



ZFS STORAGE
APPLIANCE

September 2014

Using Oracle ZFS Storage Appliance Fibre Channel and iSCSI with Symantec Storage Foundation/DMP Quickstart Guide

Overview	3
Hardware and Software Requirements	3
Configuring Oracle Solaris with Symantec Storage Foundation for Fibre Channel and iSCSI	4
Configuring Linux with Symantec Storage Foundation for Fibre Channel.....	5
Setting the Red Hat Compatible Kernel as Default	5
Loading the <code>scsi_dh_alua</code> Module During Boot on Linux	5
Creating a UDEV Rule on Linux.....	6
Installing Required Storage Foundation Patches.....	6
Required Storage Foundation 5.x/6.x Patches	6
Required Storage Foundation Patch for Oracle Solaris and Red Hat Linux Operating Systems	6
Required Storage Foundation Patch for Microsoft Windows Operating Systems	7
Required Storage Foundation 5.x/6.x ASL/APM Package	7
Required Storage Foundation 5.x/6.x DMP Parameter Settings.....	7
Configuring Fibre Channel Mode on the Oracle ZFS Storage Appliance	7
Configuring Within the Storage Area Network (SAN) BUI Screen.....	8
Setting Target Ports.....	8
Enabling LUN Masking / Access Restrictions	9
Configuring Initiators and Initiator Group.....	10
Oracle ZFS Storage Appliance Clustering Considerations	12
Creating Projects to Group Your FC Target LUNs	13
Creating and Mapping FC Target LUNs.....	14
Mapping FC Target LUNs with Multiple Projects	15

Configuring iSCSI Devices	16
Clustering Considerations	16
Building the OS Recognizable Devices.....	17
Oracle Solaris	17
Linux operating systems.....	17
Microsoft Windows	17
Fibre Channel Topology	18
Example DMP/Veritas Volume Manager (VxVM) Views of a Clustered Oracle ZFS Storage ZS3-2	19
View of Oracle ZFS Storage ZS3-2 FC LUNs with both heads active	19
View of a clustered Oracle ZFS Storage Appliance array with one head down	22
View of a clustered Oracle ZFS Storage Appliance array ready for failback	26
Acronyms / Definitions	30
References	32

Overview

The purpose of this guide is to provide a quick overview on how to configure the Oracle ZFS Storage Appliance to provide Fibre Channel and iSCSI storage for use with Symantec Storage Foundation / Dynamic Multi-Pathing (DMP). It contains details on the required Storage Foundation patches, Array Support Library/Array Policy Module (ASL/APM) packages, and DMP settings to support the Oracle ZFS Storage Appliance.

The following installation and configuration steps are required and subsequently detailed in this guide:

- Verify that your Oracle ZFS Storage Appliance array is running software revision 2011.1.x or greater. Running the latest release version of the Oracle ZFS Storage Appliance Software is always recommended.
- Disable MPxIO on your Oracle Solaris hosts or do some configuration changes on your Linux hosts (see Configuring Linux for Symantec Storage Foundation for FC part in this guide) attached to the Oracle ZFS Storage Appliance array.
- Install the required Storage Foundation VRTSvxvm patches.
- Install the latest Storage Foundation ASL/APM package supporting the Oracle ZFS Storage Appliance series.
- For clustered Oracle ZFS Storage Appliance, set DMP parameters to support takeover / failback commands.
- Enable LUN masking by creating FC target and initiator groups.
- Create FC target LUNs in Oracle ZFS Storage Appliance projects.

This document assumes you have reviewed the Oracle ZFS Storage Appliance documentation, which can be found at the web site listed in the Resources section at the end of this document. The Oracle ZFS Storage Appliance has detailed online help, accessible by clicking the Help button in its browser user interface (BUI) on your system or simulator.

To benefit from this document, you should have already performed installation of the Oracle ZFS Storage Appliance (Software version 2011.1.x or greater) and created storage pools.

Hardware and Software Requirements

Be sure to reference Symantec's Hardware Compatibility List (HCL) and Hardware Tech Notes for the latest Symantec requirements. These reference checks include:

- Symantec's 5.1 HCL for the latest qualified UNIX OS platforms/protocols:
<http://www.symantec.com/business/support/index?page=content&id=TECH74012>
- Symantec's 6.x HCL for the latest qualified UNIX OS platforms/protocols:
<http://www.symantec.com/business/support/index?page=content&id=TECH170013>
- Symantec's HW tech notes for any parameter setting needed:
<http://www.symantec.com/business/support/index?page=content&id=TECH47728>

In order to support DMP, ensure that your Oracle ZFS Storage Appliance array is running Oracle ZFS Storage Appliance Software version 2011.1.x or greater.

Configuring Oracle Solaris with Symantec Storage Foundation for Fibre Channel and iSCSI

To configure Symantec Storage Foundation on Oracle Solaris, ensure that Oracle Solaris's MPxIO (Sun StorEdge Traffic Manager) is disabled. This allows Symantec Storage Foundation DMP kernel modules to perform the multipathing.

For Oracle Solaris versions 10 and 11, if MPxIO is not disabled, use the `stmsboot -d` command to disable it. It is recommended that you do not edit the `/kernel/drv/scsi_vhci.conf` file to disable MPxIO.

Refer to the following documentation for further details on enabling and disabling MPxIO:

- For Oracle Solaris 10:
http://docs.oracle.com/cd/E19957-01/819-0139/ch_4_config_multi_SW.html
- For Oracle Solaris 11:
http://docs.oracle.com/cd/E23824_01/html/E23097/gfpvv.html#scrolltoc.

Check the FC HBA vendor documentation of your Fibre Channel (FC) host bus adapter (HBA) for other required configuration settings specific to the HBA you are using.

Configuring Linux with Symantec Storage Foundation for Fibre Channel

To configure Symantec Storage Foundation on Linux operating systems, perform the following three operations:

- Ensure that your system uses the Red Hat Compatible Kernel by default.
- Make the needed SCSI device handlers available for the storage driver.
- Create a UDEV rule on Linux to ensure that devices are properly recognized with DMP in operation.

Setting the Red Hat Compatible Kernel as Default

Perform the following steps to make your system use the Red Hat Compatible Kernel by default:

1. Edit `/etc/grub.conf` and change the value of the default parameter to indicate the Red Hat Compatible Kernel. (Each entry for a bootable kernel in the file starts with a title definition. The entries are numbered from 0 upwards, where 0 corresponds to the first entry in the file, 1 to the second entry, and so on. To view the GRUB manual, use the `info grub` command.)
2. Edit `/etc/sysconfig/kernel` and change the setting for the default kernel package type from `DEFAULTKERNEL=kernel-uek` to `DEFAULTKERNEL=kernel`.
3. Reboot.

Be sure to check the FC HBA vendor documentation for other required configuration settings specific to the HBA you are using.

Loading the `scsi_dh_alua` Module During Boot on Linux

If appropriate SCSI device handlers (`scsi_dh` modules) are not available when the storage driver (for example, `lpfc`) is first loaded, I/O operations may be issued to SCSI multipath devices that are not ready for those I/O operations. This can result in significant delays during system boot and excessive I/O error messages in the kernel log.

To ensure the appropriate SCSI device handlers (`scsi_dh` modules) are available, specify the following kernel command line parameters:

Edit `/etc/grub.conf` to add `rdloaddriver=scsi_dh_alua`.

Creating a UDEV Rule on Linux

With dynamic multipathing, there can be problems detecting re-enabled OS devices on Linux that the udev device manager may have previously removed, or have been renamed when re-enabled. To avoid this situation, use the following workaround (for details, refer to [Red Hat 6](#) and [SLES 11](#)):

1. Download and upgrade to Red Hat Linux kernel 2.6.32-71.18.1.el6 or later, or Novell Linux kernel 2.6.27.45-0.1.1 or later.
2. Create file `/etc/udev/rules.d/40-rport.rules` with the following content line:

```
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports",  
ACTION=="add",RUN+="/bin/sh -c 'echo 20  
> /sys/class/fc_remote_ports/%k/fast_io_fail_tmo;echo 864000  
>/sys/class/fc_remote_ports/%k/dev_loss_tmo' "
```

3. Reboot the system.

If new LUNs are dynamically assigned to the host, users should run the following:

```
udevadm trigger --action=add --subsystem-match=fc_remote_ports
```

Installing Required Storage Foundation Patches

The latest released Storage Foundation 5.x/6.x patches and Oracle ZFS Storage Appliance array ASL/APM packages must be installed on the system.

Required Storage Foundation 5.x/6.x Patches

Ensure you are running the latest released Storage Foundation 5.x/6.x patches by checking the Symantec website at <https://vos.symantec.com/patch/matrix>.

