An Oracle Technical White Paper
June 2015

Configuring an Oracle ZFS Storage ZS3-BA with an Oracle SuperCluster for Oracle Database Backup and Recovery

ORACLE

Executive Overview

The Oracle SuperCluster is an innovative solution which incorporates technologies from other Oracle engineered systems such as Oracle Exadata, Oracle Exalogic, and the Oracle ZFS Storage Appliance. As with Oracle Exadata, the Oracle SuperCluster requires a fast backup and recovery solution to protect its Oracle Real Application Clusters (RAC) 11*g*R2 and 12*c* Oracle Database services.

The Oracle ZFS Storage ZS3-BA is a high-performance backup and recovery system that is optimized for Oracle's engineered systems. It is able to connect directly into the Oracle SuperCluster's high-speed InfiniBand infrastructure and operates with the Oracle Database Direct NFS Client (dNFS) feature. The Oracle ZFS Storage ZS3-BA connects differently into each of the Oracle SuperCluster model's InfiniBand fabrics.

This paper provides instructions, recommendations, and examples for how to configure and tune the Oracle ZFS Storage ZS3-BA and Oracle SuperCluster, including:

- Instructions for physically connecting the Oracle ZFS Storage ZS3-BA to each Oracle SuperCluster model's InfiniBand infrastructure
- Instructions for configuring Oracle ZFS Storage ZS3-BA storage and networking
- Recommendations for tuning the Oracle Solaris 11 operating system on the Oracle SuperCluster
- Lastly, recommendations and examples on how to configure the Oracle Database feature Oracle Recovery Manager (RMAN) for best performance

NOTE: References to Oracle SPARC SuperCluster models may reflect legacy naming conventions for models still actively in use. These products are all part of the Oracle SuperCluster family.

Introduction

The configuration of the Oracle ZFS Storage ZS3-BA with the Oracle SuperCluster is tightly coupled. Only certain ports of the Oracle SuperCluster InfiniBand infrastructure are available for connectivity. Once the connections are made, the Oracle ZFS Storage ZS3-BA InfiniBand ports must be activated and then configured on the InfiniBand switches.

Properly configuring the Oracle ZFS Storage ZS3-BA requires multiple steps, including setting up InfiniBand networking, creating storage pools, creating Oracle RMAN projects and shares, and setting up DTrace Analytics on the Oracle ZFS Storage ZS3-BA for performance monitoring. All of these steps are similar to procedures for configuring an Oracle ZFS Storage ZS3-BA for an Oracle Exadata system, and they adhere to current best practices.

The remaining steps are all performed on the Oracle Solaris 11-based Oracle Database 11gR2 (Database) Server nodes. These include creating the Oracle ZFS Storage ZS3-BA mount points, tuning the Oracle Solaris 11 operating system and networking parameters, configuring dNFS for the database, tuning Oracle RMAN parameters, and creating the final Oracle RMAN scripts. These steps can either be performed manually, or by using the Oracle Engineered Solutions Backup Utility for Oracle ZFS Storage Appliance. The manual steps are provided in this paper.

Even though this document is tailored for the Oracle ZFS Storage ZS3-BA, it is also relevant for the original Sun ZFS Backup Appliance. The computing heads and attached disk differ between the two models, but the InfiniBand host bus adapter (HBA) placement is identical in both.

Connecting to an Oracle SuperCluster InfiniBand Fabric

Connection of the Oracle ZFS Storage ZS3-BA to the Oracle SuperCluster InfiniBand network depends on the Oracle SuperCluster model. In each case, both heads of the Oracle ZFS Storage ZS3-BA must have their primary and failover paths configured. This allows for data availability even when one of the system heads fails its resources over to the other head.

Oracle SPARC SuperCluster T4-4 InfiniBand Fabric

There are two methods for connecting the Oracle ZFS Storage ZS3-BA to the Oracle SPARC SuperCluster T4-4, depending on the configured environment:

- Connected directly to the SuperCluster InfiniBand switches – This option is used if the Oracle ZFS Storage ZS3-BA is the only other appliance or device (other than SuperCluster expansion) that will be connected to the infrastructure.
- Connected to external InfiniBand leaf switches – This option is used if more appliances or devices will be connected. In this case, two additional leaf switches are introduced into the InfiniBand network, but are not located within the SuperCluster or SuperCluster expansion racks. A typical scenario for this configuration is when both an Oracle ZFS Storage ZS3-BA and backup application media servers (connected to tape drives) are needed.

The Oracle ZFS Storage ZS3-BA performance is identical between these two configurations.

Connecting to Internal InfiniBand Leaf Switches

There are four ports available on each of the InfiniBand leaf switches that can be used to connect to the Oracle ZFS Storage ZS3-BA. Connect the cables to these ports as follows:

- Oracle ZFS Storage ZS3-BA Head 1:
  - PCIe 3 Port 1 to Upper IB Leaf Switch (U24) Port 2A
  - PCIe 3 Port 2 to Lower IB Leaf Switch (U18) Port 2B
  - PCIe 6 Port 1 to Upper IB Leaf Switch (U24) Port 7B
  - PCIe 6 Port 2 to Lower IB Leaf Switch (U18) Port 12A
- Oracle ZFS Storage ZS3-BA Head 2:
  - PCIe 3 Port 1 to Lower IB Leaf Switch (U18) Port 2A
  - PCIe 3 Port 2 to Upper IB Leaf Switch (U24) Port 2B
  - PCIe 6 Port 1 to Lower IB Leaf Switch (U18) Port 7B
  - PCIe 6 Port 2 to Upper IB Leaf Switch (U24) Port 12A

Figure 1 shows this direct connection configuration, including the primary resource paths (thick blue connections) and the failover resource paths (thinner red connections).

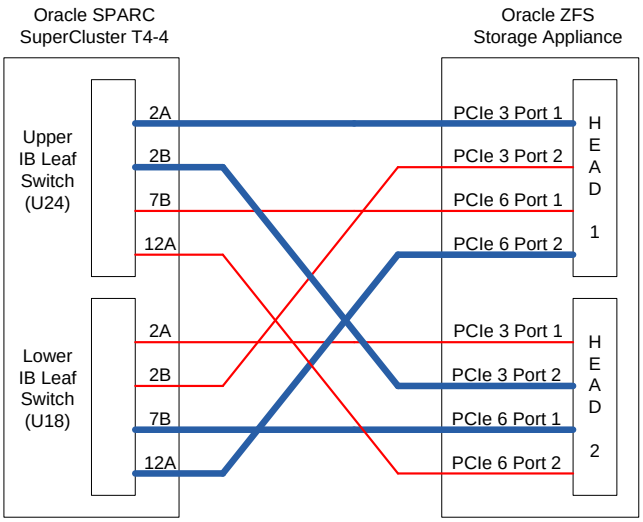

Figure 1. Oracle ZFS Storage ZS3-BA InfiniBand ports connected directly to the Oracle SPARC SuperCluster T4-4 InfiniBand leaf switches

Connecting to External InfiniBand Leaf Switches

When external InfiniBand switches are used, the eight open ports on the Oracle SPARC SuperCluster T4-4 infrastructure (Ports 2A, 2B, 7B, and 12A of each InfiniBand switch) are used for connecting to the external leaf switches. An example of external leaf switch connectivity is shown in Figure 2.

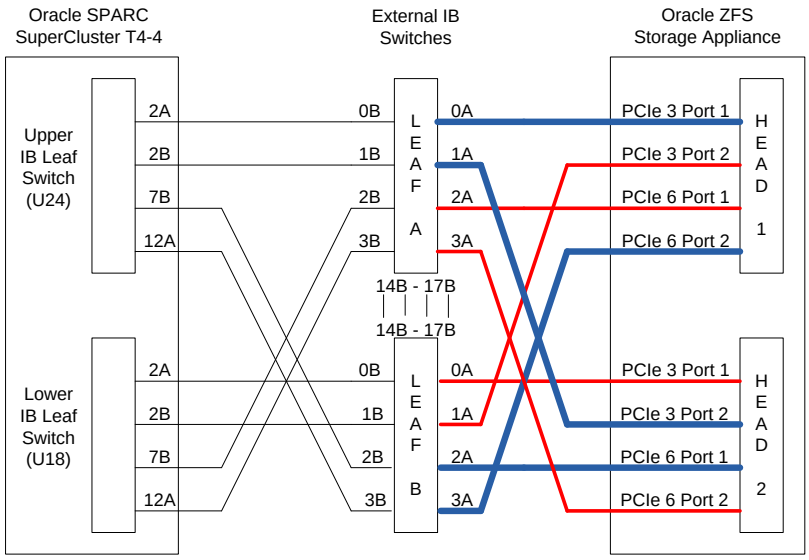

Figure 2. Oracle ZFS Storage ZS3-BA and Oracle SPARC SuperCluster T4-4 InfiniBand ports connected to external InfiniBand leaf switches

The thick blue connections indicate the Oracle ZFS Storage ZS3-BA primary resource paths and the thinner red connections indicate the failover resource paths. There are also four connections between

the two external InfiniBand switches (ports 14B, 15B, 16B, and 17B), providing for maximum performance and availability.

Oracle SuperCluster T5-8 InfiniBand Fabric

There are two methods for connecting the Oracle ZFS Storage ZS3-BA to the Oracle SuperCluster T5-8, depending on the configured environment:

- Connected directly to the Oracle SuperCluster T5-8 InfiniBand switches – This option is used if the Oracle ZFS Storage ZS3-BA is the only other appliance or device (other than Oracle SuperCluster T5-8 expansion) that will be connected to the infrastructure. This configuration does not utilize all of the InfiniBand connections available on the Oracle ZFS Storage ZS3-BA. Therefore, it cannot achieve the highest possible system performance.
- Connected to external InfiniBand leaf switches – This option is used to increase connection fault tolerance and increase system throughput. It also provides additional InfiniBand ports to attach other systems, such as backup application media servers. ***This is the preferred connection method.***

Connecting to Internal InfiniBand Leaf Switches

There are two ports available on each of the Oracle SuperCluster T5-8 InfiniBand leaf switches that can be used to connect to the Oracle ZFS Storage ZS3-BA. Connect the cables to these ports as follows:

- Oracle ZFS Storage ZS3-BA Head 1:
  - PCIe 3 Port 1 to Upper IB Leaf Switch (U32) Port 2A
  - PCIe 6 Port 1 to Upper IB Leaf Switch (U32) Port 2B
- Oracle ZFS Storage ZS3-BA Head 2:
  - PCIe 3 Port 1 to Lower IB Leaf Switch (U26) Port 2A
  - PCIe 6 Port 1 to Lower IB Leaf Switch (U26) Port 2B

Figure 3 shows the direct connection configuration, including the primary resource paths (thick blue connections) and the failover resource paths (thinner red connections).

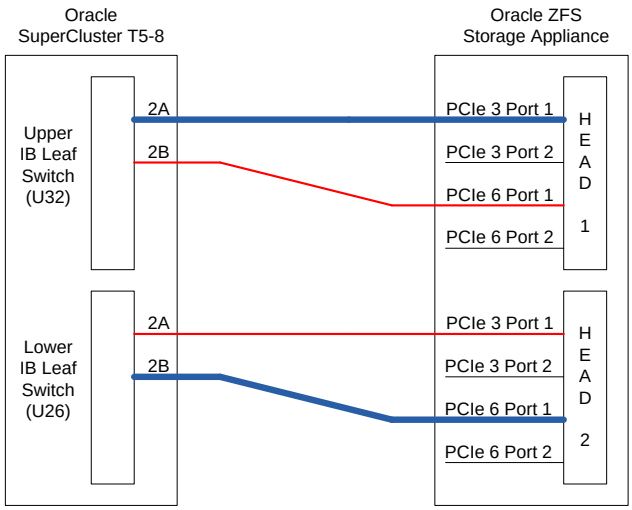

Figure 3. Oracle ZFS Storage ZS3-BA InfiniBand ports connected directly to the Oracle SuperCluster T5-8 InfiniBand leaf switches

Connecting to External InfiniBand Leaf Switches

When external InfiniBand switches are used, the four open ports on the Oracle SuperCluster T5-8 infrastructure (ports 2A and 2B of each InfiniBand switch) are used for connecting to the external leaf switches. An example of external leaf switch connectivity is shown in Figure 4.

