Oracle® Exalogic Elastic Cloud

Backup and Recovery Guide

Release EL X2-2 and X3-2

E40226-05

April 2014

This document describes how to manually back up and recover Exalogic components.

ORACLE®

Oracle Exalogic Elastic Cloud Backup and Recovery Guide, Release EL X2-2 and X3-2

E40226-05

Primary Author: Bharath K Reddy

Contributing Authors: Scott Balfour, Neeraj Gupta, Rachid Benkreira, Jeremy Hoyland

Contributors: Kumar Dhanagopal, Ashish Thomas

Contents

Preface

This document covers backing up and restoring the configuration of the Exalogic infrastructure components, backing up and restoring the management components (the Exalogic Control Stack), and the repositories of the cloud infrastructure when the Exalogic machine is deployed in a virtual configuration.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

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

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1 Overview and Scope

Oracle Exalogic is an integrated hardware and software system designed to provide a complete platform for a wide range of application types and a wide variety of workloads. Exalogic is intended for large-scale, performance-sensitive, mission-critical application deployments.

This chapter contains the following sections:

1.1 Overview of Backup and Recovery Concepts

Every Exalogic system includes some redundant components, to ensure that the failure of a single component does not affect the overall availability of the system. However, redundancy within a system does not provide sufficient protection in the following situations:

- Full Exalogic system failure due to site outage
- Incorrect data or configuration due to user error while making updates
- Data corruption

Disaster recovery (DR) technologies can be useful in some circumstances. For example, one way to safeguard against a full site outage would be to replicate an Exalogic system at another physical location. The replicated system would need to be kept up-to-date with changes made on the primary system. Such a system can provide hot-standby capability in the event of a failure of the primary system.

While hot-standby systems are extremely useful for ensuring business continuity, they are expensive to maintain and often require additional infrastructure. Further, they do not provide a convenient and optimal solution in situations where only certain segments of data need to be restored. For instance, if user error or data corruption leads to errors in certain files, these errors would also reflect on the hot-standby system. In such situations, a previous version of the files can be recovered by using a backup and restore functionality.

This document describes the backup and recovery solution for the Oracle Exalogic machine, which builds upon well-established solutions for Oracle Database and Oracle Fusion Middleware. This document also describes how to leverage the built-in features of the ZFS Storage Appliance for disk-based backups.

1.2 Scope

This document covers backing up and restoring the configuration of the Exalogic infrastructure components, and backing up and restoring the management components (the Exalogic Control Stack) of the cloud infrastructure when the Exalogic machine is deployed in a virtual configuration.

The Exalogic Control Stack consists of the Oracle VM Manager, the Oracle Enterprise Manager Ops Center and their repositories deployed to an Oracle Database.

The repository containing the disk images—for Exalogic Control vServers and for the application vServers—is located on the local ZFS Storage Appliance of the Exalogic machine.

1.3 Supported Platforms

This document is intended for use with the following platforms:

- Physical (Linux): Exalogic Elastic Cloud Software 2.0.3.x.x and earlier
- Physical (Solaris): Exalogic Elastic Cloud Software 2.0.4.x.x and earlier
- Virtual (Oracle VM): Exalogic Elastic Cloud Software 2.0.4.x.x and earlier

> **Note:** For all later releases, Oracle recommends using the ExaBR tool to back up and recover Exalogic data. For more information, see the *Backup and Recovery Guide Using ExaBR*.

1.4 Backup and Recovery Recommendations

This section provides general backup and recovery recommendations for Exalogic.

- For additional protection, it is strongly recommended that the backups be stored on an external storage appliance or on tape, depending on your current backup practices.
- It is strongly recommended that regular backups of your data be scheduled. The backup schedule should be based on the nature of your data.
- Schedule daily backups to the local ZFS storage appliance and weekly backups to the external ZFS Storage appliance.
- It is recommended that a *gold copy* be created for the operating system of the compute nodes in your environment.
- It is recommended that the components of the Exalogic Control stack and their repositories be backed up after lifecycle operations, such as adding and removing accounts, users, and vServers.

For more information, see the *Exalogic Backup and Recovery Best Practices White Paper* at `http://www.oracle.com/technetwork/database/features/availability/maa-exalogic-br-1529241.pdf`.

2

Backup and Recovery Locations

Oracle recommends creating all the initial backups of all the components, excluding the Exalogic Repository, on the local ZFS storage appliance. These backups can then be moved to an external storage device or to tape depending on existing backup policies.

To create a project and shares on the local ZFS storage, do the following:

1. Log in to the BUI of the storage appliance as the `root` user

   `https://storage_appliance_IP:215`

2. Under Shares click **Projects** to bring up the Projects page.
3. Click the plus button next to Projects.
4. Enter the name of the new project.

   Example: Exalogic_Backup.

5. From the list of projects, select the **Exalogic_Backup** project.
6. Navigate to the general properties page and update the mount point.

   Example: `/export/Exalogic_Backup`.

7. Click **Apply**.
8. Navigate to the Protocols properties page and add NFS exceptions by clicking the plus next to **NFS Exceptions**.

   Enter the following information:

Type	Entry	Access Mode	Charset	Root Access
Network	CIDR for the IPoIB default	Read/Write	Default	Yes
Network	CIDR for the IPoIB storage partition	Read/Write	Default	Yes

9. Click **Apply** to save the settings.
10. Navigate to the Shares page to create shares for each component. This is the location for the initial backup of each component.
11. Create the following shares by clicking plus symbol next to file system and providing a name for the share.

Share Name	Mount Point	Stores
`compute_nodes`	`/export/Exalogic_Backup/compute_nodes`	OS backups of compute nodes
`ib_gw_switches`	`/export/Exalogic_Backup/ib_gw_switches`	Configuration backups of IB gateway switches

Share Name	Mount Point	Stores
ib_spine_switches	/export/Exalogic_Backup/ib_spine_switches	Configuration backups of IB spine switches
management_switches	/export/Exalogic_Backup/management_switches	Cisco switch backup
control_metadata	/export/Exalogic_Backup/control_metadata	Metadata backups of the Exalogic control stack
cust_vservers	/export/Exalogic_Backup/cust_vservers	Backups of customer vServers
control_vservers	/export/Exalogic_Backup/control_vservers	Backups of Exalogic Control vServers
ecu	/export/Exalogic_Backup/ecu	ECU backup files

3

Backup and Recovery of Hardware Components

This chapter describes how to back up and recover the components of the Exalogic infrastructure.

> **Note:** Prior to recovering a component, ensure that the component is not in use.

It contains the following sections:

- Section 3.1, "Exalogic Configuration Utility"
- Section 3.2, "Exalogic Compute Nodes"
- Section 3.3, "InfiniBand Switches"
- Section 3.4, "Cisco Management Switch"
- Section 3.5, "ZFS Storage Heads"

3.1 Exalogic Configuration Utility

The Exalogic Configuration Utility (ECU) is used to configure an Exalogic machine during initial deployment. It is strongly recommended that you back up the configuration and the runtime files generated by the ECU after the initial deployment is complete.

1. Mount the NFS location defined in Chapter 2, "Backup and Recovery Locations" on the master compute node.

   The master compute node is the node in the Exalogic rack on which the ECU was run.

2. Create `tar`balls of the following directories:

   - `/opt/exalogic/ecu`: Exalogic configuration directory
   - `/var/tmp/exalogic/ecu`: Exalogic runtime directory
   - (optional) `/var/log/exalogic/ecu`: Contains ECU log files

3. To recover the ECU files, extract the `tar`ball containing the ECU configuration files to the `/opt/exalogic/ecu` directory, and extract the `tar`ball containing the runtime files to the `/var/tmp/exalogic/ecu` directory. There is no need to restore the log files.

3.2 Exalogic Compute Nodes

This section contains the following subsections:

- Section 3.2.1, "Backing Up Exalogic Compute Nodes"
- Section 3.2.2, "Reimaging and Bare Metal Restore"
- Section 3.2.3, "Recovering Exalogic Compute Nodes in a Virtual Environment"

3.2.1 Backing Up Exalogic Compute Nodes

Backing up the compute node consists of backing up the following:

- The ILOM of the compute node
- The operating system of the compute node

3.2.1.1 Backing Up the ILOM of a Compute Node

To back up the ILOM of a compute node, do the following:

> **Note:** You cannot back up the ILOM of the compute node you are using to back up and restore components. To back up the ILOM of that compute node, run the steps from a different compute node.

1. Mount the NFS location (Chapter 2, "Backup and Recovery Locations") on one of the compute nodes.
2. Log in to the compute node as the `ilom-admin` user.
3. Encode the backup by running the following command:

   ```
   set /SP/config passphrase=phrase
   ```

   Example:

   ```
   set /SP/config passphrase=mypassword1
   set 'passphrase to 'mypassword1'
   ```

   `mypassword1` is the password chosen by the user. Provide the password used when creating the backup.
4. Back up the configuration of the ILOM by running the following command:

   ```
   set /SP/config dump_uri=URI
   ```

   `URI` is the command used to perform the backup.

   Example:

   ```
   set /SP/config dump_uri=scp://root:rootpwd@hostIP/export/Exalogic_
   Backup/compute_nodes/computenode.backup
   ```

   `hostIP` is the IP address of the target host for the backup file.

   `/export/Exalogic_Backup/compute_nodes/computenode.backup` is the absolute path and the name of the backup file on the remote host.

3.2.1.2 Backing Up the Operating System of a Compute Node

The operating system of an Exalogic machine is installed on the local disk of each compute node.

If the official Exalogic base image was customized, it is recommended that you create a backup of the root file system and the customizations by using standard operating system utilities: `tar`, `dump` and so on, while excluding the `/var`, `/tmp`, `/tree`, `/proc`, `/dev`, `/poolfsmnt`, and the NFS mounted file systems. If you are running the Exalogic virtual stack, then, in addition, you should exclude the `poolfs`, `ExalogicPool`, `ExalogicRepo` file systems. The `ExalogicPool` and the `ExalogicRepo` file systems are mounted over NFS.

Save the backup to the NFS location you created for the compute nodes, as described in Chapter 2, "Backup and Recovery Locations" (for example, `/export/Exalogic_Backup/compute_nodes`).

Note: Run the following command to list the NFS file systems mounted on your compute node.

```
mount -t nfs | awk '{print $3}'
```

For more information, see the *Exalogic Backup and Recovery Best Practices White Paper* at `http://www.oracle.com/technetwork/database/features/availability/maa-exalogic-br-1529241.pdf`.

3.2.2 Reimaging and Bare Metal Restore

A compute node should be reimaged when it has been irretrievably damaged, or multiple disk failures cause local disk failure with no existing backup for the compute node. During the reimaging procedure, the other compute nodes in the Exalogic machine are available. You should restore any scripting, CRON jobs, maintenance actions, and other customizations performed on top of the Exalogic base image.

Bare-metal restore is the process of restoring a new compute node to the same state as one on which a backup was taken. To perform a bare-metal restore, the new compute node must be reimaged by performing the steps described in this section. After the node is reimaged with the Exalogic base image, it should be restored to its original state by using a previously taken backup.

When an Exalogic machine is deployed in either a physical or virtual configuration, do the following to reimage the compute node. The procedure in Section 3.2.3, "Recovering Exalogic Compute Nodes in a Virtual Environment" must be performed when the Exalogic machine is deployed in a virtual configuration.

1. Open an Oracle support request with Oracle Support Services.

   The support engineer will identify the failed server and send a replacement. Provide the support engineer the output of the `imagehistory` and `imageinfo` commands run from a surviving compute node. This output provides the details about the correct image and the patch sets that were used to image and patch the original compute node, and it provides a means to restore the system to the same level.

2. Restore the ILOM of the compute node.

   Note: You cannot restore the ILOM of the compute node you are using to back up and restore components. To restore the ILOM of that compute node, run the steps from a different compute node.

   To restore the ILOM of the compute node, do the following:

a. Mount the NFS location (Chapter 2, "Backup and Recovery Locations") on one of the compute nodes.

b. Log in to the repaired compute node as the `ilom-admin` user.

c. Encode the backup by running the following command:

```
set /SP/config passphrase=phrase
```

Example:

```
set /SP/config passphrase=mypassword1
set 'passphrase to 'mypassword1'
```

`mypassword1` is the password chosen by the user. Provide the password used when creating the backup.

d. Restore the configuration of the ILOM by running the following command:

```
set /SP/config load_uri=URI
```

`URI` is the command used to perform the backup.

Example:

```
set /SP/config load_uri=scp://root:rootpwd@hostIP/export/Exalogic_
Backup/compute_nodes/computenode.backup
```

`hostIP` is the IP address of the target host for the backup file.