Required Storage Foundation Patch for Oracle Solaris and Red Hat Linux Operating Systems

In addition, for clustered Oracle ZFS Storage Appliance configurations, you must install the hot fix VRTSvxvm 5.1RP2 or greater for Oracle Solaris and for Red Hat Linux. This patch installation requirement addresses an issue for clustered Oracle ZFS Storage Appliance fault-injected failovers and takeover/failback commands (Symantec issue/etrack e2046696) and `dmp_lun_retry_timeout` handling problems.

Required Storage Foundation Patch for Microsoft Windows Operating Systems

For Oracle ZFS Storage ZS3-2 support on Storage Foundation Windows 6.x, you must install the hot fix win_solutions-win_x64-Hotfix_6_1_0_454_3481564.

Required Storage Foundation 5.x/6.x ASL/APM Package

With Storage Foundation, all array ASLs/APMs are bundled into one package. The ASL/APM tech note is located at: <https://sort.symantec.com/asl>.

Required Storage Foundation 5.x/6.x DMP Parameter Settings

In order to support clustered Oracle ZFS Storage Appliance takeover/failback commands, the following DMP tunables must be set:

- `vxddmpadm settune dmp_health_time=0`
- `vxddmpadm settune dmp_path_age=0`
- `vxddmpadm settune dmp_lun_retry_timeout=200`

In normal situations, a clustered Oracle ZFS Storage Appliance takes a few seconds to fail over a ZFS pool/resource (FC target LUNs) to its second node, but to account for any unexpected scenarios in which a takeover takes more time, set the `dmp_lun_retry_timeout` to 200. If for some reason you are seeing 'Disabled Dmpnode' messages for the Oracle ZFS Storage Appliance FC LUNs in the `dmpevents.log` file, you may need to increase the lun retry timeout value. The 0 settings of the `dmp_health_time` and `dmp_path_age` tunables are required. This allows DMP to re-enable the path within the short period of time that it takes the Oracle ZFS Storage Appliance to complete takeover/failback.

Configuring Fibre Channel Mode on the Oracle ZFS Storage Appliance

To configure Fibre Channel mode on the Oracle ZFS Storage Appliance, at least one Fibre Channel port on the Oracle ZFS Storage Appliance must be configured to use target mode. After the desired ports are configured as targets, you need to create and configure target groups, FC initiators, and initiator groups. Then you can map a LUN to the client(s) on the SAN.

Consult the *Oracle ZFS Storage Appliance Administration Guide* (820–4167–11) or the Oracle ZFS Storage Appliance online help for more details. The following sections present examples for configuring the Oracle ZFS Storage Appliance.

Configuring Within the Storage Area Network (SAN) BUI Screen

Use the Oracle ZFS Storage Appliance BUI and navigate from Configuration to the Storage Area Network (SAN) screen, where you can configure and check the following:

- Setting target ports
- Enabling LUN masking and any access restrictions
- Configuring FC initiators and initiator groups

Setting Target Ports

By default, all Oracle ZFS Storage Appliance FC ports are defined as targets.

From the BUI Configuration-->SAN screen

Storage Area Network (SAN) Fibre Channel iSCSI SRP

To share LUNs only via particular targets or to particular initiators, build Target Groups and Initiator Groups, respectively. To create a group or add to an existing one, drag the entity from the left to the table on the right. REVERT APPLY

Ports		Initiators		Target Groups	
PCle 5		Port 1 8 Gbps	3 discovered ports	NAME	TARGETS
		21:00:00:24:ff:49:e8:c0		default	[ALL PORTS]
		Port 2 8 Gbps	3 discovered ports	fc-tg	21:00:00:24:ff:49:e8:c0 21:00:00:24:ff:49:e8:c1 21:00:00:24:ff:49:e8:e2 21:00:00:24:ff:49:e8:e3
		21:00:00:24:ff:49:e8:c1			PCle 5: Port 1 PCle 5: Port 2 aie-7330e-h2: PCle 5: Port 1 aie-7330e-h2: PCle 5: Port 2
PCle 6		Port 1 link down			
		21:00:00:24:ff:49:e8:94			
		Port 2 link down			
		21:00:00:24:ff:49:e8:95			

Figure 1. Configuring FC target ports in the Oracle ZFS Storage Appliance BUI

Click the “i” icon for a display of Port Details. It is recommended that you alias each HBA – for example, setting the alias to a hostname and its port – as seen in the following figure.

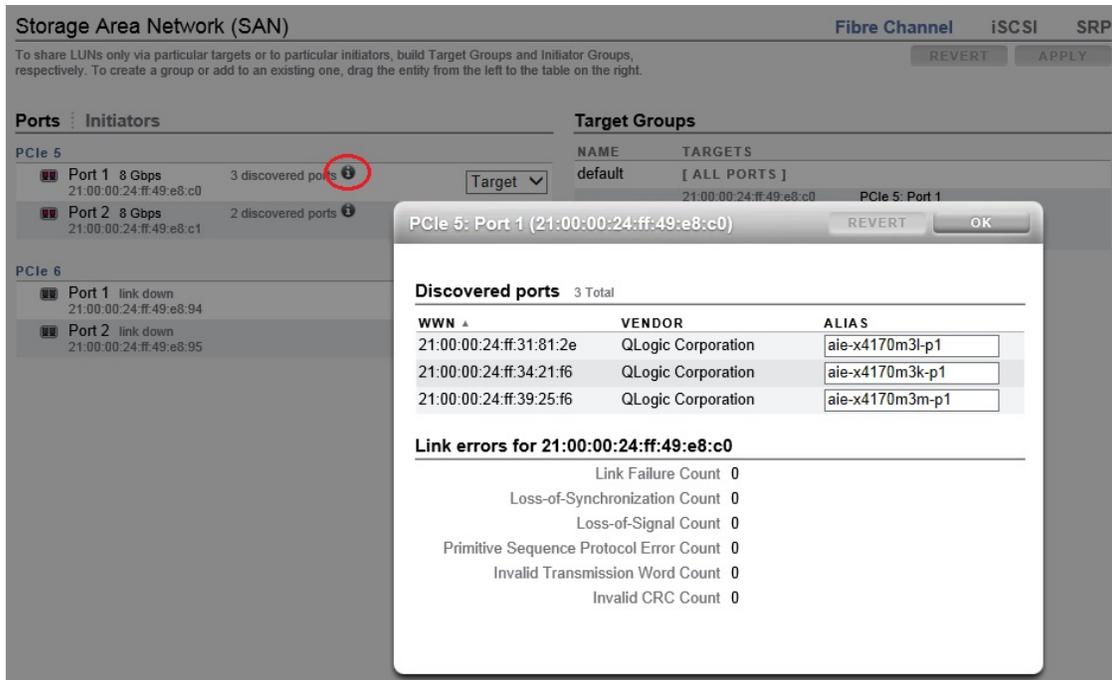


Figure 2. FC target port details including alias

Enabling LUN Masking / Access Restrictions

The Configuration > SAN screen allows for specific access restrictions to be set based on initiators and targets. This is also known as LUN masking. Alternatively, a LUN may be assigned to the default groups (initiator and target), giving all initiators access. This practice is not recommended as it greatly increases the risk of non-authorized initiators gaining access to an FC target, thereby increasing the chance of data corruption or unwanted data deletion.

If DMP is going to be used on the initiator (client side), multiple targets will be used when redundant connectivity is used between the Oracle ZFS Storage Appliance and a switch or a redundant point-to-point connection between the Oracle ZFS Storage Appliance and the server.

After an FC target port has been defined, drag and drop it into the FC Target Groups section of the BUI and edit the new FC target group to rename the alias. Select all other FC target ports to be used for this group.

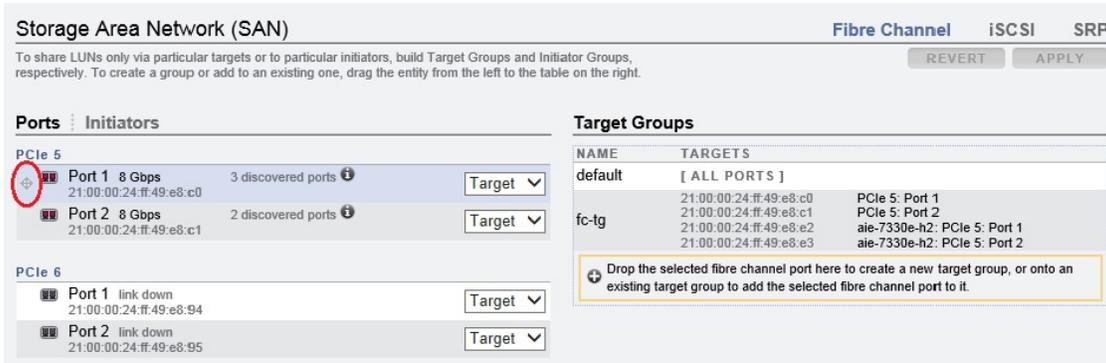


Figure 3. Creating a Fibre Channel target group

See the following illustrations for the completed setup of a target group.

