

Managing Network-Attached Storage Devices in an Oracle® Solaris Cluster 4.4 Environment

ORACLE®

Part No: E69326
April 2019

Part No: E69326

Copyright © 2000, 2019, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

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

Access to Oracle Support

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

Référence: E69326

Copyright © 2000, 2019, Oracle et/ou ses affiliés. Tous droits réservés.

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

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

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

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

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

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

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

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

Accès aux services de support Oracle

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

Contents

- Using This Documentation** 7

- 1 Installing and Maintaining Oracle ZFS Storage Appliances as NAS Devices in an Oracle Solaris Cluster Environment** 9
 - Requirements, Recommendations, and Restrictions for Oracle ZFS Storage Appliance NAS Devices 10
 - Requirements for Oracle ZFS Storage Appliance NAS Devices 10
 - Restrictions for Oracle ZFS Storage Appliance NAS Devices 13
 - Installing an Oracle ZFS Storage Appliance NAS Device in an Oracle Solaris Cluster Environment 13
 - ▼ How to Install an Oracle ZFS Storage Appliance in a Cluster 13
 - Maintaining an Oracle ZFS Storage Appliance NAS Device in an Oracle Solaris Cluster Environment 18
 - ▼ How to Prepare the Cluster for Oracle ZFS Storage Appliance NAS Device Maintenance 19
 - ▼ How to Restore Cluster Configuration After Oracle ZFS Storage Appliance NAS Device Maintenance 20
 - ▼ How to Remove an Oracle ZFS Storage Appliance NAS Device From a Cluster 21
 - ▼ How to Add Oracle ZFS Storage Appliance Directories and Projects to a Cluster 22
 - ▼ How to Remove Oracle ZFS Storage Appliance Directories and Projects and From a Cluster 27

- Index** 31

Using This Documentation

- **Overview** – Describes how to install and configure Network-Attached Storage (NAS) systems in an Oracle Solaris Cluster environment
- **Audience** – Technicians, system administrators, and authorized service providers
- **Required knowledge** – Advanced experience troubleshooting and replacing hardware

Product Documentation Library

Documentation and resources for this product and related products are available at http://www.oracle.com/pls/topic/lookup?ctx=E69294_01.

Feedback

Provide feedback about this documentation at <http://www.oracle.com/goto/docfeedback>.

Installing and Maintaining Oracle ZFS Storage Appliances as NAS Devices in an Oracle Solaris Cluster Environment

This chapter contains procedures about installing and maintaining Oracle ZFS Storage Appliances as NAS devices using NFS file systems in an Oracle Solaris Cluster environment. Before you perform any of the procedures in this chapter, read the entire procedure.

This chapter contains the following procedures:

- [“How to Install an Oracle ZFS Storage Appliance in a Cluster” on page 13](#)
- [“How to Prepare the Cluster for Oracle ZFS Storage Appliance NAS Device Maintenance” on page 19](#)
- [“How to Restore Cluster Configuration After Oracle ZFS Storage Appliance NAS Device Maintenance” on page 20](#)
- [“How to Remove an Oracle ZFS Storage Appliance NAS Device From a Cluster” on page 21](#)
- [“How to Add Oracle ZFS Storage Appliance Directories and Projects to a Cluster” on page 22](#)
- [“How to Remove Oracle ZFS Storage Appliance Directories and Projects and From a Cluster” on page 27](#)

For conceptual information about multihost storage devices, see the [Concepts for Oracle Solaris Cluster 4.4](#).

Note - The instructions for installing and maintaining Oracle ZFS Storage Appliance NAS devices are the same for both global cluster and zone cluster environments, unless stated otherwise. When you perform a procedure or step on a *cluster node*, the node is the global cluster node for a global cluster, or the zone cluster node for a zone cluster.

Requirements, Recommendations, and Restrictions for Oracle ZFS Storage Appliance NAS Devices

This section includes only restrictions and requirements that have a direct impact on the procedures in this chapter. For general support information, contact your Oracle service provider.

Requirements for Oracle ZFS Storage Appliance NAS Devices

This section describes the following requirements.

- [“Requirements When Configuring Oracle ZFS Storage Appliances” on page 10](#)
- [“Requirements When Configuring Oracle ZFS Storage Appliance NAS Devices for the Oracle RAC or HA for Oracle Database Data Service” on page 11](#)
- [“Requirements When Configuring Oracle ZFS Storage Appliance NAS Devices as Quorum Devices” on page 12](#)

Requirements When Configuring Oracle ZFS Storage Appliances

When you configure an Oracle ZFS Storage Appliance, you must meet the following requirements.

