



Deploying F5 ARX NFS Virtualization with Oracle's Sun ZFS Storage Appliance 7000 Series for NFS

Table of Contents

Configuring ARX NFS virtualization for Oracle's Sun ZFS Storage Appliance 7000 Series	
Prerequisites and configuration notes	1
Product versions and revision history	2
Configuration example	2
Configuring Oracle's Sun ZFS Storage Appliance 7000 Series for NFS	3
Performing the initial configuration of the ZFS Storage Appliance 7000 Series	3
Oracle's Sun ZFS Storage Appliance 7000 Series Setup	5
Defining the NIS Domain details	5
Creating a project	6
Creating Oracle's Sun ZFS Storage Appliance 7000 Series NFS Exports	8
Configuring the Root Directory ACL	10
Configuring the ARX	11
Creating the NFS Namespace	11
Adding the External Filers	12
Creating a Volume	12
Adding the root level export	14
Adding Oracle's Sun ZFS Storage Appliance 7000 Series NFS Share	16
Creating the Virtual Service	18
Creating the File Migration Policy	20
Removing the Legacy Storage NFS export	22
Conclusion	24

Configuring ARX NFS virtualization for Oracle's Sun ZFS Storage Appliance 7000 Series

Oracle's Sun ZFS Storage Appliance 7000 Series supports NFS exports and the F5 ARX is capable of virtualizing these NFS exports. This guide provides the procedures needed to deploy the ARX with Oracle's Sun ZFS Storage Appliance 7000 Series.

This deployment guide demonstrates how the ARX can be used to perform an NFS file system migration from the incumbent legacy storage to a new Oracle storage array. A migration policy is defined to move the files and directories to the new Oracle platform.

For more information on Oracle's Sun ZFS Storage Appliance 7000 Series, see <http://www.oracle.com/us/products/servers-storage/index.html>

For more information on the ARX series, see <http://www.f5.com/products/arx-series/>

Prerequisites and configuration notes

The following are prerequisites and configuration notes for this deployment:

- ◆ The F5 ARX is configured for network access, and the initial switch interview has been completed.
- ◆ The Oracle filer is physically installed, powered on, with network attributes reserved and available.
- ◆ The ARX platform is deployed in redundant pairs. The secondary switch is a Hot Standby switch. This guide addresses the configuration steps in order to integrate the Sun ZFS Storage Appliance with the ARX platform. Redundant switch configuration steps within the product documentation should be followed in order to deploy a high available configuration.
- ◆ Some screen shots in this guide may be different from your screens, depending on the software versions deployed.

The following ARX NFS Features are not supported

- ◆ Direct Volumes are not supported. A *direct volume* is a series of mount points to various back-end filers, with no metadata. Direct volumes are less complex than managed volumes, but do not support policies.
- ◆ NFS version 2 is not supported.

Product versions and revision history

Product and versions tested for this deployment guide:

Product Tested	Versions
Oracle's Sun ZFS Storage Appliance 7000 Series	S7310
Oracle's Sun ZFS Storage Appliance 7000 Series Software release	2010.08.17.4.0,1-1.31
F5 ARX	v5.2.2 and later

Document Version	Description
1.0	New deployment guide
1.1	-Correct <i>Adding the External Filers</i> , on page 12, to repeat the procedure to add Sun 7000 as an External Filer -Corrected incorrect references to CIFS and replaced with NFS. This did not effect the configuration procedures, only the file system name was incorrect.
1.2	Updated the guide to reflect Oracle's purchase of Sun and the latest versions of software

Configuration example

In the following diagram, we show basic connectivity between clients, ARX and Oracle's Sun ZFS Storage Appliance 7000 Series.

Sun Storage 7000 - ARX Lab Network Topology




Figure 1 Logical configuration example

Configuring Oracle's Sun ZFS Storage Appliance 7000 Series for NFS

This section addresses the configuration of the Sun ZFS Storage Appliance. Oracle's Sun ZFS Storage Appliance 7000 Series is shipped with a pre-installed kernel and an initial setup is performed on the console port. You complete the configuration using the GUI via a web browser. We configure the array for RAID volumes. The volume is NFS exported to the network.

There are several levels of authentication and access control that are available to the system administrator. This deployment guide demonstrates a very relaxed stance on access credentials for demonstration purposes only. The access rights for a live deployment most likely would not leverage such relaxed security policies.

This section contains the following procedures:

- *Performing the initial configuration of Oracle's Sun ZFS Storage Appliance 7000 Series*, on page 3
- *Oracle's Sun ZFS Storage Appliance 7000 Series Setup*, on page 5
- *Defining the NIS Domain details*, on page 5
- *Creating a project*, on page 6
- *Creating Oracle's Sun ZFS Storage Appliance 7000 Series NFS Exports*, on page 8
- *Configuring the Root Directory ACL*, on page 10

Performing the initial configuration of Oracle's Sun ZFS Storage Appliance 7000 Series

The initial configuration of Oracle's Sun ZFS Storage Appliance 7000 Series is accomplished through the console port on the array controller. Connect the management workstation to the Oracle Serial management Port. Use the Null Modem cable provided with the array hardware and a terminal emulation application.