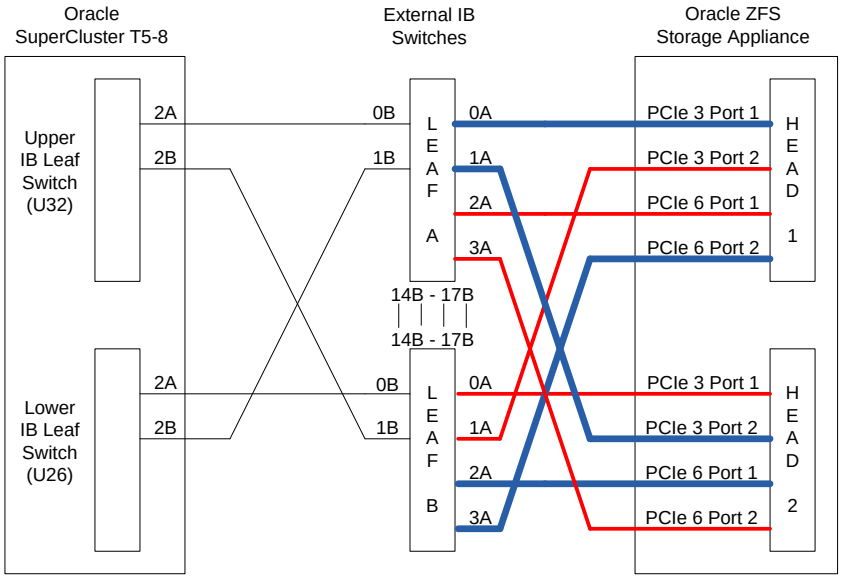

Figure 4. Oracle ZFS Storage ZS3-BA and Oracle SuperCluster T5-8 InfiniBAnd ports connected to external InfiniBand leaf switches

The thick blue connections indicate the Oracle ZFS Storage ZS3-BA primary resource paths and the thinner red connections indicate the failover resource paths. There are also four connections between

the two external InfiniBand switches (ports 14B, 15B, 16B, and 17B), providing for maximum performance and availability.

Oracle SuperCluster M6-32 InfiniBand Fabric

There is only one port on each of the Oracle SuperCluster M6-32 InfiniBand leaf switches available for additional storage connectivity. Due to this constraint, the performance of the Oracle ZFS Storage ZS3-BA will be diminished as compared to other SuperCluster models. The two available ports are:

- IB Leaf Switch (U22) Port 12A
- IB Leaf Switch (U20) Port 12A

Therefore, additional external leaf switches are required for connecting an Oracle ZFS Storage ZS3-BA to the Oracle SuperCluster M6-32 InfiniBand infrastructure. An example of external leaf switch connectivity is shown in Figure 5.

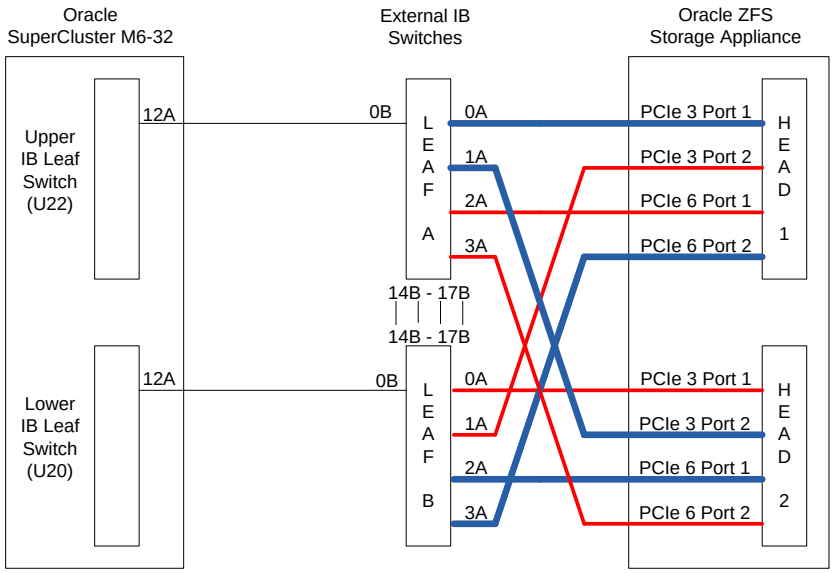

Figure 5. Oracle ZFS Storage ZS3-BA and Oracle SuperCluster M6-32 InfiniBand ports connected to external InfiniBand leaf switches

The thick blue connections indicate the Oracle ZFS Storage ZS3-BA primary resource paths and the thinner red connections indicate the failover resource paths. There are also four connections between the two external InfiniBand switches (ports 14B, 15B, 16B, and 17B), providing for maximum performance and availability.

Configuring the Oracle ZFS Storage ZS3-BA InfiniBand Datalinks

Each of the Oracle ZFS Storage ZS3-BA InfiniBand connections must be configured, using the following directions, before use.

1. Log on to the Oracle ZFS Storage ZS3-BA BUI of Head 1 and navigate to **Configuration** -> **Network**.
2. Click on the plus (+) icon next to Datalinks to open the Network Datalink dialog window.
3. Complete the Network Datalink dialog entries as follows (refer to Figure 6):
    a. Check the `IB Partition` box.
    b. Provide a meaningful name for the datalink name.
    c. Set the `Partition Key` to **`8503`**.
    d. Select **`Connected Mode`** for the `Link Mode`.
    e. Do not check the `LACP Aggregation` box.
    f. Select `Partition Device` **`ibp0`**.
    g. Record the GUID number (for example, `21280001ef43bb`) and click **APPLY**.

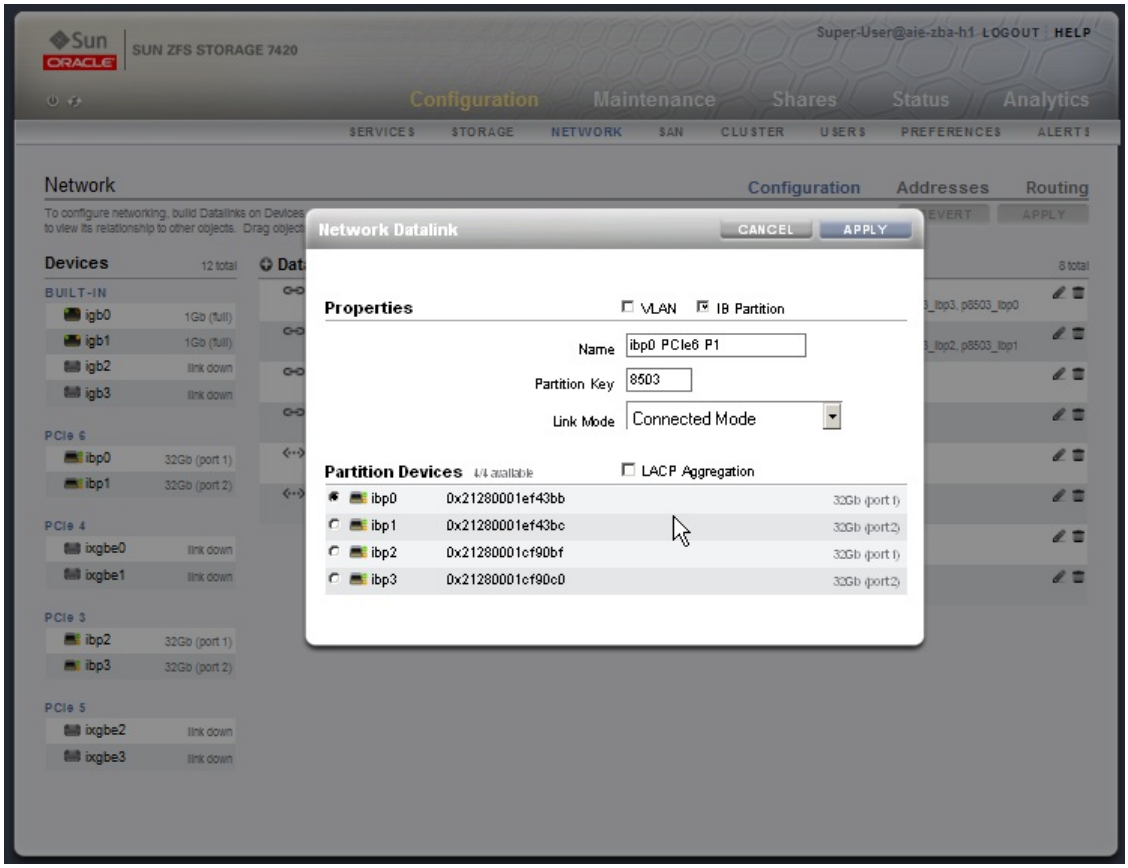

Figure 6. Network Datalink dialog box

4. Repeat Steps 2 and 3 for each remaining InfiniBand interface (`ibp1`, `ibp2`, and `ibp3`). (NOTE: Only repeat Steps 2 and 3 for InfiniBand interface `ibp2` if not connecting to an external InfiniBand leaf switch on an Oracle SuperCluster T5-8 system.)

5. Repeat Steps 1 through 4 for Oracle ZFS Storage ZS3-BA Head 2.

Reconfiguring the Oracle SuperCluster InfiniBand Switches

The GUIDs of the Oracle ZFS Storage ZS3-BA InfiniBand HBA ports must be added to the existing Oracle SuperCluster InfiniBand configuration with the partition key of `8503` in order for the components to communicate with each other. Use the following steps to do so.

1. In the CLI, log on to the Oracle SuperCluster InfiniBand spine switch as `root` and verify that it is running subnet manager by executing the `enablesm` command.

```
login as: root
root@aiescsw-ibs0's password:
Last login: Tue Sep 25 08:19:01 2012 from dhcp-brm-bl5-204-3e-east-10-135-75-
254.usdhcp.oraclecorp.com
[root@aiescsw-ibs0 ~]# enablesm
opensm (pid  15906) is already running...
Starting partitiond daemon
```

```
/usr/local/util/partitiond is already running
(You may also perform a 'restart' if wanted)
[root@aiescsw-ibs0 ~]#
```

By default, the spine switch is given a hostname of <*scid*>sw-ibs0, where <*scid*> is the prefix name given to the entire Oracle SuperCluster system. In these examples, the <*scid* > is aiesc.

The enablesm command will report that the subnet manager partitiond is already running, and will attempt to start it if it is not.

2. Enter the command getmaster to verify that this is the master switch of the configuration. If the master switch is not running on the spine switch, log out and log in to the designated master switch for the remainder of this procedure.

```
[root@aiescsw-ibs0 ~]# getmaster
Local SM enabled and running
20120913 10:16:51 Master SubnetManager on sm lid 13 sm guid 0x2128e8ac27a0a0 :
SUN DCS 36P QDR aiescsw-ibs0.us.oracle.com
[root@aiescsw-ibs0 ~]#
```

3. Back up the switch configuration according to the documented backup procedures. See the documentation for the Sun Datacenter InfiniBand Switch 36 Firmware Version 2.0 at: http://docs.oracle.com/cd/E26698_01/index.html. The doc *Oracle Integrated Lights Out Manager (ILOM) Supplement* at http://docs.oracle.com/cd/E26698_01//pdf/E26432.pdf provides specific instruction for both CLI- and Web-based methods.

Enter the command smpartition list active to verify that partition key 0x0503 is assigned to partition name sto.

NOTE: The partition key was set to 8503 on the Oracle ZFS Storage ZS3-BA datalinks, but the InfiniBand switch reports 0503. This is intentional because the InfiniBand protocol reserves the most significant bit (0x8000) of the hexadecimal partition key (pkey) for its own use. Therefore, pkeys of 0x8503 and 0x0503 are the same.

```
[root@aiescsw-ibs0 ~]# smpartition list active
# Sun DCS IB partition config file
# This file is generated, do not edit
#! version_number : 11
Default=0x7fff, ipoib : ALL_CAS=full, ALL_SWITCHES=full, SELF=full;
SUN_DCS=0x0001, ipoib : ALL_SWITCHES=full;
ic1s10 = 0x0501,ipoib,defmember=full:
0x0021280001ef30f7,
0x0021280001ef33bf,
0x0021280001ef30b7,
0x0021280001ef314b;
ic2s10 = 0x0502,ipoib,defmember=full:
0x0021280001ef30f8,
0x0021280001ef33c0,
0x0021280001ef30b8,
0x0021280001ef314c;
sto = 0x0503,ipoib,defmember=full:
0x0021280001ef43f8,
0x0021280001ef43b7,
0x0021280001cf90c0,
0x0021280001ef43bb,
… more …
```

4. Add the Oracle ZFS Storage ZS3-BA to the InfiniBand configuration.

   a. Enter the command `smpartition start` to start a reconfiguration session.

   ```
   [root@aiescsw-ibs0 ~]# smpartition start
   [root@aiescsw-ibs0 ~]#
   ```

   b. Enter the command `smpartition add` to add the eight new GUIDs to the configuration. (NOTE: Only four GUIDs will be added for the connection to the internal InfiniBand switches of an Oracle SuperCluster T5-8 system.)