- Do not use the default project. File systems created in the default project will not be fenced from failing cluster nodes.
- If you wish for a file system within a project used by the cluster to be protected by cluster fencing, ensure that the Protocols tab of the file system has the `Inherit from project` property selected for NFS. If you do not want the file system to be protected (for example, to allow file system access when cluster nodes are in non-cluster mode), ensure that the `Inherit from project` setting is not selected. This setting can be changed as needed to allow some file systems within a project to be fenced, and other file systems within the same project to not be fenced. When unselecting the `Inherit from project` setting, verify that the file system has the desired NFS exception settings for the IP address of each cluster node. Ensure that NFS file systems created within a project for cluster use with fencing control are set to inherit NFS properties from their parent project.
- For any projects that have file systems to be protected by cluster fencing, perform the following actions:

- In the Protocols tab of the project, set the Share Mode to None or Read only.
- If there are systems outside the cluster that will access the file system in this project, allow them access with the NFS exception entries created for each system. Use the Host entries or Network entries for these systems. Use the Network entry only if the system is on a different subnet than the subnet used by the cluster to access the NAS device.
- The IP address must use the format of xxx.xxx.xxx.xxx/32. Set the Access Mode for the entry to Read/Write and select Root Access for the entry. Explicitly grant access to projects to all nodes in the cluster. Use only network exceptions when granting cluster access. Add exceptions for each public IP address within the cluster that might be used to access the storage, using the format of xxx.xxx.xxx.xxx/32. If a node has multiple active public network adapters, add the IP address of each one.

Note - At run time, Oracle Solaris Cluster dynamically changes the NFS exception entries for the cluster nodes to control when a node has Read/Write access in a cluster membership. The software also controls when a node has Read access because it was evicted from the cluster membership.

- An Oracle ZFS Storage Appliance NAS device must be directly connected (through the same subnet) to all nodes of the cluster.
- The cluster can be connected to multiple public networks to communicate with external systems. However, only one network can directly access a specific Oracle ZFS Storage Appliance device. The IP addresses configured in the NFS exception entries for the cluster nodes exist in that subnet, and the Oracle ZFS Storage Appliance device's network interface that is connected to that network is also configured with an IP address in that subnet. In the Network tab of the Oracle ZFS Storage Appliance Configuration panel, ensure that the subnet's network interface has Allow Administration selected.
- Ensure that the Oracle ZFS Storage Appliance is running a qualified firmware release.

Requirements When Configuring Oracle ZFS Storage Appliance NAS Devices for the Oracle RAC or HA for Oracle Database Data Service

When you configure your Oracle ZFS Storage Appliance NAS device for use with the HA for Oracle Database data service, you must meet the following requirements:

- To guarantee data integrity, configure the Oracle ZFS Storage Appliance NAS device with fencing support.
- You can install Oracle Database and Oracle Clusterware software, as well as place files used by these installations, onto NFS shares from the NAS device, but you must ensure that

the NFS shares used to store the files are mounted with the required mount options. These mount options must be appropriate for the file type. Do not mix file types with different mount requirements on the same NFS share.

Consult your Oracle Database guide or log into My Oracle Support for the most current list of supported files and mount options. After you log into [My Oracle Support](#), click the Knowledge tab and search for Bulletin 359515.1.

When you configure your Oracle ZFS Storage Appliance NAS device for use with Oracle RAC, you must also comply with the requirements listed above.

Note - The ZFS Storage Appliance has been qualified as NAS NFS storage for Oracle RAC with Oracle Solaris Cluster, without NFS fencing support and in a configuration where the Oracle ZFS Storage Appliance and Oracle Solaris Cluster configuration tasks are not required. See [“Requirements When Configuring Oracle ZFS Storage Appliances” on page 10](#). However, starting with Oracle Solaris Cluster 4.0, you should enable the fencing support that is based on the requirements and configuration tasks described in this chapter.

Requirements When Configuring Oracle ZFS Storage Appliance NAS Devices as Quorum Devices

The administrator has the option to create and use iSCSI LUNS on the Oracle ZFS Storage Appliance to be used as quorum devices.

Note - You do not need to configure a NAS device as a quorum device in a zone cluster environment. A quorum device is not required for a zone cluster, because the zone cluster node role in the zone cluster membership directly depends on the global cluster node role on the machine that is hosting the zone cluster node.

When you use an Oracle ZFS Storage Appliance NAS device as a quorum device, you must meet the following requirements:

- When booting the cluster, always boot the Oracle ZFS Storage Appliance NAS device before you boot the cluster nodes.
If you boot the quorum device after booting the cluster nodes, your nodes cannot find the quorum device, and thus cannot count the quorum votes of the quorum device. This lack of quorum votes may result in the partition failing to form a cluster. If that situation occurs, reboot the cluster nodes.
- The Oracle ZFS Storage Appliance NAS device must be located on the same network as the cluster nodes. If an Oracle ZFS Storage Appliance NAS quorum device is not located on the same network as the cluster nodes, the quorum device is at risk of not responding at boot time. This lack of response could cause the cluster nodes to be unable to form a cluster if

they do not acquire enough votes (for example, when one node cannot be booted up). This risk is also present at quorum acquisition time for the cluster to resolve split brain situations, in which case it could cause the cluster to fail to stay up.

