ORACLE

ZFS STORAGE APPLIANCE

# Best Practices for Upgrading the Oracle ZFS Storage Appliance

ORACLE WHITE PAPER | APRIL 2018

ORACLE

## Table of Contents

Introduction 1

About System Updates 2

System Update Version 2

Retrieving a System Update 2

Remove Older System Updates 3

Understanding Release Notes and Upgrade Requirements 4

Planning an Upgrade 5

Data Services for Clustered Controllers 5

Disruptive Protocols 5

Non-disruptive Protocols 5

Maintenance Window 6

Network Interfaces 6

Data Interfaces 6

Management Interfaces 7

Backup Configuration 8

Disk Events 8

Health Check 9

Deferred Updates 11

Upgrading 12

Standalone Controller 12

Upgrade the Controller 12

Monitor Firmware Updates 12

Run Health Check 13

Clustered Controllers 13

Upgrade First Controller 13

Monitor Firmware Updates on First Controller 14

Issue Failback on the Second Controller 14

Upgrade the Second Controller 15

Post Upgrade 17

Deferred Updates 17

Restart Data Services 17

Rollback 18

Standard Rollback 18

Fail-safe Rollback 19

Appendix A – Oracle ZFS Storage Appliance Standalone Planning Sheet 20

Appendix B – Oracle ZFS Storage Appliance Clustered Planning Sheet 21

## Introduction

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.

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, iSCSI, and Fibre Channel connectivity to allow storage and network administrators to design a deployment strategy that seamlessly integrates into an existing datacenter.

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

### 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. To understand this format, look at the following browser-user interface example:

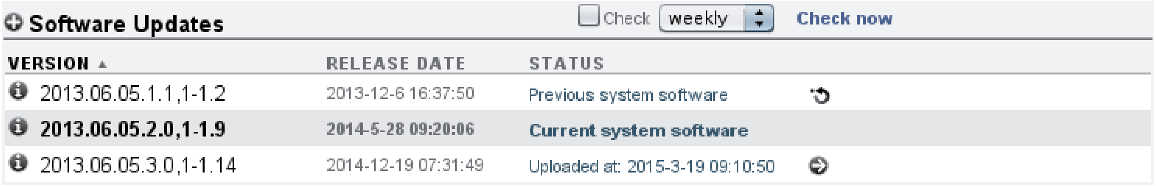

| VERSION ▲ | RELEASE DATE | STATUS |
|---|---|---|
| 2013.06.05.1.1,1-1.2 | 2013-12-6 16:37:50 | Previous system software |
| **2013.06.05.2.0,1-1.9** | **2014-5-28 09:20:06** | **Current system software** |
| 2013.06.05.3.0,1-1.14 | 2014-12-19 07:31:49 | Uploaded at: 2015-3-19 09:10:50 |

The current running version on this system is 2013.06.05.2.0,1-1.9. Here is a breakdown of this update string:

**2013.06.05** Refers to the date the underlying software was synchronized with the base Oracle Solaris operating system. This indicates a major release, but does not indicate when the software was initially released. This is stated in the RELEASE DATE field.

**2.** Refers to the minor version. A minor release is usually issued a few times a year.

**0,** Refers to the micro version. The micro release is issued as often as every two weeks.

**1-1.9** Refers to the build number. This is only relevant to Oracle Engineering.

### Retrieving a System Update

Next to Software Updates, you can click “Check now,” or you can schedule the checks by selecting the checkbox and an interval of daily, weekly, or monthly. When a new update is found, “Update available for download” is displayed under STATUS, which is also a direct download link to My Oracle Support. Click the download icon next to “Update available for download” and specify a pathname on your desktop for the update media to reside.

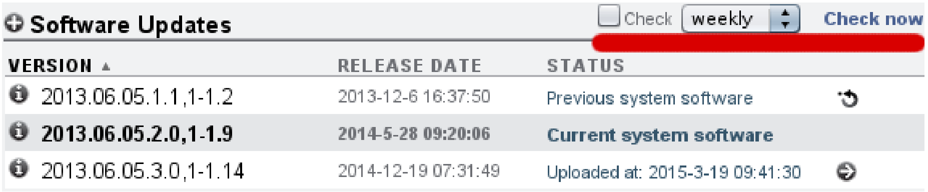

| VERSION ▲ | RELEASE DATE | STATUS |
|---|---|---|
| 2013.06.05.1.1,1-1.2 | 2013-12-6 16:37:50 | Previous system software |
| **2013.06.05.2.0,1-1.9** | **2014-5-28 09:20:06** | **Current system software** |
| 2013.06.05.3.0,1-1.14 | 2014-12-19 07:31:49 | Uploaded at: 2015-3-19 09:41:30 |

Use the following browser user interface or command-line interface procedures for uploading a system update onto the Oracle ZFS Storage Appliance. 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.