The initial configuration defines the IP network interfaces. You define an IP address, Mask, Default gateway (router), DNS Server, and root password in these steps.

◆ Note

The Oracle documentation provided with the hardware is more comprehensive than the following example. For more in depth details review the Installation documentation provided by Oracle.

To initially configure the Sun ZFS Storage Appliance System

1. Using a terminal application connect to the console port. Using the following parameters: **9600 Baud, N,8,1, no flow control**.
2. Type the default Password which is **changeme**.
3. At the command prompt, type **start /SP/console**. You see a warning to run the console, type **Y**.
4. Press any key to begin configuring the appliance.
5. Type a host name. In our example, we type **S7310**.
6. Type a network interface. In our example, we type **NET-0**.
7. Type an IP address. In our example, we type **10.51.105.34**.
8. Type a Network Mask. In our example, we type **255.255.255.0**
9. Type a Default Router. In our example, we type **10.51.105.1**
10. Type a DNS Server IP. In our example, we type **192.168.1.5**
11. Enter a password for the administration account.
12. Press **Enter** to confirm the password.

The next part of the configuration uses the Sun ZFS Storage Appliance GUI.

To configure the Sun ZFS Storage Appliance using the GUI

1. Using a web browser, log on to the Sun ZFS Storage Appliance device as root. By default the web interface is listening on port 215. Use the following URL syntax: **http://<ip address>:215**

If this is your first time configuring the Oracle device, the initial setup wizard opens.
If you need to restart the wizard, it is available from the **Maintenance** menu by clicking **System**. Refer to the Oracle documentation for more information on re-running this wizard.
2. In the first 5 steps of the wizard, verify the appropriate settings (and edit if necessary), such as the network setup, DNS, and NTP.
3. In Step 6 of 6, click the **Configure** button to configure the storage pool. The Storage wizard opens.
In this example, we create a single storage pool named **pool-0**. This pool consists of all the disks with a RAID data profile.
4. On the **Verify Storage** page (Step 1 of 2), click the **Commit** button.
5. On the **Configure Storage** page (Step 2 of 2), you select the Data Profile to apply to the storage pool. In our example, we click **Double Parity RAID**, and then click **Commit**. This creates the storage pool.

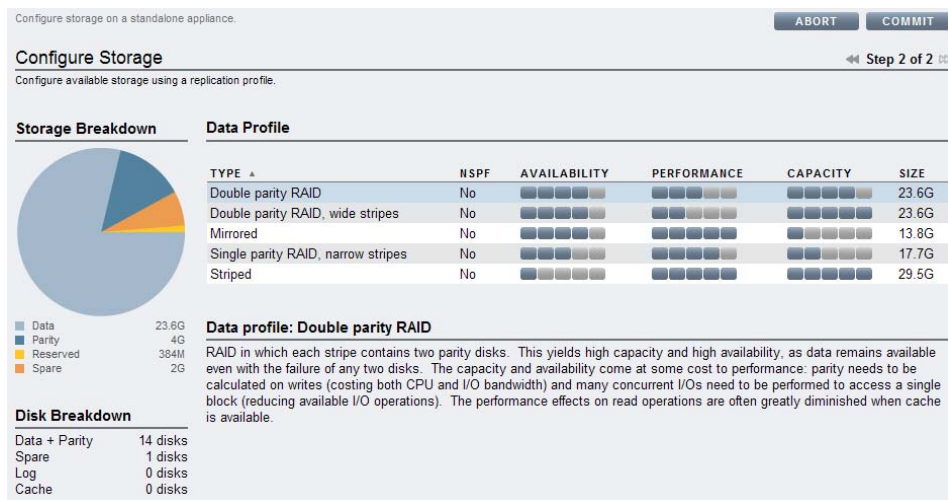


Figure 2 Configure Storage Data Profile

6. You return to Step 6 of the Configure Storage wizard. Click **Commit** to complete Step 6 of 6.

Oracle's Sun ZFS Storage Appliance 7000 Series Setup

In this section, we perform the Oracle system setup. In our configuration example the storage pools are already created. This deployment guide demonstrates how to configure a new NFS Export.

Defining the NIS Domain details

The Sun ZFS Storage Appliance can be integrated into a NIS Domain. This is an optional step. If the deployment requires NIS authentication.

To define the NIS domain

1. From the Sun ZFS Storage Appliance Status page, under Services, click **NIS**.
2. In the **Domain** box, the appropriate domain. In our example, we type **siterequest.com**.
3. Click the **Use listed servers** button.
4. In the box, type the IP address for the NIS server.
5. Click the **Apply** button. If the service was not already enabled, a you are prompted to enable the service. Click **Yes** to enable NIS.