   ```
   [root@aiescsw-ibs0 ~]# smpartition add –n sto –port 21280001ef43bb
   21280001ef43bc 21280001cf90bf 21280001cf90c0 21280001ef43f7 21280001ef43f8
   21280001ef43b7 21280001ef43b8
   [root@aiescsw-ibs0 ~]#
   ```

c. Enter the command `smpartition list modified` to verify the new GUIDs have been added correctly.

```
[root@aiescsw-ibs0 ~]# smpartition list modified
# Sun DCS IB partition config file
# This file is generated, do not edit
#! version_number : 11
Default=0x7fff, ipoib : ALL_CAS=full, ALL_SWITCHES=full, SELF=full;
SUN_DCS=0x0001, ipoib : ALL_SWITCHES=full;
ic1s10 = 0x0501,ipoib,defmember=full:
0x0021280001ef30f7,
0x0021280001ef33bf,
0x0021280001ef30b7,
0x0021280001ef314b;
ic2s10 = 0x0502,ipoib,defmember=full:
0x0021280001ef30f8,
0x0021280001ef33c0,
0x0021280001ef30b8,
0x0021280001ef314c;
sto = 0x0503,ipoib,defmember=full:
0x0021280001ef43f8,
0x0021280001ef43b7,
0x0021280001cf90c0,
0x0021280001ef43bb,
0x0021280001ef43bc,
0x0021280001cf90bf,
0x0021280001ef43b8,
0x0021280001ef43f7,
0x0021280001ef3048,
0x0021280001ef30af,
0x0021280001ef30f8,
0x0021280001ef30f7,
0x0021280001ef33c0,
0x0021280001ef33bf,
0x0021280001ef30cc,
0x0021280001ef342b,
0x0021280001ef30b8,
0x0021280001ef30b7,
0x0021280001ef314c,
0x0021280001ef314b,
0x0021280001efec65,
0x0021280001efec66,
0x0021280001efecb1,
0x0021280001efecb2;
```

d. Enter the command `smpartition commit` to apply the new configuration and propagate configuration changes to all InfiniBand switches in the configuration.

```
[root@aiescsw-ibs0 ~]# smpartition commit
[root@aiescsw-ibs0 ~]#
```

6. Log off of the InfiniBand switch.

7. Back up the switch configuration according to the documented backup procedures previously noted. (See the documentation for the Sun Datacenter InfiniBand Switch 36 Firmware Version 2.0 at: http://docs.oracle.com/cd/E26698_01/index.html).

Configuring the Oracle ZFS Storage ZS3-BA Networking for Single IP Connection

This configuration is only for an Oracle SuperCluster T5-8 environment with no external leaf switches. Use the Active-Active IPMP Configuration instructions for all other configurations (next section.)

Configure the Oracle ZFS Storage ZS3-BA InfiniBand ports for network connectivity and simple cluster failover by using the following procedure to configure Port1 with the desired IP address.

1. Log on to the Head 1 BUI and navigate to **Configuration -> Network**.
2. Click on the plus (+) icon next to **Interfaces** to open the Network Interface dialog box.
3. Complete the dialog box as follows (refer to Figure 7):
    a. Enter a meaningful name for the network interface.
    b. Verify that `Enable Interface` is checked.
    c. Verify that `Allow Administration` is checked.
    d. Verify that `Use IPv4 Protocol` is checked.
    e. Verify that the `Configure with` menu selection is `Static Address List`.
    f. In the box below that, enter the desired IP address with the appropriate netmask.
    g. Verify that `Use IPv6 Protocol` is not checked.
    h. Select the datalink for `ibp0` and click **APPLY**.
4. Repeat steps 1 through 3 on Head 2 using ibp2 as the datalink.

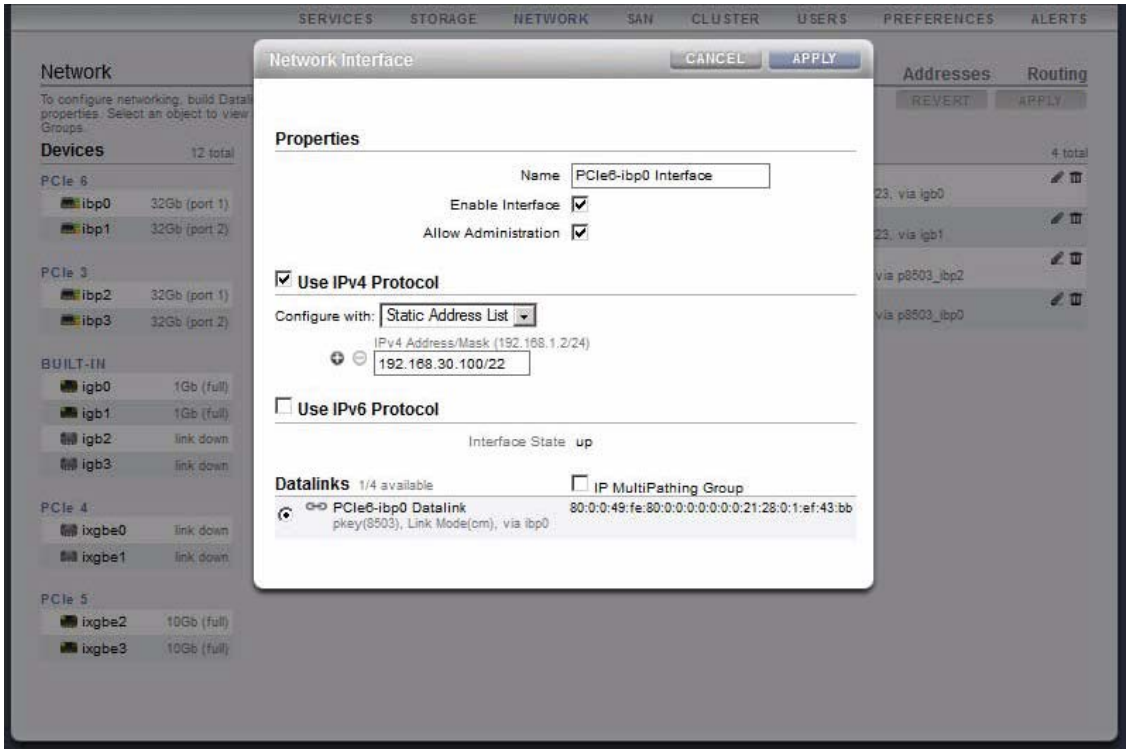

Figure 7. Configuring an IP address for an InfiniBand datalink

Configuring the Oracle ZFS Storage ZS3-BA Networking for Active-Active IPMP Connection

The InfiniBand ports on the Oracle ZFS Storage ZS3-BA must be configured for IP multipathing. Since the interfaces will be running in an active-active IPMP configuration, four IP addresses will be required for each Oracle ZFS Storage ZS3-BA head (therefore, eight addresses total). These IP addresses must be on the same private storage subnet as the <*HOST*>-storIB host addresses (see the Oracle SuperCluster host table [`/etc/inet/hosts`] on one of the Oracle SuperCluster nodes.)

To configure each InfiniBand datalink as its own network interface:

5. Log on to the Head 1 BUI and navigate to **Configuration -> Network**.
6. Click on the plus (+) icon next to **Interfaces** to open the Network Interface dialog box.
7. Complete the dialog box as follows (refer to Figure 8):
    i. Enter a meaningful name for the network interface.
    j. Verify that `Enable Interface` is checked.
    k. Verify that `Allow Administration` is checked.
    l. Verify that `Use IPv4 Protocol` is checked.

m. Verify that the `Configure with` menu selection is `Static Address List`.

n. In the box below that, enter **`0.0.0.0/8`**.

o. Verify that `Use IPv6 Protocol` is not checked.

p. Select the datalink for `ibp0` and click **APPLY**.

8. Repeat steps 2 and 3 for the remaining datalinks (`ibp1`, `ibp2`, and `ibp3`).
9. Repeat steps 1 through 4 for the datalinks on Head 2.

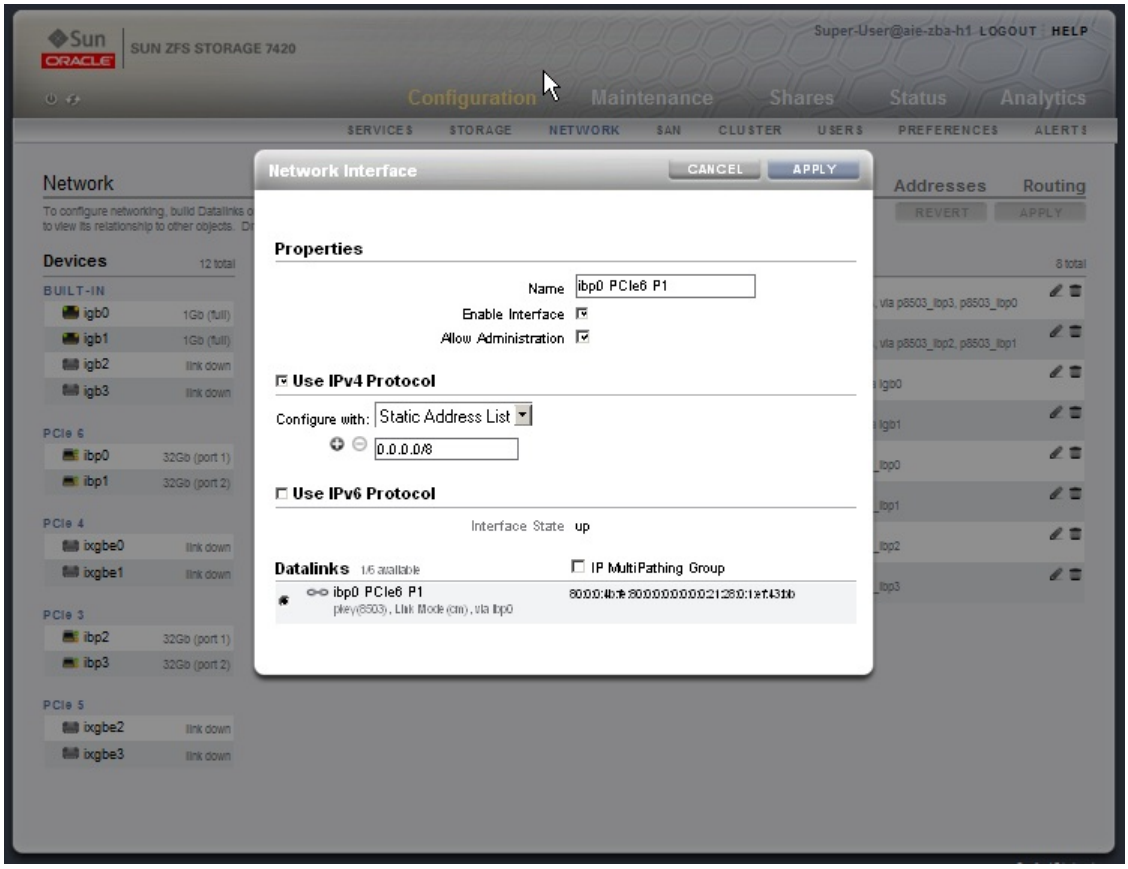

Figure 8. Creating network interfaces for each InfiniBand datalink

To configure the IPMP interface on the Oracle ZFS Storage ZS3-BA Head 1:

1. Using the Oracle ZFS Storage ZS3-BA BUI of Head 1, navigate to **Configuration -> Network**.
2. Click on the plus (+) icon next to Interfaces to open the Network Interface dialog window.
3. Complete the dialog box as follows (refer to Figure 9):

a. Enter a meaningful name for the IPMP network interface.

b. Verify that `Enable Interface` is checked.

c. Verify that `Allow Administration` is checked.

d. Verify that `Use IPv4 Protocol` is checked.

e. Verify that the `Configure with` menu selection is `Static Address List.`

f. Click the plus (+) sign next to the empty box three times, so that a total of four empty boxes are displayed.

g. In each empty box, enter one of the IP addresses reserved for the InfiniBand connections with its respective /24 netmask designation. As a best practice, do not use consecutive IP addresses from the block, but rather every other one (that is, all odd or all even).

h. Verify that `Use IPv6 Protocol` is not checked.

i. Check the `IP MultiPathing Group` box.

j. Check the boxes next to the interfaces corresponding with datalinks **`ibp0`** and **`ibp3`**.

k. Verify that each of the two interfaces is set to **`Active`** and click **APPLY**.

