



ZFS STORAGE
APPLIANCE

An Oracle Technical White Paper
April 2014

Exadata Landing Pad: Migrating a Fibre Channel-based Oracle Database Using an Oracle ZFS Storage Appliance

Introduction	2
Data Migration Architecture Using Oracle ZFS Storage Appliance.....	3
Overview of Example System Components.....	4
Configuring the Oracle ZFS Storage Appliance for Database Migration	5
Configuring the Fibre Channel Connection to the Oracle ZFS Storage Appliance	5
Creating and Assigning Fibre Channel LUNs on the Oracle ZFS Storage Appliance	5
Configuring Oracle Linux Database Server Fibre Channel LUNs.....	9
Oracle RMAN Duplication Procedures.....	10
Configuring InfiniBand iSCSI LUNs for Database Migration	11
Installing the iSCSi package on the Exadata Database Node(s).....	11
Configuring Oracle ZFS Storage Appliance InfiniBand iSCSi LUNS ...	12
Configuring iSCSi LUNs on the Exadata Database Node(s)	15
Database Migration with Oracle Data Guard.....	17
Conclusion	18
Appendix: References.....	19

Introduction

Many entities which currently run Oracle Linux-based database servers to hold their Oracle Databases are interested in upgrading to an Oracle Exadata environment to take advantage of its internal InfiniBand architecture and the impressive database processing power it provides. An interesting challenge arises when planning the database migration from this Oracle Linux server to the Oracle Exadata Database Machine. Historically, many Oracle Databases in the Oracle Linux environment are configured with direct-attached Fibre Channel disk arrays. To complicate matters further, oftentimes these Oracle Linux-based Oracle Databases are large production systems that cannot easily be taken offline. When a database administrator needs to migrate a Fibre attached database to an InfiniBand based Oracle Exadata, the administrator needs to do it in an efficient and timely manner, without disruption or downtime.

By inserting the Oracle ZFS Storage Appliance into the architecture as an intermediary, there is no need to bring down and reconfigure the production database hosts to upgrade to InfiniBand capabilities. Moreover, once the data migration is complete, the Oracle ZFS Storage Appliance can remain in the architecture and act as the RMAN backup target for both the existing legacy Oracle Linux databases and the newly attached Exadata environment.

This white paper provides best practices and recommendations for database migration of an Oracle Database from an Oracle Linux environment to an Exadata by utilizing the unified environment capabilities of the Oracle ZFS Storage Appliance.

Highlighted in this paper are:

- Architecture and best practices for an Oracle Linux-based Oracle Database migration to an Oracle Exadata with an Oracle ZFS Storage Appliance
- Guides and procedures for attaching the Oracle ZFS Storage Appliance using Fibre Channel to the host Oracle Linux host database server
- Guides and procedures for assigning Oracle ZFS Storage Appliance iSCSI LUNs to an Exadata in order to facilitate the data migration
- Procedure recommendations and references for Oracle Database duplication using the Oracle ZFS Storage Appliance

Data Migration Architecture Using Oracle ZFS Storage Appliance

This section describes the data migration architecture used in the testing.

The test environment consisted of a clustered Oracle ZFS Storage Appliance, each head containing two dual-port 8 Gbps Fibre Channel host channel adapters (HCAs) and two dual-port Infiniband HCAs, an Oracle Linux database host and a target Exadata x2-2 environment.

The following diagram demonstrates the architectural overview of the testing effort.