`/export/Exalogic_Backup/compute_nodes/computenode.backup` is the absolute path and the name of the backup file on the remote host.

3. Download the Oracle Exalogic base image and patch-set updates (PSUs).

   Download the appropriate Oracle Exalogic base image from `https://edelivery.oracle.com` and the appropriate PSU from My Oracle Support `https://support.oracle.com`.

4. Image the replacement compute node.

   The compute node being replaced can be imaged using a PXE boot server or through the web-based ILOM of the compute node. This document does not cover the steps to configure a PXE boot server, however it provides the steps to enable the compute node to use a PXE boot server.

   If a PXE boot server is being used to re-image the compute node, log in to the ILOM of the compute node through SSH, set the `boot_device` to `pxe` and reboot the compute node.

   If the web-based ILOM is being used instead, ensure that the image downloaded earlier is on the local disk of the host from which the web-based ILOM interface is being launched and then do the following:

   a. Open a web browser and bring up the ILOM of the compute node, such as http://host-ilom.mycompany.com/

   b. Log in to the ILOM as the `root` user.

   c. Navigate to Redirection under the Remote Control tab, and click the **Launch Remote Console** button. The remote console window is displayed.

Note: Do not close this window until the entire imaging process is completed. You will need to return to this window to complete the network configuration at the end of the imaging process.

d. In the remote console window, click on the Devices menu item and select:

- Keyboard (selected by default)
- Mouse (selected by default)
- CD-ROM Image

In the new dialog box that is displayed, select the Linux base image `iso` file that you downloaded.

e. On the ILOM window, navigate to the Host Control tab under the Remote Control tab.

f. Select CDROM from the drop-down list and then click **Save**.

g. Navigate to the Remote Power Control tab in the Remote Control tab.

h. Select **Power Cycle** from the drop-down list, and then click **Save**.

i. Click **OK** to confirm that you want to power cycle the machine.

This starts the imaging of the compute node. Once the imaging is complete, the first boot scripts prompt the user to provide the network configuration.

5. Configure the replacement compute node.

- If you have a valid backup, restore the `/etc` directory and customizations you made, if any, to the replacement compute node.
- If you do not have a valid backup, configure the replacement compute node with the appropriate DNS, time zone, and NTP settings. These settings should be the same on all the compute nodes in the Exalogic machine.

Note: If the compute node being replaced is the master node of the Exalogic machine, restore the ECU configuration that was backed up earlier, as described in Section 3.1, "Exalogic Configuration Utility.". The master node in an Exalogic machine is the node on which the Exalogic Configuration Utility (ECU) is run.

6. If the Exalogic machine was deployed in a physical configuration, you may have to update the VNIC configuration on the IB switches with the new IB port GUIDs.

To validate the existing VNIC configuration on both the IB switches attached to the compute node, do the following:

a. Get the port GUIDs of the replacement compute node by running the `ibstat` command on the compute node.

b. Log in to the IB switches attached to the compute node and run the `showvnics` command to view the VNICs created on the switch.

c. For the VNICs associated with the replacement compute node, verify that the port GUIDs are displayed in the output of the `ibstat` command that you ran in step a.

3.2.3 Recovering Exalogic Compute Nodes in a Virtual Environment

> **Note:** In an Exalogic virtual configuration, do not attempt to **replace** a failed compute node with an entirely new one. Contact Oracle Support for the procedure to perform such a replacement. An improperly replaced component might not be discovered correctly by Exalogic Control. You **can** use the procedures described in this document to restore a failed component after **repairing** it.

When the Exalogic machine is deployed in a virtual configuration, do the following to replace the compute node:

1. If the first compute node is down, you must migrate Oracle VM Manager and the control database vServer to a different compute node. To migrate them, do the following:
   a. Migrate the Database vServer by performing the steps in Section 4.1.1 on a running compute node.
   b. Migrate the Oracle VM Manager vServer by performing the steps in Section 4.1.2 on a running compute node
   c. Stop the components of the Exalogic Control stack. For more information, see Section 5.2.1, "Stopping the Components of the Exalogic Control Stack."
   d. Start the components of the Exalogic Control stack. For more information, see Section 5.2.5, "Starting the Exalogic Control Stack."
2. Migrate the virtual machines running on the compute node:
   a. Log in to the Oracle VM Manager BUI.
   b. Navigate to **Home**, and then to **Server Pools**.
   c. Select and expand the server pool to list the compute nodes in the pool.
   d. Select and expand the compute node to list the virtual machines running on the selected node.
   e. Migrate the virtual machines, one at a time, by doing the following:
      i. Select the virtual machine to be migrated.
      ii. Select **Migrate** under Actions to bring up the Migration Assistant.
      iii. Select the **Unassigned Virtual Machine** folder.
      iv. Select **OK**.
3. Remove the compute node from the Oracle VM server pool:
   a. Log in to Exalogic Control as the `root` user.
   b. Expand **Assets** in the navigation pane on left.
   c. Under Servers, expand the compute node that is being replaced.
   d. Select the Oracle VM Server asset.
   e. Click **Remove from Server Pool** in the Actions pane on the right.

Note: The job to remove the compute node may fail. If it does fail, examine the job in the jobs pane. The job consists of the following tasks:

- RemoveOvmServerFromPool
- OvmRefreshDomainModelTask

If the first tasks succeeds, the failure of the second task can be ignored.

f. Verify that the node has been removed from the pool by logging in to Oracle VM Manager.

g. Delete the now unassigned compute node from Oracle VM Manager.

4. Remove the compute node from the assets:

   a. Log into Exalogic Control as the `root` user.

   b. Expand **Assets** in the navigation pane on left.

   c. Expand **Servers** to list all the compute nodes.

   d. Select and expand the compute node that is being replaced.

   e. Select the operating system and place it in maintenance mode, by clicking **Place in Maintenance Mode** in the Actions pane.

   f. Select the server and place it in maintenance mode, by clicking **Place in Maintenance Mode** in the Actions pane.

   g. Delete the operating system by clicking **Delete Asset** in the Actions pane.

   h. Delete the server by clicking **Delete Asset** in the Actions pane.

5. Replace the failed compute node by following the standard replacement process.

6. Perform the steps in Section 3.2.2, "Reimaging and Bare Metal Restore" of this document to re-image the compute node and restore the previous configuration from a backup.

7. Considering that the replacement server has the same IP address but a different MAC address, you may need to flush the ARP cache of cn01 (ECU master node).

   Ping the new node from ECU master node.

   If the ping fails, but connectivity is good, flush the ARP cache on cn01. You may have to wait for some time for the cache on the Cisco switch to be cleared.

   - To look at the cache, run `arp -n`
   - To flush the cache, run `ip -s neigh flush all`

8. After the node is reimaged, log in to the compute node as `root`, and set the `ovs-agent` password for the `oracle` user:

   ```
   ovs-agent-passwd oracle password
   ```

   Note: For information on the default password, contact Oracle Support.

9. On the master compute node, go to the `/opt/exalogic/ecu` directory, and set the `ECU_HOME` environment variable, as follows:

```
export ECU_HOME=/opt/exalogic/ecu
```

Run `cd $ECU_HOME` to verify whether the `ECU_HOME` environment variable is set correctly.

10. Set up password-less SSH over IP and IB by running the `/opt/exalogic.tools/tools/setup-ssh.sh` script as follows:

```
./setup-ssh.sh -H IP-of-xenbr0-on-replaced-node
./setup-ssh.sh -H IP-of-bond1-on-replaced-node
```

11. Identify the GUIDs of the IB ports of the failed compute node.

    The GUIDs of the IB ports of the failed compute node are located in the ECU log files.

    a. Log in to the master compute node, and go to the `/var/tmp/exalogic/ecu/cnodes` directory.

       The IB port GUIDs of all the compute nodes in the machine are stored in files named `ibstat.node.`*`NodeIndex`*, where `NodeIndex` is the compute node number (1–30).

    b. Using a text editor, open the `ibstat.node.`*`NodeIndex`* file corresponding to the failed compute node.

       For example, if compute node 15 is the failed node, open `ibstat.node.15`, as shown in the following example:

```
root@exlcn15 cnodes]# cat ibstat.node.15
   CA 'mlx4_0'
           CA type: MT26428
           Number of ports: 2
           Firmware version: 2.9.1000
           Hardware version: b0
           Node GUID: 0x0021280001a122a8
           System image GUID: 0x0021280001a122ab
           Port 1:
                   State: Active
                   Physical state: LinkUp
                   Rate: 40
                   Base lid: 158
                   LMC: 0
                   SM lid: 95
                   Capability mask: 0x02510868
                   Port GUID: 0x0021280001a122a9
                   Link layer: IB
           Port 2:
                   State: Active
                   Physical state: LinkUp
                   Rate: 40
                   Base lid: 159
                   LMC: 0
                   SM lid: 95
                   Capability mask: 0x02510868
                   Port GUID: 0x0021280001a122aa
                  Link layer: IB
```

    c. Note the IB port GUIDs for the compute node.

       They are indicated for each port with the keyword `Port GUID`,

In this example, the GUID for IB port 1 is `0x0021280001a122a9` and the GUID for port 2 is `0x0021280001a122aa`.

12. Configure the networks and IB partitions on the compute node by running the Exalogic Configuration Utility (ECU):

    Note: If this node is the master node in the Exalogic rack—that is, the node from which the ECU was run initially, then, before this step, restore the ECU configuration files, run time files, and log files by performing the steps in Section 3.1.

    a. Discover the switches.

    ```
    ./ecu.sh ib_switches discover
    ```

    b. Apply the configuration to the new compute node.

    Note: Before you run the following command, verify that the `/var/tmp/ecu/cnodes_current.json` file contains the current IP addresses on the `eth-admin` and `IPoIB-default` interfaces of the node being replaced.

    ```
    ./ecu.sh apply_cnode_config node_index
    ```

    `node_index` is the compute node number in the rack.

    c. Reboot the node.

    ```
    ./ecu.sh reboot_cnode current node_index
    ```

    d. Test that the IP addresses are as expected.

    ```
    ./ecu.sh test_cnode_network target cnode_number
    ```

    The following part of the output will indicate the interfaces and IP addresses required for the next steps:

    ```
    Network IP Ping Status
    ------------------------------------------------------
    ILOM 10.196.17.152 OK
    eth-admin 10.196.17.122 OK
    IPoIB-default 192.168.17.122 OK
    IPoIB-admin 192.168.30.2 OK
    IPoIB-storage 192.168.31.2 OK
    IPoIB-virt-admin 172.36.0.2 OK
    IPoIB-ovm-mgmt 192.168.33.2 OK
    IPoIB-vserver-shared-storage 172.37.0.2 OK

    INFO:netutils:Ping to all IP addresses succeeded
    ```

13. For the virtual-machine networks to be plumbed correctly, the customer EoIB and private vNet IB partitions must be updated with the IB port GUIDs of the replacement compute node.

    a. Log in to the replacement compute node.

    b. Identify the GUIDs of the IB ports of the replacement compute node, by running the `ibstat` command, as shown in the following example:

```
root@exlcn15 ~]# ibstat
  CA 'mlx4_0'
          CA type: MT26428
          Number of ports: 2
          Firmware version: 2.9.1000
          Hardware version: b0
          Node GUID: 0x0021280001eface6
          System image GUID: 0x0021280001eface9
          Port 1:
                  State: Active
                  Physical state: LinkUp
                  Rate: 40
                  Base lid: 98
                  LMC: 0
                  SM lid: 1
                  Capability mask: 0x02510868
                  Port GUID: 0x0021280001eface7
                  Link layer: IB
          Port 2:
                  State: Active
                  Physical state: LinkUp
                  Rate: 40
                  Base lid: 99
                  LMC: 0
                  SM lid: 1
                  Capability mask: 0x02510868
                  Port GUID: 0x0021280001eface8
                Link layer: IB
```

The GUIDs for each port are indicated by the `Port GUID`.

In this example, the GUID for IB port 1 is `0x0021280001a122a9` and the GUID for port 2 is `0x0021280001eface8`.

c. Log in to the IB switch running the master subnet manager and run `smpartition start`.

This command creates a temporary file `partitions.conf.tmp` in the `/conf` directory. This file can be updated using regular Linux commands.

d. In the `/conf/partitions.conf.tmp` file, replace the failed compute node's IB port GUIDs, which you identified in step 11, with the GUIDs of the replacement node, as determined in step 13.b.