Creating a project

Oracle suggests creating a new project to manage the configuration. This facilitates inheritance of default settings for all shares and eliminate user errors and reduces the setup complexity. Several common attribute settings are needed for each export. It is best to set these values at a project level. As new exports are created they inherit these attributes by default.

The Sun ZFS Storage Appliance needs to have the **no_root_squash** NFS attribute configured. Normally this attribute is configured in a text file on the server. Oracle has created the equivalent configuration via the Web GUI interface.

To create a new project

1. From the top menu, click **Shares**.
2. From the left pane, click the arrow next to **Projects**. The **Projects List** opens.
3. Click the **Add** button. The **Create Project** dialog opens.
4. In the **Name** box, type a name for this project. In our example, we type **ARX**.
5. Click the **Apply** button.




Figure 3 Naming the Project

Once you have created the Project, you need to set the default settings for the project.

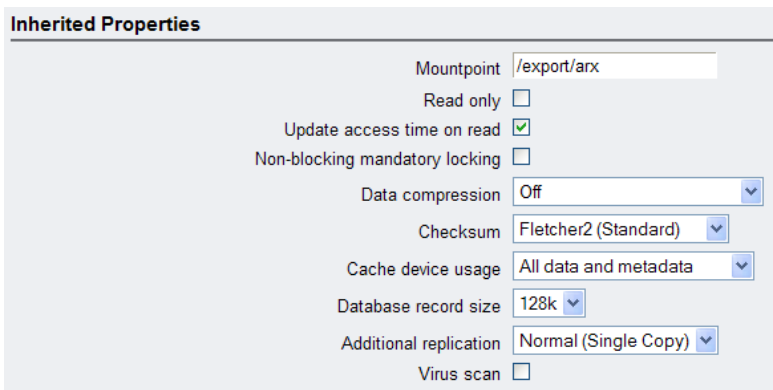
To set the default settings for the project

1. From the top menu, click **Shares**, and then click **Projects** in the sub menu.
2. From the **Project** list, point your mouse at the project you just created; a small pencil and trash can icon appear on the right. Click the pencil to edit the project (see Figure 4). A list of submenu options appear.

Projects 2 Total	
NAME ▲	SIZE
ARX ⓘ	442G
default	18K

Figure 4 List of Projects

- From the submenu bar, click **General**.
- From the Inherited Properties section, in the **Mountpoint** box, type the directory name. In our example, we type `/export/arx`.
- Click **Apply** to commit the changes.



Inherited Properties

Mountpoint

Read only

Update access time on read

Non-blocking mandatory locking

Data compression

Checksum

Cache device usage

Database record size

Additional replication

Virus scan

Figure 5 Project Inherited Properties

- From the submenu bar, click **Protocols**.
- In the **NFS Exceptions** section, click the Add (+) button to create a new default exception. The ARX requires root access to the NFS exports. By default root users are not allowed to mount NFS Exports.

Important: The root access privilege is controlled on a host by host basis. It is possible to open up root access to every host in an IP Domain, Address range, or unique IP address. For production deployments the root access should be limited to the ARX Proxy IP Addresses. This will ensure a rogue network client cannot mount the NFS export as root.

In a secured deployment scenario only the ARX Proxy IP addresses would be allowed root read/write access. For demonstration purposes, we specified the private lab storage network IP address range. From the **Type** list, select **Network**. In the **Entity** box, type the appropriate IP address. Click **Apply** to continue.

TYPE	ENTITY	ACCESS MODE	ROOT ACCESS
Network	10.0.0.0/8	Read/write	<input checked="" type="checkbox"/>

Figure 6 NFS Exceptions

The next task is to modify the permissions for the project.

8. Click the General tab.
9. In the **Filesystem Creation Defaults** section, modify the permissions as applicable to your configuration. In our example, we enable read/write/execute privileges for User, Group, and Other by clicking **R**, **W**, and **X**
10. Click the **Apply** button.

Filesystem Creation Defaults

User: nobody

Group: other

Permissions: **R** **W** **X** **R** **W** **X** **R** **W** **X**

 User Group Other

Figure 7 Filesystem Creation Defaults

The Project default values are now specified. These values are inherited by each share created as part of this project. These attributes can be overwritten within each individual export definition as necessary.

Creating Oracle's Sun ZFS Storage Appliance 7000 Series NFS Exports

The next task is to create the Sun ZFS Storage Appliance NFS exports. These exports are imported into the ARX managed volume in later steps. The exports are created within the context of the project we created in the previous section.

To create the NFS export

1. From the top menu, click Shares.
2. On the left, click the arrow next to Projects to expand the list, and then click the name of the project you just created.
3. On the sub menu, click **Shares** (if the page did not open to the Shares page).

