circular arrow icon next to the previous software update.



Software Updates				Check for updates on My Oracle Support
VERSION ▲	RELEASE DATE	RELEASE NAME	STATUS	
2013.06.05.8.62,1-1.3	2023-10-5 14:41:46	OS8.8.62	Previous system software	
2013.06.05.8.63,1-1.1	2023-10-30 09:50:17	OS8.8.63	Current system software	

3. Click OK to confirm.



Command Line Interface

1. zfs:> **maintenance system updates**
2. zfs:maintenance system updates> **show**
 Updates:

UPDATE	DATE	RELEASE NAME	STATUS
ak-nas@2013.06.05.8.61,1-1.2	2023-9-14 23:37:50	OS8.8.61	previous
ak-nas@2013.06.05.8.62,1-1.3	2023-10-5 18:41:46	OS8.8.62	previous
ak-nas@2013.06.05.8.63,1-1.1	2023-10-30 13:50:17	OS8.8.63	current
3. zfs:maintenance system updates> **select ak-nas@2013.06.05.8.62,1-1.3**
4. zfs:maintenance system updates ak-nas@2013.06.05.8.62,1-1.3> **rollback**
 You have requested a rollback of the system software to an earlier version. To complete the rollback of the system to this earlier version, the system will be rebooted, and a snapshot of the system configuration at the time that image was last running will be restored. This process will erase any system settings applied since that time. Do you wish to reboot the system, proceed with rollback, and erase all configuration changes made since you updated from this software version?

 Are you sure? (Y/N) **Y**

Fail-safe Rollback

In the event of a failed update, a fail-safe rollback may be necessary to bring your storage controller back online. This can only be done from a serial console. Refer to the [Upgrading](#) section for instructions on how to access the console.

During a reboot sequence, you can access previous updates from the GRUB menu. Use the arrow keys on your keyboard to select the desired update and then press Return to initiate the rollback.

```

+-----+
| Oracle ZFS Storage ZS3-4 2013.06.05.2.0,1-1.9
| Oracle ZFS Storage ZS3-4 2013.06.05.1.1,1-1.2
+-----+
  
```

Appendix A – Oracle ZFS Storage Appliance Standalone Planning Sheet

Storage Controller Information	
Controller Type	
Hostname	
Serial Number	
Current Software Version	
Updated Software Version	

Upgrade Sequence	Start Time	End Time	Total	Comments
1 Pre-Upgrade				
1.1 Upload Latest System Update				
1.2 Remove Older System Updates				
1.3 Download Backup Configuration				
1.4 Verify No Disk Events				
1.5 Run Health Check				
1.6 Prepare Environment for Downtime				
2 Upgrade				
2.1 Issue Upgrade on Storage Controller				
2.2 Monitor Firmware Updates After Reboot				
3 Post-Upgrade				
3.1 Run Health Check				
3.2 Apply Deferred Updates (optional)				
3.3 Restart Environment Data Services				

Appendix B – Oracle ZFS Storage Appliance Clustered Planning Sheet

Storage Controller Information	
Controller Type	
Controller 1 Hostname	
Controller 1 Serial Number	
Controller 2 Hostname	
Controller 2 Serial Number	
Current Software Version	
Updated Software Version	

Upgrade Sequence	Start Time	End Time	Total	Comments
1 Pre-Upgrade				
1.1 Upload Latest System Update				
1.2 Remove Older System Updates				
1.3 Download Backup Configuration				
1.4 Check Network Interfaces				
1.5 Verify No Disk Events				
1.6 Run Health Check				
1.7 Prepare Environment				
2 Upgrade				
2.1 Upgrade Controller 1				
2.2 Run Health Check on Controller 1				
2.3 Monitor Firmware Updates on Controller 1				
2.4 Issue Failback on Controller 2				
2.5 Upgrade Controller 2				
2.6 Run Health Check on Controller 2				
2.7 Monitor Firmware Updates on Controller 2				
2.8 Issue Failback on Controller 1				
3 Post-Upgrade				
3.1 Final Health Check (both controllers)				
3.2 Apply Deferred Updates (optional)				
3.3 Restart Environment Data Services				

Connect with us

Call +1.800.ORACLE1 or visit oracle.com. Outside North America, find your local office at: oracle.com/contact.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2023, Oracle and/or its affiliates. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

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

Best Practices for Upgrading Oracle ZFS Storage Appliance v2.0 December 2023

Author: Joe Hartley, Paul Johnson, Ulrich Conrad