You can do this by using a text editor, or by using the `sed` command, as shown in the following example:

```
sed 's/0x0021280001a122a9/0x0021280001a122a9/g' /conf/partitions.conf.tmp
```

e. Propagate the configuration to all the IB switches in the fabric by running `smpartition commit`.

14. Update the credentials for the ILOM and the compute node:

a. Log in to the Exalogic Control BUI.

b. Navigate to **Credentials** in the **Plan Management** section.

c. Enter the host name of the compute node in the search box, and click **Search**.

The IPMI and SSH credential entries for the ILOM and the compute node are displayed.

d. To update all four credentials, do the following:

i. Select the entry for the credentials and click **Edit**. The Update Credentials dialog box is displayed.

ii. Update the password and confirm the password fields.

iii. Click **Update**.

15. Rediscover and add the asset:
    a. Log in to the Exalogic Control BUI.
    b. In the navigation pane on the left, expand **Plan Management**, and under **Profiles and Policies**, expand **Discovery**.
    c. Select the appropriate **Server OS @ *host*** discovery profile.
    d. In the Actions pane on the right, click **Add Assets**.
    e. On the resulting screen, verify whether the correct discovery profile is displayed.
    f. Click **Add Now**.
    g. Wait until the discovery process succeeds.
    h. Select the appropriate **Server ILOM @ *host*** discovery profile.
    i. In the Actions pane on the right, click **Add Assets**.
    j. On the resulting screen, verify whether the correct discovery profile is displayed.
    k. Click **Add Now**.
    l. Wait until the discovery process succeeds.
    m. In the left navigation pane, expand the **Assets** section to display all the assets.
    n. Verify whether the replaced server is displayed in the Assets section and positioned correctly in the photo-realistic view.
16. Log in to Oracle VM Manager using the `admin` user credentials and discover the new compute node. Use the IP address of the `IPoIB-ovm-mgmt` partition.

    Note: The default partition key and network CIDR for the IPoIB-ovm-mgmt-partition are `0x8004` and `192.168.23.0/24` respectively.

    This information is also available in the `/opt/exalogic/ecu/config/cnode_ipoib_networks.json` file on the master compute node.

17. After the new compute node is discovered, ensure that it is added to the required pool.

    If the compute node is in the unassigned-server group, you should add it manually in Oracle VM Manager.

    a. In the Hardware tab of the left navigation pane, expand **Resources**, and right-click on the name of the server pool to which you want to add the compute node.
    b. From the resulting context menu, select **Add/Remove Servers**.

       The Add/Remove Servers from the Server Pool dialog box is displayed.

c. Select the server that you want to add from the **Available Servers** list and move it to the **Selected Servers** list.

d. Click **OK**.

18. Refresh the repository in Oracle VM Manager by doing the following:

a. Log in to the Oracle VM Manager console.

b. Click **Home** under the View menu.

c. Select **Servers Pools** in the left pane.

d. Select **Repositories** in the right pane.

e. Click the **Refresh Repositories** icon. This is the icon with curved blue arrows.

19. Present the repository to the compute nodes in Oracle VM Manager:

a. Log in to the Oracle VM Manager console.

b. Click on **Home** under the View menu.

c. Select **Servers Pools** in the left pane.

d. Select **Repositories** in the right pane.

e. Select the entry with a forward slash under the Repositories table.

f. Click the **Present-Unpresent Selected Repository** icon. This is the icon with green up and down arrows.

g. In the **Present this Repository to Server(s)** dialog box, select the compute nodes listed under the **Servers** column and move them to the **Present to Server(s)** column.

h. Click **OK**.

i. Verify whether the compute node has been added by monitoring the Oracle VM Manager job.

20. Add the compute node as an admin server to the repository.

a. Log in to the Oracle VM Manager console.

b. Click on **Hardware** under the View menu.

c. Select the **Storage** tab in the left pane.

d. Expand **File Servers**.

e. Expand **Generic Network File System**.

f. Select the **Generic Network File System**, and then select **Add/Remove Admin Servers** from the menu.

g. In the **Present this Repository to Server(s)** dialog box, select the compute nodes listed under the **Servers** column and move them to the **Present to Server(s)** column.

h. Click **OK**.

i. Verify whether the compute node has been added by monitoring the Oracle VM Manager job.

21. Log in to the Exalogic Control BUI as the `Cloud Admin` user, and the start the virtual machines that were migrated to the **Unassigned Virtual Machine** folder in the Oracle VM Manager BUI in step 2. The virtual machines will be started up on the replacement compute node.

Note:

- If you migrated the virtual machines to another compute node within the pool, use Oracle VM Manager to migrate them back to the replaced compute node. Follow the instructions in step 2.e, but instead of selecting the **Unassigned Virtual Machine** folder, select the replacement compute node.
- Currently, migrating virtual machines between pools is not supported.

3.3 InfiniBand Switches

The InfiniBand switches are a core part of an Exalogic machine and the configurations of all the Infiniband switches must be backed up regularly. The configuration backups of the Service Processor can be created either using the ILOM BUI or CLI.

This section contains the following subsections:

3.3.1 Backing Up the InfiniBand switches

Save the IB switch backups to the NFS locations you created for the IB switches as described in Chapter 2, "Backup and Recovery Locations" (for example, `/export/Exalogic_Backup/ib_gw_switches` and `/export/Exalogic_Backup/ib_spine_switches`). Backups must be created for all the switches in the fabric. Create separate directories under the NFS share for each switch in the fabric.

To back up the Service Processor configuration of an IB switch by using the ILOM CLI, do the following:

1. Mount the NFS location (Chapter 2, "Backup and Recovery Locations") on one of the compute nodes.
2. Log in to the InfiniBand switch as the `ilom-admin` user.
3. Encode the backup by running the following command:

   ```
   set /SP/config passphrase=phrase
   ```

   Example:

   ```
   set /SP/config passphrase=mypassword1
   set 'passphrase to 'mypassword1'
   ```

   `mypassword1` is the password chosen by the user. Provide the password used when creating the backup.
4. Back up the configuration of all the Infiniband switches by running the following command:

   ```
   set /SP/config dump_uri=URI
   ```

   `URI` is the command used to perform the backup.

   Example:

```
set /SP/config dump_uri=scp://root:rootpwd@hostIP/export/Exalogic_Backup/ib_
type_switches/switch.backup
```

`hostIP` is the IP address of the target host for the backup file.

`type` is either `gw` or `spine` depending on the type of IB switch.

`/export/Exalogic_Backup/ib_type_switches/switch.backup` is the absolute path and the name of the backup file on the remote host.

5. To back up the user settings of the InfiniBand switch, manually back up the `/etc/opensm/opensm.conf` file to the same location you backed up Service Processor configuration.
6. To back up the partitions of the InfiniBand switch, manually back up the `/conf/partitions.current` file to the same location you backed up Service Processor configuration.

   After the files are transferred to the NFS location, they can be backed up to more permanent storage as part of the operating system backup.

3.3.2 Recovering the InfiniBand Switches in a Physical Environment

Ensure that no configuration changes are being made while performing the restore. Configuration changes include vServer creation and vNet creation.

Note: During the restore, there will be a temporary disruption of traffic.

Restore the configuration of an IB switch through the ILOM CLI, by doing the following:

1. Mount the NFS location (Chapter 2, "Backup and Recovery Locations") on one of the compute nodes.
2. Log in to the Infiniband switch as the `ilom-admin` user.
3. Encode the backup by running the following command:

   ```
   set /SP/config passphrase=phrase
   ```

   Example:

   ```
   set /SP/config passphrase=mypassword1
   set 'passphrase to 'mypassword1'
   ```

   `mypassword1` is the password chosen by the user. Provide the password used when creating the backup.
4. Restore the configuration of all the Infiniband switches by doing the following:
   a. Run the following command:

      ```
      set /SP/config load_uri=URI
      ```

      `URI` is the command used to perform the backup.

      Example:

      ```
      set /SP/config load_uri=scp://root:rootpwd@hostIP/export/Exalogic_
      Backup/ib_type_switches/switch.backup
      ```

      `hostIP` is the IP address of the target host for the backup file.

`type` is either `gw` or `spine` depending on the type of IB switch.

`/export/Exalogic_Backup/ib_type_switches/switch.backup` is the absolute path and the name of the backup file on the remote host.

b. If the failed switch is replaced with a new one, as opposed to being repaired and reinstalled, add the GUIDs of the BridgeX ports to the EoIB partitions on the switch.

Identify the GUIDs of the BridgeX ports by running `showgwports` on the switch:

```
showgwports
INTERNAL PORTS:
---------------
Device Port Portname PeerPort PortGUID LID IBState GWState
-----------------------------------------------------------------------
Bridge-0 1 Bridge-0-1 4 0x002128f4832ec001 0x0007 Active Up
Bridge-0 2 Bridge-0-2 3 0x002128f4832ec002 0x0006 Active Up
Bridge-1 1 Bridge-1-1 2 0x002128f4832ec041 0x000a Active Up
Bridge-1 2 Bridge-1-2 1 0x002128f4832ec042 0x000e Active Up
```

Log in to the switch running the master subnet manager and add the Bridgex ports as full members to each of the EoIB partitions by running the following command:

```
smpartition add -pkey PKEY -port BridgeXGUID -m full
```

Example:

```
smpartition add -pkey 0x8006 -port 0x002128f4832ec002 -m full
```

c. Restore the user settings that you backed up earlier as described in Section 3.3.1, "Backing Up the InfiniBand switches."

d. Restore the partitions that you backed up earlier as described in Section 3.3.1, "Backing Up the InfiniBand switches." Before restoring the partitions, review the current partitions and the `partitions.current` file.

Restore the `partitions.current` file to the master switch and propagate the subnet manager configuration, by running the following commands:

Note: To ensure that the new switch joins the IB fabric, enable the subnet manager by running the `enablesm` command on both switches. However, only one of the gateway switches should be set as the master sm.

```
smpartition start
smpartition commit
```

3.3.3 Recovering the InfiniBand Switches in a Virtual Environment

---

Note: In an Exalogic virtual configuration, do not attempt to **replace** a failed InfiniBand switch with an entirely new one. Contact Oracle Support for the procedure to perform such a replacement. An improperly replaced component might not be discovered correctly by Exalogic Control. You **can** use the procedures described in this document to restore a failed component after **repairing** it.

When Exalogic is deployed in a virtual configuration, do the following to replace a failed InfiniBand switch.

1. Remove the InfiniBand switch from the assets:
   a. Log in to the Exalogic Control BUI as the `root` user.
   b. From the Assets accordion in the navigation pane on the left, expand **Switches**.
   c. Select the switch being replaced.
   d. Place the switch in maintenance mode, by clicking **Place in Maintenance Mode** in the Actions pane.
   e. Select the switch being replaced.
   f. In the Actions pane on the right, click **Delete Assets**.
2. Replace the failed switch by following the standard replacement procedure.

   Note: Before connecting the switch to the IB fabric, disable the subnet manager by running the `disablesm` command on the switch.

3. Restore the switch from the latest backup, by performing the procedure described in Section 3.3.2, "Recovering the InfiniBand Switches in a Physical Environment."

   Note: For the NM2-36P switches in a virtual configuration, do not restore the `partitions.current` file.

4. Identify the BridgeX port GUIDs of the failed IB switch.

   The port GUIDs are created when running the ECU, and can be retrieved from the runtime ECU configuration files.

   a. Log in to the master compute node, and go to the `/var/tmp/exalogic/ecu/switches` directory.

      The BridgeX port GUIDs of all the switches in the machine are stored in files named *`switchHostname`*`_showgwports.out`, where `switchHostname` is the hostname or IP address of the IB switch.

   b. Using a text editor, open the file corresponding to the failed IB switch. For example, if `elswib02` is the failed IB switch, open `elswib02_showgwports.out`, as shown in the following example:

```
cat elswib02_showgwports.out
showgwports

INTERNAL PORTS:
---------------

Device   Port Portname  PeerPort PortGUID           LID    IBState  GWState
-------------------------------------------------------------------------------
Bridge-0  1   Bridge-0-1    4    0x002128deb28ac001 0x004b Active   Up
Bridge-0  2   Bridge-0-2    3    0x002128deb28ac002 0x004f Active   Up
Bridge-1  1   Bridge-1-1    2    0x002128deb28ac041 0x0055 Active   Up
Bridge-1  2   Bridge-1-2    1    0x002128deb28ac042 0x0059 Active   Up

CONNECTOR 0A-ETH:
```

```
-----------------

Port      Bridge      Adminstate Link  State       Linkmode       Speed
-------------------------------------------------------------------------
0A-ETH-1  Bridge-0-2  Enabled    Up    Up          XFI            10Gb/s
0A-ETH-2  Bridge-0-2  Enabled    Up    Up          XFI            10Gb/s
0A-ETH-3  Bridge-0-1  Enabled    Up    Up          XFI            10Gb/s
0A-ETH-4  Bridge-0-1  Enabled    Up    Up          XFI            10Gb/s

CONNECTOR 1A-ETH:
-----------------

Port      Bridge      Adminstate Link  State       Linkmode       Speed
-------------------------------------------------------------------------
1A-ETH-1  Bridge-1-2  Enabled    Up    Up          XFI            10Gb/s
1A-ETH-2  Bridge-1-2  Enabled    Up    Up          XFI            10Gb/s
1A-ETH-3  Bridge-1-1  Enabled    Up    Up          XFI            10Gb/s
1A-ETH-4  Bridge-1-1  Enabled    Up    Up          XFI            10Gb/s
```

c. Note the four BridgeX port GUIDs, which are displayed in the `PortGUID` column of the `INTERNAL PORTS` section.