-
4. Click the Add (+) button next to **Filesystems**. The Create Filesystem wizard opens.
 5. In the **Name** box, type a name for this Filesystem. In our example, we type **share**.
 6. In our example, we leave the remaining settings at the default levels.
 7. Click the **Apply** button.
The new share appears in the list of shares.

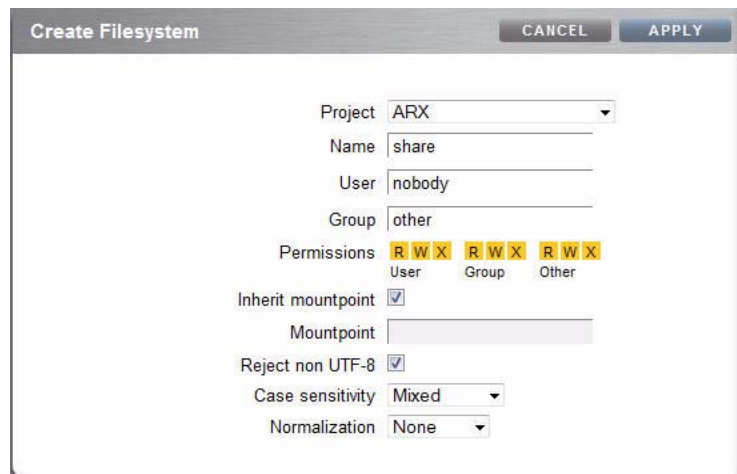


Figure 8 Create File system

Configuring the Root Directory ACL

The next task is to configure the Root directory ACL.

To configure the root directory ACL

1. From the Filesystems list, point your mouse at the filesystem you just created; a small pencil and trash can icon appear on the right. Click the pencil to edit the filesystem.
2. Click the **Access** tab to edit the Root Directory ACL. By default the Root Directory ACL has an allow and deny entry for Owner, Group, and Everyone as per the Unix style of for file permissions.
3. In the Root Directory ACL section, delete all but one of the rows by hovering over the row, and clicking the trash can icon for each.
4. In the last row, leave **Everyone** selected. Hover on the row and click the Pencil button to edit permissions. Click the **Full Control** check box and then click **OK** to continue.

Verify the ACL entry matches the permissions and inheritance settings in figure 9.

Root Directory ACL			
TYPE	TARGET	MODE	PERMISSIONS : INHERITANCE
Everyone	not applicable	Allow	rw-xpdaARWcCo:fd--

Figure 9 Root Directory ACL Permissions and Inheritance settings

5. Click **Apply** to commit the changes.
6. The previous procedures demonstrated the process to create a new NFS export with the appropriate access credentials to allow the ARX to virtualize the export into a managed volume.

Return to *Creating Oracle's Sun ZFS Storage Appliance 7000 Series NFS Exports*, on page 8 and repeat both procedures to create a second NFS export. This export is used by the ARX to store the metadata for the managed volume. In our example we created an export and named it **metaData2**.

Configuring the ARX

In this section, we configure the ARX to access the Incumbent Legacy Storage, and the Sun ZFS Storage Appliance. We create a NFS namespace with two shares added to it. The shares are added from the Incumbent Legacy NFS Export and the new Oracle NFS Export. A file migration policy is applied in order to migrate the file contents from the Legacy Storage to the new ZFS Storage Appliance 7000 Series.

This section contains the following tasks:

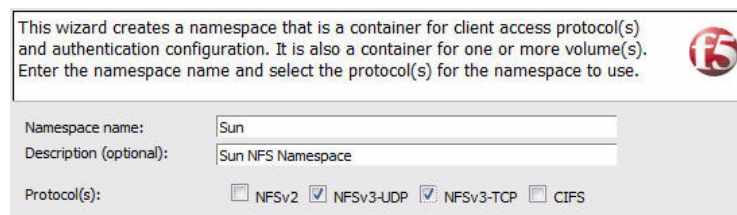
- *Creating the NFS Namespace*, on page 11
- *Adding the External Filers*, on page 12
- *Creating a Volume*, on page 12
- *Adding the root level export*, on page 14
- *Adding Oracle's Sun ZFS Storage Appliance 7000 Series NFS Share*, on page 16
- *Creating the Virtual Service*, on page 18
- *Creating the File Migration Policy*, on page 20
- *Removing the Legacy Storage NFS export*, on page 22

Creating the NFS Namespace

The first task is to create the NFS Namespace on the ARX device.

To create the NFS namespace

1. Open the ARX web-based Configuration utility, and then, from the navigation pane, click **Common Operations**.
2. Click **Create Namespace**. The Create Namespace wizard opens.
3. In the **Namespace name** box, type a name. In our example, we type **Sun**. You can optionally type a description.
4. From the **Protocol** list, check the appropriate NFS box(es). In our example, we check **NFS Version 3 TCP and UDP protocol**.
5. Click the **Next** button.