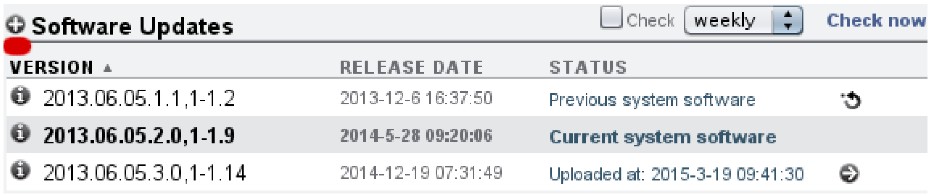

| VERSION ▲ | RELEASE DATE | STATUS |
|---|---|---|
| 2013.06.05.1.1,1-1.2 | 2013-12-6 16:37:50 | Previous system software |
| **2013.06.05.2.0,1-1.9** | **2014-5-28 09:20:06** | **Current system software** |
| 2013.06.05.3.0,1-1.14 | 2014-12-19 07:31:49 | Uploaded at: 2015-3-19 09:41:30 |

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

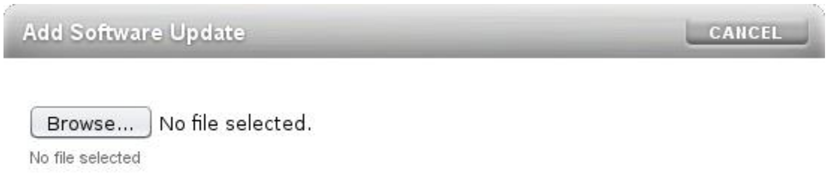

No file selected.

No file selected

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.gz`
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, maintain no more than three updates at any given time.

Browser User Interface

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

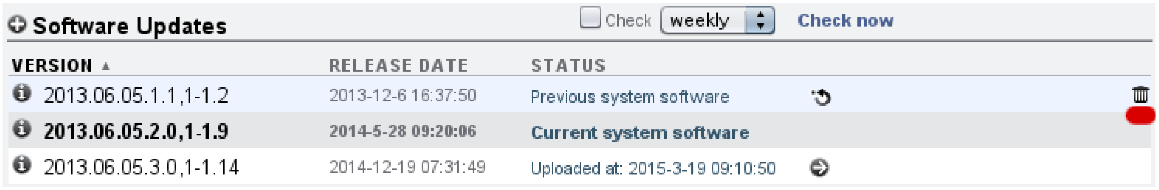

| VERSION ▲ | RELEASE DATE | STATUS |
|---|---|---|
| 2013.06.05.1.1,1-1.2 | 2013-12-6 16:37:50 | Previous system software |
| **2013.06.05.2.0,1-1.9** | **2014-5-28 09:20:06** | **Current system software** |
| 2013.06.05.3.0,1-1.14 | 2014-12-19 07:31:49 | Uploaded at: 2015-3-19 09:10:50 |

3. Click OK to confirm.

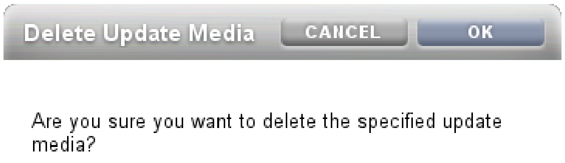

Are you sure you want to delete the specified update media?

Command Line Interface

1. `zfs:> **maintenance system updates**`

```
zfs:> maintenance system updates
```

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

UPDATE                                  DATE                      STATUS
ak-nas@2013.06.05.1.1,1-1.2             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.
```
zfs:maintenance system updates> destroy ak-nas@2013.06.05.1.1,1-1.2
This will destroy the update "ak-nas@2013.06.05.1.1,1-1.2". Are you sure? (Y/N) Y
```

## Understanding Release Notes and Upgrade Requirements

Each system update comes with a set of release notes that discuss new features and bug fixes. These can be obtained by navigating to Maintenance → System and clicking on the information icon next to each update.

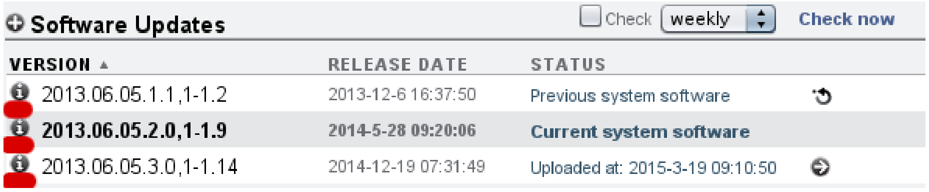

**Software Updates** □ Check weekly **Check now**

| VERSION ▲ | RELEASE DATE | STATUS |
|---|---|---|
| 2013.06.05.1.1,1-1.2 | 2013-12-6 16:37:50 | Previous system software |
| **2013.06.05.2.0,1-1.9** | **2014-5-28 09:20:06** | **Current system software** |
| 2013.06.05.3.0,1-1.14 | 2014-12-19 07:31:49 | Uploaded at: 2015-3-19 09:10:50 |

A link is provided to My Oracle Support, as well as instructions for finding the appropriate release notes.