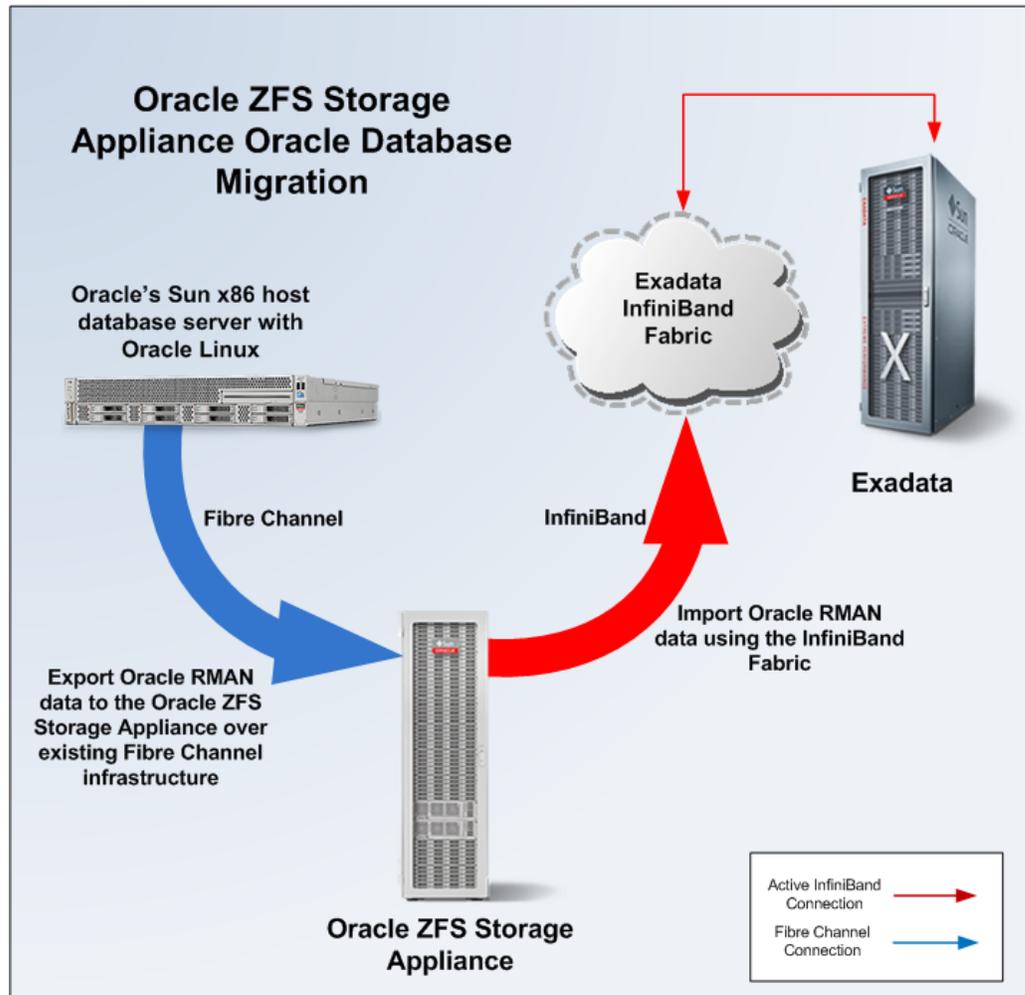


Figure 1. Data Migration with Oracle ZFS Storage Appliance intermediary – architecture overview

Overview of Example System Components

The following tables describe the hardware configuration, operating systems, and software releases utilized by this white paper and its reference architecture.

Table 1 shows the hardware of the Oracle ZFS Storage Appliance used.

TABLE 1. ORACLE ZFS STORAGE APPLIANCE HARDWARE USED IN REFERENCE ARCHITECTURE

EQUIPMENT	QUANTITY	CONFIGURATION
Storage	1 cluster (2 controllers)	Sun ZFS Storage 7420 cluster 256 GB DRAM per controller 2 x 512 GB read cache solid state drives (SSDs) per controller 2 x 20 2 TB SAS-2 (serial attached SCSI-2) disk trays 2 x dual port 10GbE Network Interface Card (NIC) 2 x dual port 8 Gbps FC HBA 2 x 17 GB log device

Table 2 shows hardware of the Oracle Linux host and the Exadata target.

TABLE 2. ORACLE DATABASE HOST SERVER AND EXADATA TARGET COMPONENTS USED IN REFERENCE ARCHITECTURE

EQUIPMENT	QUANTITY	CONFIGURATION
Oracle Linux Database Server	1	Sun Fire X4440 * Note that this server used in testing is now End of Life; a similar newer generation server is the Sun Server X2-4 256 GB DRAM 2 internal HDDs 2 x dual port 8Gbps FC HCA 2 x dual port InfiniBand HBA
Exadata	1	Oracle Exadata Database Machine X2-2 half rack Four x Sun Fire X4170 M2 Servers with 2 six-core Intel Xeon X5675 processors (3.06 GHz); 96 GB memory (expandable to 144 GB with optional memory expansion kit); 4 x 300 GB 10,000 RPM SAS disks; 2 x QDR (40Gbit/s) ports, and 2 x 10 Gb Ethernet ports based on the Intel 82599 10GbE controller Seven x Oracle Exadata Storage Servers

Configuring the Oracle ZFS Storage Appliance for Database Migration

This section provides best practices and recommendations for configuring the Oracle ZFS Storage Appliance as the intermediary for database migration.