In this example, the BridgeX ports are `0x002128deb28ac001`, `0x002128deb28ac002`, `0x002128deb28ac041` and `0x002128deb28ac042`.

5. Add the gateway port GUIDs of the switch to the existing EoIB partitions:

a. Log in to the master compute node of the rack—that is, the compute node on which the Exalogic Configuration Utility (ECU) was run.

b. Set the `ECU_HOME` variable in your shell, as shown in the following example:

```
export ECU_HOME=/opt/Exalogic/ecu
```

c. Discover all the IB switches in the fabric, by running the following command:

```
./ecu.sh ib_switches discover
```

d. Discover and add the gateway bridge GUIDs of all the IB switches to the system EoIB partitions of the switches, by running the following commands:

`./ecu.sh ib_switch_gw_ports discover` — to discover the GUIDs

`./ecu.sh ib_switch_gw_ports add_ports` — to add the GUIDs to the `0x8006` partition (the EoIB partition for the Exalogic Control stack)

`./ecu.sh ib_switch_gw_ports show` — to display the ports GUIDs

e. Add the gateway port GUIDs of the switch being replaced to the custom EoIB partitions, as **full** members.

i. Log in to the switch running the master subnet manager.

ii. Run `smpartition start` to start editing the partitions. This command creates a temporary file `partitions.conf.tmp` in the `/conf` directory. This file can be updated using regular Linux commands.

iii. In the `/conf/partitions.conf.tmp` file, replace the failed switch's BridgeX port GUIDs, which you identified in step 4, with the BridgeX port GUIDs of the replacement switch, as identified in step 5.d.

You can do this by using a text editor, or by using the `sed` command, as shown in the following example:

```
sed 's/0x002128deb28ac001/0x002128fe54f6c001/g' /conf/partitions.conf.tmp
```

iv. Run `smpartition commit` to commit and propagate the configuration to all the switches in the fabric.

6. Verify whether the output of the `smnodes list` command on all the InfiniBand switches is correct. The command must display the IP addresses of the switches in the fabric.

   Note:

   - If the output of `smnodes list` does not contain the IP addresses of all the IB switches intended to run the subnet manager, use the `smnodes add` command to update the SM nodes across all the switches in the fabric.

     ```
     smnodes add IP_address_of_IB_switch
     ```

   - To delete the IP address of a switch from the smnodes list, use the `smnodes delete` command.

     ```
     smnodes delete IP_address_of_IB_switch
     ```

7. Propagate the current subnet manager configuration to the switch:

   Note: To ensure that the new switch joins the IB fabric, enable the subnet manager by running the `enablesm` command on both switches. However, only one of the gateway switches should be set as the master sm.

   a. Identify the switch running the master subnet manager, by logging into any one of the switches and running `getmaster`.

   b. Log in to the switch running the master subnet manager.

   c. Run `smpartition start` to edit the subnet manager configuration.

   d. Run `smpartition commit` to save and propagate the subnet manager configuration.

   e. Log in to the switch being replaced, and verify whether the subnet manager configuration has been propagated by running `smpartition list active`.

8. Update the credentials for the switch:

   a. Log in to the Exalogic Control BUI.

   b. Select **Credentials** in the **Plan Management** accordion.

   c. Enter the host name of the switch in the search box and click **Search**.

      The IPMI and SSH credential entries for the switch are displayed.

   d. To update all four credentials, do the following:

      i. Select the entry for the credentials and click **Edit**. The Update Credentials dialog box is displayed.

      ii. Update the password and confirm the password fields.

      iii. Click **Update**.

9. Rediscover the asset:

a. Log in to the Exalogic Control BUI.

b. In the navigation pane on the left, expand **Plan Management**, and under **Profiles and Policies**, expand **Discovery**.

c. Select the appropriate **Infiniband @ *host*** discovery profile.

d. In the Actions pane on the right, click **Add Assets**.

e. On the resulting screen, verify whether the correct discovery profile is displayed.

f. Click **Add Now**.

g. Wait until the discovery process succeeds.

h. In the left navigation pane, expand **Assets** to display all the assets.

i. Verify whether the replaced server is displayed in the Assets section and positioned correctly in the photo-realistic view.

10. Add the switch as an asset:

a. Log in to the Exalogic Control BUI.

b. Expand the **Assets** accordion.

c. Select the appropriate rack, and then select **Place/Remove Assets** from the **Actions** accordion.

d. In **Place/Remove Assets in the Oracle Exalogic Rack** dialog box, select the switch, and then click **Submit**.

After the job is complete, the switch will be visible in the Assets tab.

3.4 Cisco Management Switch

The Cisco Management switch provides connectivity on the management interface and must be backed up regularly.

This section contains the following subsections:

- Section 3.4.1, "Backing Up the Management Switch"
- Section 3.4.2, "Recovering the Management Switch in a Physical Environment"
- Section 3.4.3, "Recovering the Management Switch in a Virtual Environment"

3.4.1 Backing Up the Management Switch

Save the Cisco switch backups to the NFS locations you created for the Cisco switch as described in Chapter 2, "Backup and Recovery Locations" (for example, `/export/Exalogic_Backup/management_switches`).

1. Enable the FTP service on the local ZFS storage appliance:

a. Log in to the storage BUI, `https://storageIP:215/`, as the `root` user.

b. Click the **Shares** tab.

c. Double click the `management_switches` share.
The properties page of the `management_switches` share is displayed.

d. Click the **Protocols** tab.

e. Under the FTP section, deselect **Inherit from project**.

f. Set the Share mode as **Read/write**.

g. Click **Apply**.

h. Navigate to **Services** under Configuration, and select the **FTP** service.

i. On the FTP page, in the General Settings section, set the **Default Login** root to the NFS share location created for the Cisco switch as described in Chapter 2, "Backup and Recovery Locations."

Example: `/export/Exalogic_Backup/management_switches`

j. Under Security Settings, permit `root` login.

k. Click **Apply**.

2. Back up the configuration of the Cisco switch:

a. Log in to the Cisco switch and at the `Router>` prompt, issue the `enable` command.

Provide the required password when prompted. The prompt changes to `Router#`, which indicates that the router is now in privileged mode.

b. Configure the FTP user name and password:

```
Router#config terminal
Router (config)#ip ftp username root
Router (config)#ip ftp password password for root
Router (config)#end
Router#
```

c. Copy the configuration to the FTP server.

```
Router#copy running-config ftp:
Address or name of remote host []? IP address of your storage
Destination filename [Router-confg]? backup_cfg_for_router
Writing backup_cfg_for_router !
1030 bytes copied in 3.341 secs (308 bytes/sec)
Router#
```

d. Open the configuration file using a text editor. Search for and remove any line that starts with `AAA`.

Note: This step is performed to remove any security commands that can lock you out of the router.

3.4.2 Recovering the Management Switch in a Physical Environment

To recover the Cisco switch, do the following:

Note: Ensure that no configuration changes are being made while performing the restore.

---

1. Log in to the Cisco switch.

At the `Router>` prompt, issue the `enable` command, provide the required password when prompted. The prompt changes to `Router#`, which indicates that the router is now in privileged mode.

2. Configure the FTP user name and password:

```
Router#config terminal
Router (config)#ip ftp username root
Router (config)#ip ftp password password for root
Router (config)#end
Router #
```

3. Copy the configuration file from the FTP server to a router in privileged (enable) mode which has a basic configuration.
4. Run the following commands:

```
Router# copy ftp: running-config
Address or name of remote host [IP address]?
Source filename [backup_cfg_for_router]?
Destination filename [running-config]?
Accessing ftp://storageIP/backup_cfg_for_router...
Loading backup_cfg_for_router !
[OK - 1030/4096 bytes]
1030 bytes copied in 13.213 secs (78 bytes/sec)
Router#
```

3.4.3 Recovering the Management Switch in a Virtual Environment

> **Note:** In an Exalogic virtual configuration, do not attempt to **replace** a failed Cisco switch with an entirely new one. Contact Oracle Support for the procedure to perform such a replacement. An improperly replaced component might not be discovered correctly by Exalogic Control. You **can** use the procedures described in this document to restore a failed component after **repairing** it.

When Exalogic is deployed in a virtual configuration, do the following to replace a failed Cisco switch.

1. Remove the Cisco switch from the assets:
   a. Log in to the Exalogic Control BUI as the `root` user.
   b. Navigate to the **Assets** section on the left side of the page.
   c. Expand **Switches** to list all the switches associated with the vDC.
   d. Select the switch being replaced.
   e. Click **Delete Assets** in the Actions pane.
2. Replace the failed Cisco switch by following the standard replacement process.
3. After the switch has been replaced, perform the steps in Section 3.4.2, "Recovering the Management Switch in a Physical Environment" to restore the switch from the latest backup.
4. Update the credentials for the switch:
   a. Log in to the Exalogic Control BUI.
   b. Select **Credentials** in the **Plan Management** accordion.
   c. Enter the host name of the switch in the search box and click **Search**.

      The IPMI and SSH credential entries for the switch are displayed.
   d. To update all four credentials, do the following:

i. Select the entry for the credentials and click **Edit**. The Update Credentials dialog box is displayed.

ii. Update the password and confirm the password fields.

iii. Click **Update**.

5. Rediscover the Cisco switch.
   a. Log in to the Exalogic Control BUI.
   b. In the navigation pane on the left, expand **Plan Management**, and under **Profiles and Policies**, expand **Discovery**.
   c. Select the **Cisco Switch @ Cisco-switch** discovery profile.
   d. In the Actions pane on the right, click **Add Assets**.
   e. On the resulting screen, verify whether the correct discovery profile is displayed.
   f. Click **Add Now**.
   g. Wait until the discovery process succeeds.
   h. In the left navigation pane, expand the **Assets** section to display all the assets.
   i. Verify whether the replaced switch is displayed in the Assets section and positioned correctly in the photo-realistic view.
6. If the switch is not displayed in the **Assets** section, add the Cisco switch manually by doing the following:
   a. Log in to the Exalogic Control BUI.
   b. Expand the **Assets** section.
   c. Select the appropriate rack, and then select **Place/Remove Assets** in the **Actions** section on the left side of the page.
   d. In the **Place/Remove assets in the Oracle Exalogic Rack** dialog box, select the switch, and then click **Submit**.

   After the job is complete, the switch is shown in the Assets tab.

3.5 ZFS Storage Heads

The ZFS storage appliance in an Exalogic machine has two heads that are deployed in a clustered configuration. At any given time, one head is active and the other is passive. If a storage head fails and has to be replaced, the configuration from the surviving active node is pushed to the new storage head. It is not required to backup the configuration of the ZFS Storage appliance.

Note: Ensure that the passive head is being restored and not the active head. While performing the restore, ensure that no configuration changes are being made

---

Recovering the ZFS Storage Head in a Virtual Configuration

> **Note:** In an Exalogic virtual configuration, do not attempt to **replace** a failed ZFS storage head with an entirely new one. Contact Oracle Support for the procedure to perform such a replacement. An improperly replaced component might not be discovered correctly by Exalogic Control. You **can** use the procedures described in this document to restore a failed component after **repairing** it

When the Exalogic machine is deployed in a virtual configuration, do the following to add a ZFS storage head.

1. Remove the failed storage head from the assets:
   a. Log in to the Exalogic Control BUI as the `root` user.
   b. Navigate to the **Assets** section on the left side of the page.
   c. Expand **Storage** to list all the storage heads associated with the vDC.
   d. Select the storage head being replaced.
   e. Click **Delete Assets** in the Actions pane.
2. Replace the failed storage by following the standard replacement process.
3. After the storage head has been replaced, update the credentials for the switch:
   a. Log in to the Exalogic Control BUI.
   b. Select **Credentials** under the **Plan Management** section.
   c. Enter the host name of the switch in the search box and click **Search**.