When you use an iSCSI LUN from an Oracle ZFS Storage Appliance NAS device as a cluster quorum device, the device appears to the quorum subsystem as a regular SCSI shared disk. The iSCSI connection to the NAS device is completely invisible to the quorum subsystem.

For instructions on adding an Oracle ZFS Storage Appliance NAS quorum device, see [“How to Add an Oracle ZFS Storage Appliance NAS Quorum Device”](#) in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

Restrictions for Oracle ZFS Storage Appliance NAS Devices

The Oracle Solaris Cluster interface for configuring the NFS file systems from the Oracle ZFS Storage Appliance does not support the configuration at the individual file system level. The configuration of such file systems is restricted to the projects in the Oracle ZFS Storage Appliance that contain the file systems.

Installing an Oracle ZFS Storage Appliance NAS Device in an Oracle Solaris Cluster Environment

▼ How to Install an Oracle ZFS Storage Appliance in a Cluster

Note - You can also add an Oracle ZFS Storage Appliance using the Oracle Solaris Cluster Manager (Cluster Manager) browser interface. For Cluster Manager log-in instructions, see [“How to Access Oracle Solaris Cluster Manager”](#) in *Administering an Oracle Solaris Cluster 4.4 Configuration*. After you install the appliance, you can also use Cluster Manager to edit the `Export List` property.

Before You Begin This procedure relies on the following assumptions:

- Your cluster nodes have the operating system and Oracle Solaris Cluster software installed.

- You have administrative access to the Oracle ZFS Storage Appliance.

To perform this procedure, assume the `root` role or a role that provides `solaris.cluster.read` and `solaris.cluster.modify` authorization.

1. Set up the Oracle ZFS Storage Appliance.

You can set up the appliance at any point in your cluster installation. Follow the instructions in your Oracle ZFS Storage Appliance's documentation. You can also click Help in the Oracle ZFS Storage Appliance GUI to access information specific to the device you are installing.

When setting up your Oracle ZFS Storage Appliance, follow the standards that are described in [“Requirements, Recommendations, and Restrictions for Oracle ZFS Storage Appliance NAS Devices” on page 10.](#)

2. On each cluster node, add the Oracle ZFS Storage Appliance name to the `/etc/inet/hosts` file.

Add a hostname-to-address mapping for the device in the `/etc/inet/hosts` file on all cluster nodes, as shown in the following example:

```
192.0.2.19 zfssa-123
```

3. In the `/etc/nsswitch.conf` file on every cluster node, ensure that files follows `cluster` and precedes any directory or name service.

For example:

```
hosts:      cluster files nis
```

a. Display the current setting for the host and netmask lookup.

```
# /usr/sbin/svccfg -s svc:/system/name-service/switch listprop config/host
config/host astring "cluster files nis dns"
```

```
# /usr/sbin/svccfg -s svc:/system/name-service/switch listprop config/netmask
config/netmask astring "cluster files nis"
```

b. If the `cluster` lookup is not included at the beginning of either lookup list, set the correct lookup list and refresh the name-service switch.

```
# /usr/sbin/svccfg \  
-s svc:/system/name-service/switch setprop config/host =astring: \  
"cluster files nis \  
\  
" \  
# /usr/sbin/svccfg \  
-s svc:/system/name-service/switch setprop config/netmask =astring: \  
"cluster files \  
nis\  
" \  
# /usr/sbin/svcadm refresh svc:/system/name-service/switch
```

- c. **Verify that the lookup lists now have the cluster lookup at the beginning of the list.**

```
# /usr/sbin/svccfg -s svc:/system/name-service/switch listprop config/host
# /usr/sbin/svccfg -s svc:/system/name-service/switch listprop config/netmask
```

4. **Configure the filer workflow for Oracle Solaris Cluster NFS.**

- a. **In the Oracle ZFS Storage Appliance GUI, select Maintenance, select Workflows, and click the workflow called Configure for Oracle Solaris Cluster NFS.**

- b. **Provide a password for this workflow.**

This same password will be used again in [Step 7](#).

Perform the workflow configuration from only one head in a dual-head configuration.

Note - If the workflow of the specified name is not present, it is likely that the filer is not running the correct software release. See [“Requirements for Oracle ZFS Storage Appliance NAS Devices” on page 10](#) for an example of a supported software release.

5. **In the global zone, install the zfssa-client package on all cluster nodes.**

On all nodes within the global zone, install the `zfssa-client` package from the repository. You can use the `pkg publisher` command to check that the publisher is already set for the `zfssa-client` package. For example, the `pkg publisher` command might return the following location: `https://pkg.oracle.com/ha-cluster/release`. For more information about setting the publisher, see [“How to Install Oracle Solaris Cluster Software \(pkg\)” in *Installing and Configuring an Oracle Solaris Cluster 4.4 Environment*](#).