Configuring the Fibre Channel Connection to the Oracle ZFS Storage Appliance

The following configurations for the Oracle ZFS Storage Appliance are recommended to help facilitate database migration.

Creating and Assigning Fibre Channel LUNs on the Oracle ZFS Storage Appliance

To configure Fibre Channel initiators of the Oracle Linux database server on the Oracle ZFS Storage Appliance, perform the following steps:

1. Find the WWN from the Oracle Linux database server.

Log on to the Oracle Linux database server with root credentials and list the World Wide Name (WWN) for the Fibre Channel initiators assigned by the operating system's Fibre Channel driver.

```
#cat /sys/class/fc_host/host<XX>/port_name
```

Note: `host<XX>` represents the FC host number assigned by the operating system's Fibre Channel driver.

The following is an example of the output:

```
#ls -l /sys/class/fc_hosts
```

```
lrwxrwxrwx 1 root root 0 Dec 17 10:27 host13 ->
../../../../devices/pci0000:80/0000:80:0b.0/0000:81:00.0/host13/fc_host/host13

lrwxrwxrwx 1 root root 0 Dec 17 10:27 host14 ->
../../../../devices/pci0000:80/0000:80:0b.0/0000:81:00.1/host14/fc_host/host14
```

For Port 0 (pci0000:80/0000:80:0b.0/0000:81:00.0):

```
#cat /sys/class/fc_host/host13/port_name
0x21000024ff39300e
```

For Port 1 (pci0000:80/0000:80:0b.0/0000:81:00.1):

```
#cat /sys/class/fc_host/host14/port_name  
0x21000024ff39300f
```

2. On the Oracle ZFS Storage Appliance, add the FC initiators from the Oracle Linux database server.

To add FC initiators through the Oracle ZFS Storage Appliance BUI:

Select the tabs: Configuration - SAN - Fibre Channel - Initiators and the “+” button to add an FC initiator. This will bring up the “Identify FC Initiator” window seen in the following figure. Add the WWN obtained from the Oracle Linux database server and create an Alias for the WWN port.

WWN	VENDOR
21:00:00:24:ff:39:30:0f	QLogic Corporation

Figure 2. Adding an FC initiator on the Oracle ZFS Storage Appliance

Add the newly identified initiator to Initiator Groups by dragging the initiator to an empty slot in the Initiator Group panel.

Edit initiator Group name. In this example the initiator group is edited to reflect the Oracle Linux database server FC port number: aie-4440b-port0(1).

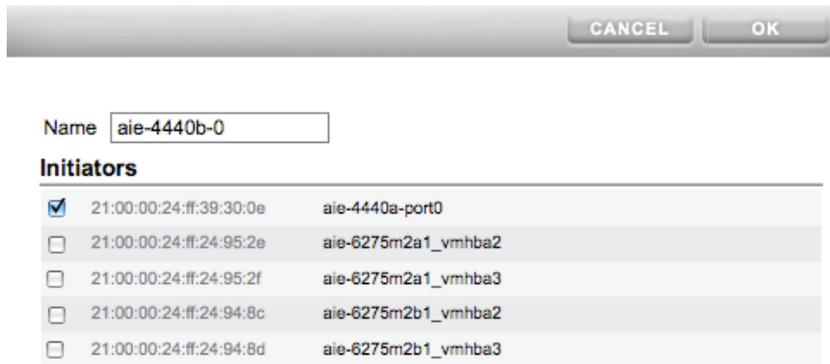


Figure 3. Adding an FC initiator group on the Oracle ZFS Storage Appliance

To add FC initiators and initiator groups using the CLI:

```
configuration san fc initiators groups> create
configuration san fc initiators group (uncommitted)> set name=aie444b-0
                        name = aie444b-0 (uncommitted)
configuration san fc initiators group (uncommitted)> set
initiators=wwn.21000024FF39300E
                        initiators = wwn.21000024FF39300E (uncommitted)
configuration san fc initiators group (uncommitted)> commit
configuration san fc initiators groups> list
GROUP      NAME
group-001  aie-4440b-0
          |
          +--> INITIATORS
                wwn.21000024FF39300E
configuration san fc initiators> create
configuration san fc initiators initiator (uncommitted)> set alias=aie-4440b-
port0
                        alias = aie-4440b-port0 (uncommitted)
configuration san fc initiators initiator (uncommitted)> commit
configuration san fc initiators> list
NAME      ALIAS
initiator-000  aie-4440a-port0
          |
          +--> INITIATOR
```

```
wwn.21000024FF39300E
```