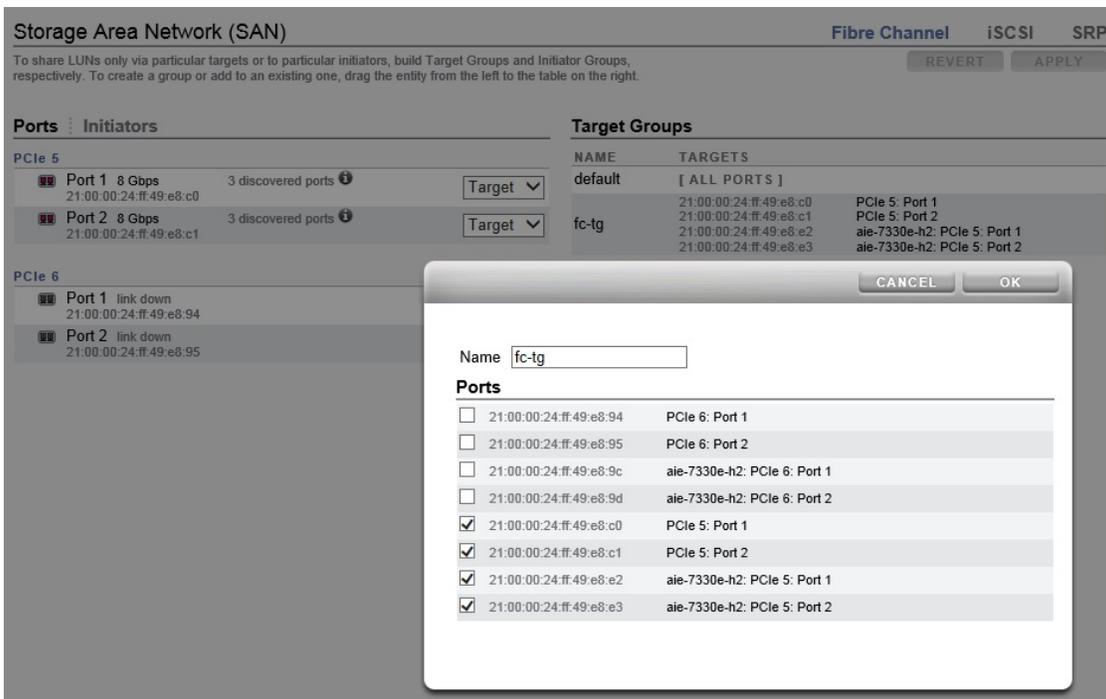


Figure 4. Assigning FC ports to a target group

Configuring Initiators and Initiator Group

Get the WWPN from the client initiator and create a new FC initiator on the Oracle ZFS Storage Appliance.

From the BUI, navigate to Configuration>SAN>Fibre Channel Initiators>Initiators “+” add icon.

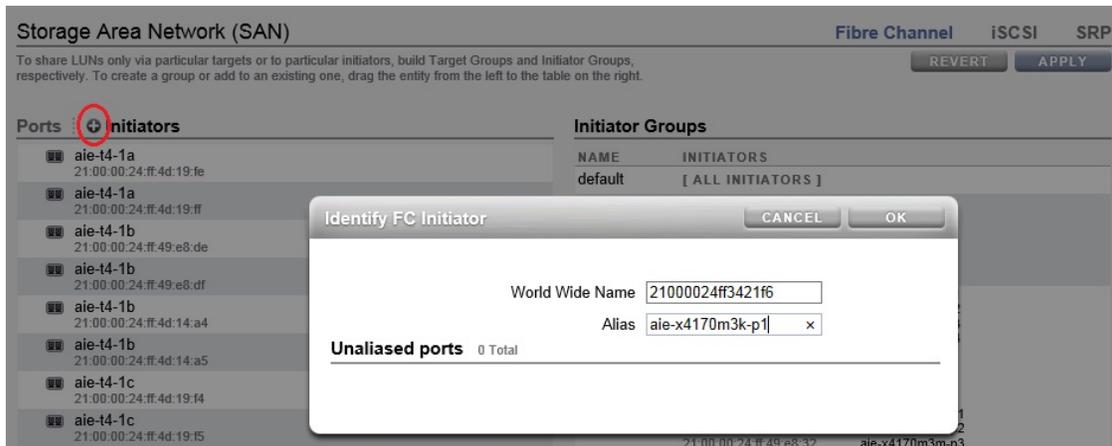


Figure 5. Adding a World Wide Name for a new FC initiator

Next, drag and drop the newly created FC initiator into the Initiator Group section of the BUI.

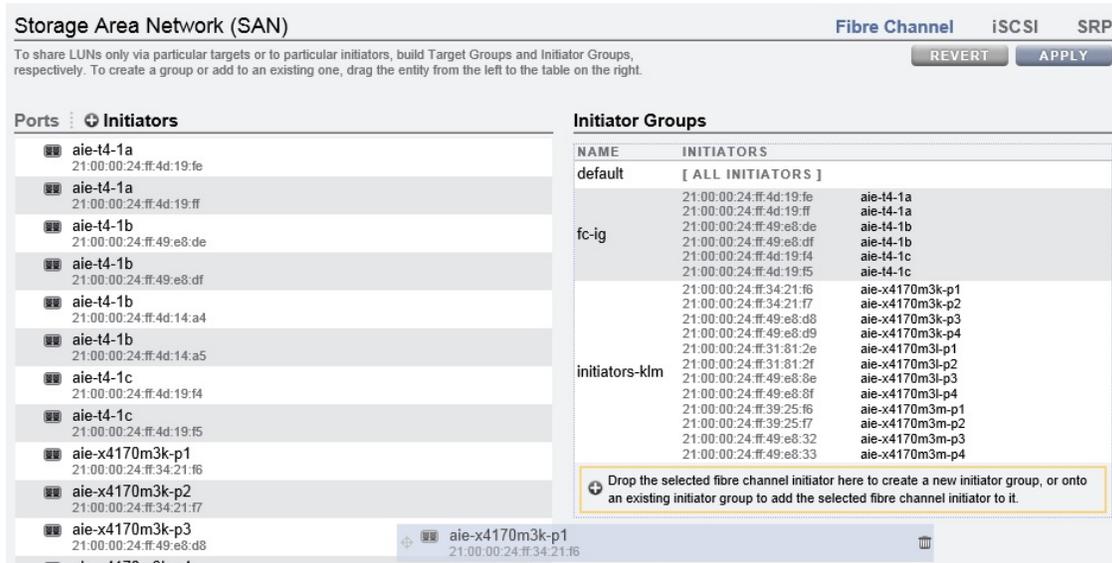


Figure 6. Creating a new initiator group

Edit the new Initiator Group alias and add any additional initiators that are required for DMP to this group.

Storage Area Network (SAN) Fibre Channel iSCSI SRP

To share LUNs only via particular targets or to particular initiators, build Target Groups and Initiator Groups, respectively. To create a group or add to an existing one, drag the entity from the left to the table on the right. REVERT APPLY

Ports Initiators

- aie-t4-1a
21:00:00:24:ff:4d:19:fe
- aie-t4-1a
21:00:00:24:ff:4d:19:ff
- aie-t4-1b
21:00:00:24:ff:49:e8:de
- aie-t4-1b
21:00:00:24:ff:49:e8:df
- aie-t4-1b
21:00:00:24:ff:4d:14:a4
- aie-t4-1b
21:00:00:24:ff:4d:14:a5
- aie-t4-1c
21:00:00:24:ff:4d:19:f4
- aie-t4-1c
21:00:00:24:ff:4d:19:f5
- aie-x4170m3k-p1
21:00:00:24:ff:34:21:f6
- aie-x4170m3k-p2
21:00:00:24:ff:34:21:f7
- aie-x4170m3k-p3
21:00:00:24:ff:49:e8:d8
- aie-x4170m3k-p4
21:00:00:24:ff:49:e8:d9
- aie-x4170m3l-p1
21:00:00:24:ff:31:81:2e
- aie-x4170m3l-p2
21:00:00:24:ff:31:81:2f
- aie-x4170m3l-p3
21:00:00:24:ff:49:e8:8e
- aie-x4170m3l-p4
21:00:00:24:ff:49:e8:8f
- aie-x4170m3m-p1
21:00:00:24:ff:39:25:f6
- aie-x4170m3m-p2
21:00:00:24:ff:39:25:f7
- aie-x4170m3m-p3
21:00:00:24:ff:49:e8:32