For example:

```
# pkg publisher
# pkg list -af zfssa-client
# pkg install zfssa-client

=====
#
Packages to install: 1
Create boot environment: No
Create backup boot environment: No

DOWNLOAD          PKGS      FILES    XFER (MB)
Completed         1/1       7/7      0.2/0.2$<3>
```

```

PHASE                ACTIONS
Install Phase        17/17

PHASE                ITEMS
Package State Update Phase  1/1
Image State Update Phase  2/2
#
=====

```

6. **If you have a zone cluster, log in to the zone cluster and install the `zfssa-client` package on the zone cluster non-global zone.**

```
# pkg install zfssa-client
```

7. **Configure Oracle Solaris Cluster fencing support for the Oracle ZFS Storage Appliance.**

Note - If you skip this step, Oracle Solaris Cluster will not provide fencing support for the appliance.

- a. **Add the device and provide the cluster network addresses used to access the appliance.**

Perform this command from any cluster node:

```
# clnasdevice add -t sun_uss -p userid=osc_agent -p "nodeIPs{node_name}"=ip_address
myfiler
```

For example:

```
# clnasdevice add -t sun_uss -p userid=osc_agent
-p "nodeIPs{node1}"=10.111.11.111
-p "nodeIPs{node2}"=10.111.11.112 device1.us.example.com
Please enter password
```

<code>-t sun_uss</code>	Specifies <code>sun_uss</code> as the type of device you are adding.
<code>ip_address</code>	Specifies the IP address used to perform I/O to the appliance from this node.
<code>myfiler</code>	Specifies the name of the Oracle ZFS Storage Appliance that you are adding.
<code>node_name</code>	Specifies the name of cluster node whose IP addresses is being added.

This step allows the cluster fencing framework to restrict access to the filer for the specified IP address to read-only when nodes leave the cluster.

Note - The IP addresses configured for the cluster nodes should match the ones configured in the Oracle ZFS Storage Appliance as described in [“Requirements for Oracle ZFS Storage Appliance NAS Devices”](#) on page 10.

Note - If you want to add an appliance and provide the cluster network addresses used to access the appliance for a zone cluster but you need to issue the command from the global zone, use the `clnasdevice` command with the `-Z` option:

```
# clnasdevice add -t sun_uss -p userid=osc_agent -Z zcname \
-p "nodeIPs{node_name}"=ip_address myfiler
Please enter password
```

`-Z zcname` Specifies the name of the zone cluster where the Oracle ZFS Storage Appliance is being added.

b. At the prompt, type the same password that you used in [Step 4](#).

c. Confirm that the device has been added to the cluster.

Perform this command from any cluster node:

```
# clnasdevice show
===NAS Devices===
Nas Device:                device1.us.example.com
Type:                      sun_uss
userid:                    osc_agent
nodeIPs{node1}             10.111.11.111
nodeIPs{node2}             10.111.11.112
nodeIPs{node3}             10.111.11.113
nodeIPs{node4}             10.111.11.114
```

For more information about the `clnasdevice` command, see the [clnasdevice\(8CL\)](#) man page.

Note - If you are checking for the device for a zone cluster but you need to issue the command from the global zone, use the `clnasdevice show` command with the `-Z` option:

```
# clnasdevice show -Z zcname
```

You can also perform zone cluster-related commands inside the zone cluster by omitting the `-Z` option. For more information about the `clnasdevice` command, see the [clnasdevice\(8CL\)](#) man page.

- 8. To enable fencing support for the NFS file systems used by the cluster nodes, add the associated projects to the cluster configuration.**

Follow the directions in [“How to Add Oracle ZFS Storage Appliance Directories and Projects to a Cluster”](#) on page 22.

- 9. Configure a LUN on the Oracle ZFS Storage Appliance NAS device as a quorum device.**

Note - You can skip this step if the cluster does not require a quorum device or if it has been configured with quorum services from other devices or quorum servers.

See [“How to Add an Oracle ZFS Storage Appliance NAS Quorum Device”](#) in *Administering an Oracle Solaris Cluster 4.4 Configuration* for instructions for configuring an Oracle ZFS Storage Appliance NAS quorum device.

Maintaining an Oracle ZFS Storage Appliance NAS Device in an Oracle Solaris Cluster Environment

This section contains procedures about maintaining Oracle ZFS Storage Appliance NAS devices that are attached to a cluster. If a device's maintenance procedure might jeopardize the device's availability to the cluster, you must always perform the steps in [“How to Prepare the Cluster for Oracle ZFS Storage Appliance NAS Device Maintenance”](#) on page 19 before performing the maintenance procedure. After performing the maintenance procedure, perform the steps in [“How to Restore Cluster Configuration After Oracle ZFS Storage Appliance NAS Device Maintenance”](#) on page 20 to return the cluster to its original configuration.