3. On the Oracle ZFS Storage Appliance, create and assign FC LUN(s).

To create and assign a Fibre Channel LUN using the BUI:

Select Tabs: Shares – LUNs and the “+” button to add a LUN. This will bring up the “Create LUN” dialog window. Enter the FC LUN and define volume size, volume block size, Target group (fc_target) and (FC) Initiator group as well as fields indicated in the following figure.

The screenshot shows the 'Create LUN' dialog box with the following configuration:

- Project: Mule
- Name: Lun0
- Volume size: 120 G
- Thin provisioned:
- Volume block size: 128k
- Online:
- Target group: fc_target
- Initiator group(s): OracleVM, VMware_6275m2ab0, aie-4440b-0, vmware_ipoib_init_group
- LU number: 0 Auto-assign

Figure 4. Creating a Fibre Channel LUN

To create and assign a Fibre Channel LUN using the CLI:

Select the Project name to add the LUN to:

```
shares> select Mule
```

Create a new FC LUN and define the volume size, FC initiator group, volume block size and FC target group:

```
shares Mule> lun lun0
shares Mule/lun0 (uncommitted)> set volsize=120g
                               volsize = 120G (uncommitted)
shares Mule/lun0 (uncommitted)> set initiatorgroup=aie-4440b-0
                               initiatorgroup = aie-4440b-0 (uncommitted)
shares Mule/lun0 (uncommitted)> set volblocksize=128k
                               volblocksize = 128K (uncommitted)
```

```
shares Mule/lun0 (uncommitted)> set targetgroup=fc_target
                                targetgroup = fc_target (uncommitted)
shares Mule/lun0 (uncommitted)> commit
```

Configuring Oracle Linux Database Server Fibre Channel LUNs

After creating, defining and assigning the Oracle ZFS Storage Appliance LUNs as outlined in the previous section, configure the new LUNS on the Oracle Linux database server(s). The following sections describe this process.

1. Scan for the new FC LUNs on the Oracle Linux database server.

The Fibre Channel driver will already be installed and configured on the Oracle Linux database server. Once the Oracle ZFS Storage Appliance LUNS are configured and assigned, rescan for Fibre Channel devices on the Oracle Linux database server. Root user access is required for these procedures:

```
# for host in `ls /sys/class/fc_host`; do
> echo "- - -" > /sys/class/scsi_host/${HOST}/scan;
> done;
```

2. Validate that the database server recognizes the new Fibre Channel LUNs.

Use the `fdisk` command to validate that the new `/dev/sd*` device files were created. In the following example, `/dev/sdj` is the newly discovered LUN.

```
#fdisk -l |egrep '^Disk' |egrep -v 'dm-'
.
.
.
Disk identifier: 0x00000000
Disk /dev/sdj (Sun disk label): 255 heads, 63 sectors, 15665 cylinders
```

Note: The `lsscsi` command can be used interchangeably to accomplish the same task as well.

3. Format the FC LUN with `fdisk`.

Using `fdisk`, add a new partition. In the following example, using `/dev/sdj` from the discovery in the previous example, `fdisk` is used to create a new partition with a size of 120 GB on partition 1.

```
#fdisk /dev/sdj
n (new partition)
p 1
size +120G
w
```

4. Add the ext3 filesystem to the newly formatted LUN.

Run the `mksfs.ext3` command to add the Linux ext3 filesystem.

```
#mkfs.ext3 /dev/sdj1
```

Run a consistency check of the new ext3 filesystem.

```
#fsck -t ext3 /dev/sdj1
```

5. Add an fstab entry.

Edit the fstab and add a mount point entry for an ext3 filesystem.

```
#vi /etc/fstab  
  
/dev/sdj1          /zfssa/mule/backup1      ext3    defaults  
1 2
```

6. Mount the FC LUN.

After editing the `/etc/fstab`, mount the FC LUN and change permissions for RMAN/Oracle use.