This wizard creates a namespace that is a container for client access protocol(s) and authentication configuration. It is also a container for one or more volume(s). Enter the namespace name and select the protocol(s) for the namespace to use.

Namespace name:

Description (optional):

Protocol(s): NFSv2 NFSv3-UDP NFSv3-TCP CIFS

Figure 10 Create NFS namespace details

6. Review the summary, and then click **Finish**.

Adding the External Filers

The next task is to add the Legacy Storage and the Sun ZFS Storage Appliance as external filer entries in the ARX. These entries are referenced later when we add the filer shares to the managed volume.

To add the External Filers

1. From the navigation pane, click **File Servers**.
2. Click the **Add** button. The Add File Server screen opens.
3. In the **Name** box, type a name for this File Server. In our example, we type **Legacy**.
4. In the **Primary IP Address** box, type the primary IP address.
5. In the **Secondary IP Address** box, type any secondary IP addresses, and click the **Add** button. In our example, we do not include any secondary IP addresses.
6. In the **Description** box, you can optionally type a description.
7. In the **Ignore Directories (optional)** box, type any snapshot directories the ARX should ignore on the backend file shares, and click the **Add** button. In our example, this filer does not support snapshots so that feature was not enabled.
8. Click the **OK** button. You return to the File Server Summary page.
9. Click **OK** to complete the file server addition.
10. Repeat this entire procedure to add the Oracle 7000 as an external filer.

Creating a Volume

The backend filer NFS exports are incorporated into an ARX Managed Volume. File placement policy is managed at the volume level.

ARX best practices state Metadata should be created on an NFS export if one is available. This deployment guide is demonstrating how to migrate file systems from a legacy storage platform to the new Oracle Unified Storage. After the migration the incumbent product maybe removed or re-deployed. For this reason, the ARX MetaData repository is configured on the new storage platform. In this example, we place the ARX Managed Volume Metadata onto the Oracle Unified Storage platform.

To create a volume on the ARX

1. From the left navigation pane, click **Common Operations**.
2. Click the **Create Volume** button. The Create Volume wizard opens.
3. From the **Namespace** list, select the name of the namespace you created in *Creating the NFS Namespace*, on page 11. In our example, we select **Sun**.

-
- In the **Volume name** box, type the volume name. In our example, we type **/NFS**. You can optionally add a description. See Figure 11.
 - Click the **Next** button.

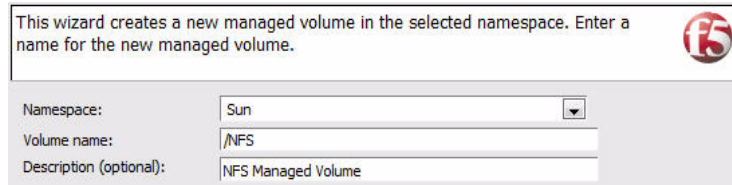


Figure 11 New Managed Volume

- From the **Metadata file server protocol** list, select an appropriate protocol. In our example, we select **NFSv3-UDP**.
- From the **Metadata file server** list, select the file server if it has already been configured. If it has not, you can click **Add** to add one. In our example, we select **Sun**.
- In the **Metadata CIFS share/NFS path** box, type the path. In our example, we type **/export/metaData2**.
- Click the **Next** button.

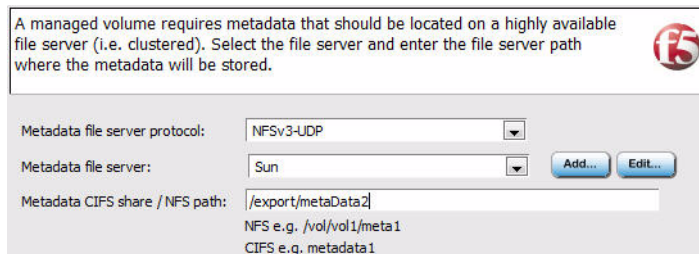


Figure 12 New managed volume metadata file server

- From the Volume Parameters page, in the Performance Tuning section, from the **VPU ID** list, select **Dynamic**. Ensure the **Auto Reserve files** box is checked.
- In the Import Conflict Resolution section, click the **Files and directories can be renamed during import and a re-import**. **Makes the volume read write** button.
- In the Enable Volume section, click the **Enable the volume when finished** button.
- Click the **Next** button.
- Review the summary, and click **Finish**.

Set the volume parameters or accept the defaults. Set the maximum file count to the approximate number of files that will be in the volume.

Performance Tuning

VPU ID:

Auto reserve files

Maximum files:

Import Conflict Resolution

Files and directories are not renamed during import. Makes the volume read only.

Files and directories can be renamed during import and a re-import. Makes the volume read write.
(Rename files and rename directories should also be set when adding shares to this volume.)

Shadow Target

Make this volume a shadow-copy target.

Enable Volume

Enable the volume when finished.