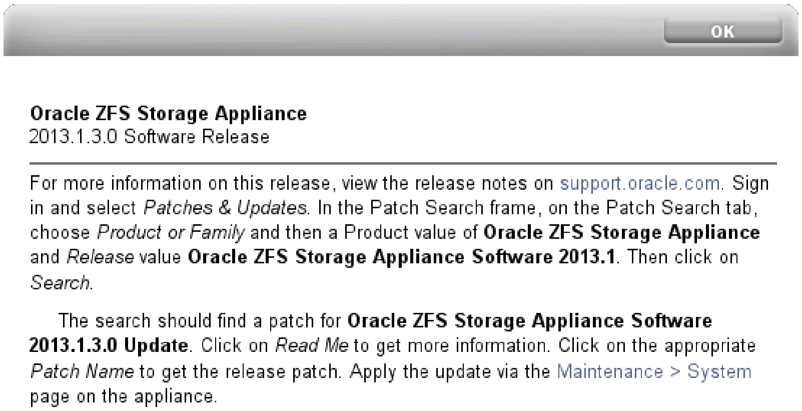

**Oracle ZFS Storage Appliance**
2013.1.3.0 Software Release

For more information on this release, view the release notes on support.oracle.com. Sign in and select *Patches & Updates*. In the Patch Search frame, on the Patch Search tab, choose *Product or Family* and then a Product value of **Oracle ZFS Storage Appliance** and *Release* value **Oracle ZFS Storage Appliance Software 2013.1**. Then click on *Search*.

The search should find a patch for **Oracle ZFS Storage Appliance Software 2013.1.3.0 Update**. Click on *Read Me* to get more information. Click on the appropriate *Patch Name* to get the release patch. Apply the update via the Maintenance > System page on the appliance.

Read the release notes 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.

# Planning an Upgrade

Oracle ZFS Storage Appliance 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 Appliance 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 `nfsstat -m` on all attached devices. |
| Fibre Channel | Verify all attached devices have multiple active and standby paths. This can be verified with `multipath -ll` on Linux, `mpathadm show LU` 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 `multipath -ll` on Linux, `mpathadm show LU` 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 `multipath -ll` on Linux and `mpathadm show LU` on Solaris. |
| SRP | Verify all attached devices have multiple active and standby paths. This can be verified with `multipath -ll` 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 data interfaces for clustered controllers are 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.

Browser User Interface

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

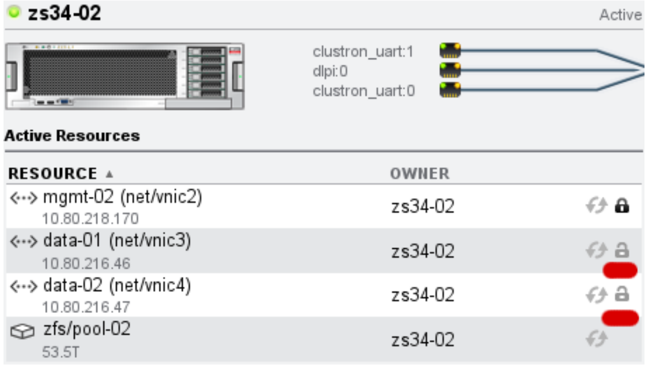

zs34-02 Active

clustron_uart:1
dlpi:0
clustron_uart:0

Active Resources

| RESOURCE ▲ | OWNER |
|---|---|
| mgmt-02 (net/vnic2)<br>10.80.218.170 | zs34-02 |
| data-01 (net/vnic3)<br>10.80.216.46 | zs34-02 |
| data-02 (net/vnic4)<br>10.80.216.47 | zs34-02 |
| zfs/pool-02<br>53.5T | zs34-02 |

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

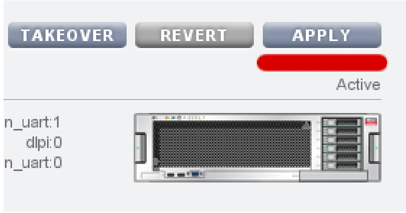

Command Line Interface

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

RESOURCE        OWNER           TYPE       LABEL         CHANGES  DETAILS
net/vnic2       zs34-02         private    mgmt-02       no       10.80.218.170
net/vnic3       zs34-02         singleton  data-01       no       10.80.216.46
net/vnic4       zs34-02         singleton  data-02       no       10.80.216.47
zfs/pool-01     zs34-01         singleton                no
zfs/pool-02     zs34-02         singleton                no       53.5T
```

2.
```
zfs:configuration cluster resources> select net/vnic3 set type=singleton
                                          type = singleton
```

3.
```
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.

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.