Initiator Groups

NAME	INITIATORS
default	[ALL INITIATORS]
fc-ig	21:00:00:24:ff:4d:19:fe aie-t4-1a
	21:00:00:24:ff:4d:19:ff aie-t4-1a
	21:00:00:24:ff:49:e8:de aie-t4-1b
	21:00:00:24:ff:49:e8:df aie-t4-1b
	21:00:00:24:ff:4d:19:f4 aie-t4-1c
initiators-klm	21:00:00:24:ff:4d:19:f5 aie-t4-1c
	21:00:00:24:ff:34:21:f6 aie-x4170m3k-p1
	21:00:00:24:ff:34:21:f7 aie-x4170m3k-p2
	21:00:00:24:ff:49:e8:d8 aie-x4170m3k-p3
	21:00:00:24:ff:49:e8:d9 aie-x4170m3k-p4
	21:00:00:24:ff:31:81:2e aie-x4170m3l-p1
	21:00:00:24:ff:31:81:2f aie-x4170m3l-p2
	21:00:00:24:ff:49:e8:8e aie-x4170m3l-p3

Name

Initiators

<input type="checkbox"/>	21:00:00:24:ff:4d:19:fe	aie-t4-1a
<input type="checkbox"/>	21:00:00:24:ff:4d:19:ff	aie-t4-1a
<input type="checkbox"/>	21:00:00:24:ff:49:e8:de	aie-t4-1b
<input type="checkbox"/>	21:00:00:24:ff:49:e8:df	aie-t4-1b
<input type="checkbox"/>	21:00:00:24:ff:4d:14:a4	aie-t4-1b
<input type="checkbox"/>	21:00:00:24:ff:4d:14:a5	aie-t4-1b
<input type="checkbox"/>	21:00:00:24:ff:4d:19:f4	aie-t4-1c
<input type="checkbox"/>	21:00:00:24:ff:4d:19:f5	aie-t4-1c
<input checked="" type="checkbox"/>	21:00:00:24:ff:34:21:f6	aie-x4170m3k-p1
<input checked="" type="checkbox"/>	21:00:00:24:ff:34:21:f7	aie-x4170m3k-p2
<input checked="" type="checkbox"/>	21:00:00:24:ff:49:e8:d8	aie-x4170m3k-p3
<input checked="" type="checkbox"/>	21:00:00:24:ff:49:e8:d9	aie-x4170m3k-p4

CANCEL OK

Figure 7. Editing assignees to an FC initiator group

Oracle ZFS Storage Appliance Clustering Considerations

In an Oracle ZFS Storage Appliance cluster, initiators will have two paths (or sets of paths) to each LUN: one path (or set of paths) will be to the head that has imported the storage associated with the LUN; the other path (or set of paths) will be to that head's clustered peer. The first path (or set of paths) is active; the second path (or set of paths) is standby; in the event of a takeover, the active paths will become unavailable, and the standby paths will (after a short time) be transitioned to be active, after which I/O will continue. This approach to multipathing is known as asymmetric logical unit access (ALUA) and – when coupled with an ALUA-aware initiator – allows cluster takeover to be transparent to higher-level applications.

When sizing an Oracle ZFS Storage series system for use in a cluster, two additional considerations gain importance. Perhaps the most important decision is whether all storage pools will be assigned ownership to the same head, or split between them. Please reference the “Clustering Consideration for Storage” section in the *Oracle ZFS Storage Appliance Administration Guide* for more details.

Creating Projects to Group Your FC Target LUNs

Once you have created initiator and target groups, it is recommended that you group your FC target LUNs by creating a project. To create a project from the BUI, select Shares>PROJECTS and then click on the “+” symbol to the left of Projects. Projects are grouped by storage pools.

The screenshot displays the Oracle ZFS Storage BUI interface. At the top, the system is identified as 'ORACLE ZFS STORAGE ZS3-2'. The user is logged in as 'Super-User@aie-7330e-h1'. The main navigation bar includes 'Configuration', 'Maintenance', 'Shares', 'Status', and 'Analytics'. The 'Shares' section is active, showing 'PROJECTS' and 'SCHEMA' sub-sections. The 'Projects' section is expanded, showing 'All Projects'. On the left, usage statistics are displayed: 1.7% of 7.20T referenced data, 89.4M snapshot data, and 127G total space. The main area shows a table of projects with columns for NAME, SIZE, and CREATION. A 'Create Project' dialog box is open, with the name 'FC-target-luns-for-klm' entered in the text field. The dialog box has 'CANCEL' and 'APPLY' buttons.

NAME	SIZE	CREATION
7120_LUNs	42.0G	2013-12-23 16:04:42
7330_LUNs	42.3G	2014-1-18 20:25:58
7410_LUNs	42.7G	2013-12-23 15:53:56
default	106K	2013-12-9 12:28:39

Figure 8. Creating a project to group FC target LUNs

Creating and Mapping FC Target LUNs

Storage that is offered through the FC protocol requires the creation of a LUN, not a filesystem. From the Configuration>SAN>Initiators BUI, you can create storage for the initiator group by clicking on the Add LUN icon on the right side of the initiator group as you pass your mouse over it, as illustrated in the following figure.

The screenshot shows the 'Storage Area Network (SAN)' configuration page. It has tabs for 'Fibre Channel', 'iSCSI', and 'SRP'. Below the tabs, there are instructions: 'To share LUNs only via particular targets or to particular initiators, build Target Groups and Initiator Groups, respectively. To create a group or add to an existing one, drag the entity from the left to the table on the right.' There are 'REVERT' and 'APPLY' buttons.

On the left, under 'Ports' and 'Initiators', there is a list of initiators with their WWPNs:

- aie-t4-1a (21:00:00:24:ff:4d:19:fe)
- aie-t4-1a (21:00:00:24:ff:4d:19:ff)
- aie-t4-1b (21:00:00:24:ff:49:e8:de)
- aie-t4-1b (21:00:00:24:ff:49:e8:df)
- aie-t4-1b (21:00:00:24:ff:4d:14:a4)
- aie-t4-1b (21:00:00:24:ff:4d:14:a5)
- aie-t4-1c (21:00:00:24:ff:4d:19:f4)
- aie-t4-1c (21:00:00:24:ff:4d:19:f5)
- aie-x4170m3k-p1 (21:00:00:24:ff:34:21:86)

On the right, under 'Initiator Groups', there is a table:

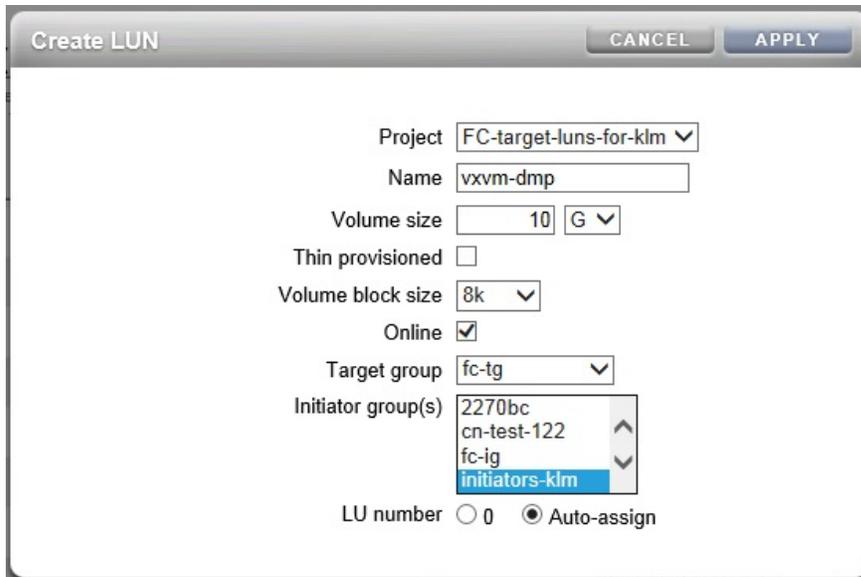
NAME	INITIATORS
default	[ALL INITIATORS]
fc-ig	21:00:00:24:ff:4d:19:fe aie-t4-1a
	21:00:00:24:ff:4d:19:ff aie-t4-1a
	21:00:00:24:ff:49:e8:de aie-t4-1b
	21:00:00:24:ff:49:e8:df aie-t4-1b
initiators-klm	21:00:00:24:ff:4d:19:f4 aie-t4-1c
	21:00:00:24:ff:4d:19:f5 aie-t4-1c
	21:00:00:24:ff:34:21:86 aie-x4170m3k-p1
	21:00:00:24:ff:34:21:f7 aie-x4170m3k-p2
	21:00:00:24:ff:49:e8:d8 aie-x4170m3k-p3
	21:00:00:24:ff:49:e8:d9 aie-x4170m3k-p4
	21:00:00:24:ff:31:81:2e aie-x4170m3l-p1
	21:00:00:24:ff:31:81:2f aie-x4170m3l-p2
	21:00:00:24:ff:49:e8:8e aie-x4170m3l-p3
	21:00:00:24:ff:49:e8:8f aie-x4170m3l-p4
	21:00:00:24:ff:39:25:86 aie-x4170m3m-p1
	21:00:00:24:ff:39:25:f7 aie-x4170m3m-p2
	21:00:00:24:ff:49:e8:32 aie-x4170m3m-p3
	21:00:00:24:ff:49:e8:33 aie-x4170m3m-p4

A tooltip 'Add a LUN for this initiator group to the current pool' is shown over the 'initiators-klm' group, with a red circle highlighting the 'Add LUN' icon.