▼ How to Prepare the Cluster for Oracle ZFS Storage Appliance NAS Device Maintenance

Follow the instructions in this procedure whenever the Oracle ZFS Storage Appliance NAS device maintenance you are performing might affect the device's availability to the cluster nodes.



Caution - If your cluster requires a quorum device (for example, a two-node cluster) and you are maintaining the only shared storage device in the cluster, your cluster is in a vulnerable state throughout the maintenance procedure. Loss of a single node during the procedure causes the other node to panic and your entire cluster becomes unavailable. Limit the amount of time for performing such procedures. To protect your cluster against such vulnerability, add a shared storage device to the cluster.

Before You Begin To perform this procedure, assume the root role or a role that provides `solaris.cluster.read` and `solaris.cluster.modify` authorization.

1. Stop I/O to the Oracle ZFS Storage Appliance NAS device.

If you have data services using NFS file systems from the Oracle ZFS Storage Appliance, bring the data services offline and disable the resources for the applications using those file systems. On each node, ensure that no existing processes are still using any of the NFS file systems from the device.

2. On each cluster node, unmount the NFS file systems from the Oracle ZFS Storage Appliance NAS device.

If you have a resource of type `SUNW.ScalMountPoint` managing the file system, disable that resource.

Note - For more information on disabling a resource, see [“How to Disable a Resource and Move Its Resource Group Into the UNMANAGED State”](#) in *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*.

If that resource is not configured, use the Oracle Solaris `umount(8)` command. If the file system cannot be unmounted because it is still busy, check for applications or processes that are still on that file system, as explained in [Step 1](#). You can also force the unmount by using the `-f` option with the `umount` command.

3. Determine whether a LUN on this Oracle ZFS Storage Appliance NAS device is a quorum device.

```
# clquorum show
```

4. **If the LUNs on this NAS device are not quorum devices, you are finished with this procedure.**
5. **If a LUN is a quorum device, perform the following steps:**
 - a. **If your cluster uses other shared storage devices or a quorum server, select and configure another quorum device.**

- b. **Remove this quorum device.**

See [Chapter 6, “Administering Quorum”](#) in *Administering an Oracle Solaris Cluster 4.4 Configuration* for instructions on adding and removing quorum devices.

▼ **How to Restore Cluster Configuration After Oracle ZFS Storage Appliance NAS Device Maintenance**

Follow the instructions in this procedure after performing any Oracle ZFS Storage Appliance NAS device maintenance that might affect the device's availability to the cluster nodes.

1. **Mount the NFS file systems from the Oracle ZFS Storage Appliance NAS device.**

If you have configured a resource of type `SUNW.ScalMountPoint` for the file system, enable the resource and bring its resource group online.
2. **Determine whether you want an iSCSI LUN on this Oracle ZFS Storage Appliance NAS device to be a quorum device.**

If you do, configure the LUN as a quorum device by following the steps in [“How to Add an Oracle ZFS Storage Appliance NAS Quorum Device”](#) in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

Remove any extraneous quorum device that you configured in [“How to Prepare the Cluster for Oracle ZFS Storage Appliance NAS Device Maintenance”](#) on page 19.
3. **I/O to the NFS file systems from the Oracle ZFS Storage Appliance NAS device can be resumed by bringing up the applications using the file systems.**

If the application is managed by a data service, enable the corresponding resources and bring their resource group online.

▼ How to Remove an Oracle ZFS Storage Appliance NAS Device From a Cluster

Note - You can also remove an Oracle ZFS Storage Appliance by using the Oracle Solaris Cluster Manager (Cluster Manager) browser interface. For Cluster Manager log-in instructions, see [“How to Access Oracle Solaris Cluster Manager”](#) in *Administering an Oracle Solaris Cluster 4.4 Configuration*.

Before You Begin This procedure relies on the following assumptions:

- Your cluster is operating.
- You have prepared the cluster by performing the steps in [“How to Prepare the Cluster for Oracle ZFS Storage Appliance NAS Device Maintenance”](#) on page 19.
- You have removed any device directories from the cluster by performing the steps in [“How to Remove Oracle ZFS Storage Appliance Directories and Projects and From a Cluster”](#) on page 27.

Note - When you remove the device from cluster configuration, the data on the device is not available to the cluster. Ensure that other shared storage in the cluster can continue to serve the data when the Oracle ZFS Storage Appliance NAS device is removed. When the device is removed, change the following items in the cluster configuration:

- Change the NFS file system entries in the `/etc/vfstab` file for that device, and unconfigure any `SUNW.ScalMountPoint` resources.
 - Reconfigure applications or data services with dependencies on these file systems to use other storage devices, or remove them from the cluster.
-

To perform this procedure, assume the root role or a role that provides `solaris.cluster.read` and `solaris.cluster.modify` authorization.