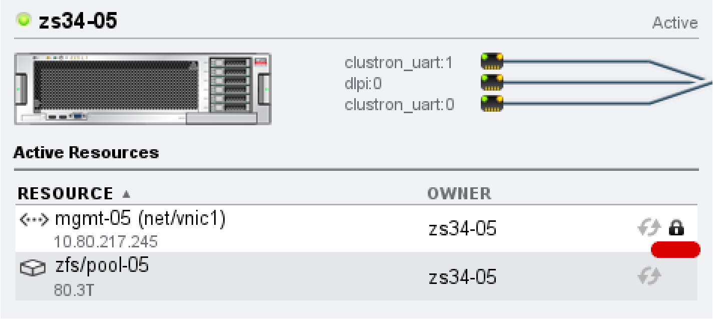

| RESOURCE ▲ | OWNER |
|---|---|
| mgmt-05 (net/vnic1)<br>10.80.217.245 | zs34-05 |
| zfs/pool-05<br>80.3T | zs34-05 |

3. Click APPLY for the change to take effect.

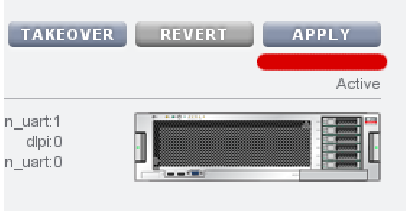

Command Line Interface

1.
```
zfs:> configuration cluster resources
```

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

RESOURCE        OWNER           TYPE       LABEL         CHANGES  DETAILS
net/vnic1       zs34-05         private    mgmt-05       no       10.80.217.245
zfs/pool-05     zs34-05         singleton                no       80.3T
```

3.
```
zfs:configuration cluster resources> select net/vnic1 set type=private
                                          type = private
```

4. 
   ```
   zfs:configuration cluster resources> commit
   ```

## 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.

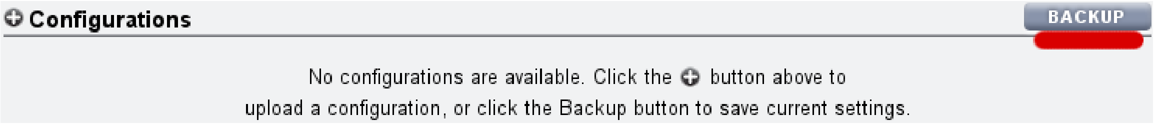

No configurations are available. Click the ⊕ button above to upload a configuration, or click the Backup button to save current settings.

3. Create a comment for the configuration.

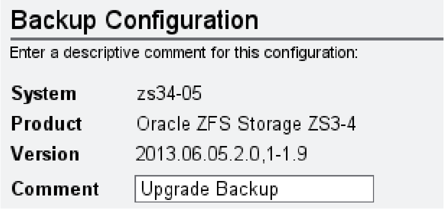

Backup Configuration

Enter a descriptive comment for this configuration:

| | |
|---|---|
| System | zs34-05 |
| Product | Oracle ZFS Storage ZS3-4 |
| Version | 2013.06.05.2.0,1-1.9 |
| Comment | Upgrade Backup |

4. Click APPLY to create the backup configuration.
5. Download the configuration locally to your computer by clicking on the download icon.

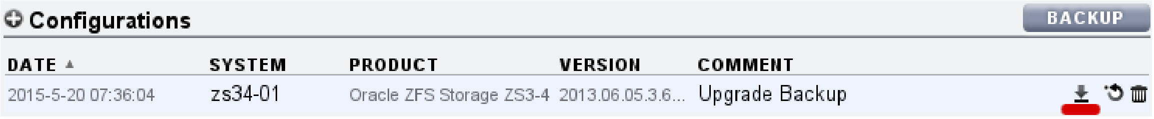

| DATE ▲ | SYSTEM | PRODUCT | VERSION | COMMENT |
|---|---|---|---|---|
| 2015-5-20 07:36:04 | zs34-01 | Oracle ZFS Storage ZS3-4 | 2013.06.05.3.6... | Upgrade Backup |

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 disks resilvering events.
3. Verify there are no scrub activities.

Command Line Interface

1.
```
zfs:> configuration storage
```

2.
```
zfs:configuration storage (pool-01)> show
Pools:

POOL             OWNER         DATA PROFILE  LOG PROFILE  STATUS   ERRORS
pool-01          zs34-01       mirror_nspf   log_stripe   online   0

Properties:
                          pool = pool-01
                        status = online
                        errors = 0
                         owner = zs34-01
                       profile = mirror_nspf
                   log_profile = log_stripe
         cache_profile = cache_stripe
                 scrub = scrub in progress for 0h0m, 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.