Create a directory for the mount point if it does not exist. In the following example the directory `/zfssa/mule/backup1` is created.

```
#mount /zfssa/mule/backup1  
#chown -R oracle:oinstall /zfssa/mule
```

7. Initiate an Oracle RMAN backup or duplication procedure.

Issue an `rman backupset`, `copy` or `duplication` here. More reference information regarding these procedures is located in the following section.

8. Upon completion of the RMAN backup/duplication, unmount the LUN.

```
# umount /zfssa/mule/backup1
```

Oracle RMAN Duplication Procedures

Database administrators typically have their own procedures regarding backup, recovery and duplication. Given these variations, detailing them is outside the scope of this document. There are several procedures documented and supported by Oracle. More information on the RMAN duplication process can be found on the Oracle Support website at <https://support.us.oracle.com>, including:

- ID 1375864.1 – Perform Backup Based RMAN DUPLICATE Without Connecting To Target Database For Both Disk & Tape Backups
- ID 300062.1 - How To Clone An Existing RDBMS Installation Using OUI
- ID 562556.1 - How to Manually Backup - Restore or Clone a Database to Another Node
- ID 732624.1 - Clone a Database With RMAN Without Connecting To Target Database
- ID 1338233.1 - Clone/Refresh Non-prod database from production

Configuring InfiniBand iSCSI LUNs for Database Migration

The following configurations for the Oracle ZFS Storage Appliance are recommended to help facilitate database migration.

Installing the iSCSI package on the Exadata Database Node(s)

The Exadata image distribution does not provide or update the Exadata database nodes with the iSCSI packages. The iSCSI package must be installed separately and added to the Exadata database operating system as a service. The following sections describe this procedure.

1. On the Exadata database node(s), install the iSCSI package and start the iSCSI service.

Install the `iscsi` initiator utility on the Exadata database nodes used for the iSCSI LUN attachments. Root user access is required for these procedures.

```
#yum install iscsi-initiator-utils
#yum list iscsi-initiator-utils
```

```
Installed Packages
iscsi-initiator-utils.x86_64 6.2.0.872-13.0.1.el5 installed
```

To manage the software, use the `chkconfig` and `service` commands. The service for Red Hat Enterprise Linux systems is `iscsid`.

Add and validate the service with the `chkconfig` command.

```
#chkconfig --add iscsid
#chkconfig --list iscsid
```

```
iscsi 0:off 1:off 2:off 3:on 4:on 5:on 6:off
[not set for reboot]
```

```
Start the iscsid with
#service iscsi start
#service iscsid start
```

2. Display the Exadata database server iSCSI initiator identifier.

Look at the `initiatorname.iscsi` file to determine the initiator identifier. The identifier (or the `iqn`) is located near the bottom of the file.

```
#cat /etc/iscsi/initiatorname.iscsi  
  
.br/>.br/>InitiatorName=iqn.1988-12.com.oracle:f964b675e9a 1020 2013-12-19.09:50:44
```

3. Configure CHAP Authentication.

If you are using CHAP, you will need to edit your `iscsid.conf` file and update it with correct authentication info:

```
#vi /etc/iscsi/iscsid.conf  
  
node.startup = automatic  
node.session.auth.username = username  
node.session.auth.password = password  
discovery.sendtargets.auth.username = username  
discovery.sendtargets.auth.password = password
```

4. Save and close the file.

Configuring Oracle ZFS Storage Appliance InfiniBand iSCSi LUNS

Once the iSCSI package has been installed and activated as a service on the Exadata database server(s), the RMAN backup LUN(s) that are currently configured on the Oracle ZFS Storage Appliance will need to be reconfigured as iSCSI LUN(s). The following sections describe the process of defining iSCSI groups and iSCSI targets to be discovered on the Exadata database servers.

1. On the Oracle ZFS Storage Appliance, create an iSCSI target from the Exadata database server.

Select Tabs: Configuration – SAN – iSCSI – Targets and the “+” button to add a LUN. This will bring up the “Create iSCSI Target” dialog window seen in the following figure.