1. Remove the device.

Perform this command from any cluster node:

```
# clnasdevice remove myfiler
```

myfiler Specifies the name of the Oracle ZFS Storage Appliance NAS device that you are removing.

For more information about the `clnasdevice` command, see the [clnasdevice\(8CL\)](#) man page.

Note - If you want to remove a NAS device from a zone cluster but you need to issue the command from the global zone, use the `clnasdevice` command with the `-Z` option:

```
# clnasdevice remove -Z zcname myfiler
```

`-Z zcname` Specifies the name of the zone cluster where the Oracle ZFS Storage Appliance NAS device is being removed.

2. Confirm that the device has been removed from the cluster.

Perform this command from any cluster node:

```
# clnasdevice list
```

Note - If you want to check the NAS device for a zone cluster but you need to issue the command from the global zone, use the `clnasdevice` command with the `-Z` option:

```
# clnasdevice list -Z zcname
```

You can also perform zone cluster-related commands inside the zone cluster by omitting the `-Z` option. For more information about the `clnasdevice` command, see the [clnasdevice\(8CL\)](#) man page.

▼ How to Add Oracle ZFS Storage Appliance Directories and Projects to a Cluster

Before You Begin The procedure relies on the following assumptions:

- Perform the steps in this procedure only if the directory or project is meant to be protected by cluster fencing, restricting access to read-only for nodes that leave the cluster.
- Your cluster is operating.
- The Oracle ZFS Storage Appliance NAS device is properly configured.
See [“Requirements, Recommendations, and Restrictions for Oracle ZFS Storage Appliance NAS Devices” on page 10](#) for the details about required device configuration.
- You have added the device to the cluster by performing the steps in [“How to Install an Oracle ZFS Storage Appliance in a Cluster” on page 13](#).

An NFS file system or directory from the Oracle ZFS Storage Appliance is already created in a project, which is itself in one of the storage pools of the device. It is important that in order

for a directory (for example, the NFS file system) to be used by the cluster, to perform the configuration at the project level, as described below.

To perform this procedure, assume the root role or a role that provides `solaris.cluster.read` and `solaris.cluster.modify` authorization.

1. Use the Oracle ZFS Storage Appliance GUI to identify the project associated with the NFS file systems for use by the cluster.

After you have identified the appropriate project, click `Edit` for that project.

2. If read/write access to the project has not been configured, set up read/write access to the project for the cluster nodes.

a. Access the NFS properties for the project.

In the Oracle ZFS Storage Appliance GUI, select the `Protocols` tab in the `Edit Project` page.

b. Set the Share Mode for the project to None or Read only, depending on the desired access rights for nonclustered systems.

The `Share Mode` can be set to `Read/Write` if it is required to make the project world-writable, but doing so renders the data less secure.

c. Add a read/write NFS Exception for each cluster node by performing the following steps.

- Under `NFS Exceptions`, click `+`.
- Use the pull-down menu to select a `Type`.

Note - If you are using AK 2013.1.6 or later versions, you must select `IPv4 Subnet` as the `Type`. If you are using versions earlier than AK 2013.1.6, then select `Network` as the `Type`.

■ **Enter the public IP address the cluster node will use to access the appliance as the Entity.**

Use a CIDR mask of `/32`. For example, `192.168.254.254/32`.

■ **Select Read/Write as the Access Mode.**

■ **If desired, select Root Access.**

`Root Access` is required when configuring applications, such as Oracle RAC or HA for Oracle Database.

- **Add exceptions for all cluster nodes.**
 - **Click Apply after the exceptions have been added for all IP addresses.**
3. **Ensure that the directory being added is set to inherit its NFS properties from its parent project.**
 - a. **Navigate to the Shares tab in the Oracle ZFS Storage Appliance GUI.**
 - b. **Click Edit Entry to the right of the Share that will have fencing enabled.**
 - c. **Navigate to the Protocols tab for that share, and ensure that the Inherit from project property is set in the NFS section.**

If you are adding multiple directories within the same project, verify that each directory that needs to be protected by cluster fencing has the Inherit from project property set.

4. **If the project has not already been configured with the cluster, add the project to the cluster configuration.**
 - a. **Use the `clnasdevice show -v` command to determine whether the project has already been configured with the cluster.**

```
# clnasdevice show -v

===NAS Devices===
Nas Device:                device1.us.example.com
Type:                      sun_uss
userid:                    osc_agent
nodeIPs{node1}             10.111.11.111
nodeIPs{node2}             10.111.11.112
nodeIPs{node3}             10.111.11.113
nodeIPs{node4}             10.111.11.114
Project:                   pool-0/local/projecta
Project:                   pool-0/local/projectb
```

- b. **If you need to add a project to the cluster configuration, perform this command from any cluster node:**

```
# clnasdevice add-dir -d project1,project2 myfiler
```


-d project1, project2 Specifies the project or projects that you are adding. Specify the full path name of the project, including the pool. For example, `pool-0/local/projecta`.

myfiler Specifies the name of the NAS device containing the projects.