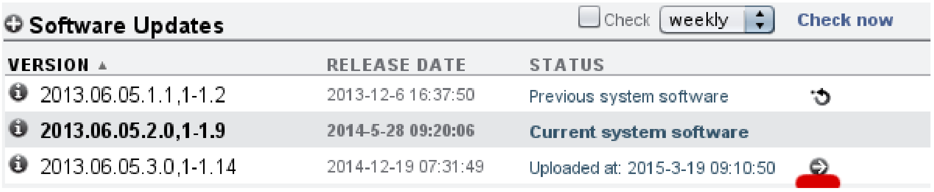

| VERSION ▲ | RELEASE DATE | STATUS |
|---|---|---|
| 2013.06.05.1.1,1-1.2 | 2013-12-6 16:37:50 | Previous system software |
| **2013.06.05.2.0,1-1.9** | **2014-5-28 09:20:06** | **Current system software** |
| 2013.06.05.3.0,1-1.14 | 2014-12-19 07:31:49 | Uploaded at: 2015-3-19 09:10:50 |

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

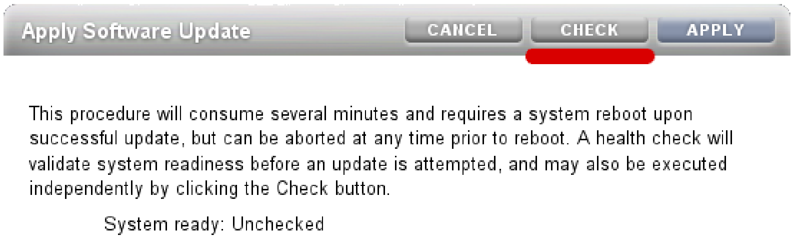

This procedure will consume several minutes and requires a system reboot upon successful update, but can be aborted at any time prior to reboot. A health check will validate system readiness before an update is attempted, and may also be executed independently by clicking the Check button.

System ready: Unchecked

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.

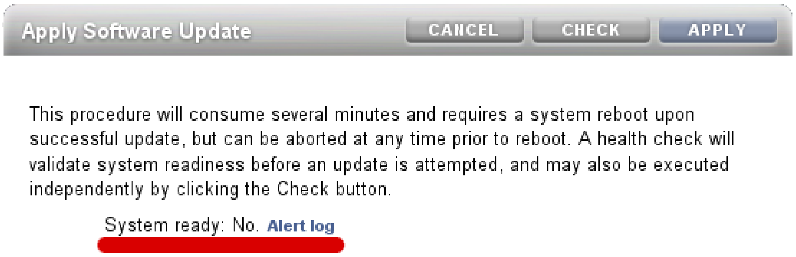

This procedure will consume several minutes and requires a system reboot upon successful update, but can be aborted at any time prior to reboot. A health check will validate system readiness before an update is attempted, and may also be executed independently by clicking the Check button.

System ready: No. Alert log

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

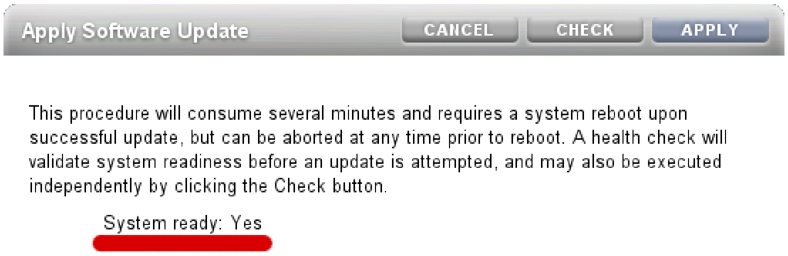

This procedure will consume several minutes and requires a system reboot upon successful update, but can be aborted at any time prior to reboot. A health check will validate system readiness before an update is attempted, and may also be executed independently by clicking the Check button.

System ready: Yes

### Command Line Interface

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

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

UPDATE                                   DATE                      STATUS
ak-nas@2013.06.05.1.1,1-1.2              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. 
```
zfs:maintenance system updates> select ak-nas@2013.06.05.3.0,1-1.14
```

4. 
```
zfs:maintenance system updates ak-nas@2013.06.05.3.0,1-1.14> 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)
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.

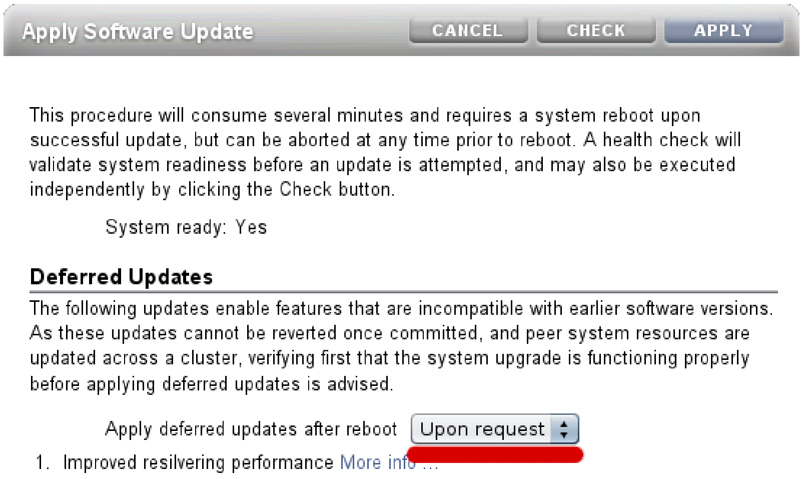

This procedure will consume several minutes and requires a system reboot upon successful update, but can be aborted at any time prior to reboot. A health check will validate system readiness before an update is attempted, and may also be executed independently by clicking the Check button.

System ready: Yes

**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.

Apply deferred updates after reboot: Upon request

1. Improved resilvering performance More info ...

### Command Line Interface

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

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

UPDATE                                     DATE                       STATUS
ak-nas@2011.04.24.6.0,1-0.41               2013-3-11 18:43:29         current
ak-nas@2013.06.05.1.1,1-1.1                2013-12-4 02:53:01         waiting
```