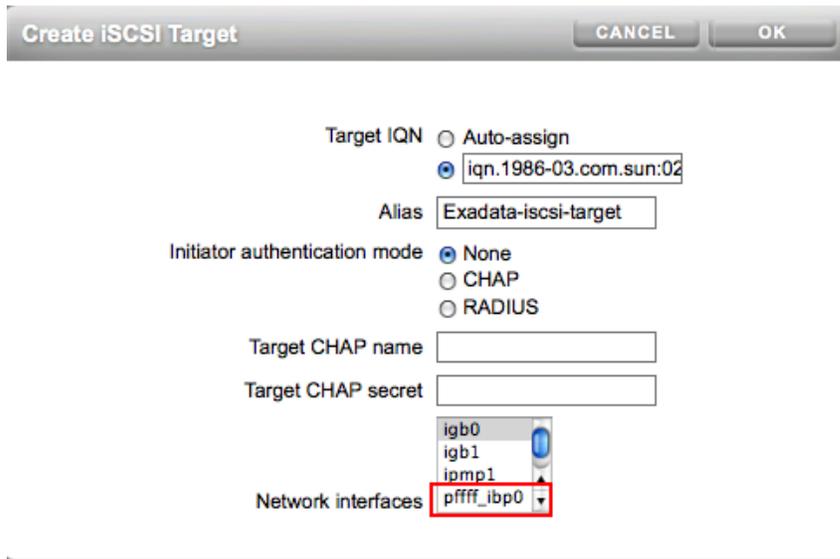


Figure 5. Creating an iSCSI target to the Exadata database server

Enter the iSCSI Target IQN, Alias, and Initiator authentication mode (if applicable), noting the following:

- Target IQN – Enter the initiator identifier (iqn) for the Exadata database server you wish to make a target. To find the iqn, reference step 2, "Display Exadata database server iSCSI initiator identifier," in the previous section "Installing the iSCSi package on the Exadata Database Node(s)."
- Target CHAP name (and secret) – Enter if CHAP authentication is used. To determine if CHAP authentication is currently being used by the Exadata database server, reference the configuration step 3, "Configure CHAP Authentication," in the previous section "Installing the iSCSi package on the Exadata Database Node(s)."
- Network interfaces – Select the InfiniBand interface. In this example, select `pffff_ibp0`.

2. Create the iSCSI Group on the Oracle ZFS Storage Appliance.

Add the iSCSI initiator to iSCSI Initiator Groups by dragging the initiator to an empty slot in the Initiator Groups panel.

Edit the iSCSI Initiator Group name; enter the name and select the target iqn of the Exadata database server you wish to assign as the target:

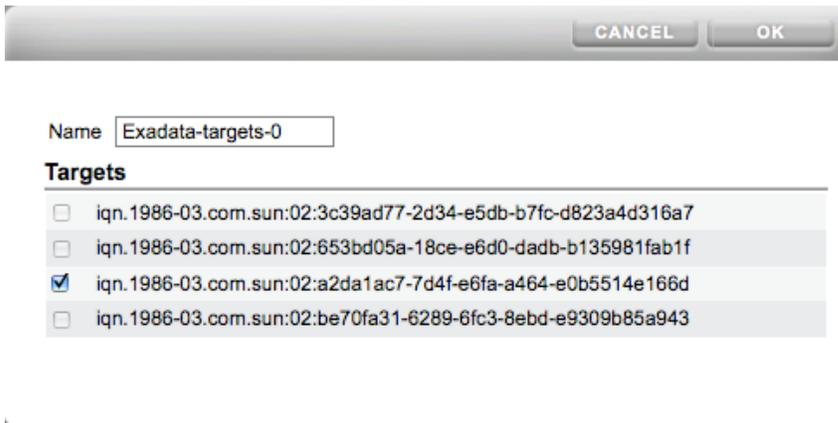


Figure 6. Creating an iSCSI Target Group

3. Assign the original FC LUN to an iSCSI Target Group and Initiator.

Select Tabs: Shares– LUN.

Edit the LUN(s) used to perform the Oracle RMAN backup/duplication process.

Select Tab: Protocols.