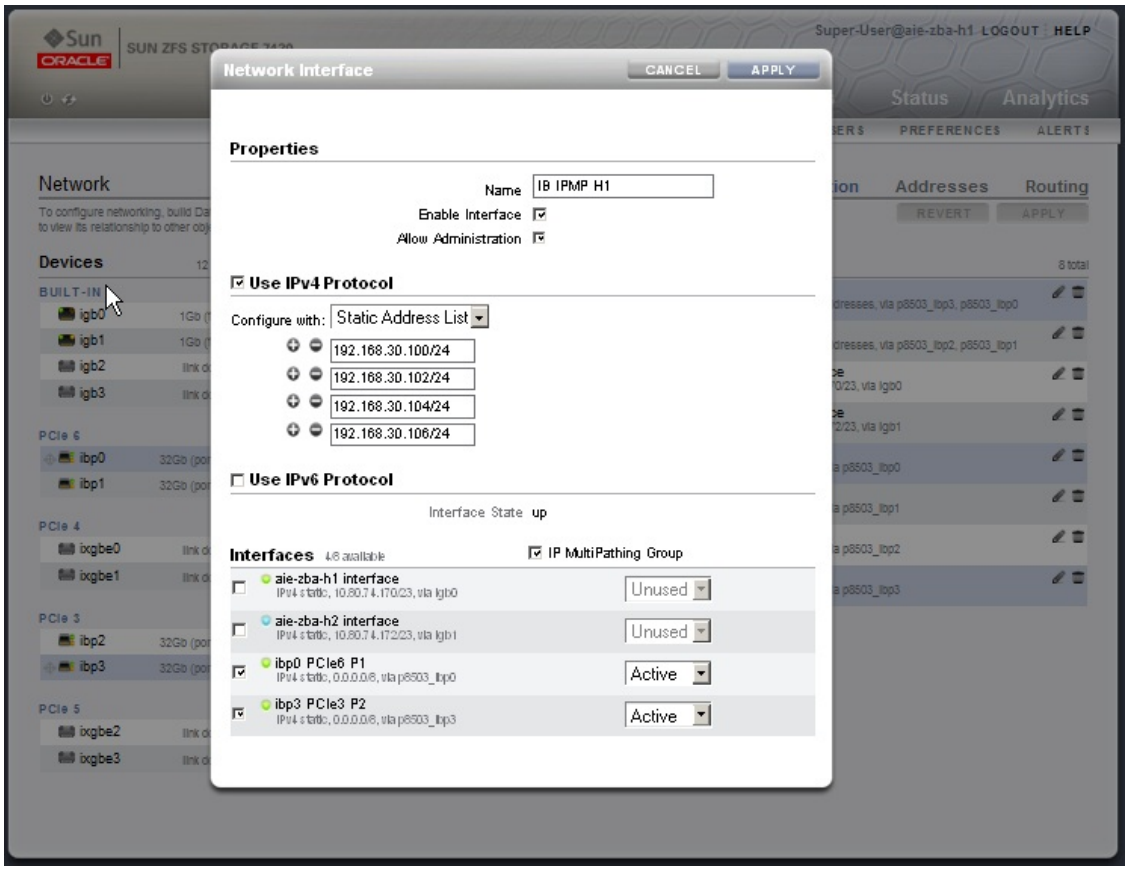

Figure 9. IPMP InfiniBand group for Oracle ZFS Storage ZS3-BA Head 1

4. From **Configuration** -> **Network**, click **Routing**.

5. Click on the `Multihoming model` button corresponding with `Adaptive` (refer to Figure 10).

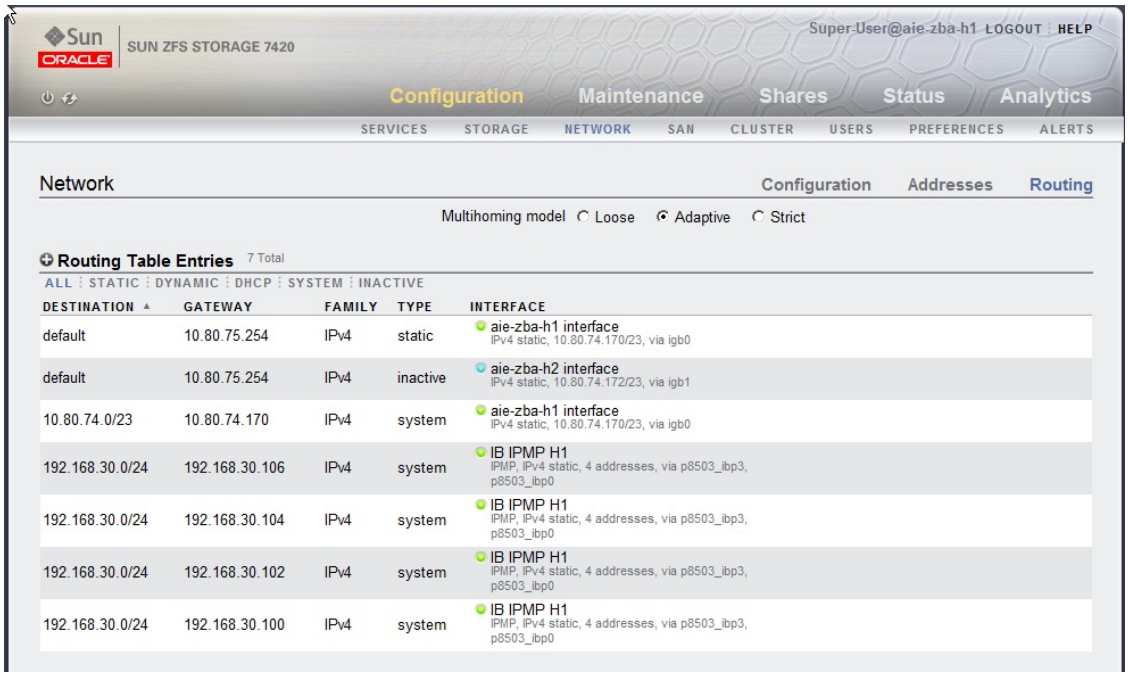

Figure 10. Selecting “Adaptive” for Multihoming model

Next, configure the IPMP interface on Head 2:

1. Log on to the BUI of Head 2 and navigate to **Configuration -> Network**.
2. Click the plus (+) icon next to **Interfaces** to open the Network Interface dialog box.
3. Complete the dialog box as follows (refer to Figure 11):
    a. Enter a meaningful name for the IPMP network interface.
    b. Verify that `Enable Interface` is checked.
    c. Verify that `Allow Administration` is checked.
    d. Verify that `Use IPv4 Protocol` is checked.
    e. Verify that the `Configure with` menu selection is `Static Address List`.
    f. Click the plus (+) sign next to the empty box three times so that a total of four empty boxes are displayed.
    g. In each empty box, enter one of the remaining four IP addresses reserved for the InfiniBand connections with their respective /24 netmask designation. These IP addresses should be the ones not used on Head 1.
    h. Verify that `Use IPv6 Protocol` is not checked.

i. Check the `IP MultiPathing Group` box.

j. Check the boxes next to the interfaces corresponding with datalinks **`ibp1`** and **`ibp2`**.

k. Verify that each of the two interfaces are set to `Active` and click **APPLY**.

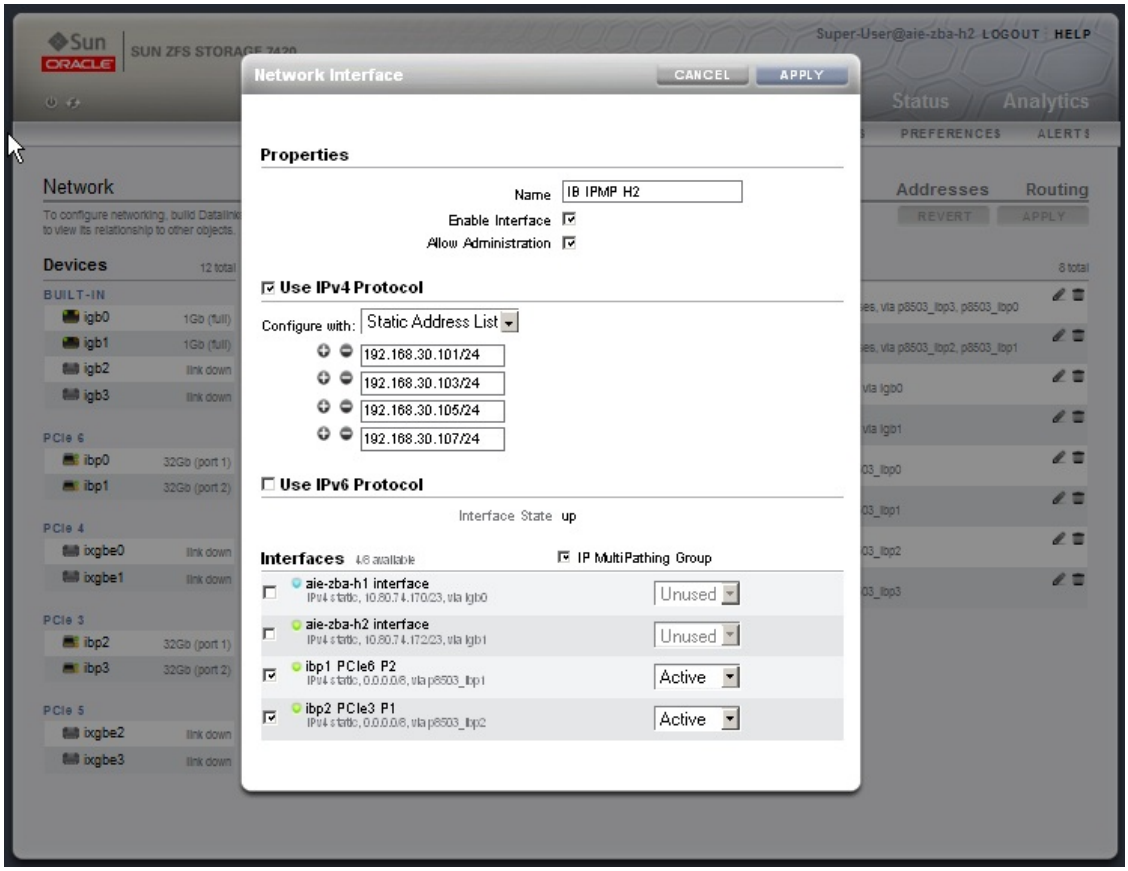

Figure 11. IPMP InfiniBand group for Oracle ZFS Storage ZS3-BA Head 2

4. From **Configuration** -> **Network**, click **Routing**.

5. Click on the `Multihoming model` button corresponding with `Adaptive`.

Finally, verify connectivity with the Oracle RAC Database SuperCluster nodes.

Verify that each of the Oracle Database nodes can ping each of the eight addresses used in the IPMP groups on the Oracle ZFS Storage ZS3-BA. Add these IP addresses to the `/etc/inet/hosts` table of each Oracle Database node.

Configuring the Oracle ZFS Storage ZS3-BA Storage Pools

Pool configuration assigns physical disk drive resources to logical storage pools for backup data storage. To maximize system throughput, configure two equally sized storage pools by assigning half of the physical drives in each drive tray to each storage pool as shown in Figure 12.