      The IPMI and SSH credential entries for the ILOM and the compute node are displayed.
   d. To update all four credentials do the following:

      i. Select the entry for the credentials and click **Edit**. The Update Credentials dialog box is displayed.

      ii. Update the password and confirm the password fields.

      iii. Click **Update**.
4. Rediscover the storage appliance.
   a. Log in to the Exalogic Control BUI.
   b. In the navigation pane on the left, expand **Plan Management**, and under **Profiles and Policies**, expand **Discovery**.
   c. Select the appropriate **Storage Appliance @ *host*** discovery profile.
   d. In the Actions pane on the right, click **Add Assets**.
   e. On the resulting screen, verify whether the correct discovery profile is displayed.
   f. Click **Add Now**.
   g. Wait until the discovery process succeeds.
   h. In the left navigation pane, expand the **Assets** section to display all the assets.

i. Verify whether the replaced storage is displayed in the Assets section and positioned correctly in the photo-realistic view.

5. Add the replaced storage as an asset:

   a. Log in to the Exalogic Control BUI.

   b. Expand the **Assets** section.

   c. Select the appropriate rack, and then select **Place/remove Assets** in the **Actions** section on the right side of the page.

   d. In the **Place/Remove assets in Oracle Exalogic Rack** dialog box, select the storage head, and then click **Submit**.

   After the job is complete, the storage head is shown in the Assets tab.

6. Add the IB port GUIDs of the replaced storage head to the `IPoIB-admin`, `IPoIB-storage`, and the `IPoIB-vserver-shared-storage` partitions. The default keys for these partitions are `0x8001`, `0x8002`, and `0x8005` respectively.

   a. Log in to the replaced storage head by using SSH.

   i. Identify the IB port GUID for the first port as follows:

```
storagehead:> configuration net devices

storagehead:configuration net devices> select ibp0

storagehead sn02:configuration net devices ibp0> show

Properties:

            speed = 32000 Mbit/s

               up = true

           active = false

            media = Infiniband

      factory_mac = not available

             port = 1

             guid = 0x212800013f279b


storagehead:configuration net devices ibp0>
```

   The IB port GUID for port 1 is shown by the `guid` entry. In this example, the GUID is `0x212800013f279b`.

   ii. Identify the IB port GUID for the second port as follows:

```
storagehead:> configuration net devices

storagehead:configuration net devices> select ibp1

storagehead sn02:configuration net devices ibp1> show

Properties:

            speed = 32000 Mbit/s

               up = true

           active = false

            media = Infiniband
```

```
factory_mac = not available
        port = 1
        guid = 0x212800013f279c
```

`storagehead:configuration net devices ibp0>`

iii. The IB port GUID for port 1 is shown by the `guid` entry. In this example, the GUID is `0x212800013f279c`.

b. Add the IB port GUIDs as full members of the `IPoIB-admin` partition with the default `pkey` of `0x8001`:

i. Log in to the switch running the master subnet manager.

ii. Run `smpartition start` to edit the partitions.

iii. Add the GUIDs to partition `0x8001` using the following commands:

```
smpartition add -pkey 8001 -port GUID_for_port1 -m full
smpartition add -pkey 8001 -port GUID_for_port2 -m full
```

Example:

```
smpartition add -pkey 8001 -port 0x212800013f279b -m full
smpartition add -pkey 8001 -port 0x212800013f279c -m full
```

iv. Run `smpartition commit` to update and the propagate the configuration to all the switches in the fabric.

c. Repeat step (b) to add the IB port GUIDs identified in step (a) to the `IPoIB-storage` network with a default `pkey` of `0x8002` and to the `IPoIB-vserver-shared-storage` network with a default `pkey` of `0x8005`.

4

Recovery of the Exalogic Control Stack from Hardware Failures and Corruption

Exalogic Control is a comprehensive software management stack providing onboarded capabilities for Exalogic machine, vDC management, and monitoring.

This chapter contains the following sections:

4.1 Recovering from a Hardware Failure

This section describes how you can recover each component of the Exalogic Control stack if a compute node running an Exalogic Control vServer crashes.

4.1.1 Recovering the Database vServer from a Hardware Failure

The database repositories for Oracle VM Manager (`ovs`) and Enterprise Manager Ops Center (`emoc`, `emoc_ro`) are deployed to a database vServer. By default, the database vServer is deployed to the first compute node in the first pool. If the database vServer crashes, both Oracle VM manager and Enterprise Manager Ops Center will stop being operational.

After the compute node is restored to its previous state, do the following to start the database vServer:

1. Log in to the compute node as the `root` user.
2. Change directory to `/OVS/Repositories/*/VirtualMachines`.

   ```
   hostname#cd /OVS/Repositories/*/VirtualMachines
   ```

3. Find the absolute path to the virtual machine configuration file for the database vServer.

   Run the following `grep` command to identify the correct configuration file corresponding to the Exalogic Control database vServer:

   ```
   hostname# grep -i ExalogicControlDB */vm.cfg
   ```

   The output is similar to:

   ```
   0004fb00000600002c18bee8647fb8f7/vm.cfg:OVM_simple_name = 'ExalogicControlDB'
   ```

4. Start the database vServer by using the `xm create` command.

   ```
   xm create absolute_path_to_vm.cfg
   ```

   Example:

   ```
   xm create 0004fb00000600002c18bee8647fb8f7/vm.cfg
   ```

5. Verify whether the database vServer and all the database processes are running by logging in to the vServer.
6. Restart the Oracle VM Manager process:
   a. Log in to the Oracle VM Manager vServer as `root`.
   b. Run `service ovmm stop` to stop the Oracle VM Manager process.
   c. Run `service ovmm start` to start the Oracle VM Manager process.
   d. Verify whether Oracle VM Manager started correctly by logging in to the Oracle VM Manager console.
7. Restart the Proxy Controller and Enterprise Controller processes:
   a. Log in to each of the Proxy Controller vServers as `root`.
   b. Run `proxyadm stop` to stop the Proxy Controller.
   c. Run `proxyadm start` to start the Proxy Controller.
   d. Log in to the Enterprise Controller vServer as `root`.
   e. Run `satadm stop` to stop the Enterprise Controller.
   f. Run `satadm start` to start the Enterprise Controller.
   g. Verify whether the Proxy Controller and Enterprise Controller vServers restarted successfully by logging in to the Exalogic Control BUI.

4.1.2 Recovering the Oracle VM Manager vServer from a Hardware Failure

The Oracle VM Manager is deployed to the ovmm vServer. By default, it is deployed to the first compute node in the first pool. If the Oracle VM Manager vServer crashes, Enterprise Manager Ops Center functionality will be affected.

After the compute node is restored to its previous state, do the following to start the Oracle VM Manager vServer:

1. Log in to the compute node as the `root` user.
2. Change directory to `/OVS/Repositories/*/VirtualMachines`.

   ```
   cd /OVS/Repositories/*/VirtualMachines
   ```

3. Find the absolute path to virtual machine configuration file for the Oracle VM Manager vServer.

   Run the following `grep` command to identify the correct configuration file:

   ```
   Hostname# grep -i ExalogicControlOVMM */vm.cfg
   ```

   The output is similar to:

   ```
   0004fb000006000088afde54f9794d32/vm.cfg:OVM_simple_name = 'ExalogicControlOVMM'
   ```

4. Start the Oracle VM Manager vServer using the `xm create` command.

   ```
   xm create absolute_path_to_vm.cfg
   ```

Example:

```
xm create 0004fb000006000088afde54f9794d32/vm.cfg
```

5. Verify whether the Oracle VM Manager vServer and all the Oracle VM processes are running by logging in to the vServer.
6. Verify whether Enterprise Manager Ops Center is fully functional by logging in to the Exalogic Control BUI.

4.1.3 Recovering the Proxy Controller vServer from a Hardware Failure

The Proxy Controller component is deployed as two vServers. By default, the second vServer (pc2) is deployed to the second compute node in the first pool and the first proxy controller (pc1) vServer is deployed to the third compute node in the first pool. Enterprise Manager Ops Center functionality will be affected if either of the Proxy Controller vServers crash.

After the compute nodes are recovered to their previous state, do the following to start the Proxy Controller vServer:

1. Log in to the compute node as the `root` user.
2. Change directory to `/OVS/Repositories/*/VirtualMachines`.

   ```
   cd /OVS/Repositories/*/VirtualMachines
   ```

3. Find the absolute path to virtual machine configuration file for the Proxy Controller vServer.

   Run the following `grep` command to identify the correct configuration file:

   ```
   Hostname# grep -i ExalogicControlOpsCenterPC* */vm.cfg
   ```

   The output is similar to:

   ```
   0004fb0000060000821f3e60a6d3502d/vm.cfg:OVM_simple_name =
   'ExalogicControlOpsCenterPC2'

   0004fb000006000084a183dbe7c3dba0/vm.cfg:OVM_simple_name =
   'ExalogicControlOpsCenterPC1'
   ```

4. Start the Proxy Controller vServer by using the `xm create` command.

   ```
   xm create absolute_path_to_vm.cfg
   ```

   Example:

   ```
   xm create 0004fb0000060000cf01f02c2fb5adaf/vm.cfg
   ```

5. Verify whether Enterprise Manager Ops Center is fully functional by logging in to the Exalogic Control BUI.

4.1.4 Recovering the Enterprise Controller vServer from a Hardware Failure

The Enterprise Controller component of the Enterprise Manager Ops Center is deployed to the enterprise controller vServer. By default, the enterprise controller vServer is deployed to the fourth compute node in the first pool. All provisioning and lifecycle management functionality is affected if the enterprise controller vServer is unavailable. Once the compute node is recovered to its previous state, follow the steps to start the enterprise controller vServer:

1. Log in to the compute node as the `root` user.
2. Change the directory to `/OVS/Repositories/*/VirtualMachines`:

   ```
   cd /OVS/Repositories/*/VirtualMachines
   ```

3. Find the absolute path to virtual machine configuration file for the Proxy Controller vServer.

   Run the following `grep` command to identify the correct configuration file:

   ```
   Hostname# grep -i ExalogicControlOpsCenterEC1 */vm.cfg
   ```

   The output is similar to:

   ```
   0004fb0000060000cf01f02c2fb5adaf/vm.cfg:OVM_simple_name =
   'ExalogicControlOpsCenterEC1'
   ```

4. Manually start the enterprise controller vServer using the `xm create` command. The syntax for the command is:

   ```
   xm create absolute_path_to_vm.cfg
   ```

   Example:

   ```
   xm create 0004fb0000060000cf01f02c2fb5adaf/vm.cfg
   ```

5. Validate that Ops Center is fully functional by logging in to the Ops Center BUI.

4.2 Backing Up and Recovering Oracle VM Manager

This section provides the backup and recovery steps for the Oracle VM Manager and its repository.

Save the backup to the NFS location you created for the Exalogic Control stack, as described in Chapter 2, "Backup and Recovery Locations" (for example, `/export/Exalogic_Backup/control_metadata`).

Create a directory for storing the backups of the Oracle VM manager under this NFS share (for example, `/export/Exalogic_Backup/control_metadata/ovmm`).

4.2.1 Backing Up Oracle VM Manager

To back up Oracle VM Manager, you should back up the Oracle VM Manager configuration file, and the Oracle VM Manager database schema. By default, this schema is named `ovs`, and this name is used in the backup example; when you run this procedure, replace the schema name with your own.

The Oracle VM Manager configuration file is stored at the following location on the Oracle VM manager vServer:

```
/u01/app/oracle/ovm-manager-3/.config
```

This configuration file contains database connection information, ports, and the UUID used by Oracle VM Manager.

The following is an example of this configuration file:

```
DBHOST=<hostname of database server>
SID=<oracle SID>
LSNR=<listener port number defaults 1521>
APEX=<application express port number defaults 8080>
OVSSCHEMA=<database schema name for oracle vm manager defaults ovs>
```

```
WLSADMIN=<weblogic server admin defaults weblogic>
OVSADMIN=<oracle vm manager administrator name defaults admin>
COREPORT=<oracle vm manager core port defaults 54321>
UUID=<oracle vm manager uuid>
```

To back up Oracle VM Manager, do the following:

1. Back up or copy the Oracle VM Manager configuration file located at:

   ```
   /u01/app/oracle/ovm-manager-3/.config
   ```

2. As the `root` user, shut down Oracle VM Manager.

   ```
   # /sbin/service ovmm stop
   ```

3. Back up the Oracle VM Manager database schema.
   a. Log in to any compute node on the Exalogic machine.
   b. Mount the NFS location created for the exalogic control stack.
   c. Log in to the Oracle database vServer from the compute node using the IP address of the IPoIB-admin interface.
   d. Log in to the operating system as the `oracle` user.