Figure 9. Creating and mapping FC target LUNs

The Create LUN dialog window provides the entry fields to specify a project, give the LUN a name, set the Volume size, allow for Thin provisioning, set the Volume block size, select the Target group and Initiator group, assign an LU number (if you do not want to let the system do this automatically), and set the status to online or offline.

Target group and initiator group setup allows for assigning the LUN a specific target or targets and allowing specific initiators to access those targets, as shown in the following screenshot.



The image shows a 'Create LUN' dialog window with the following fields and options:

- Project:** FC-target-luns-for-klm (dropdown)
- Name:** vxvm-dmp (text input)
- Volume size:** 10 G (text input with dropdown)
- Thin provisioned:**
- Volume block size:** 8k (dropdown)
- Online:**
- Target group:** fc-tg (dropdown)
- Initiator group(s):** A list box containing 2270bc, cn-test-122, fc-ig, and initiators-klm (with 'initiators-klm' selected).
- LU number:** 0 Auto-assign

Figure 10. Create LUN dialog window

FC target LUNs can be created from the BUI, choosing Shares>Shares>LUNs windows or the CLI interface. Review the Oracle ZFS Storage Appliance online help or the administration guide for more details.

Mapping FC Target LUNs with Multiple Projects

The following BUI illustration (Shares>Shares) shows four projects with 64 FC LUNs. Click the “+” sign to the left of the LUNs to create more FC target LUNs and map those to specific projects and initiator/target groups.

The screenshot displays the Oracle ZFS Storage Appliance GUI. The top navigation bar includes 'Configuration', 'Maintenance', 'Shares' (highlighted), 'Status', and 'Analytics'. Below this, there are tabs for 'POOLS', 'POOL2', 'SHARES', 'PROJECTS', and 'SCHEMA'. The main content area shows a 'Projects' sidebar on the left and a table of LUNs on the right. The table has columns for 'NAME', 'SIZE', and 'GUID'. The LUNs listed are:

NAME	SIZE	GUID
7410_LUNs / 7410_lun7	2G	600144F0A1A8E7EE000052983DEB0017
7410_LUNs / 7410_lun8	2G	600144F0A1A8E7EE000052983DEC0018
7410_LUNs / 7410_lun9	2G	600144F0A1A8E7EE000052983DEE0019
FC-target-luns-for-klm / vxvm-dmp	10G	600144F0A9E3B5DC000052FB1D450004

Figure 11. BUI view of FC target LUNs with multiple projects

Configuring iSCSI Devices

It is recommended that you consult the online help iSCSI configuration guide. A few important points that follow which have been extracted from the online help iSCSI help page:

When using the iSCSI protocol, the target portal refers to the unique combination of an IP address and TCP port number by which an initiator can contact a target.

When using the iSCSI protocol, a target portal group is a collection of target portals. Target portal groups are managed transparently; each network interface has a corresponding target portal group with that interface's active addresses. Binding a target to an interface advertises that iSCSI target using the portal group associated with that interface.

An IQN (iSCSI qualified name) is the unique identifier of a device in an iSCSI network. iSCSI uses the form `iqn.date.authority:uniqueid` for IQNs. For example, the Oracle ZFS Storage Appliance may use the IQN: `iqn.1986-03.com.sun:02:c7824a5b-f3ea-6038-c79d-ca443337d92c` to identify one of its iSCSI targets. This name shows that this is an iSCSI device built by a company registered in March of 1986. The naming authority is just the DNS name of the company reversed, in this case, "com.sun". Everything following is a unique ID that Sun uses to identify the target.

Clustering Considerations

On clustered platforms, targets that have at least one active interface on that cluster node will be online. Take care when assigning interfaces to targets; a target may be configured to use portal groups on disjointed head nodes. In that situation, the target will be online on both heads yet will export

different LUNs, depending on the storage owned by each head node. As network interfaces migrate between cluster heads as part of takeover/failback or ownership changes, iSCSI targets will move online and offline as their respective network interfaces are imported and exported.

Targets that are bound to an IPMP interface will be advertised only through the addresses of that IPMP group. The target will not be reachable through that group's test addresses. Targets bound to interfaces built on top of an LACP aggregation will use the address of that aggregation. If an LACP aggregation is added to an IPMP group, a target can no longer use that aggregation's interface, as that address will become an IPMP test address.

When planning your iSCSI client configuration, you will need answers for the following:

- What initiators (and their IQNs) will be accessing the SAN?
- If you plan on using CHAP authentication, what CHAP credentials does each initiator use?
- How many iSCSI disks (LUNs) are required, and how big should they be?
- Do the LUNs need to be shared between multiple initiators?

Building the OS Recognizable Devices

To make the newly defined LUNs accessible by the operating system, use the following instructions for the corresponding OS.

Oracle Solaris

For Oracle Solaris, run the `devfsadm -Cv` as user root. This command creates the related special device path files.

Linux operating systems

For Linux, use the following command to scan storage interconnects:

```
echo "1" > /sys/class/fc_host/hostX/issue_lip
```

Microsoft Windows

On the Windows operating system, launch Veritas Enterprise Administrator (VEA) > Actions > Rescan to get new devices.

The following example shows output for the Oracle Solaris `format` command for an Oracle Solaris host that has four paths to 10 Oracle ZFS Storage Appliance 7310 FC LUNs (five FC target LUNs off head 1 and five FC target LUNs off head 2).

```
isv-2000a# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:

    0. clt2101001B32345466d0 <SUN-SunStorage7310-1.0 cyl 1785 alt 2 hd 254 sec 254>
        /pci@7c0/pci@0/pci@1/pci@0,2/SUNW,qlc@1/fp@0,0/ssd@w2101001b32345466,0
    1. clt2101001B32345466d1 <SUN-SunStorage7310-1.0 cyl 3248 alt 2 hd 254 sec 254>
        /pci@7c0/pci@0/pci@1/pci@0,2/SUNW,qlc@1/fp@0,0/ssd@w2101001b32345466,1
    2. clt2101001B32345466d2 <SUN-SunStorage7310-1.0 cyl 3248 alt 2 hd 254 sec 254>
        /pci@7c0/pci@0/pci@1/pci@0,2/SUNW,qlc@1/fp@0,0/ssd@w2101001b32345466,2
        ...
    40. c3t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
        /pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2/sd@0,0

Specify disk (enter its number):
```

Fibre Channel Topology

The Oracle ZFS Storage Appliance series with software version 2011.1.x or later supports direct and switched-attached fabrics. The following illustration is an example of a Veritas Cluster Server (VCS) setup with redundant paths to all FC connects. Ensure you have redundant FC connections. An FC failure outside of the Oracle ZFS Storage Appliance will not trigger an automatic path failover.