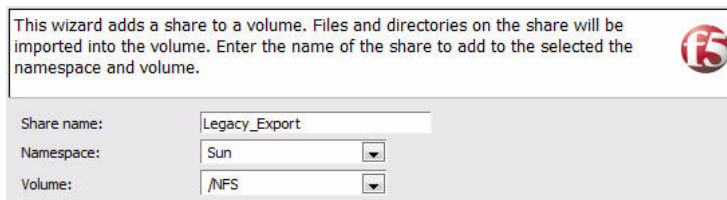
Figure 13 Volume Parameters

Adding the root level export

The next task is to create the root level export. The first file export we add is the root level export. This is the incumbent legacy storage volumes with file content. Subsequent exports adapt to the root volume permissions.

To add a root level share

1. From the left navigation pane, click **Common Operations**.
2. Click the **Add Share** button. The Add Share Wizard opens.
3. In the **Share Name** box, type a name for this Share. In our example, we type **Legacy_Export**.
4. From the **Namespace** list, select the name of the namespace you created in *Creating the NFS Namespace*, on page 11. In our example, we select **Sun**.
5. From the **Volume** list, select the name of the Volume you created in *Creating a Volume*, on page 12. In our example, we select **/NFS**.
6. Click the **Next** button.



This wizard adds a share to a volume. Files and directories on the share will be imported into the volume. Enter the name of the share to add to the selected the namespace and volume.

Share name: Legacy_Export

Namespace: Sun

Volume: /NFS

Figure 14 *New Legacy Storage Share*

7. From the **File Server** list, select the name of the file server you created in *Adding the External Filers*, on page 12. In our example, we select **Legacy**.
8. In the **NFS Path** box, type the path. In our example, we type **/volume**.
9. In the Import Conflict Resolution section, check the **Rename files with naming collisions on import** and **Rename directories with naming collisions on import** boxes (see Figure 15). In our example, we also check the **Disable the managed file system check on import box**. This is optional.
10. Check the **Enable this share when finished** box.
11. Optional: You can select the ability for the ARX to take ownership if the export is owned by another ARX and if the export had been previously managed by an ARX by checking the **Allow this switch to import this share, even if it's owned by another Acopia switch box**
12. Click the **Next** button.
13. Review the summary, and click the **Finish** button.

Select the file server and file server share to add to the volume. If the volume is a managed volume, select any import options and indicate whether the share contains access control lists with local groups.

Share name: Legacy_Export

File Server: Legacy Add...

NFS Path: /volume

Import Conflict Resolution

- Rename files with naming collisions on import.
- Rename directories with naming collisions on import.
- Synchronize directory attributes between shares on import.
- Disable the managed file system check on import.

Enable Share

- Enable this share when finished.
- Allow this switch to import this share, even if it is owned by another Acopia switch.

Figure 15 File Server and Share name

Adding Oracle's Sun ZFS Storage Appliance 7000 Series NFS Share

The Sun ZFS Storage Appliance share needs to be added to the volume.

To add the NFS Shares to the volume

1. From the left navigation pane, click **Common Operations**.
2. Click the **Add Share** button.
3. In the **Share Name** box, type a name for this Share. In our example, we type **Sun_Export**.
4. From the **Namespace** list, select the name of the namespace you created in *Creating the NFS Namespace*, on page 11. In our example, we select **Sun**.
5. From the **Volume** list, select the name of the Volume you created in *Creating a Volume*, on page 12. In our example, we select **NFS**.
6. Click the **Next** button.
7. From the **File Server** list, select the name of the file server you created in *Adding the External Filers*, on page 12. In our example, we select **Sun**.
8. In the **NFS path** box, type the path.
In our example, we type **/export/share**.
9. In the Import Conflict Resolution section, check the **Rename files with naming collisions on import**, **Rename directories with naming collisions on import**, and **Synchronize directory**

- attributes between shares on import** boxes. In our example, we also check the **Disable the managed file system check on import** box. This is optional.
- Optional: You can select the ability for the ARX to take ownership if the share is owned by another ARX by clicking the **Allow this switch to import this share, even if it's owned by another Acopia switch** box.
 - Check the **Enable this share when finished** box.

Select the file server and file server share to add to the volume. If the volume is a managed volume, select any import options and indicate whether the share contains access control lists with local groups.

Share name: Sun_Export

File Server:

NFS Path:

Import Conflict Resolution

- Rename files with naming collisions on import.
- Rename directories with naming collisions on import.
- Synchronize directory attributes between shares on import.
- Disable the managed file system check on import.

Enable Share

- Enable this share when finished.
- Allow this switch to import this share, even if it is owned by another Acopia switch.

Figure 16 File Server and NFS Export Path

- Click the **Next** button.
- Review the summary, and click the **Finish** button.

From the *Managed Volume Shares* tab, you should see both NFS exports with a status of online.