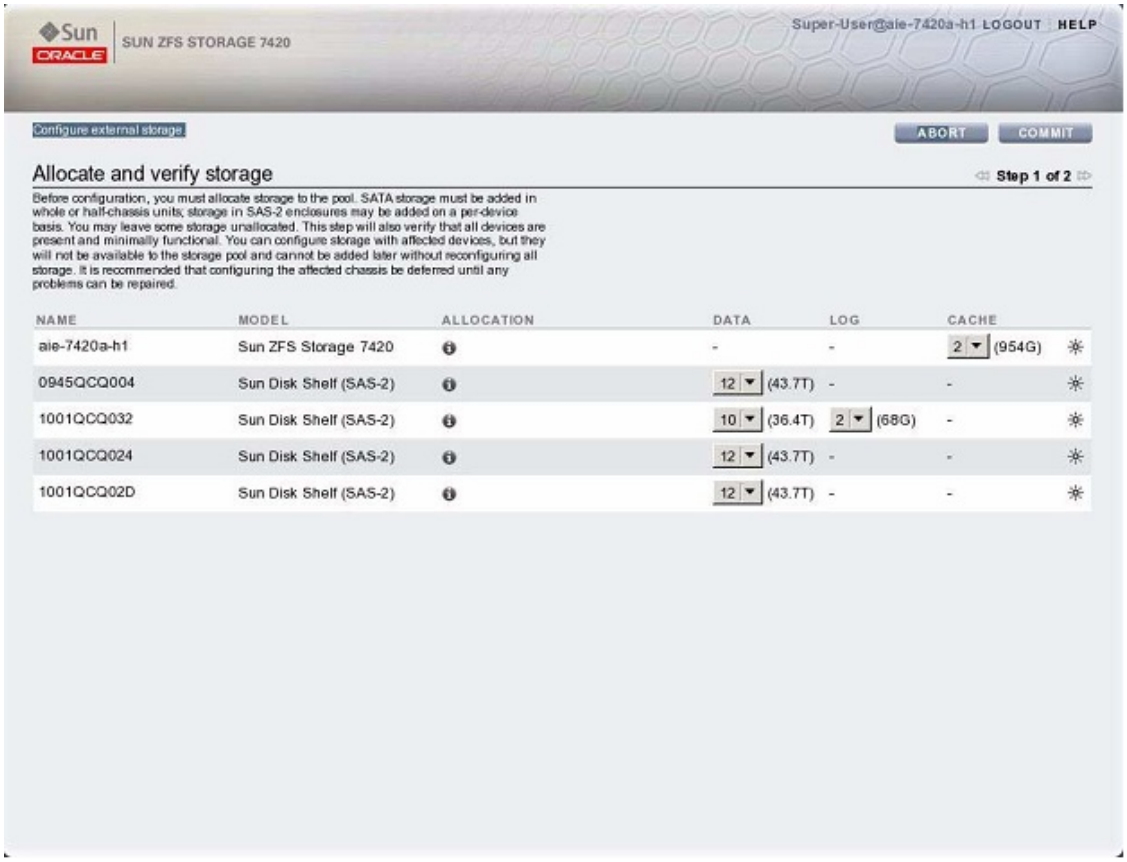

Figure 12. Storage pool configured based on half of the drives in each tray

Configuring the Oracle Engineered Systems Backup Utility for Oracle ZFS Storage Appliance

The remaining configuration tasks can be performed by the Oracle Engineered Systems Backup Utility for Oracle ZFS Storage Appliance, v2.0 and higher. This utility verifies the system configuration, sets up the requested projects and shares, configures the NFS mounts, enables dNFS, configures `oranfstab`, creates backup listener services, and even generates sample Oracle RMAN backup and recovery scripts. It can be downloaded from the Oracle ZFS Storage Appliance plug-in downloads web page found in References at the end of this document. A sample execution of the utility is also provided in Appendix A.

Otherwise, the following "Manually Configuring/Tuning …" sections must be performed.

Manually Configuring the Oracle ZFS Storage ZS3-BA RMAN Projects and Shares

Share configuration is the process of setting up and running NFS mount points for client access. Two projects should be created for the Oracle Database SuperCluster configuration; one project per pool. A project is an entity that provides a higher-level management interface point for a collection of shares. To optimize share management, update the default mount point for shares contained in the project to reference the database name, such as `/export/dbname`. For a performance-optimized system, create four shares for each project in each pool, for a total of eight shares (four on each head). To configure a project, perform the following:

1. Log on to the BUI of Head 1 and navigate to **Shares -> Projects**.
2. Click on the plus (+) icon next to **Projects**, enter a meaningful name for the project, and click on **APPLY**. Since a similar project will be created on the other head, uniquely name the project for Head 1, such as `H1-dbname`.
3. Click the Pencil icon next to the new Project Name to edit the project.
4. Click **General** and complete the properties as follows (refer to Figure 13):
    a. Change the `Mountpoint` to include the database name (for example, `/export/dbname`).
    b. Change `Synchronous write bias` from `Latency` to **`Throughput`** and click **APPLY**.

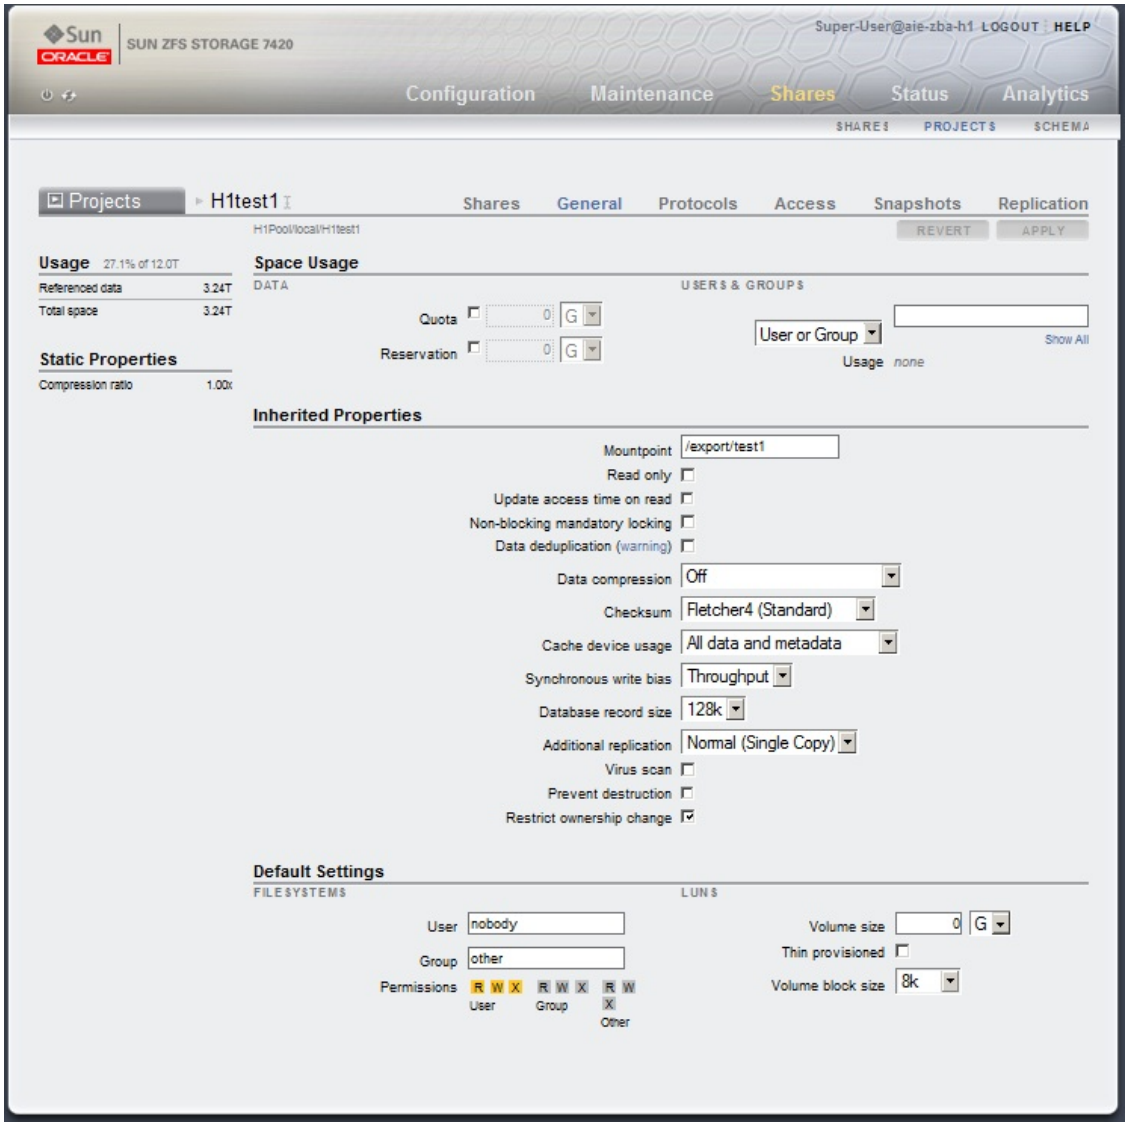

Figure 13. Project general parameter settings

5. Click Protocols and add an NFS Exception as follows (refer to Figure 14):

   a. Click on the plus (+) icon next to **NFS Exceptions**.

   b. Change `Type` to **`Network`**.

   c. Enter the subnet and netmask (for example, `192.168.30.0/24`) of the InfiniBand network.

   d. Change `Access Mode` to **`Read/write`**.

   e. Verify that `Charset` is set to **`default`**.

   f. Check the `Root Access` box and click **APPLY**.

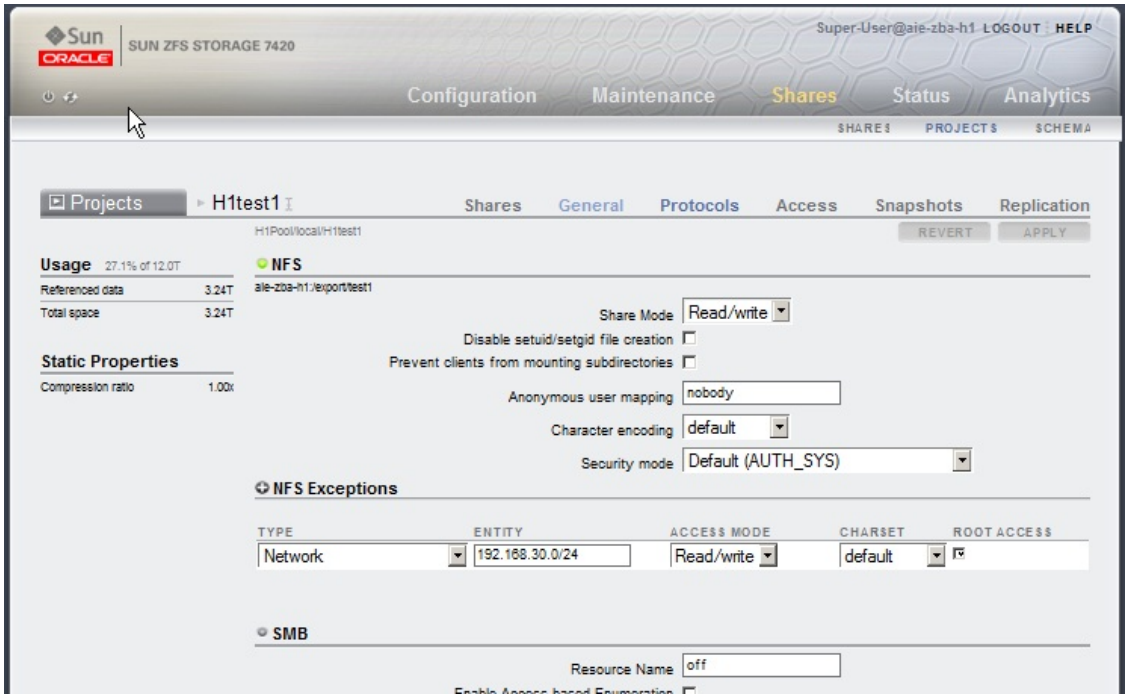

Figure 14. Setting up project NFS exceptions

6. Next to **General**, click **Shares**.

7. Create four filesystems for Head 1 and uniquely name them so they will be different from the names for Head 2. The names of the filesystems on Head 1 should be `backup1`, `backup3`, `backup5`, and `backup7`. The names of the filesystems on Head 2 will be `backup2`, `backup4`, `backup6`, and `backup8`. This is done for interleaving of the backup streams to distribute the data across the two heads, thereby providing better performance.

To create the filesystems, click the plus (+) icon next to **Filesystems**, enter the name of the filesystem (`backup1`), and click **APPLY**. Repeat this step to create the remaining three filesystems (`backup3`, `backup5`, and `backup7`). The filesystem listing should be similar to Figure 15.

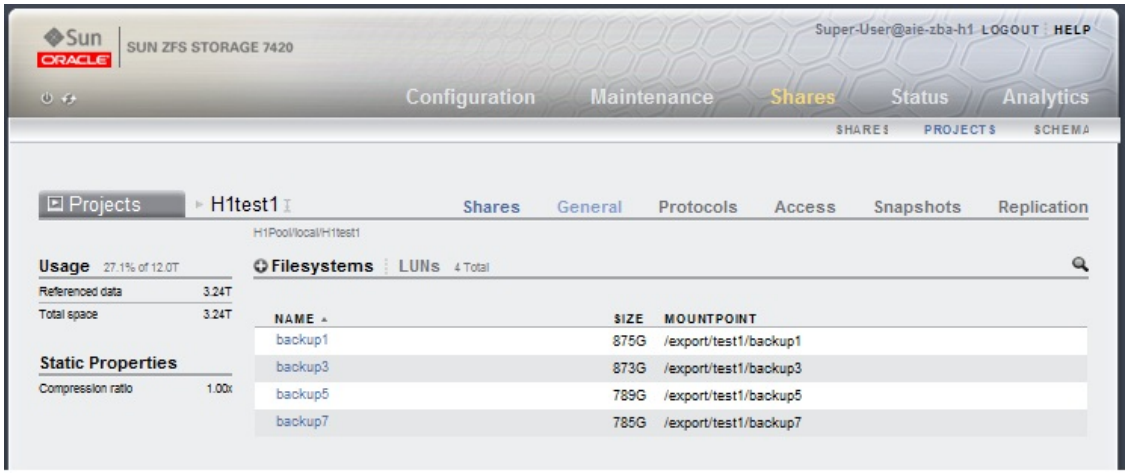

Figure 15. Filesystem listing for Head 1

8. Repeat steps 1 through 7 using Head 2. Remember to use a unique project name (for example, `H2-dbname`) for the project, and specify the even-numbered backup IDs (`backup2`, `backup4`, `backup6`, and `backup8`) for the filesystem names. The filesystem listing should be similar to the example in figure 16.