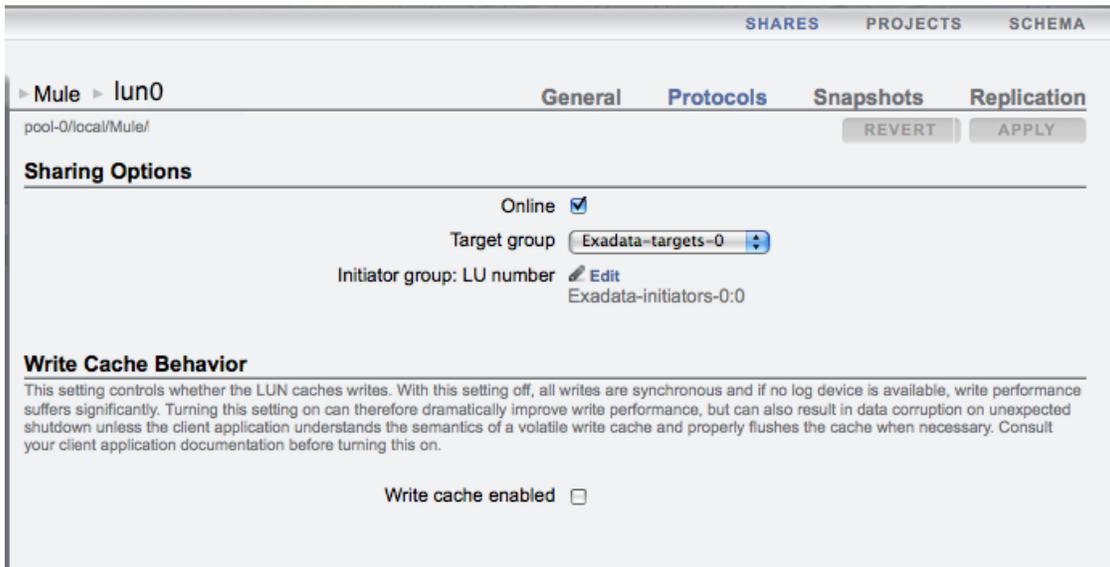


Figure 7. Assigning the original FC LUN to an iSCSI Target Group and Initiator

Target group – Change the new Exadata iSCSI target group.

Initiator group: LU number – Select Edit and add the new iSCSI initiator alias.

Configuring iSCSI LUNs on the Exadata Database Node(s)

The following sections describe the procedures to discover, build and mount the iSCSI LUN in order to complete the Oracle Database migration.

1. Discover iSCSI targets on the Exadata database server node.

Execute an iSCSI discovery against the Oracle ZFS Storage Appliance:

```
#iscsiadm -m discovery -t sendtargets -p aie-7420e-h1
10.80.74.204:3260,2 iqn.1986-03.com.sun:02:653bd05a-18ce-e6d0-dadb-b135981fab1f
192.168.50.1:3260,3 iqn.1986-03.com.sun:02:653bd05a-18ce-e6d0-dadb-b135981fab1f
192.168.36.208:3260,2 iqn.1986-03.com.sun:02:a2dalac7-7d4f-e6fa-a464-e0b5514e166d
192.168.36.210:3260,2 iqn.1986-03.com.sun:02:a2dalac7-7d4f-e6fa-a464-e0b5514e166d
192.168.36.208:3260,2 iqn.1986-03.com.sun:02:be70fa31-6289-6fc3-8ebd-e9309b85a943
192.168.36.210:3260,2 iqn.1986-03.com.sun:02:be70fa31-6289-6fc3-8ebd-e9309b85a943
```

2. Build the iSCSI device file.

Following the iSCSI discovery, build the new iSCSI LUN devices.

```
#iscsiadm -m node -T"iqn.1986-03.com.sun:02:a2dalac7-7d4f-e6fa-a464-e0b5514e166d" --
login
#service iscsi stop
#service iscsi start
```

3. Validate that the Exadata database server recognizes the new iSCSI LUN(s).

Execute the `fdisk` command to validate that the `/dev/sd*` device files were created.

In the following example, `/dev/sdc` is the LUN with the RMAN backup files.

```
#fdisk -l |egrep '^Disk' |egrep -v 'dm-'
Disk /dev/dm-0 doesn't contain a valid partition table
Disk /dev/dm-1 doesn't contain a valid partition table
Disk /dev/dm-2 doesn't contain a valid partition table
Disk /dev/dm-3 doesn't contain a valid partition table
Disk /dev/sda: 896.9 GB, 896998047744 bytes
Disk /dev/sdb (Sun disk label): 255 heads, 63 sectors, 15665 cylinders
```

4. Mount the iSCSI LUN.

List the device to be mounted.