For example:

```
# clnasdevice add-dir -d pool-0/local/projecta device1.us.example.com
# clnasdevice add-dir -d pool-0/local/projectb device1.us.example.com
# clnasdevice find-dir -v
=== NAS Devices ===
```

```
Nas Device:                device1.us.example.com
Type:                      sun_uss
Unconfigured Project:      pool-0/local/projecta
File System:                /export/projecta/filesystem-1
File System:                /export/projecta/filesystem-2
Unconfigured Project:      pool-0/local/projectb
File System:                /export/projectb/filesystem-1
```

For more information about the `clnasdevice` command, see the [clnasdevice\(8CL\)](#) man page.

Note - If you want to add the project from an Oracle ZFS Storage Appliance to a zone cluster but you need to issue the command from the global zone, use the `clnasdevice` command with the `-Z` option:

```
# clnasdevice add-dir -d project1,project2 -Z zcname myfiler
```

-Z zcname Specifies the name of the zone cluster where the NAS projects are being added.

5. Confirm that the directory and project have been configured.

Perform this command from any cluster node:

```
# clnasdevice show -v -d all
```

For example:

```
# clnasdevice show -v -d all
```

```
===NAS Devices===
Nas Device:                device1.us.example.com
```

```
Type: sun_uss
nodeIPs{node1} 10.111.11.111
nodeIPs{node2} 10.111.11.112
nodeIPs{node3} 10.111.11.113
nodeIPs{node4} 10.111.11.114
userid: osc_agent
Project: pool-0/local/projecta
File System: /export/projecta/filesystem-1
File System: /export/projecta/filesystem-2
Project: pool-0/local/projectb
File System: /export/projectb/filesystem-1
```

Note - If you want to check the projects for a zone cluster but you need to issue the command from the global zone, use the `clnasdevice` command with the `-Z` option:

```
# clnasdevice show -v -Z zcname
```

You can also perform zone cluster-related commands inside the zone cluster by omitting the `-Z` option. For more information about the `clnasdevice` command, see the [clnasdevice\(8CL\)](#) man page.

After you confirm that a project name is associated with the desired NFS file system, use that project name in the configuration command.

6. If you do not use the automounter, mount the directories manually.

- a. **On each node in the cluster, create a mount-point directory for each Oracle ZFS Storage Appliance NAS project that you added.**

```
# mkdir -p /path-to-mountpoint
```

path-to-mountpoint Name of the directory on which to mount the project.

- b. **On each node in the cluster, add an entry to the `/etc/vfstab` file for the mount point.**

If you are using your Oracle ZFS Storage Appliance NAS device for Oracle RAC or HA for Oracle Database, consult your Oracle Database guide or log into My Oracle Support for a current list of supported files and mount options. After you log into [My Oracle Support](#), click the Knowledge tab and search for Bulletin 359515.1.

When mounting Oracle ZFS Storage Appliance NAS directories, select the mount options appropriate to your cluster applications. Mount the directories on each node that will access the directories. Oracle Solaris Cluster places no additional restrictions or requirements on the options that you use.

7. **To enable file system monitoring, configure a resource of type SUNW.ScalMountPoint for the file systems.**

For more information, see [“Configuring Failover and Scalable Data Services on Shared File Systems”](#) in *Planning and Administering Data Services for Oracle Solaris Cluster 4.4*.

▼ How to Remove Oracle ZFS Storage Appliance Directories and Projects and From a Cluster

Before You Begin This procedure relies on the following assumptions:

- Your cluster is operating.
- You have prepared the cluster by performing the steps in [“How to Prepare the Cluster for Oracle ZFS Storage Appliance NAS Device Maintenance”](#) on page 19.

Note - When you remove the directories, the data on those directories is not available to the cluster. Ensure that other device projects or shared storage in the cluster can continue to serve the data when these directories are removed. When the directory is removed, change the following items in the cluster configuration:

- Change the NFS file system entries in the `/etc/vfstab` file for that device, and unconfigure any SUNW.ScalMountPoint resources.
 - Reconfigure applications or data services with dependencies on these file systems to use other storage devices, or remove them from the cluster.
-

To perform this procedure, assume the `root` role or a role that provides `solaris.cluster.read` and `solaris.cluster.modify` authorization.

1. **If you are using hard mounts or the automounter, unconfigure the NFS file system.**
 - a. **On each node in the cluster, unmount the file system you are removing.**

```
# umount /mount-point
```
 - b. **On each node in the cluster, remove the entries in the `/etc/vfstab` file for the projects you are removing.**
Skip this step if you are using the automounter.