      You can use the `su - oracle` command as the `root` user if you do not have the password for the `oracle` user.

   e. Find the version of Oracle Database installed by navigating to the following directory:

      ```
      /u01/app/oracle/product/
      ```

   f. Set the ORACLE_HOME, PATH and ORACLE_SID environment variables by running the following commands:

      ```
      export ORACLE_HOME=/u01/app/oracle/product/ProductVersion/dbhome_1
      ```

      `ProductVersion` is the version of Oracle Database installed on your Exalogic rack that you found in step e.

      ```
      export PATH=$ORACLE_HOME/bin:$PATH

      export ORACLE_SID=elctrldb
      ```

   g. Navigate to the following directory:

      ```
      /u01/app/oracle/product/ProductVersion/dbhome_1/bin
      ```

   h. Export the schema using the using the `exp` command:

      ```
      exp ovs/password grants=y compress=y file=/tmp/ovsbackup.dmp
      ```

   i. FTP the backup file to the compute node from step a. Use the mount point of the NFS share as the destination in the `ftp` command.
   j. Store the Oracle VM Manager database schema backup along with the Oracle VM Manager configuration file.

4.2.2 Restoring Oracle VM Manager

To restore Oracle VM Manager, and the Oracle VM Manager database schema from a backup, you must have performed the steps to back up Oracle VM Manager as

described in Section 4.2.1, "Backing Up Oracle VM Manager."

> **Note:** The OVS repository and the Oracle VM Manager must always be restored together. They cannot be restored as individual components.

1. In certain cases it may be necessary to either re-install or upgrade the Oracle VM Manager. For more information, see the following documentation:

   - Installing Oracle VM Manager: `http://docs.oracle.com/cd/E27300_01/E27308/html/vmiug-manager-install.html`
   - Upgrading Oracle VM Manager: `http://docs.oracle.com/cd/E27300_01/E27308/html/vmiug-manager-upgrading.html`

   Log in to the Oracle VM Manager vServer and then perform the install using the `runInstaller.sh --uuid` *`uuid`* command and provide the UUID from the previous manager installation you created a backup from. The UUID can be found in the Oracle VM Manager configuration file.

   > **Note:** The Oracle VM Manager UUID is also persisted in the `/etc/sysconfig/ovmm` file. If the system disk of the server on which you are installing or restoring Oracle VM Manager was not wiped entirely, the existing UUID is still present and will be detected when running the installer.
   >
   > - The `--uuid` option overrides this existing UUID.
   > - If no UUID is present in `/etc/sysconfig/ovmm`, the `--uuid` option adds the UUID to the file.

   Example:

   ```
   # ./runInstaller.sh --uuid 0004FB000000100002CB7F2DFFA8D8
   ```

   When the Oracle VM Manager installer prompts for installation information other than passwords, reuse the same user names for the Oracle Database schema, Oracle WebLogic Server and Oracle VM Manager administration user, as set out in the backup of the Oracle VM Manager configuration file. You must set the passwords again as the passwords are not backed up and cannot be restored.

2. After installation, reinstallation or upgrade, stop Oracle VM Manager before you restore the backup:

   ```
   # /sbin/service ovmm stop
   ```

3. Log in to any compute node in the Exalogic machine.
4. Mount the NFS location created for the Exalogic control stack.
5. Log in to the Oracle Database vServer from the compute node using the IP address of the IPoIB-admin interface.
6. Log in to the operating system as the oracle user. You can `su - oracle` as the root user if you do not have the password for the oracle user.
7. Find the version of Oracle Database installed by navigating to the following directory:

   ```
   /u01/app/oracle/product/
   ```

8. Set the ORACLE_HOME, PATH and ORACLE_SID environment variables by running the following commands:

   ```
   export ORACLE_HOME=/u01/app/oracle/product/ProductVersion/dbhome_1
   ```

   `ProductVersion` is the version of Oracle Database installed on your Exalogic rack that you found in step e.

   ```
   export PATH=$ORACLE_HOME/bin:$PATH

   export ORACLE_SID=elctrldb
   ```

9. Log in to the Oracle Database as the `sys` or `system` user to delete the Oracle VM Manager administration user. The default Oracle VM Manager administration user is ovs.

   ```
   $ sqlplus system/password
   SQL> drop user ovs cascade;
   ```

   Then, re-create the Oracle VM Manager administration user, with the necessary grants:

   ```
   SQL> create user ovs identified by password;
   SQL> grant connect, resource to ovs;
   SQL> exit;
   ```

10. FTP the backup file of the OVS schema from the NFS location mounted on the compute node from Step 1 to a temporary location on the database vServer.

11. Restore the OVS schema by importing it from the backup file.

    ```
    # imp ovs/password file=/tmp/ovmmovsbackup.dmp full=y
    ```

12. Restart Oracle VM Manager.

    ```
    # /sbin/service ovmm start
    ```

4.3 Recovering Oracle VM Manager After Database Corruption

This section provides the procedure for recovering from corruption of the Oracle VM Manager data in the database.

Save the backups for the Oracle VM Manager and Enterprise Manager Ops Center repositories to the NFS locations created as described in Chapter 2, "Backup and Recovery Locations" (for example, `/export/Exalogic_Backup/control_metadata`).

This procedure has the following steps:

Step 1: Shut Down the Enterprise Manager Ops Center VMs

Step 2: Cleanup the Oracle VM Manager RDBMS Schema

Step 3: Rediscover the Environment

Step 4: Start the Enterprise Manager Ops Center Control VMs

Step 5: Synchronize the Clocks of the Control vServers

Step 6: Verify Whether Enterprise Manager Ops Center is Running

Step 1: Shut Down the Enterprise Manager Ops Center VMs

When the Oracle VM Manager database gets corrupted, the Oracle VM Manager can no longer operate. As a result Enterprise Manager Ops Center cannot perform any of the Exalogic cloud management functions.

The Enterprise Manager Ops Center control VMs must be shut down in the following order:

1. ExalogicControlOpsCenterEC1 VM
2. ExalogicControlOpsCenterPC1 VM
3. ExalogicControlOpsCenterPC2 VM

First, determine the OVS server on which the Enterprise Manager Ops Center control VMs are running. Note that the VMs must be restarted on the same OVS servers later.

Note: By default, the VMs run on the following OVS servers:

- ExalogicControlOpsCenterEC1 VM: cn04
- ExalogicControlOpsCenterPC1 VM: cn03
- ExalogicControlOpsCenterPC2 VM: cn02

---

1. Launch a web browser and log in to the Oracle VM Manager web console.
2. Expand **Server Pools** under the **Home** section.
3. Under the **Serverpool1** pool, expand the OVS servers one by one till you find the three VMs listed earlier.

There are several ways to shut down the Enterprise Manager Ops Center control VM instances. This section documents two of them.

- Using Oracle VM Manager web console

  1. Select the **ExalogicControlOpsCenterEC1** VM and click **Stop**.
  2. Select the **ExalogicControlOpsCenterPC1** VM and click **Stop**.
  3. Select the **ExalogicControlOpsCenterPC2** VM and click **Stop**.

  Wait for each VM to be shut down before proceeding.

- Using Oracle VM Manager CLI

  Do the following for each of the Exalogic Control VMs, in the order listed earlier:

  1. SSH to the Oracle VM Manager VM by using the `EoIB-external-mgmt` IP address assigned to that VM.
  2. From inside the Oracle VM Manager VM, SSH to the OVS server (`dom0`) where the VM is running using the `IPoIB-ovm-mgmt` IP address assigned to that OVS server. If you used the default ECU settings, then the IP address will be:

     `192.168.23.`*`N`*

     `N` is the index of the OVS server.

     Example:

     For cn01, the IP will be:

     `192.168.23.1`

3. Find the UUID of the VM that must be shut down.

   This can be done using the Oracle VM Manager web console: Select the VM instance. The ID will be shown in the window on the right. You can confirm that the VM runs on this OVS server by running the following command:

   ```
   xm list
   ```

   If the control VMs were never stopped, then they will have index 1 on the respective OVS server.

4. To gracefully shut down the VM, run the following command:

   ```
   xm shutdown UUID
   ```

   Optionally, you can monitor the VM while it is shutting down by running this command:

   ```
   xm console UUID
   ```

5. Verify that the VM instance has shut down by running the following command:

   ```
   xm list
   ```

   The UUID of the VM that was stopped should not be listed.

Note: The VM can be shut down forcefully. Oracle does not recommend this process. Use it only if the previous steps do not work:

```
xm destroy UUID
```

Step 2: Cleanup the Oracle VM Manager RDBMS Schema

1. **Stop the Oracle VM Manager service**.

   1. SSH as `root` to the Oracle VM Manager VM using its EoIB-external-mgmt IP address.
   2. Run the following command:

      ```
      service ovmm stop
      ```

2. **Fix any RDBMS corruptions**.

   For example, if there is a table containing a corrupted row, that row must be deleted using normal RDBMS tools such as SQL*Plus. If the RDBMS error ORA-1555 is present in any of the log files, then there is possibly a corrupted BLOB value. The section below covers how to confirm and delete that row.

   Note: If more than one table contains corrupted BLOBs you should repeat this procedure for each table. You should drop the `corrupted_lob_data` table by using the following statement:

   ```
   SQL> drop table corrupted_lob_data;
   ```

   The following are the steps to find the corrupted BLOB value and to delete it:

   1. SSH as `root` to the Oracle VM Manager VM using its EoIB-external-mgmt IP address.

2. From the Oracle VM Manager VM, SSH as `root` to the DB VM instance using its IPoIB-admin IP address. The default IP address is 192.168.20.10.

   Run the following commands to connect to the Oracle VM Manager database schema:

```
# ssh root@192.168.20.10
The authenticity of host '192.168.20.10 (192.168.20.10)' can't be
established.
RSA key fingerprint is f6:14:37:f9:ef:45:ba:48:73:76:35:7f:a9:e0:99:ab.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '192.168.20.10' (RSA) to the list of known
hosts.
root@192.168.20.10's password:
Last login: Wed May 16 13:39:58 2012
[root@elir-db ~]# su - oracle
[oracle@elir-db root]$ export ORACLE_
HOME=/u01/app/oracle/product/DatabaseVersion/dbhome_1
```

   `DatabaseVersion` is the version of Oracle DB installed on your Exalogic machine.

```
[oracle@elir-db root]$ export ORACLE_SID=elctrldb
[oracle@elir-db root]$ cd /u01/app/oracle/product/DatabaseVersion/dbhome_
1/bin
[oracle@elir-db bin]$ ./sqlplus ovs@elctrldb

SQL*Plus: Release 11.2.0.1.0 Production on Wed May 16 13:39:58 2012

Copyright (c) 1982, 2009, Oracle.  All rights reserved.

Enter password:

Connected to:
Oracle Database 11g Enterprise Edition Release 11.2.0.1.0 - 64bit
Production
With the Partitioning, OLAP, Data Mining and Real Application Testing
options

SQL>
```

3. Create a new temporary table for storing all row IDs of the corrupted LOBs. In this example, we will call it `corrupted_lob_data`.

```
SQL> create table corrupted_lob_data (corrupted_rowid rowid);
```

4. Execute the following SQL script to detect BLOB data corruptions:

```
SQL> set concat off

declare
  error_1555 exception;
  pragma exception_init(error_1555,-1555);
  num number;
begin
  for cursor_lob in (select rowid r, &&lob_column from &table_owner.&table_
with_lob) loop
    begin
      num := dbms_lob.instr (cursor_lob.&&lob_column, hextoraw ('889911'))
;
    exception
      when error_1555 then
```

```
        insert into corrupted_lob_data values (cursor_lob.r);
        commit;
    end;
  end loop;
end;
/
```

After this step, the following prompts are displayed:

```
Enter value for lob_column    : m_data
Enter value for table_owner   : ovs
Enter value for table_with_LOB: You need to get this information from your
application log file, for example, Mgr_Eventlog, run this for each table
you saw in the log prior to the exception
```

In the end all row IDs of the corrupted LOBs are inserted into the `corrupted_lob_data` table created earlier.