3.
```
zfs:maintenance system updates> select ak-nas@2013.06.05.1.1,1-1.1
```

4.
```
zfs:maintenance system updates ak-nas@2013.06.05.1.1,1-1.1> 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.

### Upgrade the Controller

Browser User Interface

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

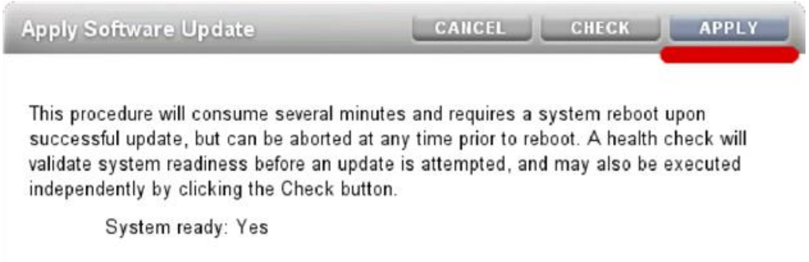

This procedure will consume several minutes and requires a system reboot upon successful update, but can be aborted at any time prior to reboot. A health check will validate system readiness before an update is attempted, and may also be executed independently by clicking the Check button.

System ready: Yes

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

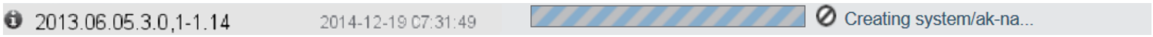

2013.06.05.3.0,1-1.14 2014-12-19 07:31:49 Creating system/ak-na...

Command Line Interface

```
zfs:> maintenance system updates ak-nas@2013.06.05.6.3,1-2.1> 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@2011.04.24.6.0,1-0.41
Loading media  … done.
Selecting alternate product … SUNW,ankimo
```

### Monitor Firmware Updates

The following hardware components can to 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

Navigate to `maintenance system updates firmware show`.

### Run Health Check

Refer to 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.

### Upgrade First Controller

Browser User Interface

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

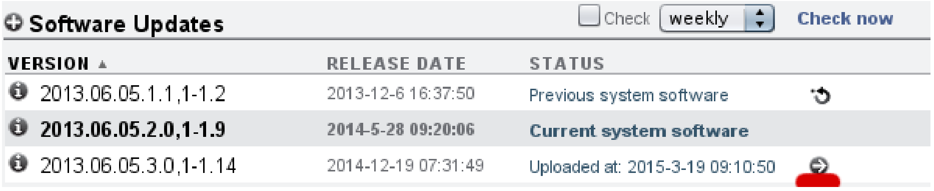

| VERSION ▲ | RELEASE DATE | STATUS |
|---|---|---|
| 2013.06.05.1.1,1-1.2 | 2013-12-6 16:37:50 | Previous system software |
| **2013.06.05.2.0,1-1.9** | **2014-5-28 09:20:06** | **Current system software** |
| 2013.06.05.3.0,1-1.14 | 2014-12-19 07:31:49 | Uploaded at: 2015-3-19 09:10:50 |

3. (Optional) Run Health Check on the first controller. Refer to section Health Check.
4. Click APPLY to begin the upgrade process.

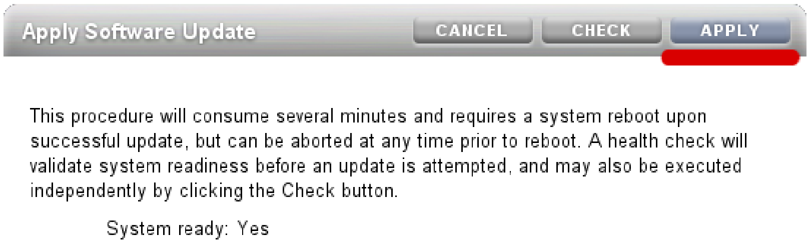

This procedure will consume several minutes and requires a system reboot upon successful update, but can be aborted at any time prior to reboot. A health check will validate system readiness before an update is attempted, and may also be executed independently by clicking the Check button.

System ready: Yes

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

Command Line Interface

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