2. **(Optional) Perform the remaining steps in this procedure only if you want to remove the project containing this directory from the cluster configuration.**
 - a. **Ensure that no directories within the project are in use within the cluster.**
 - b. **Remove the projects.**

Perform this command from any cluster node:

```
# clnasdevice remove-dir -d project1 myfiler
```

-d project1 Specifies the project or projects that you are removing. To remove all of this device's projects, specify `all` for the `-d` option.

myfiler Specifies the name of the Oracle ZFS Storage Appliance NAS device containing the projects.

For more information about the `clnasdevice` command, see the [clnasdevice\(8CL\)](#) man page.

Note - If you want to remove a project from an Oracle ZFS Storage Appliance device from a zone cluster but you need to issue the command from the global zone, use the `clnasdevice` command with the `-Z` option:

```
# clnasdevice remove-dir -d project1 -Z zcname myfiler
```

-Z zcname Specifies the name of the zone cluster where the Oracle ZFS Storage Appliance NAS projects are being removed.

For more information about the `clnasdevice` command, see the [clnasdevice\(8CL\)](#) man page.

3. **Confirm that the projects have been removed.**

Perform this command from any cluster node:

```
# clnasdevice show -v
```

Note - If you want to check the NAS projects for a zone cluster but you need to issue the command from the global zone, use the `clnasdevice` command with the `-Z` option:

```
# clnasdevice show -v -Z zcname
```

You can also perform zone cluster-related commands inside the zone cluster by omitting the `-Z` option. For more information about the `clnasdevice` command, see the [clnasdevice\(8CL\)](#) man page.

See Also To remove the device, see [“How to Remove an Oracle ZFS Storage Appliance NAS Device From a Cluster”](#) on page 21.

Index

A

adding

- Oracle ZFS Storage Appliance directories, 22
- Oracle ZFS Storage Appliance NAS devices, 13
- Oracle ZFS Storage Appliance projects, 22

B

- boot order, 12

C

- clnasdevice add command, 16
- clnasdevice add-dir command, 24
- clnasdevice list command, 22
- clnasdevice remove command, 21
- clnasdevice remove-dir command, 28
- clnasdevice show command, 17, 25
- clquorum show command, 19
- cluster preparation, 19
- cluster restoration, 20
- commands
 - clnasdevice add, 16
 - clnasdevice add-dir, 24
 - clnasdevice list, 22
 - clnasdevice remove, 21
 - clnasdevice remove-dir, 28
 - clnasdevice show, 17, 24, 25
 - clquorum show, 19
- configuration requirements
 - Oracle ZFS Storage Appliance, 11

D

- devices *See* Oracle ZFS Storage Appliance

E

- /etc/vfstab file
 - adding mount points, 26

F

- fencing
 - configuring for Oracle ZFS Storage Appliance NAS devices, 13
- filers *See* Oracle ZFS Storage Appliance
- files
 - fstab, adding mount points, 26

H

- hosts file, 14

I

- installing Oracle ZFS Storage Appliance NAS devices, 13

M

- maintenance
 - preparing the cluster, 19
 - restoring the cluster, 20

mount options
 Oracle ZFS Storage Appliance NAS devices, 11
mount points
 modifying the `/etc/vfstab` file, 26

N

NAS devices
 adding Oracle ZFS Storage Appliance, 13
 adding Oracle ZFS Storage Appliance directories, 22
 adding Oracle ZFS Storage Appliance projects, 22
 installing Oracle ZFS Storage Appliance, 13
 removing Oracle ZFS Storage Appliance directories from a cluster, 27
 removing Oracle ZFS Storage Appliance projects from a cluster, 27
network-attached storage devices *See* as a NAS device
`nsswitch` file, 14

O

Oracle Solaris Cluster Manager
 tasks you can perform
 add an Oracle ZFS Storage Appliance, 13
 edit `Export List` property for an Oracle ZFS Storage Appliance, 13
 remove an Oracle ZFS Storage Appliance, 21
Oracle ZFS Storage Appliance
 as a NAS device, 9
 mount options, 11
Oracle ZFS Storage Appliance NAS devices
 configuring fencing support, 13
 removing, 21
 requirements, 10
 restrictions, 13

P

packages
 `zfssa-client`, 13
preparing the cluster for maintenance, 19

Q

quorum devices, 12, 12

R

removing
 Oracle ZFS Storage Appliance directories from a cluster, 27
 Oracle ZFS Storage Appliance NAS devices, 21, 21
 Oracle ZFS Storage Appliance projects from a cluster, 27
requirements
 for configuring Oracle ZFS Storage Appliance, 11
requirements for Oracle ZFS Storage Appliance NAS device use, 10
restoring the cluster after maintenance, 20
restrictions of Oracle ZFS Storage Appliance NAS device use, 13

V

`vfstab` file
 adding mount points, 26

W

workflow
 configuring for Oracle ZFS Storage Appliance, 15