5. Check whether any corrupted BLOBs were found by executing the following query:

```
SQL> select * from corrupted_lob_data;

CORRUPTED_ROWID
--------------------
AAEWBsAAGAAACewAAC
AAEWBsAAGAAACewAAF
AAEWBsAAGAAACewAAG

3 rows selected
```

6. Remove the corrupted BLOB value, by doing *one of the following*:

Empty the affected LOBs:

```
SQL> update table_name_with_corrupted_blob set m_data = empty_blob()
     where rowid in (select corrupted_rowid from corrupted_lob_data);
```

Delete the rows with the corrupted BLOB value:

```
SQL> delete from table_name_with_corrupted_blob where rowid in (select
corrupted_rowid from corrupted_lob_data);
```

Notes: The SQL*Plus tool is available only on the RDBMS VM. You can access that VM through SSH from the Oracle VM Manager VM to the RDBMS VM using the IPoIB-admin IP address assigned to that VM.

The following My Oracle Support documents describe how to find corrupted BLOB segments:

- 833635.1: Export Fails with ORA-2354 ORA-1555 ORA-22924 and How to Confirm LOB Segment Corruption Using Export Utility?
- 787004.1: Export Receives the Errors ORA-1555 ORA-22924 ORA-1578 ORA-22922

---

3. Delete all tables in the Oracle VM Manager RDBMS schema.

1. SSH as `root` to the Oracle VM Manager VM using its EoIB-external-mgmt IP address.

2. Gather some metadata information about the current Oracle VM Manager instance by running the following commands:

```
[root@elir-ovmm ovm-manager-3]# cat /u01/app/oracle/ovm-manager-3/.config
DBHOST=192.168.20.10
SID=elctrldb
LSNR=1521
APEX=8080
OVSSCHEMA=ovs
WLSADMIN=weblogic
OVSADMIN=admin
COREPORT=54321
UUID=0004fb0000010000e224fcdfc21df2d2
BUILDID=3.0.3.240
```

   You can find the UUID assigned to the Oracle VM Manager instance by running the following command. This UUID is required if Oracle VM Manager needs to be reinstalled.

```
[root@elir-ovmm ovm-manager-3]# cat /etc/sysconfig/ovmm
UUID=0004fb0000010000e224fcdfc21df2d2
RUN_OVMM=YES
```

3. Delete all tables in the Oracle VM Manager RDBMS schema by running the following commands and getting the exact values for the parameters from file `/u01/app/oracle/ovm-manager-3/.config`.

```
# cd /u01/app/oracle/ovm-manager-3/ovm_upgrade/bin/
# bash ./ovm_upgrade.sh --dbuser=OVSSCHEMA --dbpass=OVSSCHEMA_PASSWORD
--dbhost=DBHOST --dbport=LSNR --dbsid=SID --deletedb
```

   Example, using default values:

```
# cd /u01/app/oracle/ovm-manager-3/ovm_upgrade/bin/
# bash ./ovm_upgrade.sh --dbuser=ovs --dbpass=default_password
--dbhost=192.168.20.10 --dbport=1521 --dbsid=elctrldb --deletedb
```

   After running this command, the Oracle VM Manager RDBMS schema should still exist but it should be empty (no tables in it).

Step 3: Rediscover the Environment

1. Start up the Oracle VM Manager service.

```
# service ovmm start
```

2. Discover all the OVS servers as follows:

   a. Launch a web browser and access the Oracle VM Manager web console by using the following URL:

```
http://OVMM_VM_EoIB-external-mgmt_IP_Address:7002/ovm/console
```

   b. Expand the **Hardware** section.

   c. Select the **Hardware** tab.

   d. Click the **Discover Servers** button

      The Discover Servers window is displayed.

   e. From inside the Oracle VM Manager VM, SSH to the OVS server (`dom0`) where the VM is running using the IPoIB-ovm-mgmt IP address assigned to that OVS server. If you used the default ECU settings, then that IP address will be:

```
192.168.23.N
```

where N is the index of the OVS server

Example:

For cn01, the IP will be:

```
192.168.23.1
```

f. Enter the Oracle VM Agent password.

g. Click **OK**.

h. Wait till all the servers are discovered. At this point, all OVS servers should be OK.

If any errors occur, you should select the OVS server that has errors. Then, go to the **Events** tab and acknowledge all events.

3. Register the File Server.

a. While still in the **Hardware** pane, select the **Storage** tab.

b. Click on the **Discover a File Server** button.

The Discover a File Server wizard is displayed.

c. Enter the following values:

Name: Generic Network File System

Access Host: *IPoIB-storage_IP_address_of_the_storage*

Note: In a quarter rack at default settings, the above IP address is 192.168.21.9

---

d. Click **Next**.

e. Select all the OVS servers and move them to the Selected Servers section.

f. Click **Next**.

g. Select only the `nfs:/export/ExalogicRepo` file system.

h. Click **Finish**.

4. Present the repository to all servers

a. Expand the **Home** pane.

b. Select **Server Pools**.

c. Click on the **Repositories** tab on the right side of the page.

d. Select the `exlcontrol_repo` repository (it will be the only one listed).

e. Click on **Present-Unpresent Selected Repository** icon (represented by up and down green arrows).

f. Present the repository to all OVS servers by moving all servers to the Present to Server(s) side.

g. Click **OK**.

5. Refresh the repository by selecting it and clicking **Refresh Selected Repository Content** in the toolbar.

6. Click on the **Servers and VMs** tab.
7. Rediscover all the servers by repeating the following procedure for every OVS server in each server pool:
   a. Select the OVS server.
   b. Click the **Rediscover Server** button in the toolbar.

      Alternatively, right-click and select **Rediscover Server**.

   After this step, all running VM instances should be shown under their respective OVS servers. Only the stopped VM instances should be listed under Unassigned Virtual Machines, which should include the three Enterprise Manager Ops Center control VMs.

Note: The storage volumes assigned to each VM will be discovered but instead of displaying their names, Oracle VM Manager will display their UUIDs. This is not a problem.

Step 4: Start the Enterprise Manager Ops Center Control VMs

The Enterprise Manager Ops Center VMs should be started in the following order:

1. ExalogicControlOpsCenterPC2
2. ExalogicControlOpsCenterPC1

   Wait for five minutes before proceeding.
3. ExalogicControlOpsCenterEC1

Refer to where the VMs were running, as noted in the Step 1: Shut Down the Enterprise Manager Ops Center VMs step.

Do the following for each VM:

1. Select the VM from the Unassigned Virtual Machines list.
2. Click **Migrate** in the toolbar.

   Alternatively, right-click on the VM name, and select **Migrate**.
3. Select the OVS server in which the VM should run.
4. Select the migrated VM. Expand the OVS server in which the VM should run.
5. Click **Start** in the toolbar.

   Alternatively, right-click on the VM name and select **Start**.

Step 5: Synchronize the Clocks of the Control vServers

1. Stop the components of the Exalogic Control stack, as described in Section 5.2.1, "Stopping the Components of the Exalogic Control Stack."
2. Log in as `root` on the first compute node.
3. Set up SSH for all the control vServers by running the following commands:

```
#/opt/exalogic.tools/tools/setup-ssh.sh -H 192.168.23.10 -P password
#/opt/exalogic.tools/tools/setup-ssh.sh -H 192.168.23.11 -P password
#/opt/exalogic.tools/tools/setup-ssh.sh -H 192.168.23.12 -P password
#/opt/exalogic.tools/tools/setup-ssh.sh -H 192.168.23.13 -P password
#/opt/exalogic.tools/tools/setup-ssh.sh -H 192.168.23.14 -P password
```

4. Create a file with the IP addresses of all the control vServers as follows:

```
> 192.168.23.10
> 192.168.23.11
> 192.168.23.12
> 192.168.23.13
> 192.168.23.14
```

5. Ensure that the time matches for all control vServers by running the following commands:

```
# /opt/exalogic.tools/tools/dcli -g IPAddressesFile 'date'

# /opt/exalogic.tools/tools/dcli -g IPAddressesFile 'service ntpd stop'

# /opt/exalogic.tools/tools/dcli -g IPAddressesFile 'ntpd -gq'

# /opt/exalogic.tools/tools/dcli -g IPAddressesFile 'service ntpd start'

# /opt/exalogic.tools/tools/dcli -g IPAddressesFile 'date'
```

> **Note:** To synchronize the time between control vServers, you may have to repeat the `ntpd` related commands.

6. Start the components of the Exalogic Control stack, as described in Section 5.2.5, "Starting the Exalogic Control Stack."

Step 6: Verify Whether Enterprise Manager Ops Center is Running

After the ExalogicControlOpsCenterEC1 VM is up, it may take 10 to 15 minutes for the Enterprise Manager Ops Center to completely start and become accessible from a web browser.

Launch a web browser and access the Exalogic Control BUI. Enterprise Manager Ops Center should be fully functional at this point.

5

Backup and Recovery of the Exalogic Control Repository and Stack

This chapter provides the steps for backing up and recovering the Exalogic Control repository and stack. It contains the following sections:

- Section 5.1, "Recovering After Hardware Failure"
- Section 5.2, "Backing Up the Exalogic Control Repository"
- Section 5.3, "Restoring the Exalogic Control Repository"
- Section 5.4, "Restoring the Exalogic Control Stack"

5.1 Recovering After Hardware Failure

If the compute node on which a vServer is running crashes, recover the compute node to its previous state, log in to the Exalogic Control BUI, and start the vServer.

If the compute node on which a vServer is running needs to be replaced, follow the steps in Section 3.2.2, "Reimaging and Bare Metal Restore" to recover the compute node. After the compute node is restored to its previous state, log in to the Exalogic Control BUI, and start the vServer.

5.2 Backing Up the Exalogic Control Repository

The Exalogic Control repository contains virtual-machine disk images, templates, and virtual-machine configuration files for all the virtual machines running on the Exalogic machine. The Exalogic Control repository resides in the ZFS storage appliance in the Exalogic machine. It is recommended that you use ZFS snapshots to back up the Exalogic Control repository.

To ensure data consistency, the virtual machines running the Exalogic Control stack must be shut down before taking a ZFS snapshot. Note that the customer virtual machines do not need to be stopped as long as all the applications are running off an NFS mount.

To back up the Exalogic Control repository, do the following:

1. Stop the components of the Exalogic Control stack as described in Section 5.2.1.
2. Shut down the Exalogic Control virtual machines as described in Section 5.2.2.
3. Create a ZFS snapshot of the Exalogic Control repository as described in Section 5.2.3.
4. Restart the components of the Exalogic Control stack as described in Section 5.2.5.

5.2.1 Stopping the Components of the Exalogic Control Stack

To stop the components of the Exalogic Control stack, do the following:

1. Log in to the Proxy Controller-2 virtual machine as the `root` user.
2. Stop the Proxy Controller process.

   ```
   # /opt/sun/xvmoc/bin/proxyadm stop -w
   ```

3. Log in to the Proxy Controller-1 virtual machine as the `root` user.
4. Stop the Proxy Controller process.

   ```
   # /opt/sun/xvmoc/bin/proxyadm stop -w
   ```

5. Log in to the Enterprise Controller virtual machine as the `root` user.
6. Stop the Enterprise Controller process.

   ```
   # /opt/sun/xvmoc/bin/ecadm stop -w
   ```

7. Log in to the Oracle VM Manager virtual machine as the `root` user.
8. Stop the Oracle VM Manager process.

   ```
   # service ovmm stop
   ```

9. Log in to the database virtual machine as the `root` user.
10. Stop the database process.

    ```
    # service oracle-db stop
    ```

5.2.2 Shutting Down the Exalogic Control Virtual Machines

Gracefully shut down all Exalogic Control virtual machines by doing the following:

1. Log in to the Proxy Controller-2 virtual machine as the `root` user.
2. Shut down the virtual machine, by running the following command:

   ```
   # shutdown -h now
   ```

3. Log in to the Proxy Controller-1 virtual machine as the `root` user.
4. Shut down the virtual machine, by running the following command:

   ```
   # shutdown -h now
   ```

5. Log in to the Enterprise Controller virtual machine as the `root` user.
6. Shut down the virtual machine, by running the following command:

   ```
   # shutdown -h now
   ```

7. Log in to the Oracle VM Manager virtual machine as the `root` user.
8. Shut down the virtual machine, by running the following command:

   ```
   # shutdown -h now
   ```

9. Log in to the database virtual machine as the `root` user.
10. Shut down the virtual machine, by running the following command:

    ```
    # shutdown -h now
    ```

5.2.3 Creating a ZFS Snapshot of the Exalogic Control repository

The snapshot created in this section can be used in the following restoration scenarios:

- Section 5.3, "Restoring the Exalogic Control Repository"
- Section 5.4.1, "Restoring the Exalogic Control Stack from a ZFS Snapshot"

If you want to use the snapshot to restore the Exalogic Control repository, shutdown customer vServers before creating the snapshot.