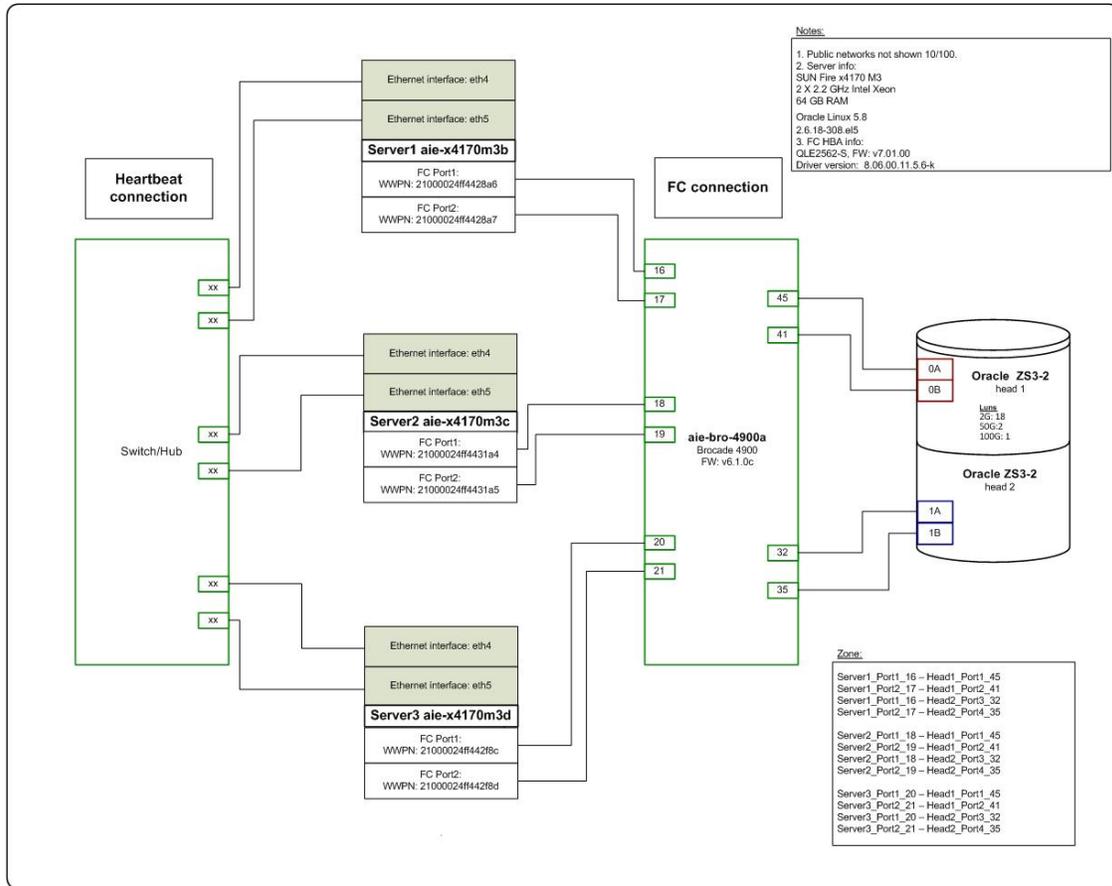


Figure 12. FC topology of VCS server connected to a clustered Oracle ZFS Storage ZS3-2

Example DMP/Veritas Volume Manager (VxVM) Views of a Clustered Oracle ZFS Storage ZS3-2

The following examples show details for various states of a clustered Oracle ZFS Storage ZS3-2 as seen through the BUI as well as CLI sessions.

View of Oracle ZFS Storage ZS3-2 FC LUNs with both heads active

The following example is from an Oracle Linux 5.8 host hooked up to an Oracle ZFS Storage ZS3-2 cluster which has 4 paths to each ZS3-2 FC LUN. The example assumes you have already run `vxdisksetup` and added the ZS3-2 FC LUNs to VxVM disk groups.

```
[root@aie-x4170m3b ~]# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
--------	------	------	-------	--------

```

disk_0      auto:none      -          -          online invalid
disk_1      auto:none      -          -          online invalid
zfs_73300_0 auto:cdsdisk    zfs_73300_0 dg1        online shared
zfs_73300_1 auto:cdsdisk    zfs_73300_1 dg2        online shared
zfs_73300_2 auto:cdsdisk    zfs_73300_2 dg1        online shared
zfs_73300_3 auto:cdsdisk    zfs_73300_3 dg3        online shared
zfs_73300_4 auto:cdsdisk    zfs_73300_4 dg1        online shared
zfs_73300_5 auto:cdsdisk    zfs_73300_5 dg2        online shared
zfs_73300_6 auto:cdsdisk    zfs_73300_6 dg3        online shared
zfs_73300_7 auto:cdsdisk    zfs_73300_7 dg1        online shared
zfs_73300_8 auto:cdsdisk    zfs_73300_8 dg2        online shared
zfs_73300_9 auto:cdsdisk    zfs_73300_9 dg3        online shared
zfs_73300_10 auto:cdsdisk    -          -          online
zfs_73300_11 auto:cdsdisk    -          -          online
zfs_73300_12 auto:cdsdisk    zfs_73300_12 dg3        online shared
zfs_73300_13 auto:cdsdisk    zfs_73300_13 dg1        online shared
zfs_73300_14 auto:cdsdisk    zfs_73300_14 dg2        online shared
zfs_73300_15 auto:cdsdisk    zfs_73300_15 dg3        online shared
zfs_73300_16 auto:cdsdisk    -          -          online
zfs_73300_17 auto:cdsdisk    zfs_73300_17 dg1        online shared
zfs_73300_18 auto:cdsdisk    zfs_73300_18 dg2        online shared
zfs_73300_19 auto:cdsdisk    zfs_73300_19 dg3        online shared
zfs_73300_20 auto:cdsdisk    zfs_73300_20 dg2        online shared

```

```
[root@aie-x4170m3b ~]# vxdisk list zfs_73300_0
```

```

Device:      zfs_73300_0
devicetag:   zfs_73300_0
type:        auto
clusterid:   clus_us
disk:        name=zfs_73300_0 id=1396091048.157.aie-x4170m3b
group:       name=dg1 id=1398139104.41.aie-x4170m3b
info:        format=cdsdisk,privoffset=256,pubslice=3,privslice=3
flags:       online ready private autoconfig shared autoimport imported

```

```
pubpaths:  block=/dev/vx/dmp/zfs_73300_0s3 char=/dev/vx/rdmp/zfs_73300_0s3
guid:      {d99aaef0-b731-11e3-8304-0f11e44e9518}
udid:
SUN%5FZFS%20Storage%207330%5FSUN%5F7x10%5FSeries%5F600144F08F8B540E00005334C7
F70016
site:     -
version:  3.1
iosize:   min=512 (bytes) max=1024 (blocks)
public:   slice=3 offset=65792 len=4120320 disk_offset=0
private:  slice=3 offset=256 len=65536 disk_offset=0
update:   time=1398140448 seqno=0.260
ssb:      actual_seqno=0.0
headers:  0 240
configs:  count=1 len=51360
logs:     count=1 len=4096
Defined regions:
config  priv 000048-000239[000192]: copy=01 offset=000000 enabled
config  priv 000256-051423[051168]: copy=01 offset=000192 enabled
log     priv 051424-055519[004096]: copy=01 offset=000000 disabled
lockrgn priv 055520-055663[000144]: part=00 offset=000000
Multipathing information:
numpaths: 4
sdas      state=enabled  type=secondary
sdf       state=enabled  type=primary
sdbf      state=enabled  type=secondary
sdn       state=enabled  type=primary
[root@aie-x4170m3b ~]#
```

```
[root@aie-x4170m3b ~]# vxdmpadm getsubpaths
```

NAME	STATE[A]	PATH-TYPE[M]	DMPNODENAME	ENCLR-NAME	CTLR	ATTRS
sdb	ENABLED(A)	-	disk_0	disk	c0	-
sda	ENABLED(A)	-	disk_1	disk	c0	-

```

sdas          ENABLED(A) SECONDARY   zfs_73300_0  zfs_73300    c2          -
sdbf          ENABLED(A) SECONDARY   zfs_73300_0  zfs_73300    c3          -
sdf           ENABLED      PRIMARY     zfs_73300_0  zfs_73300    c2          -
sdn           ENABLED      PRIMARY     zfs_73300_0  zfs_73300    c3          -
sdaq          ENABLED(A) SECONDARY   zfs_73300_1  zfs_73300    c2          -
sdbd          ENABLED(A) SECONDARY   zfs_73300_1  zfs_73300    c3          -
sde           ENABLED      PRIMARY     zfs_73300_1  zfs_73300    c2          -
sdl           ENABLED      PRIMARY     zfs_73300_1  zfs_73300    c3          -
...
sdaa          ENABLED      PRIMARY     zfs_73300_9  zfs_73300    c3          -
sdbe          ENABLED(A) SECONDARY   zfs_73300_9  zfs_73300    c2          -
sdbp          ENABLED(A) SECONDARY   zfs_73300_9  zfs_73300    c3          -
sdq           ENABLED      PRIMARY     zfs_73300_9  zfs_73300    c2          -
[root@aie-x4170m3b ~]#

```

View of a clustered Oracle ZFS Storage Appliance array with one head down

The following exhibits an example in which one head (node) of a clustered Oracle ZFS Storage ZS3-2 is turned off (or down). Review the illustration and VxVM commands that follow for details on how DMP views the standby paths in this scenario.

To recover from this issue, first power on the Oracle ZFS Storage Appliance head (node). After the node joins the Oracle ZFS Storage Appliance cluster, the 'Failback' button will be highlighted. Execute the Failback button in the Oracle ZFS Storage Appliance BUI to transfer the ZFS pool resource to the correct head (node). Then run `vxdisk scandisks` or wait 3 minutes for DMP to automatically rescan the devices.



Figure 13. Oracle ZFS Storage ZS3-2 with failed head