```
Updates:

UPDATE                                  DATE                    STATUS
ak-nas@2013.06.05.5.1,1-1.2             2015-12-6 23:37:50      previous
ak-nas@2013.06.05.2.0,1-1.9             2016-6-28 15:20:06      current
ak-nas@2013.06.05.6.3,1-2.1             2016-8-19 14:31:49      waiting
```

3. `zfs:maintenance system updates> select ak-nas@2013.06.05.6.3,1-2.1`

4. (Optional) Run Health Check on the first controller. Refer to section Health Check.

5.
   ```
   zfs:maintenance system updates ak-nas@2013.06.05.6.3,1-2.1> upgrade
   This procedure will consume several minutes and requires a system reboot upon
   successful update, but can be aborted with [Control-C] at any time prior to
   reboot. A health check will validate system readiness before an update is
   attempted, and may also be executed independently using the check command.

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

6.
   ```
   zfs:> configuration cluster show

   state = AKCS_STRIPPED

   description = Ready (waiting for failback)
   ```

### Monitor Firmware Updates on First Controller

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 Firmware Updates.

Command Line Interface

Navigate to `maintenance system updates firmware show`.

### 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 want 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.

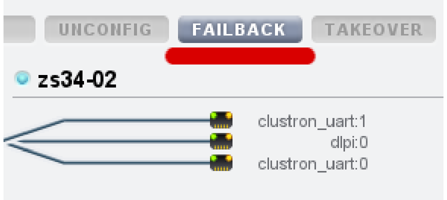

Command Line Interface

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

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

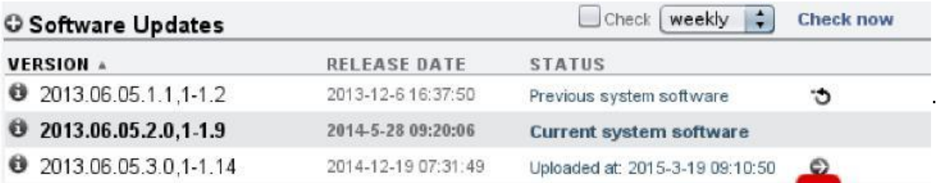

| VERSION ▲ | RELEASE DATE | STATUS |
|---|---|---|
| 2013.06.05.1.1,1-1.2 | 2013-12-6 16:37:50 | Previous system software |
| **2013.06.05.2.0,1-1.9** | **2014-5-28 09:20:06** | **Current system software** |
| 2013.06.05.3.0,1-1.14 | 2014-12-19 07:31:49 | Uploaded at: 2015-3-19 09:10:50 |

3. (Optional) Run Health Check on the second controller. Refer to section Health Check.
4. Click APPLY to begin the upgrade process.

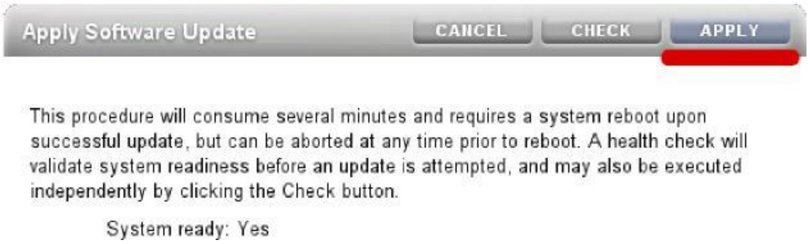

This procedure will consume several minutes and requires a system reboot upon successful update, but can be aborted at any time prior to reboot. A health check will validate system readiness before an update is attempted, and may also be executed independently by clicking the Check button.

System ready: Yes

5. Navigate to Configuration → Cluster to verify that the second controller is in the "Ready (waiting for failback)" state.

» Command-Line Interface

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

```
Updates:

         UPDATE                                    DATE                      STATUS
         ak-nas@2013.06.05.5.1,1-1.2               2015-12-6 23:37:50        previous
         ak-nas@2013.06.05.6.0,1-1.9               2016-6-28 15:20:06        current
         ak-nas@2013.06.05.6.3,1-2.1               2016-8-19 14:31:49        waiting
```

3. `zfs:maintenance system updates>` **`select ak-nas@2013.06.05.6.3,1-2.1`**
4. (Optional) Run Health Check on Second Controller. Refer to section Health Check.
5. `zfs:maintenance system updates ak-nas@2013.06.05.6.3,1-2.1>` **`upgrade`**
6. `zfs:> configuration cluster show`

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

### Monitor Firmware Updates on the Second Controller

Refer to section "Monitor Firmware Updates on First Controller."