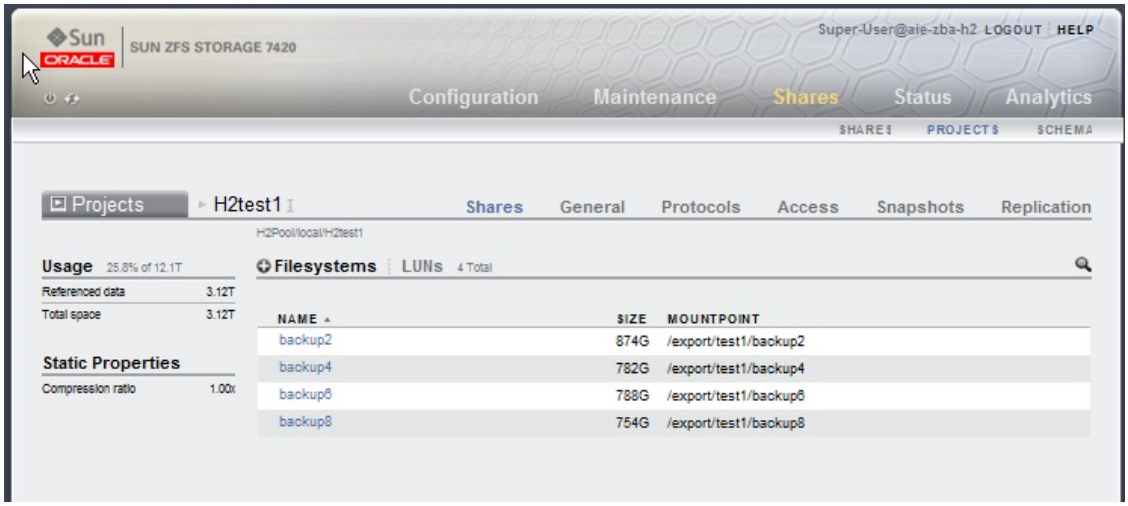

Figure 16. Filesystem listing for Head 2

Tuning the Oracle Solaris 11 Kernel and Network

Starting with the January 2013 quarterly full stack download patch (QFSDP) release of Oracle SuperCluster updates, an "ssctuner" process exists to verify that the Oracle SuperCluster is running at the optimal Network and Kernel configuration. The following settings should already be in effect on the system:

In the `/etc/system` file of each Oracle SuperCluster node:

```
set rpcmod:clnt_max_conns = 8
set nfs:nfs3_bsize = 131072
```

The Oracle Solaris 11 `ipadm` command should display the following "CURRENT" `tcp` settings for `max_buf`, `send_buf`, and `recv_buf`:

```
root:~# ipadm show-prop -p max_buf,send_buf,recv_buf tcp
PROTO PROPERTY    PERM CURRENT      PERSISTENT   DEFAULT      POSSIBLE
tcp   max_buf     rw   2097152      --           1048576      1048576-1073741824
tcp   send_buf    rw   1048576      --           49152        4096-2097152
tcp   recv_buf    rw   1048576      --           128000       2048-2097152
root:~#
```

Use the `ipadm set-prop –p <property>=<value> tcp` command to make any necessary changes.

The ssctuner process will flag warnings if any of these settings are not in effect.

NOTE: The ssctuner sets these values only in the Global Zone. If additional zones are used for databases, each zone needs to be manually set with these tcp values.

Manually Configuring the Client NFS Mount

For configuration, any server that accesses the Oracle ZFS Storage ZS3-BA, including Oracle SuperCluster nodes, is considered a client. Configuring the client NFS mount includes creating the target directory structure for access to the Oracle ZFS Storage ZS3-BA as well as the specific NFS mount options necessary for optimal system performance. Mount options for Oracle Solaris clients are:

```
rw,bg,hard,nointr,rsize=131072,wsize=131072,proto=tcp,vers=3
```

The mount points of the directories created on the Oracle ZFS Storage ZS3-BA should be created on each of the Oracle Database SuperCluster nodes and added to their `/etc/inet/hosts` table.

Manually Configuring Oracle Direct NFS (dNFS)

On each Oracle Database SuperCluster node, configure dNFS as follows:

1. Shut down the running instance of the Oracle Database software.
2. Change directory to `$ORACLE_HOME/rdbms/lib`.
3. Enable dNFS:

   ```
   make –f $ORACLE_HOME/rdbms/lib/ins_rdbms.mk dnfs_on
   ```

4. Update the `oranfstab` (located in `$ORACLE_HOME/dbs`) with the server, path, and export names specific to the configuration.

a. The `server` parameter refers to the local name of the Oracle ZFS Storage ZS3-BA head on the InfiniBand network.

b. The `path` parameters should reflect the addresses for that head specified in the IPMP groups.

c. The `export` parameters should reflect the mount points similar to the entries created in the `/etc/vfstab`.

The entries should look similar to these:

```
server: aie-ZS3-BA-h1-stor
path: 192.168.30.100
path: 192.168.30.102
path: 192.168.30.104
path: 192.168.30.106
export: /export/test1/backup1 mount: /ZS3-BA/test1/backup1
export: /export/test1/backup3 mount: /ZS3-BA/test1/backup3
export: /export/test1/backup5 mount: /ZS3-BA/test1/backup5
export: /export/test1/backup7 mount: /ZS3-BA/test1/backup7
server: aie-ZS3-BA-h2-stor
path: 192.168.30.101
path: 192.168.30.103
path: 192.168.30.105
path: 192.168.30.107
export: /export/test1/backup2 mount: /ZS3-BA/test1/backup2
export: /export/test1/backup4 mount: /ZS3-BA/test1/backup4
export: /export/test1/backup6 mount: /ZS3-BA/test1/backup6
export: /export/test1/backup8 mount: /ZS3-BA/test1/backup8
```

5. Restart the Oracle Database software instance.

Manually Tuning the Oracle Database Instance for Oracle RMAN

Optimizing high-bandwidth backup and restore operations using Oracle RMAN and the Oracle ZFS Storage ZS3-BA requires adjusting the instance parameters that control I/O buffering. For information about how to tune these parameters, see Article ID 1072545.1: *RMAN Performance Tuning Using Buffer Memory Parameters* at http://support.oracle.com.

For Oracle Database on SuperCluster, consider tuning the following four parameters:

- `_backup_disk_bufcnt` – Number of buffers used to process backup sets
- `_backup_disk_bufsz` – Size of the buffers used to process backup sets
- `_backup_file_bufcnt` – Number of buffers used to process image copies
- `_backup_file_bufsz` – Size of the buffers used to process image copies

For backup and restore operations on backup sets and image copies, set the number of buffers to 64 and the buffer size to 1 MB:

```
SQL> alter system set "_backup_disk_bufcnt"=64;
SQL> alter system set "_backup_file_bufcnt"=64;
SQL> alter system set "_backup_disk_bufsz"=1048576;
```

```
SQL> alter system set "_backup_file_bufsz"=1048576;
```

These commands may be configured persistently by adding them to the SPFILE, or they may be dynamically set in the Oracle RMAN run block used to execute the backup or restore operations.

The following code fragments show how to dynamically tune the buffer sizes and counts for backup and restore operations.

- Backup set backup:

```
run
{
       sql 'alter system set "_backup_disk_bufcnt"=64';
       sql 'alter system set "_backup_disk_bufsz"=1048576';
       allocate channel...
       ...
       backup as backupset database;
}
```

- Backup set restore:

```
run
{
       sql 'alter system set "_backup_disk_bufcnt"=64';
       sql 'alter system set "_backup_disk_bufsz"=1048576';
       allocate channel...
       ...
       restore database;
}
```

- Image copy backup:

```
run
{
       sql 'alter system set "_backup_file_bufcnt"=64';
       sql 'alter system set "_backup_file_bufsz"=1048576';
       allocate channel...
       ...
       backup as copy database;
}
```

- Image copy restore:

```
run
{
       sql 'alter system set "_backup_file_bufcnt"=64';
       sql 'alter system set "_backup_file_bufsz"=1048576';
       allocate channel...
       ...
       restore database;
}
```

Performing an incrementally applied backup requires reading an incremental backup set and writing to an image copy. To tune buffers for incrementally applied backups, run the following:

```
run
{
       sql 'alter system set "_backup_disk_bufcnt"=64';
       sql 'alter system set "_backup_disk_bufsz"=1048576';
       sql 'alter system set "_backup_file_bufcnt"=64';
       sql 'alter system set "_backup_file_bufsz"=1048576';
       allocate channel...
       ...
       recover copy of database;
```

```
}
```

Manually Configuring Oracle RMAN

Configuring Oracle RMAN channel and parallelism includes specifying the filesystem targets for the Oracle RMAN backup channels and the total number of channels used for backup and restore operations. Performance benefits can be realized by configuring 16 or 32 Oracle RMAN channels such that they are evenly distributed over the Oracle Database instances and nodes in the Oracle RAC cluster and evenly distributed over the shares exported from the Oracle ZFS Storage ZS3-BA.

The following code fragments show sample Oracle RMAN run blocks for performing backup and restore operations for backup sets and image copies as well as applying incremental merges to image copies. The sample code is based on the following database configuration:

- Database name: `mydb`
- SYSDBA login: `sys/welcome`
- Scan address: `myssc-scan`
- Service names for the backup: `mydb_bkup[1-2]`

The Oracle RMAN run blocks for backup and restore using backup sets and image copies are shown in the examples in the following sections. In these examples, the mount points for the eight shares are accessed as `/ZS3-BA/mydb/backup1` through `/ZS3-BA/mydb/backup8`. Sixteen channels are configured, two per filesystem.

Backup set level 0 backup:

```
run
{
       sql 'alter system set "_backup_disk_bufcnt"=64 scope=memory';
       sql 'alter system set "_backup_disk_bufsz"=1048576 scope=memory';
       allocate channel ch01 device type disk connect 'sys/welcome@myssc-
       scan/dbname_bkup1' format '/ZS3-BA/mydb/backup1/%U';
       allocate channel ch02 device type disk connect 'sys/welcome@myssc-
       scan/dbname_bkup2' format '/ZS3-BA/mydb/backup2/%U';
       allocate channel ch03 device type disk connect 'sys/welcome@myssc-
       scan/dbname_bkup1' format '/ZS3-BA/mydb/backup3/%U';
       allocate channel ch04 device type disk connect 'sys/welcome@myssc-
       scan/dbname_bkup2' format '/ZS3-BA/mydb/backup4/%U';
       allocate channel ch05 device type disk connect 'sys/welcome@myssc-
       scan/dbname_bkup1' format '/ZS3-BA/mydb/backup5/%U';
       allocate channel ch06 device type disk connect 'sys/welcome@myssc-
       scan/dbname_bkup2' format '/ZS3-BA/mydb/backup6/%U';
       allocate channel ch07 device type disk connect 'sys/welcome@myssc-
       scan/dbname_bkup1' format '/ZS3-BA/mydb/backup7/%U';
       allocate channel ch08 device type disk connect 'sys/welcome@myssc-
       scan/dbname_bkup2' format '/ZS3-BA/mydb/backup8/%U';
       allocate channel ch09 device type disk connect 'sys/welcome@myssc-
       scan/dbname_bkup1' format '/ZS3-BA/mydb/backup1/%U';
       allocate channel ch10 device type disk connect 'sys/welcome@myssc-
       scan/dbname_bkup2' format '/ZS3-BA/mydb/backup2/%U';
       allocate channel ch11 device type disk connect 'sys/welcome@myssc-
       scan/dbname_bkup1' format '/ZS3-BA/mydb/backup3/%U';
       allocate channel ch12 device type disk connect 'sys/welcome@myssc-
       scan/dbname_bkup2' format '/ZS3-BA/mydb/backup4/%U';
       allocate channel ch13 device type disk connect 'sys/welcome@myssc-
```

```
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup5/%U';
        allocate channel ch14 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup6/%U';
        allocate channel ch15 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup7/%U';
        allocate channel ch16 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup8/%U';
        configure snapshot controlfile name to '/ZS3-BA/mydb/backup1/snapcf_dbname.f';
        backup as backupset incremental level 0 section size 32g database
        tag 'FULLBACKUPSET_L0' plus archivelog tag 'FULLBACKUPSET_L0';
}
```