```
[root@aie-x4170m3b ~]# vxdisk scandisks
[root@aie-x4170m3b ~]# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
disk_0	auto:none	-	-	online invalid
disk_1	auto:none	-	-	online invalid
zfs_73300_0	auto:cdsdisk	zfs_73300_0	dg1	online shared
zfs_73300_1	auto:cdsdisk	zfs_73300_1	dg2	online shared
zfs_73300_2	auto:cdsdisk	zfs_73300_2	dg1	online shared
zfs_73300_3	auto:cdsdisk	zfs_73300_3	dg3	online shared
zfs_73300_4	auto:cdsdisk	zfs_73300_4	dg1	online shared
zfs_73300_5	auto:cdsdisk	zfs_73300_5	dg2	online shared
zfs_73300_6	auto:cdsdisk	zfs_73300_6	dg3	online shared
zfs_73300_7	auto:cdsdisk	zfs_73300_7	dg1	online shared
zfs_73300_8	auto:cdsdisk	zfs_73300_8	dg2	online shared
zfs_73300_9	auto:cdsdisk	zfs_73300_9	dg3	online shared
zfs_73300_10	auto:cdsdisk	-	-	online

```
zfs_73300_11 auto:cdsdisk - - online
zfs_73300_12 auto:cdsdisk zfs_73300_12 dg3 online shared
zfs_73300_13 auto:cdsdisk zfs_73300_13 dg1 online shared
zfs_73300_14 auto:cdsdisk zfs_73300_14 dg2 online shared
zfs_73300_15 auto:cdsdisk zfs_73300_15 dg3 online shared
zfs_73300_16 auto:cdsdisk - - online
zfs_73300_17 auto:cdsdisk zfs_73300_17 dg1 online shared
zfs_73300_18 auto:cdsdisk zfs_73300_18 dg2 online shared
zfs_73300_19 auto:cdsdisk zfs_73300_19 dg3 online shared
zfs_73300_20 auto:cdsdisk zfs_73300_20 dg2 online shared
[root@aie-x4170m3b ~]#
```

```
[root@aie-x4170m3b ~]# vxdisk list zfs_73300_0
```

```
Device: zfs_73300_0
devicetag: zfs_73300_0
type: auto
clusterid: clus_us
disk: name=zfs_73300_0 id=1396091048.157.aie-x4170m3b
group: name=dg1 id=1398139104.41.aie-x4170m3b
info: format=cdsdisk,privoffset=256,pubslice=3,privslice=3
flags: online ready private autoconfig shared autoimport imported
pubpaths: block=/dev/vx/dmp/zfs_73300_0s3 char=/dev/vx/rdmp/zfs_73300_0s3
guid: {d99aaef0-b731-11e3-8304-0f11e44e9518}
udid:
SUN%5FZFS%20Storage%207330%5FSUN%5F7x10%5FSeries%5F600144F08F8B540E00005334C7
F70016
site: -
version: 3.1
iosize: min=512 (bytes) max=1024 (blocks)
public: slice=3 offset=65792 len=4120320 disk_offset=0
private: slice=3 offset=256 len=65536 disk_offset=0
update: time=1398140448 seqno=0.260
ssb: actual_seqno=0.0
headers: 0 240
```

```

configs:   count=1 len=51360
logs:     count=1 len=4096
Defined regions:
  config   priv 000048-000239[000192]: copy=01 offset=000000 enabled
  config   priv 000256-051423[051168]: copy=01 offset=000192 enabled
  log      priv 051424-055519[004096]: copy=01 offset=000000 disabled
  lockrgn  priv 055520-055663[000144]: part=00 offset=000000
Multipathing information:
numpaths: 4
sdas      state=enabled  type=primary
sdf       state=disabled  type=primary
sdbf      state=enabled  type=primary
sdn       state=disabled  type=primary
[root@aie-x4170m3b ~]#

```

```
[root@aie-x4170m3b ~]# vxddmpadm getsubpaths
```

NAME	STATE[A]	PATH-TYPE[M]	DMPNODENAME	ENCLR-NAME	CTLR	ATTRS
sdb	ENABLED(A)	-	disk_0	disk	c0	-
sda	ENABLED(A)	-	disk_1	disk	c0	-
sdas	ENABLED(A)	PRIMARY	zfs_73300_0	zfs_73300	c2	-
sdbf	ENABLED(A)	PRIMARY	zfs_73300_0	zfs_73300	c3	-
sdf	DISABLED	PRIMARY	zfs_73300_0	zfs_73300	c2	-
sdn	DISABLED	PRIMARY	zfs_73300_0	zfs_73300	c3	-
sdaq	ENABLED(A)	PRIMARY	zfs_73300_1	zfs_73300	c2	-
sdbd	ENABLED(A)	PRIMARY	zfs_73300_1	zfs_73300	c3	-
sde	DISABLED	PRIMARY	zfs_73300_1	zfs_73300	c2	-
sd1	DISABLED	PRIMARY	zfs_73300_1	zfs_73300	c3	-
...						
sdaa	DISABLED	PRIMARY	zfs_73300_9	zfs_73300	c3	-
sdbe	ENABLED(A)	PRIMARY	zfs_73300_9	zfs_73300	c2	-
sdbp	ENABLED(A)	PRIMARY	zfs_73300_9	zfs_73300	c3	-

```
sdq          DISABLED  PRIMARY      zfs_73300_9  zfs_73300    c2          -
[root@aie-x4170m3b ~]#
```

View of a clustered Oracle ZFS Storage Appliance array ready for failback

When one head of clustered Oracle ZFS Storage array reboots or fails for any reason, the other head of the Oracle ZFS Storage array takes over all I/O and its resources. After the rebooted or failed Oracle ZFS Storage Appliance rejoins the cluster, the Oracle ZFS Storage Appliance administrator initiates a 'failback' to switch the ZFS resource pool to the proper head. The following illustration shows an Oracle ZFS Storage ZS3-2 waiting for a failback to be initiated. The following illustration shows how VxVM DMP commands view the Oracle ZFS Storage ZS3-2 FC LUNs in this scenario.

The screenshot shows the Oracle ZFS Storage ZS3-2 management console. At the top, there is a notification: "The cluster peer has rejoined the cluster." The main navigation bar includes "Configuration", "Maintenance", "Shares", "Status", and "Analytics". Below this, there are tabs for "SERVICES", "STORAGE", "NETWORK", "SAN", "CLUSTER", "USERS", "PREFERENCES", and "ALERTS". The "CLUSTER" tab is active, showing two nodes: "aie-zs3-2d-h1" (Active, takeover completed) and "aie-zs3-2d-h2" (Ready, waiting for failback). A diagram shows connections between the nodes. Below the diagram, there are two sections for "Active Resources". The left section shows a list of resources and their owners, while the right section shows "No resources are active on this cluster node."

RESOURCE	OWNER
↔ aie-zs3-2d-h1 (net/ixgbe2) 10.80.44.218	aie-zs3-2d-h1
↔ aie-zs3-2d-h2 (net/ixgbe3) 10.80.44.220	aie-zs3-2d-h2
↔ SRM-REPLICA (net/ixgbe4) 192.168.55.3	aie-zs3-2d-h1
↔ 10GbE_SRM (net/ixgbe6) 192.168.220.220	aie-zs3-2d-h1
zfs/new-pool2-h2 10.7T	aie-zs3-2d-h2
zfs/new-raidz2 18.7T	aie-zs3-2d-h2
zfs/newpool0-h1 8.03T	aie-zs3-2d-h1

Figure 14. Oracle ZFS Storage ZS3-2 waiting for failback

```
[root@aie-x4170m3b ~]# vxdisk scandisks
[root@aie-x4170m3b ~]# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
disk_0	auto:none	-	-	online invalid
disk_1	auto:none	-	-	online invalid
zfs_73300_0	auto:cdsdisk	zfs_73300_0	dg1	online shared

```
zfs_73300_1 auto:cdsdisk zfs_73300_1 dg2 online shared
zfs_73300_2 auto:cdsdisk zfs_73300_2 dg1 online shared
zfs_73300_3 auto:cdsdisk zfs_73300_3 dg3 online shared
zfs_73300_4 auto:cdsdisk zfs_73300_4 dg1 online shared
zfs_73300_5 auto:cdsdisk zfs_73300_5 dg2 online shared
zfs_73300_6 auto:cdsdisk zfs_73300_6 dg3 online shared
zfs_73300_7 auto:cdsdisk zfs_73300_7 dg1 online shared
zfs_73300_8 auto:cdsdisk zfs_73300_8 dg2 online shared
zfs_73300_9 auto:cdsdisk zfs_73300_9 dg3 online shared
zfs_73300_10 auto:cdsdisk - - online
zfs_73300_11 auto:cdsdisk - - online
zfs_73300_12 auto:cdsdisk zfs_73300_12 dg3 online shared
zfs_73300_13 auto:cdsdisk zfs_73300_13 dg1 online shared
zfs_73300_14 auto:cdsdisk zfs_73300_14 dg2 online shared
zfs_73300_15 auto:cdsdisk zfs_73300_15 dg3 online shared
zfs_73300_16 auto:cdsdisk - - online
zfs_73300_17 auto:cdsdisk zfs_73300_17 dg1 online shared
zfs_73300_18 auto:cdsdisk zfs_73300_18 dg2 online shared
zfs_73300_19 auto:cdsdisk zfs_73300_19 dg3 online shared
zfs_73300_20 auto:cdsdisk zfs_73300_20 dg2 online shared
[root@aie-x4170m3b ~]#
```