### 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 want an Active / Passive configuration.

» Browser User Interface

1. Navigate to Configuration → Cluster to verify that the first 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.

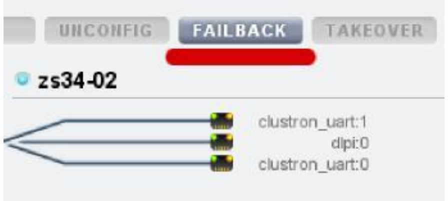

» Command-Line Interface

1. `zfs:> **maintenance system show**`

```
 state = AKCS_STRIPPED
 description = Active (takeover complete)
```

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. `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
```

# 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.

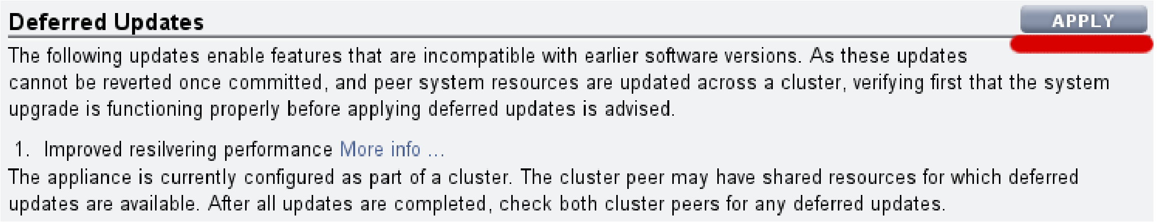

**Deferred Updates**

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                      STATUS
ak-nas@2011.04.24.6.0,1-0.41             2013-3-11 18:43:29        previous
ak-nas@2013.06.05.1.1,1-1.1              2013-12-4 02:53:01        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.

## 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.

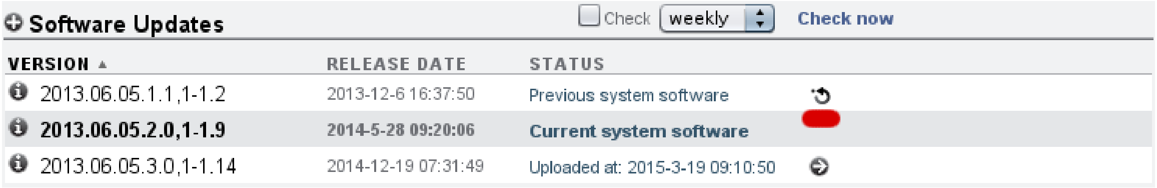

Software Updates

□ Check weekly Check now

| VERSION ▲ | RELEASE DATE | STATUS |
|---|---|---|
| 2013.06.05.1.1,1-1.2 | 2013-12-6 16:37:50 | Previous system software |
| **2013.06.05.2.0,1-1.9** | **2014-5-28 09:20:06** | **Current system software** |
| 2013.06.05.3.0,1-1.14 | 2014-12-19 07:31:49 | Uploaded at: 2015-3-19 09:10:50 |

3. Click OK to confirm.

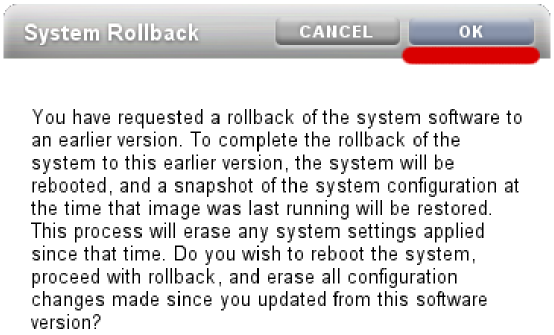

System Rollback

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?

Command Line Interface

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

    UPDATE                                     DATE                   STATUS
    ak-nas@2011.04.24.6.0,1-0.41               2013-3-11 18:43:29     previous
    ak-nas@2013.06.05.1.1,1-1.1                2013-12-4 02:53:01     current
```
3. 
```
zfs:maintenance system updates> select ak-nas@2011.04.24.6.0,1-0.41
```
4. 
```
zfs:maintenance system updates ak-nas@2011.04.24.6.0,1-0.41> 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.

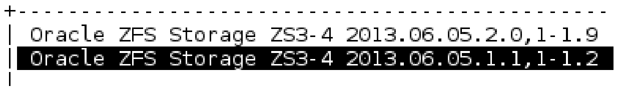

```
+---------------------------------------------
| Oracle ZFS Storage ZS3-4 2013.06.05.2.0,1-1.9
| Oracle ZFS Storage ZS3-4 2013.06.05.1.1,1-1.2
|
```

## 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 | | | | |

**Oracle Corporation, World Headquarters**
500 Oracle Parkway
Redwood Shores, CA 94065, USA

**Worldwide Inquiries**
Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

- blogs.oracle.com/oracle
- facebook.com/oracle
- twitter.com/oracle
- oracle.com

Hardware and Software, **Engineered to Work Together**

Best Practices for Upgrading Oracle ZFS Storage Appliance
April 2018
Author: Paul Johnson, Ulrich Conrad