Backup set level 1 backup:

```
run
{
        sql 'alter system set "_backup_disk_bufcnt"=64 scope=memory';
        sql 'alter system set "_backup_disk_bufsz"=1048576 scope=memory';
        allocate channel ch01 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup1/%U';
        allocate channel ch02 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup2/%U';
        allocate channel ch03 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup3/%U';
        allocate channel ch04 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup4/%U';
        allocate channel ch05 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup5/%U';
        allocate channel ch06 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup6/%U';
        allocate channel ch07 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup7/%U';
        allocate channel ch08 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup8/%U';
        allocate channel ch09 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup1/%U';
        allocate channel ch10 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup2/%U';
        allocate channel ch11 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup3/%U';
        allocate channel ch12 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup4/%U';
        allocate channel ch13 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup5/%U';
        allocate channel ch14 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup6/%U';
        allocate channel ch15 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup7/%U';
        allocate channel ch16 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup8/%U';
        configure snapshot controlfile name to
        '/ZS3-BA/mydb/backup1/snapcf_dbname.f';
        backup as backupset incremental level 1 database tag
        'FULLBACKUPSET_L1' plus archivelog tag 'FULLBACKUPSET_L1';
}
```

Image copy backup:

```
run
{
        sql 'alter system set "_backup_file_bufcnt"=64 scope=memory';
        sql 'alter system set "_backup_file_bufsz"=1048576 scope=memory';
        allocate channel ch01 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup1/%U';
        allocate channel ch02 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup2/%U';
        allocate channel ch03 device type disk connect 'sys/welcome@myssc-
```

```
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup3/%U';
        allocate channel ch04 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup4/%U';
        allocate channel ch05 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup5/%U';
        allocate channel ch06 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup6/%U';
        allocate channel ch07 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup7/%U';
        allocate channel ch08 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup8/%U';
        allocate channel ch09 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup1/%U';
        allocate channel ch10 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup2/%U';
        allocate channel ch11 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup3/%U';
        allocate channel ch12 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup4/%U';
        allocate channel ch13 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup5/%U';
        allocate channel ch14 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup6/%U';
        allocate channel ch15 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1' format '/ZS3-BA/mydb/backup7/%U';
        allocate channel ch16 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2' format '/ZS3-BA/mydb/backup8/%U';
        configure snapshot controlfile name to
        '/ZS3-BA/mydb/backup1/snapcf_dbname.f';
        backup incremental level 1 for recover of copy with tag 'IMAGECOPY'
        database;
}
```

Incremental merge to image copy:

```
run
{
        sql 'alter system set "_backup_disk_bufcnt"=64 scope=memory';
        sql 'alter system set "_backup_disk_bufsz"=1048576 scope=memory';
        sql 'alter system set "_backup_file_bufcnt"=64 scope=memory';
        sql 'alter system set "_backup_file_bufsz"=1048576 scope=memory';
        allocate channel ch01 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup1';
        allocate channel ch02 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup2';
        allocate channel ch03 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup1';
        allocate channel ch04 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup2';
        allocate channel ch05 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup1';
        allocate channel ch06 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup2';
        allocate channel ch07 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup1';
        allocate channel ch08 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup2';
        allocate channel ch09 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup1';
        allocate channel ch10 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup2';
        allocate channel ch11 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup1';
        allocate channel ch12 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup2';
        allocate channel ch13 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup1';
```

```
        allocate channel ch14 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup2';
        allocate channel ch15 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup1';
        allocate channel ch16 device type disk connect 'sys/welcome@ad01-
        scan/dbname_bkup2';
        configure snapshot controlfile name to
        '/zfssa/dbname/backup1/snapcf_dbname.f';
        recover copy of database with tag 'IMAGECOPY';
}
```

Restore validate:

```
run
{
        sql 'alter system set "_backup_disk_bufcnt"=64 scope=memory';
        sql 'alter system set "_backup_disk_bufsz"=1048576 scope=memory';
        sql 'alter system set "_backup_file_bufcnt"=64 scope=memory';
        sql 'alter system set "_backup_file_bufsz"=1048576 scope=memory';
        allocate channel ch01 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1';
        allocate channel ch02 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2';
        allocate channel ch03 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1';
        allocate channel ch04 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2';
        allocate channel ch05 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1';
        allocate channel ch06 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2';
        allocate channel ch07 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1';
        allocate channel ch08 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2';
        allocate channel ch09 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1';
        allocate channel ch10 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2';
        allocate channel ch11 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1';
        allocate channel ch12 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2';
        allocate channel ch13 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1';
        allocate channel ch14 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2';
        allocate channel ch15 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup1';
        allocate channel ch16 device type disk connect 'sys/welcome@myssc-
        scan/dbname_bkup2';
        configure snapshot controlfile name to
        '/ZS3-BA/mydb/backup1/snapcf_dbname.f';
        restore validate database;
}
```

Configuring DTrace Analytics in the Oracle ZFS Storage ZS3-BA

The Oracle ZFS Storage ZS3-BA includes a comprehensive performance analysis tool call DTrace Analytics. DTrace Analytics is a framework that monitors important subsystem performance accounting statistics. A subnet of the available accounting statistics should be monitored to provide comprehensive data on the effectiveness and performance of Oracle RMAN backup and restore

workloads. The following analytics are configured on the Oracle ZFS Storage ZS3-BA (**Configuration -> Preferences -> Enable Advanced Analytics**):

- CPU: Percent utilization broken down by CPU mode
- Disk: Average number of I/O operations broken down by state of operation
- Disk: I/O bytes per second broken down by type of operation
- Disk: I/O operations per second broken down by latency
- Disk: Disks with utilization of at least 95 percent broken down by disk
- Network: Interface bytes per second broken down by direction
- Network: Interface bytes per second broken down by interface
- Protocol: NFSv3 operations per second broken down by size
- Protocol: NFSv3 operations per second broken down by type of operation
- Protocol: NFSv3 operations per second of type read broken down by latency
- Protocol: NFSv3 operations per second of type write broken down by latency
- Protocol: NFSv3 operations per second of type read broken down by size
- Protocol: NFSv3 operations per second of type write broken down by size

Implementing these accounting statistics helps end users gain a quantitative understanding of the instantaneous and historical resource consumption used in Quality of Service (QoS) Management for their specific implementation.

Conclusion

The Oracle ZFS Storage ZS3-BA is a high-performance backup and recovery system that is optimized for Oracle's engineered systems. In the case of Oracle SuperCluster Engineered Systems, it easily integrates into the InfiniBand environment and enables high-speed backup and restore for Oracle Databases capitalizing on the dNFS features.