```
[root@aie-x4170m3b ~]# vxdisk list zfs_73300_0
Device: zfs_73300_0
devicetag: zfs_73300_0
type: auto
clusterid: clus_us
disk: name=zfs_73300_0 id=1396091048.157.aie-x4170m3b
group: name=dg1 id=1398139104.41.aie-x4170m3b
info: format=cdsdisk,privoffset=256,pubslice=3,privslice=3
flags: online ready private autoconfig shared autoimport imported
pubpaths: block=/dev/vx/dmp/zfs_73300_0s3 char=/dev/vx/rdmp/zfs_73300_0s3
guid: {d99aaef0-b731-11e3-8304-0f11e44e9518}
```

```

udid:
SUN%5FZFS%20Storage%207330%5FSUN%5F7x10%5FSeries%5F600144F08F8B540E00005334C7
F70016
site:      -
version:   3.1
iosize:    min=512 (bytes) max=1024 (blocks)
public:    slice=3 offset=65792 len=4120320 disk_offset=0
private:   slice=3 offset=256 len=65536 disk_offset=0
update:    time=1398140448 seqno=0.260
ssb:       actual_seqno=0.0
headers:   0 240
configs:   count=1 len=51360
logs:      count=1 len=4096
Defined regions:
  config  priv 000048-000239[000192]: copy=01 offset=000000 enabled
  config  priv 000256-051423[051168]: copy=01 offset=000192 enabled
  log     priv 051424-055519[004096]: copy=01 offset=000000 disabled
  lockrgn priv 055520-055663[000144]: part=00 offset=000000
Multipathing information:
numpaths:  4
sdas        state=enabled   type=secondary
sdf         state=enabled   type=primary
sdbf        state=enabled   type=secondary
sdn         state=enabled   type=primary
[root@aie-x4170m3b ~]#

```

```
[root@aie-x4170m3b ~]# vxmpadm getsubpaths
```

NAME	STATE[A]	PATH-TYPE[M]	DMPNODENAME	ENCLR-NAME	CTLR	ATTRS
sdb	ENABLED(A)	-	disk_0	disk	c0	-
sda	ENABLED(A)	-	disk_1	disk	c0	-
sdas	ENABLED(A)	SECONDARY	zfs_73300_0	zfs_73300	c2	-
sdbf	ENABLED(A)	SECONDARY	zfs_73300_0	zfs_73300	c3	-
sdf	ENABLED	PRIMARY	zfs_73300_0	zfs_73300	c2	-

```
sdn          ENABLED    PRIMARY    zfs_73300_0  zfs_73300    c3          -
sdaq         ENABLED(A) SECONDARY  zfs_73300_1  zfs_73300    c2          -
sdbd         ENABLED(A) SECONDARY  zfs_73300_1  zfs_73300    c3          -
sde          ENABLED    PRIMARY    zfs_73300_1  zfs_73300    c2          -
sdl          ENABLED    PRIMARY    zfs_73300_1  zfs_73300    c3          -
...
sdaa         ENABLED    PRIMARY    zfs_73300_9  zfs_73300    c3          -
sdbe         ENABLED(A) SECONDARY  zfs_73300_9  zfs_73300    c2          -
sdbp         ENABLED(A) SECONDARY  zfs_73300_9  zfs_73300    c3          -
sdq          ENABLED    PRIMARY    zfs_73300_9  zfs_73300    c2          -
[root@aie-x4170m3b ~]#
```

Acronyms / Definitions

APM – Array Policy Modules contain the kernel-level logic required to control path states and properly handle SCSI errors from a managed device. The base DMP packages come with APMs that implement procedures that are ‘typical’ for well-known storage system behavior such as Active/Active systems, Active/Passive systems, systems that comply with the SCSI-3 ALUA standard, and so on. And for those systems that do not fit into any of the broad categories and require specific handling, Symantec has the ability to build hardware-specific APMs to handle those specificities.

ASL – Array Support Libraries allow DMP’s Device Discovery Layer (DDL) to properly claim LUNs from the storage hardware, identify multiple paths to the same device, discover device-specific information and basically configure DMP’s kernel code to use the correct set of procedures to control the discovered devices. As such, storage hardware-specific Array Support Libraries (ASL) implement array-specific logic that allows DMP to configure itself to the connected storage and provide visibility into vendor-specific attributes for that storage.

DMP – Symantec Storage Foundation Dynamic Multi-Pathing. DMP provides automated discovery and configuration of the driver for any storage connected that is supported by the ASLs and APMs installed on the host.

Failback – On a clustered Oracle ZFS Storage Appliance, the Failback never occurs automatically. When a failed head is repaired and booted, it will rejoin the cluster (resynchronizing its view of all resources, their properties, and their ownership) and proceed to wait for an administrator to perform a failback operation. Until then, the original surviving head will continue to provide all services. This allows for a full investigation of the problem that originally triggered the takeover, validation of a new software revision, or other administrative tasks prior to the head returning to production service. Because failback is disruptive to clients, it should be scheduled according to business-specific needs and processes. There is one exception: Suppose that head A has failed and head B has taken over. When head A rejoins the cluster, it becomes eligible to take over if it detects that head B is absent or has failed. The principle is that it is always better to provide service than not, even if there has not yet been an opportunity to investigate the original problem. So while failback to a previously-failed head will never occur automatically, it may still perform takeover at any time.

FC – Fibre Channel

iSCSI – is an acronym for Internet SCSI (Small Computer System Interface), an Internet Protocol (IP)-based storage networking standard for linking data storage subsystems. This networking standard was developed by the Internet Engineering Task Force (IETF). For more information about the iSCSI technology, see RFC 3720: <http://www.ietf.org/rfc/rfc3720.txt>.

MPxIO – Sun StorEdge Traffic Manager /Oracle Solaris's multipathing software provides a multipathing solution for storage devices accessible through multiple physical paths. It abstracts physical paths to a device, providing unified access through a single virtual device. Physical path failures are transparently recovered, so failures are not exposed to the applications using the virtual device.

SF – Symantec Storage Foundation

Takeover – On a clustered Oracle ZFS Storage Appliance, the takeover command can occur at any time; please refer to the Oracle ZFS Storage Appliance online help or the administration guide for more details. Takeover is attempted whenever peer failure is detected. It can also be triggered manually using the cluster configuration CLI or BUI. This is useful for testing purposes as well as to perform rolling software upgrades (upgrades in which one head is upgraded while the other provides service running the older software, then the second head is upgraded once the new software is validated). Finally, takeover will occur when a head boots and detects that its peer is absent. This allows service to resume normally when one head has failed permanently or when both heads have temporarily lost power.

VxVM – Symantec Storage Foundation Volume Manager

References

NOTE: References to Sun ZFS Storage Appliance, Sun ZFS Storage 7000, and ZFS Storage Appliance all refer to the same family of Oracle ZFS Storage Appliance products. Some cited documentation or screen code may still carry these legacy naming conventions.

- Oracle Support Center
<http://www.oracle.com/support>
- Patches and updates downloads from My Oracle Support (MOS)
(search under the Patches and Updates tab, Oracle ZFS Storage Appliance Software Patches). Log in to My Oracle Support (MOS), which requires a registration, through the Oracle Support Center URL.

Also, you can reference the fishworks wiki for Oracle ZFS Storage Appliance Software updates

<http://wikis.oracle.com/display/FishWorks/Software+Updates>

"Oracle ZFS Storage Appliance: How to Set Up Client Multipathing"

MOS note 1628999.1.

https://support.oracle.com/epmos/faces/DocumentDisplay?_afrcLoop=2210813044875205&id=1628999.1

- Oracle ZFS Storage Appliance Documentation
<http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html>
- Additional Oracle ZFS Storage Appliance technical information and white papers
<http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html>



Using Oracle ZFS Storage Appliance Fibre
Channel and iSCSI with Symantec Storage
Foundation / DMP Quickstart Guide
September 2014, Version 2.0
Oracle Application Integration Engineering
Kitty Qi, Mark Schein

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



| Oracle is committed to developing practices and products that help protect the environment

Copyright © 2013, 2014. Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle 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 licensed through X/Open Company, Ltd. 0611

Hardware and Software, Engineered to Work Together