Managed Volume Details

Namespace: Volume:

Volume **Shares**

<input type="checkbox"/>	Share	File Server	File Server Path	Free Space	Transitions	Status
<input type="checkbox"/>	Legacy_Export	Legacy 10.20.2.61	/volume	79 G	1	● Online
<input type="checkbox"/>	Sun_Export	Sun 10.51.105.34	/export/share	0	1	● Online

Figure 17 Managed Volume Details

Creating the Virtual Service

The Virtual Service is how the ARX presents the NFS exports to the network clients. Clients send file requests through the Virtual Service and the ARX proxies these requests to the appropriate backend filer.

To create the virtual service

1. From the navigation pane, click **Virtual Services**.
The Virtual Service Summary page opens.
2. Click the **Add** button. The Add Virtual Service Wizard opens.
3. From the **Namespace** list, select the namespace you created in *Creating the NFS Namespace*, on page 11. In our example, we select **Sun**.
4. Click the **Create a new virtual service (VIP) button**.
 - a) In the **Virtual service DNS name** box, type the DNS name for the virtual service. In our example, we type **export.siterequest.com**.
 - b) In the **IP Address** box, type the IP address of the VIP. In our example, we type **10.0.0.1**.
 - c) In the **Subnet Mask** box, type the appropriate subnet mask. In our example, we type **255.0.0.0**.
 - d) From the **VLAN ID** list, select the appropriate VLAN ID. In our example, we select **301**.
5. Ensure the **Enable the virtual service when finished** box is checked.

This wizard creates either a new virtual service (virtual IP address) or adds a new export to an existing virtual service. Select a namespace and whether to add a new service or add an export to an existing service.

Namespace: Sun

Add export(s) to an existing virtual service (VIP)

Virtual service DNS name:

IP Address: -- No VIPs are defined --

Create a new virtual service (VIP)

Virtual service DNS name:

IP Address:

Subnet Mask:

VLAN ID:

[Enable Virtual Service](#)

Enable the virtual service when finished

Figure 18 Create new Virtual Service Wizard

6. Click the **Next** button. The Virtual Service Exports screen opens.

7. From the **Volume** list, select the name of the Volume you created in *Creating a Volume*, on page 12. In our example, we select **/NFS**.
8. In the **Volume Path** box, type the path. In our example, we type **/**.
9. In the **Export Name** box, type a name for the Export. In our example, we type **arx_export**. You can optionally type a description.
10. Click the **Add Export** button.
11. Click the **Next** button.
12. Review the summary and click **Finish**.

You can review the Virtual Service by clicking **Virtual Services** from the navigation pane. Notice the admin state is enabled and the status is ready.

Virtual Service Summary

Click on a virtual service to view its details, or select a virtual service and click on an action button.

<input type="checkbox"/>	Domain Name	Virtual IP	VLAN	Exports	Domain Join	Admin State	Status
<input type="checkbox"/>	export.siterequest.com	10.0.0.1 255.0.0.0	1	i		NFS: Enabled	NFS: Ready

Figure 19 Virtual Service Summary

Click the Domain Name. The Export Summary opens. Notice the status is Online.

Services | NFS Exports

Export	Virtual Vol Path	NFS Access List	State
arx_export	/NFS/	(rw)	Online

Figure 20 Virtual Service Export Summary

Creating the File Migration Policy

A fileset migration policy is assigned to a managed volume. It facilitates file movement between backend file shares based on file attributes. The fileset matching attributes include exact match, wildcard expressions, and regular expressions.

In this example, we create a Fileset Migration Policy to move files from the Legacy Storage to the Sun ZFS Storage Appliance. The files are copied to the NFS export **/export/share**.

To create the file migration policy

1. From the left navigation pane, click **Common Operations**.
2. Click the **Fileset Migration** button. The Fileset Migration Wizard opens.
3. In the **Policy Name** box, type a name for this policy. In our example, we type **Migrate_to_Sun**.
4. From the **Namespace** list, select the name of the namespace you created in *Creating the NFS Namespace*, on page 11. In our example, we select **Sun**.
5. From the **Managed volume** list, select the name of the Volume you created in *Creating a Volume*, on page 12. In our example, we select **/NFS**.
6. Click the **Next** button.
7. In the **Fileset** row, click the **Add** button. Configure the following options to create a new fileset:
 - a) In the **Fileset Name** box, type a name. In our example, we type **Migrate_All**.
 - b) From the **Fileset Type** list, select the appropriate type. In our example, we select **filename**.
 - c) In the **Filename Matching Criteria** section, click the appropriate button and then type the criteria in the box. In our example, we click **Wildcard Expression**, and type ***** in the box.
 - d) Configure the rest of the options as applicable. In our example, we leave the settings at the defaults.
 - e) Click the **Save** button to return to the Fileset Migration wizard.

Create a new fileset.