For this example, `/dev/sdb1` is the partition to be used.

```
#ls -l /dev/sd*
.
.
.
brw-r----- 1 root disk 8, 32 Dec 19 10:50 /dev/sdb
brw-r----- 1 root disk 8, 33 Dec 19 10:50 /dev/sdb1
```

Create a directory to mount the LUN and change the directory ownership to reflect the Oracle Database user and group.

```
#mkdir /zfssa/mule/backup1
#chown -R oracle:dba /zfssa/mule
```

Mount the iSCSI LUN and make sure the LUN has oracle user and group permissions.

```
#mount /dev/sdc1 /zfssa/mule/backup1
#chown -R oracle:dba /zfssa/mule
```

5. Validate the iSCSI LUN permissions.

Check to see if the Exadata database server node can access the rman files from the newly mounted iSCSI LUN.

```
#su - oracle

$rman target /

RMAN> CONFIGURE SNAPSHOT CONTROLFILE NAME TO '/zfssa/mule/backup1/snapcf_mule.f';
[no snapshot control in the directory]

RMAN> crosscheck backup;

RMAN> catalog start with '/zfssa/mule/backup1/';

searching for all files that match the pattern /zfssa/mule/backup1/

List of Files Unknown to the Database
=====
File Name: /zfssa/mule/backup1/0korr7c5_8_1
File Name: /zfssa/mule/backup1/0torr7c5_1_1
File Name: /zfssa/mule/backup1/0norr7c5_2_1
File Name: /zfssa/mule/backup1/0lorr7c5_2_1
File Name: /zfssa/mule/backup1/0lorr7c5_1_1
File Name: /zfssa/mule/backup1/0uorr7c5_2_1
File Name: /zfssa/mule/backup1/0qorr7c5_1_1
File Name: /zfssa/mule/backup1/12orr7c5_1_1
File Name: /zfssa/mule/backup1/0korr7c5_4_1
File Name: /zfssa/mule/backup1/0qorr7c5_2_1
.
.
```

Continue with the Oracle RMAN duplication process to complete the database migration. For information regarding the RMAN database duplication, reference the section "Oracle RMAN Duplication Procedures" earlier in the document.

Database Migration with Oracle Data Guard

A customer may already have deployed an Oracle Data Guard disaster protection strategy. The Oracle ZFS Storage Appliance can be easily added to the existing Data Guard environment and can provide database snapshots and clones.

The document detailing and providing best practices for such a database migration using the Oracle ZFS Storage Appliance with Oracle Data Guard can be found here:

<http://www.oracleimg.com/technetwork/database/features/availability/maa-db-clone-szfssa-172997.pdf>

Conclusion

Oracle ZFS Storage Appliance offers a fast, efficient and non-disruptive solution for Oracle Database administrators needing to migrate their Oracle Linux-based Oracle Databases with Fibre Channel-attached disk to an InfiniBand-based Oracle Exadata environment.

By adding an Oracle ZFS Storage Appliance into the middle of the architecture, the database administrator can seamlessly integrate Oracle RMAN backup/duplication procedures to the Oracle ZFS Storage Appliance by leveraging already existing Fibre Channel networks, and then complete the RMAN duplication process to the Exadata utilizing InfiniBand technology.

With its versatile unified storage networking protocol capabilities, the Oracle ZFS Storage Appliance platforms and technology are an excellent choice for your Oracle Database migration environment.

Appendix: References

See the following resources for additional information relating to the products covered in this document.

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 may still carry these legacy naming conventions.

- Oracle ZFS Storage Appliance Documentation Library, including Installation, Analytics, Customer Service, and Administration guides:
<http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html>
- The *Oracle ZFS Storage Appliance Administration Guide* is also available through the Oracle ZFS Storage Appliance help context.
The Help function in Oracle ZFS Storage Appliance can be accessed through the browser user interface.
- Oracle Support Center
<http://www.oracle.com/support>
- Patches and updates downloads from My Oracle Support (MOS)
(search under Oracle ZFS Storage Software Patches)
- Oracle ZFS Storage Appliance Plug-ins
<http://www.oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html>
- Oracle Storage Product Information
<http://www.oracle.com/us/products/storage/overview/index.html>
- Oracle ZFS Storage Appliance Technical White Papers and Solution Briefs
<http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html>



Exadata Landing Pad: Migrating a Fibre Channel-based Oracle Database Using an Oracle ZFS Storage Appliance

April 2014, v1.0

Application Integration Engineering

Author: Joseph Pichette

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 © 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