Appendix A: Oracle Engineered Systems Backup Utility for Oracle ZFS Storage Appliance, v2.0, Example Session

```
root@aiesscdbadm01:~/esbu# ./esbu

This utility configures Engineered System database backup on
                 Oracle ZFS Storage Appliance


The following types of Engineered Systems are supported for database backup:
    0. Exadata
    1. SuperCluster
    2. ODA


Enter the number corresponding to the type of Engineered System database to
 be backed up: 1


Step 1 Configure Oracle ZFS Storage Appliance

In this step a project and shares will be created in the appliance.
Prior to this step, a storage pool and a network interface must be
configured first using the BUI of the appliance
(https://appliance-mgt-ip:215).

To complete this step successfully, the following information will
be required:

- Project name to be created on Oracle ZFS Storage Appliance. All
  shares will be created under this project name
- IP addresses of the Engineered System to have NFS read/write
  access
- Number of shares per pool to be used for database backup
- Management user name and host name of the appliance

Continue? (y|n) y

Enter the project name to be created on Oracle ZFS Storage Appliance:
EXDBPROJECT

Enter the IPs of the Engineered System for NFS rw access:
List all IPs or use netmasks. Separate multiple entries by comma (,)
Example: 10.80.36.0/24,10.80.37.81/0 (netmask is required)
192.168.30.2/22,192.168.30.3/22


Enter the number of shares per pool to be used for database backup:
Use 1 share per pool for management-optimized systems.
Use 2 - 4 shares per pool for performance-optimized systems.
4

Enter the host name of the appliance:
zba-h1

Enter the user name of the appliance:
root

Obtaining system info from root@zba-h1 ...

Type:           Sun ZFS Storage 7420
Version:        ak/SUNW,otoro@2013.06.05.1.1,1-1.2
Cluster:        true
```

```
Cluster State:  AKCS_CLUSTERED

Continue? (y|n) y

Obtaining available pools from root@zba-h1 ...

Available storage pools on zba-h1:

                  POOL          PROFILE           STATUS
  1 -           H1Pool           raidz1           online
  0 - Desired pool does not exist. Quit
Enter digit to select a pool: 1

H1Pool will be used for the database backup

Obtaining available network interfaces from zba-h1 ...
Available network interfaces to access H1Pool:

       INTERFACE          ADDRESS            STATE            LABEL
            igb0  10.80.74.170/23              up H1 Admin Interface
           ipmp1 192.168.30.100/22               up H1 IB IPMP Group
           ipmp1 192.168.30.102/22               up H1 IB IPMP Group
           ipmp1 192.168.30.104/22               up H1 IB IPMP Group
           ipmp1 192.168.30.106/22               up H1 IB IPMP Group
          ixgbe2  192.168.66.1/24               up Ledbetter Interface
          ixgbe2 192.168.66.66/24               up Ledbetter Interface
      p8503_ibp0        0.0.0.0/8               up PCIe6-ibp0 Interface
      p8503_ibp3        0.0.0.0/8               up PCIe3-ibp3 Interface

Enter the IP addesses to access the selected pool:
List all IPs without netmasks. Separate multiple entries by comma (,)
Example: 10.80.36.81,10.80.37.82
192.168.30.100,192.168.30.102,192.168.30.104,192.168.30.106

192.168.30.100,192.168.30.102,192.168.30.104,192.168.30.106 will be used to access
H1Pool

Select one more pool? (y|n) y

Select the second pool from the cluster peer head? (y|n) y

Obtaining system info from root@zba-h2 ...

Type:           Sun ZFS Storage 7420
Version:        ak/SUNW,otoro@2013.06.05.1.1,1-1.2
Cluster:        true
Cluster State:  AKCS_CLUSTERED

Continue? (y|n) y

Obtaining available pools from root@zba-h2 ...

Available storage pools on zba-h2:

                  POOL          PROFILE           STATUS
  1 -           H2Pool           raidz1           online
  0 - Desired pool does not exist. Quit
Enter digit to select a pool: 1

H2Pool will be used for the database backup

Obtaining available network interfaces from zba-h2 ...
Available network interfaces to access H2Pool:

       INTERFACE          ADDRESS            STATE            LABEL
            igb1  10.80.74.172/23              up H2 Admin Interface
```

```
          ipmp2 192.168.30.101/22              up H2 IB IPMP Grouip
          ipmp2 192.168.30.103/22              up H2 IB IPMP Grouip
          ipmp2 192.168.30.105/22              up H2 IB IPMP Grouip
          ipmp2 192.168.30.107/22              up H2 IB IPMP Grouip
         ixgbe0 192.168.212.2/24              up     ZBAinterface
     p8503_ibp1        0.0.0.0/8              up PCIe6-ibp1 Interface
     p8503_ibp2        0.0.0.0/8              up PCIe3-ibp2 Interface

Enter the IP addesses to access the selected pool:
List all IPs without netmasks. Separate multiple entries by comma (,)
Example: 10.80.36.81,10.80.37.82
192.168.30.101,192.168.30.103,192.168.30.105,192.168.30.107

192.168.30.101,192.168.30.103,192.168.30.105,192.168.30.107 will be used to access
H2Pool
Verify configuration settings:

    Engineered System type: supercluster
    ZFSSA project name:     EXDBPROJECT
    ZFSSA login:            root
    ZFSSA shares per pool:  4
    ZFSSA hostname - 1:     zba-h1
    ZFSSA pool - 1:         H1Pool
    ZFSSA interface - 1:    192.168.30.100,192.168.30.102,192.168.30.104,192.168.30.106
    ZFSSA nfs exception:    192.168.30.2/22,192.168.30.3/22
    ZFSSA hostname - 2:     zba-h2
    ZFSSA pool - 2:         H2Pool
    ZFSSA interface - 2:    192.168.30.101,192.168.30.103,192.168.30.105,192.168.30.107

The following NFS shares will be used to back up database:


192.168.30.100,192.168.30.102,192.168.30.104,192.168.30.106:/export/EXDBPROJECT/backup1

192.168.30.100,192.168.30.102,192.168.30.104,192.168.30.106:/export/EXDBPROJECT/backup3

192.168.30.100,192.168.30.102,192.168.30.104,192.168.30.106:/export/EXDBPROJECT/backup5

192.168.30.100,192.168.30.102,192.168.30.104,192.168.30.106:/export/EXDBPROJECT/backup7


192.168.30.101,192.168.30.103,192.168.30.105,192.168.30.107:/export/EXDBPROJECT/backup2

192.168.30.101,192.168.30.103,192.168.30.105,192.168.30.107:/export/EXDBPROJECT/backup4

192.168.30.101,192.168.30.103,192.168.30.105,192.168.30.107:/export/EXDBPROJECT/backup6

192.168.30.101,192.168.30.103,192.168.30.105,192.168.30.107:/export/EXDBPROJECT/backup8

Correct? (y|n) y

Verifying project and shares on H1Pool ...
Verifying project and shares on H2Pool ...

Verifying IPMP properties ... - zba-h1

Creating project and shares on H1Pool ...
project created - EXDBPROJECT
share created - backup1
share created - backup3
share created - backup5
share created - backup7

Verifying IPMP properties ... - zba-h2

Creating project and shares on H2Pool ...
```

```
project created - EXDBPROJECT
share created - backup2
share created - backup4
share created - backup6
share created - backup8

Step 1 Configure Oracle ZFS Storage Appliance
Completed Successfully.

Continue to Step 2 Configure Engineered System Database Nodes? (y|n) y


Step 2 Configure SuperCluster Database Nodes

In this step the SuperCluster database nodes will be configured. Prior to this
step, the Oracle ZFS Storage Appliance must be configured by following
Step 1.

To complete this step successfully, the following information will
be required:

- Name of the database (DB_UNIQUE_NAME) to be backed up
- Host name (node name) of a SuperCluster database node
- Owner user and group name of the database and RMAN. The owner will own
  the NFS directories from the Oracle ZFS Storage Appliance where RMAN
  stores database backup.
- File permission of the NFS directories where RMAN will store backup.
- (Optional) Host IP addresses of the SuperCluster database nodes that will be
  used by Oracle Direct NFS (dNFS).

The following configuration will be performed in this step:
- Start database - DB_UNIQUE_NAME
- Create mount points /zfssa/DB_UNIQUE_NAME/backup[1-8]
- Add entries to /etc/vfstab
- Create and start /etc/init.d/zfssa_DB_UNIQUE_NAME
  to mount the NFS shares of the ZFSSA
- Add entries to ORACLE_HOME/dbs/oranfstab
- Enable Oracle Direct NFS
- Add backup services
- Restart database DB_UNIQUE_NAME

Continue? (y|n) y


Enter the database name (DB_UNIQUE_NAME) to be backed up:
exdb


Enter the host name of a SuperCluster database node:
aiesscdbadm01


Enter the owner user name of the database and RMAN. The user will
own the NFS directories where RMAN stores backup:
Hit Enter key to use the default (oracle)



Enter the owner group name of the database and RMAN. The group will
own the NFS directories where RMAN stores backup:
Hit Enter key to use the default (oinstall)



Enter the file mode of the NFS directories where RMAN stores backup:
Hit Enter key to use the default (750)
```

```
Verify configuration settings:

    Database name:          exdb
    SuperCluster hostname:  aiesscdbadm01
    DB Backup - owner:      oracle:oinstall
    DB Backup - file mode:  750

Correct? (y|n) y

About to start configuring SuperCluster database nodes
Continue? (y|n) y

Current ORACLE_HOME is set to /u01/app/oracle/product/11.2.0.3/dbhome_1_cloned

Correct? (y|n) n

Enter the ORACLE_HOME path:
/u01/app/oracle/product/11.2.0.3/dbhome_1

ORACLE_HOME is set to /u01/app/oracle/product/11.2.0.3/dbhome_1

Obtaining database status - exdb
Database is already running - exdb

Instance exdb1 is running on node aiesscdbadm01
Instance exdb2 is running on node aiesscdbadm02

Continue? (y|n) y

Performing system precheck - aiesscdbadm01
Performing system precheck - aiesscdbadm02

Creating mount points - aiesscdbadm01
/zfssa/exdb/backup1 - created
/zfssa/exdb/backup3 - created
/zfssa/exdb/backup5 - created
/zfssa/exdb/backup7 - created
/zfssa/exdb/backup2 - created
/zfssa/exdb/backup4 - created
/zfssa/exdb/backup6 - created
/zfssa/exdb/backup8 - created

Adding mount points to vfstab - aiesscdbadm01
Older version /etc/vfstab is saved to /etc/vfstab_20140507T080638
/etc/vfstab: /zfssa/exdb/backup1 - added
/etc/vfstab: /zfssa/exdb/backup3 - added
/etc/vfstab: /zfssa/exdb/backup5 - added
/etc/vfstab: /zfssa/exdb/backup7 - added
/etc/vfstab: /zfssa/exdb/backup2 - added
/etc/vfstab: /zfssa/exdb/backup4 - added
/etc/vfstab: /zfssa/exdb/backup6 - added
/etc/vfstab: /zfssa/exdb/backup8 - added

Creating /etc/init.d/zfssa_exdb - aiesscdbadm01
start service - zfssa_exdb
/etc/init.d/zfssa_exdb - created

Does aiesscdbadm01 have more than one network interface
 available for the dNFS data paths? (y|n) n

Older version /u01/app/oracle/product/11.2.0.3/dbhome_1/dbs/oranfstab is saved to
/u01/app/oracle/product/11.2.0.3/dbhome_1/dbs/oranfstab_20140507T080657

Updating oranfstab - aiesscdbadm01
```

```
/bin/chown oracle:oinstall /u01/app/oracle/product/11.2.0.3/dbhome_1/dbs/oranfstab -
aiesscdbadm01
/bin/chmod 640 /u01/app/oracle/product/11.2.0.3/dbhome_1/dbs/oranfstab - aiesscdbadm01
/usr/bin/make -f /u01/app/oracle/product/11.2.0.3/dbhome_1/rdbms/lib/ins_rdbms.mk
dnfs_on ORACLE_HOME=/u01/app/oracle/product/11.2.0.3/dbhome_1 - aiesscdbadm01
dNFS is enabled - aiesscdbadm01

Creating mount points - aiesscdbadm02
/zfssa/exdb/backup1 - created
/zfssa/exdb/backup3 - created
/zfssa/exdb/backup5 - created
/zfssa/exdb/backup7 - created
/zfssa/exdb/backup2 - created
/zfssa/exdb/backup4 - created
/zfssa/exdb/backup6 - created
/zfssa/exdb/backup8 - created

Adding mount points to vfstab - aiesscdbadm02
Older version /etc/vfstab is saved to /etc/vfstab_20140507T080712
/etc/vfstab: /zfssa/exdb/backup1 - added
/etc/vfstab: /zfssa/exdb/backup3 - added
/etc/vfstab: /zfssa/exdb/backup5 - added
/etc/vfstab: /zfssa/exdb/backup7 - added
/etc/vfstab: /zfssa/exdb/backup2 - added
/etc/vfstab: /zfssa/exdb/backup4 - added
/etc/vfstab: /zfssa/exdb/backup6 - added
/etc/vfstab: /zfssa/exdb/backup8 - added

Creating /etc/init.d/zfssa_exdb - aiesscdbadm02
start service - zfssa_exdb
/etc/init.d/zfssa_exdb - created
Older version /u01/app/oracle/product/11.2.0.3/dbhome_1/dbs/oranfstab is saved to
/u01/app/oracle/product/11.2.0.3/dbhome_1/dbs/oranfstab_20140507T080735

Updating oranfstab - aiesscdbadm02
/bin/chown oracle:oinstall /u01/app/oracle/product/11.2.0.3/dbhome_1/dbs/oranfstab -
aiesscdbadm02
/bin/chmod 640 /u01/app/oracle/product/11.2.0.3/dbhome_1/dbs/oranfstab - aiesscdbadm02
/usr/bin/make -f /u01/app/oracle/product/11.2.0.3/dbhome_1/rdbms/lib/ins_rdbms.mk
dnfs_on ORACLE_HOME=/u01/app/oracle/product/11.2.0.3/dbhome_1 - aiesscdbadm02
dNFS is enabled - aiesscdbadm02

/bin/chown oracle:oinstall /zfssa/exdb/backup1
/bin/chmod 750 /zfssa/exdb/backup1
/bin/chown oracle:oinstall /zfssa/exdb/backup2
/bin/chmod 750 /zfssa/exdb/backup2
/bin/chown oracle:oinstall /zfssa/exdb/backup3
/bin/chmod 750 /zfssa/exdb/backup3
/bin/chown oracle:oinstall /zfssa/exdb/backup4
/bin/chmod 750 /zfssa/exdb/backup4
/bin/chown oracle:oinstall /zfssa/exdb/backup5
/bin/chmod 750 /zfssa/exdb/backup5
/bin/chown oracle:oinstall /zfssa/exdb/backup6
/bin/chmod 750 /zfssa/exdb/backup6
/bin/chown oracle:oinstall /zfssa/exdb/backup7
/bin/chmod 750 /zfssa/exdb/backup7
/bin/chown oracle:oinstall /zfssa/exdb/backup8
/bin/chmod 750 /zfssa/exdb/backup8

Adding backup services ...
/u01/app/oracle/product/11.2.0.3/dbhome_1/bin/srvctl add service -d exdb -r exdb1 -a
exdb2 -s exdb_bkup1
/u01/app/oracle/product/11.2.0.3/dbhome_1/bin/srvctl start service -d exdb -s
exdb_bkup1
/u01/app/oracle/product/11.2.0.3/dbhome_1/bin/srvctl add service -d exdb -r exdb2 -a
exdb1 -s exdb_bkup2
```

```
/u01/app/oracle/product/11.2.0.3/dbhome_1/bin/srvctl start service -d exdb -s
exdb_bkup2

About to restart the database exdb
This step will interrupt the database. You may skip this step
and restart the database manually later.
Continue to restart the database? (y|n) y


Stopping database ...
Database stopped successfully - exdb
Starting database ...
Database started successfully - exdb

Stopping services ...
Services stopped successfully - exdb
Starting services ...
Services started successfully - exdb

Step 2 Configure SuperCluster Database Nodes
Completed Successfully.

Continue to Step 3 Generate RMAN Run Block Scripts? (y|n) y


Step 3 Generate RMAN Run Block Scripts

In this step the RMAN run block scripts will be generated. Prior to
this step the Oracle ZFS Storage Appliance and the Engineered System
database nodes must be configured by following the Step 1 and 2
respectively.

To complete this step successfully, the following information will
be required:

- Database login string (user/password@scan_address) or scan address

Continue? (y|n) y

Enter the login string (user/password@scan_address) or scan address of exdb:
Example: sys/oracle@testdb-scan or testdb-scan
sys/welcome1@aiessc-scan

Verify the login string for exdb:

Database login:      sys/welcome1@aiessc-scan

Correct? (y|n) y


The following RMAN run block scripts are generated:

    /root/esbu/rman/
        exdb_backupset_L0 - level 0 backupset
        exdb_backupset_L1 - level 1 backupset
        exdb_image_copy - level 0/1 image copy
        exdb_inc_merge - incremental merge
        exdb_restore_validate_backupset - restore validate for backupset
        exdb_restore_validate_image_copy - restore validate for image copy

The RMAN run block scripts are intended for samples only. Please review
and modify the scripts as necessary for your requirements.

Step 3 Generate RMAN Run Block Scripts
Completed Successfully.
```

```
This session is saved in /root/esbu/template/EXDBPROJECT

To continue next time, run
esbu -t /root/esbu/template/EXDBPROJECT

root@aiesscdbadm01:~/esbu#
```

Appendix B: 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
- Sun Datacenter InfiniBand Switch 36 Firmware Version 2.0
  http://docs.oracle.com/cd/E26698_01/index.html
- Oracle SPARC SuperCluster
  http://www.oracle.com/technetwork/server-storage/sun-sparc-enterprise/sparc-supercluster-191641.html

Configuring an Oracle ZFS Storage ZS3-BA with an Oracle SuperCluster for Oracle Database Backup and Recovery
June 2015, Version 2.1
Author: Eric Polednik

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

Hardware and Software, Engineered to Work Together