Fileset Name: Migrate_All

Fileset Type: filename

Filename Matching Criteria: Exact Match Wildcard Expression (i.e. shell style) Regular Expression

Criteria: *

Exclude Name:

Ignore Case for Name:

Path Matching Criteria: Exact Match Wildcard Expression (i.e. shell style) Regular Expression

Criteria: /

Exclude Path:

Ignore Case for Path:

Recurse Subdirectories:

Figure 21 New Fileset criteria

8. From the **Match Criteria** list, select the appropriate criteria. In our example, we select **Files and Directories**.
9. From the **Source** list, select the applicable source. In our example, we select **Legacy_Export**
10. From the **Target** list, select the appropriate target. In our example, we select **Sun_Export**.
11. Click the **Next** button.

Figure 22 Fileset Migration Criteria

12. In the Options section, configure the settings as applicable to your configuration. In our example, we leave the settings at the defaults.
13. Make sure **Enable this policy when finished** is checked.
14. Click the **Next** button.
15. Review the summary and then click the **Finish** button.

By default a report is generated for this policy. To view the report, from the navigation pane, click **Report**. A list of reports opens. The first report is the one created for the fileset migration. You can click the report to view it.

Reports

Click on a report to view it, or select a report and click on an action button

Type: All

Download Download Format: TXT Remove...

Report	Date	Size	Type	Status
Migrate to Sun_200911120919.rpt	Nov 12, 2009 9:19:22 AM	41 k	Place Rule	DONE

Figure 23 Reports Summary

In our example, the report shows that all the files have been moved from the Legacy Storage to the Oracle Export. A metadata report can confirm this. To generate a metadata report, click **Managed Volumes**. Click the **/NFS** volume to display the details. Click the **Report** button, and in the **Report Type** row, click **Metadata**. Click **OK** to generate the report.

The report is available by clicking **Reports** from the navigation pane. Open the report and review the Share column. For each file, the Share on which it is physically present is reported. In our example, we see all the files reside on the share **Sun_Export**.

Removing the Legacy Storage NFS export

The final task is to remove the legacy storage NFS export from the managed volume.

To remove the Legacy Storage NFS export

1. From the navigation pane, click **Managed Volumes**.
2. From the **Volume** list, select the volume you created in *Creating a Volume*, on page 12. In our example, we select **/NFS**.
3. Click the Shares tab.
4. Check the box next to the Share you created in *Adding the root level export*, on page 14. In our example, we click the box next to **Legacy_Export**.

Managed Volume Details

Namespace: Sun Volume: /NFS

Volume Shares

Add... Remove... Edit... Enable Disable Sync...

<input type="checkbox"/>	Share	File Server	File Server Path	Free Space	Transitions	Status
<input checked="" type="checkbox"/>	Legacy_Export	Legacy 10.20.2.61	/volume	79 G	3	● Online
<input type="checkbox"/>	Sun_Export	Sun 10.51.105.34	/export/share	7.1 T	3	● Online

Figure 24 Managed Volume Details

5. Click the **Remove** button. The Remove Share wizard opens. The removal process ensures any files which may exist on the old share are migrated prior to removal. This is important because users may have added files to the managed volume.
6. Click the **Next** button.

This wizard removes a share from a volume and optionally allows its contents to be migrated. Select the namespace, volume and the share to remove.

Namespace: Sun

Volume: /NFS

Share to remove: Legacy_Export

Figure 25 Remove share wizard

7. Click the **Remove and Migrate the contents to another target** button.
8. In the Migrate files to the selected share or share-farm section, select the share you created in *Adding Oracle's Sun ZFS Storage Appliance 7000 Series NFS Share*, on page 16. In our example, we select **Sun_Export**.
9. Click the Next button.
10. Review the summary and click **Finish** to remove the export.

You return to the Managed Volume Details screen. If you refresh your browser, the volume shows only the Oracle export exists.

The screenshot shows the 'Managed Volume Details' interface. At the top, there are dropdown menus for 'Namespace: Sun' and 'Volume: /NFS'. Below these are tabs for 'Volume' and 'Shares', with 'Shares' being the active tab. A row of buttons includes 'Add...', 'Remove...', 'Edit...', 'Enable', 'Disable', and 'Sync...'. Below the buttons is a table with the following data:

Share	File Server	File Server Path	Free Space	Transitions	Status
<input type="checkbox"/> Sun_Export	Sun 10.51.105.34	/export/share	7.1 T	3	● Online

Figure 26 *Managed Volume Details*

The Legacy Storage platform can now be decommissioned or repurposed safely.

Conclusion

This deployment guide demonstrated the way to integrate F5 ARX platform with the Sun ZFS Storage Appliance. The deployment demonstrated and enabled the migration of files from legacy storage platforms onto the Sun ZFS Storage Appliance using NFS.

For more information on configuring the F5 ARX, refer to the documentation, available on [Ask F5](#).