To create a ZFS storage appliance snapshot of the Exalogic Control repository, do the following:

1. Log in to the BUI of the ZFS storage appliance at `http://storage_host:215` as the `root` user.
2. Navigate to **Shares** and then to **Projects**, and select the **ExalogicControl** project.
3. Under the ExalogicControl project, select the **ExalogicRepo** share.
4. Navigate to Snapshots, and click **+** to create a snapshot.
5. In the Create Snapshot dialog box, provide a name, and click **Apply**.

5.2.4 Creating a Full Backup of the Exalogic Control Artifacts

It is recommended that you create a full backup of the Exalogic Control artifacts regularly.

1. Stop the components of the Exalogic Control stack as described in Section 5.2.1.
2. Shut down the Exalogic Control virtual machines as described in Section 5.2.2.
3. Mount the NFS location defined in Chapter 2, "Backup and Recovery Locations" on one of the compute nodes.
4. Go to the `/OVS/Repositories/*/VirtualMachines` directory.
5. Find the virtual-machine configuration files for the Exalogic Control virtual machines, as shown in the following example:

```
grep -i ExalogicControl */vm.cfg
0004fb000006000038a23e7d5c307e00/vm.cfg:OVM_simple_name =
 'ExalogicControlDB'
0004fb000006000056cafd150ebc34c0/vm.cfg:OVM_simple_name =
 'ExalogicControlOpsCenterPC2'
0004fb00000600009c3bf065d07004cb/vm.cfg:OVM_simple_name =
 'ExalogicControlOVMM'
0004fb0000060000a5143f2ef7d264ac/vm.cfg:OVM_simple_name =
 'ExalogicControlOpsCenterPC1'
0004fb0000060000f0ccd3391c56a578/vm.cfg:OVM_simple_name =
 'ExalogicControlOpsCenterEC1'
```

6. Identify the location and name of the virtual-disk image file for each component of the Exalogic Control stack.

   You can identify the location and name of the virtual-disk image file from the `disk` parameter in the `vm.cfg` file, as shown in the following example:

```
cat 0004fb000006000056cafd150ebc34c0/vm.cfg | grep file
disk =
['file:/OVS/Repositories/0004fb00000300007d5117500d92ae54/VirtualDisks/0004fb00
```

```
001200003d3b8b4058ee7682.img,hda,w'
```

7. Create a `tar` file with the `vm.cfg` from step 5 and the virtual-disk image file from step 6, as shown in the following example:

```
tar -czvf pc2_oct24.tgz
/OVS/Repositories/0004fb00000300007d5117500d92ae54/VirtualDisks/0004fb000012000
03d3b8b4058ee7682.img
/OVS/Repositories/0004fb00000300007d5117500d92ae54/VirtualMachines/0004fb000006
000056cafd150ebc34c0/vm.cfg
```

8. Copy the `tar` file to the backup location that you mounted in step 3.
9. Start the components of the Exalogic Control stack as described in Section 5.2.5.

5.2.5 Starting the Exalogic Control Stack

To start the Exalogic Control stack, do the following:

> **Note:** If the Exalogic Control virtual machines are not up, follow Step 4: Start the Enterprise Manager Ops Center Control VMs in Section 5.4.

1. Log in to the database virtual machine as the `root` user.
2. Start the database process.

```
# service oracle-db start
```

3. Log in to the Oracle VM Manager virtual machine as the `root` user.
4. Start the Oracle VM Manager process.

```
# service ovmm start
```

5. Log in to the Enterprise Controller virtual machine as the `root` user.
6. Start the Enterprise Controller process.

```
# /opt/sun/xvmoc/bin/ecadm start -w
```

7. Log in to the Proxy Controller-1 virtual machine as the `root` user.
8. Start the Proxy Controller process.

```
# /opt/sun/xvmoc/bin/proxyadm start -w
```

9. Log in to the Proxy Controller-2 virtual machine as the `root` user.
10. Start the Proxy Controller process.

```
# /opt/sun/xvmoc/bin/proxyadm start -w
```

5.3 Restoring the Exalogic Control Repository

The entire Exalogic Control repository can be recovered to the point in time when the last good ZFS snapshot was taken.

To restore the Exalogic Control repository, do the following:

Caution: Rolling back to a snapshot reverts all the active data, including customer vServers, to the point in time when the snapshot was created. Any vServers that were created after the snapshot was taken, will no longer be present. Also, if any vServers associated with EoIB networks were deleted after the snapshot was taken, then, when such vServers are restored from the snapshot, the VNICs associated with the vServers will not be re-created

In addition, any recent snapshots and clones will be destroyed. This operation cannot be reversed.

1. Stop the components of the Exalogic Control stack as described in Section 5.2.1, "Stopping the Components of the Exalogic Control Stack."
2. Restore the Exalogic Control repository from a ZFS snapshot as follows:
   a. Log in to the BUI of the ZFS storage appliance at `http://storage_host:215` as the `root` user.
   b. Navigate to **Shares** and then to **Projects**, and select the **ExalogicControl** project.
   c. Under the ExalogicControl project, select the **ExalogicRepo** share.
   d. Navigate to Snapshots, select the required snapshot, and click **Rollback data to this snapshot**.
   e. In the resulting dialog box, click **OK**.
3. Start the components of the Exalogic Control stack as described in Section 5.2.5, "Starting the Exalogic Control Stack."

5.4 Restoring the Exalogic Control Stack

To ensure data consistency, all the Exalogic Control virtual machines must be restored together. The components of the Exalogic Control stack can be recovered from either the ZFS snapshot or the latest full backup.

5.4.1 Restoring the Exalogic Control Stack from a ZFS Snapshot

To restore the Exalogic Control stack from a ZFS snapshot, do the following:

1. Stop the Exalogic Control stack as described in Section 5.2.1, "Stopping the Components of the Exalogic Control Stack."
2. Log in to the BUI of the ZFS storage appliance at `http://storage_host:215` as the `root` user.
3. Navigate to **Shares** and then to **Projects**, and select the **ExalogicControl** project.
4. Under the ExalogicControl project, select the **ExalogicRepo** share.
5. Navigate to Snapshots, select the required snapshot, and clone the snapshot.
6. Select the location of the `Backup` project, as defined in Chapter 2, "Backup and Recovery Locations."
7. Enter a name for the clone (example: `RepoClone_date`).
8. Note the mount point for the clone.
9. Mount the clone on one of the compute nodes, as shown in the following example:

```
mount -t storage_host:/export/Exalogic_Backup/RepoClone_Oct24 /backup
```

10. Go to the `/backup/VirtualMachines` directory.
11. Find the virtual-machine configuration files for the Exalogic Control virtual machines, as shown in the following example:

```
grep -i ExalogicControl /backup/VirtualMachines/*/vm.cfg
0004fb000006000038a23e7d5c307e00/vm.cfg:OVM_simple_name =
 'ExalogicControlDB'
0004fb000006000056cafd150ebc34c0/vm.cfg:OVM_simple_name =
 'ExalogicControlOpsCenterPC2'
0004fb00000600009c3bf065d07004cb/vm.cfg:OVM_simple_name =
 'ExalogicControlOVMM'
0004fb0000060000a5143f2ef7d264ac/vm.cfg:OVM_simple_name =
 'ExalogicControlOpsCenterPC1'
0004fb0000060000f0ccd3391c56a578/vm.cfg:OVM_simple_name =
 'ExalogicControlOpsCenterEC1'
```

12. Identify the location and name of the virtual-disk image file for each component of the Exalogic Control stack.

    You can identify the location and name of the virtual-disk image file from the `disk` parameter in the vm.cfg files, as shown in the following example:

```
cat 0004fb000006000056cafd150ebc34c0/vm.cfg | grep file
disk =
['file:/OVS/Repositories/0004fb00000300007d5117500d92ae54/VirtualDisks/0004fb00
001200003d3b8b4058ee7682.img,hda,w'
```

13. Copy the `vm.cfg` files identified in step 11 to the `/OVS/Repositories/*/VirtualMachines/`*`vmGUID`* directory.

    Note that `vmGUID` is name of the directory in which the `vm.cfg` file is located. For example, step 11 shows the following as the location of the `ExalogicControlDB` virtual machine:

```
0004fb000006000038a23e7d5c307e00/vm.cfg:OVM_simple_name = 'ExalogicControlDB'
```

    In this example, `0004fb000006000038a23e7d5c307e00` is the `vmGUID`.

14. Go to the `/backup/VirtualDisks` directory, and copy the virtual-disk image file identified in step 12 to the location identified by the `disk` parameter in the `vm.cfg` file, as shown in the following example:

```
cp /backup/VirtualDisks/0004fb00001200003d3b8b4058ee7682.img
/OVS/Repositories/0004fb00000300007d5117500d92ae54/VirtualDisks/0004fb000012000
03d3b8b4058ee7682.img
```

15. Repeat steps 12 to 14 to restore each component of the Exalogic Control stack.

5.4.2 Restoring the Exalogic Control Stack from a Full Backup

To restore the Exalogic Control stack from a full backup, do the following:

1. Stop the Exalogic Control stack as described in Section 5.2.1, "Stopping the Components of the Exalogic Control Stack."
2. Mount the backup location as defined in Section 5.2.4, "Creating a Full Backup of the Exalogic Control Artifacts."
3. Identify the backup file from which you want to restore the Exalogic Control stack.

4. Untar the backup files for each of the Exalogic Control components to a directory of your choice.
5. Copy the `vm.cfg` files identified earlier to the `/OVS/Repositories/*/VirtualMachines/`*`vmGUID`* directory.

   Note that `vmGUID` is name of the directory in which the `vm.cfg` file is located. For example, the following was the previously identified location of the `ExalogicControlDB` virtual machine:

   ```
   0004fb000006000038a23e7d5c307e00/vm.cfg:OVM_simple_name = 'ExalogicControlDB'
   ```

   In this example, `0004fb000006000038a23e7d5c307e00` is the `vmGUID`.
6. Copy the virtual-disk image file identified earlier to the location indicated by the `disk` parameter in the `vm.cfg` file
7. Start the Exalogic Control virtual machines as described in Section 4.1, "Recovering from a Hardware Failure" and the Exalogic Control processes as detailed in Step 4: Start the Enterprise Manager Ops Center Control VMs.

6

Backup and Recovery of Customer vServers

This chapter provides the steps for backing up and recovering customer vServers. It contains the following sections:

- Section 6.1, "Backing Up Customer vServers"
- Section 6.2, "Restoring Customer vServers"
- Section 6.3, "Re-creating vServers"

6.1 Backing Up Customer vServers

All the artifacts for the customer vServers are stored in the `ExalogicRepo` share on the ZFS storage appliance. These artifacts can be backed up by either creating ZFS snapshots as described in Section 5.2.3, "Creating a ZFS Snapshot of the Exalogic Control repository," or by performing a full backup to an external storage device using your existing backup strategy (for example: agent-based backup, NDMP, or ZFS replication).

For more information, see the *Exalogic Backup and Recovery Best Practices White Paper* at `http://www.oracle.com/technetwork/database/features/availability/maa-exalogic-br-1529241.pdf`.

6.2 Restoring Customer vServers

A customer vServer can be recovered either from a ZFS snapshot or from the latest full backup.

> **Note:** You can only restore a vServer if the vServer exists in Enterprise Manager Ops Center. If you deleted the vServer in Enterprise Manager Ops Center, you cannot restore the vServer.

6.2.1 Restoring a Customer vServer from a ZFS Snapshot

To restore a customer vServer from a ZFS snapshot, do the following:

1. Log in to the BUI of the ZFS storage appliance at `http://`*`storage_host`*`:215` as the `root` user.
2. Navigate to **Shares** and then to **Projects**, and select the **ExalogicControl** project.
3. Under the ExalogicControl project, select the **ExalogicRepo** share.
4. Navigate to Snapshots, select the required snapshot, and clone the snapshot.

Note: If the vServer you are restoring hosts a critical application, ensure that you are using a snapshot which was taken when either the application was stopped or the virtual machine was shutdown.

5. Select the location of the backup project, as defined in Chapter 2, "Backup and Recovery Locations."
6. Enter a name for the clone (example: `RepoClone_date`).
7. Note the mount point for the clone.
8. Mount the clone on one of the compute nodes, as shown in the following example:

   ```
   mount -t storage_host:/export/Exalogic_Backup/RepoClone_Oct24 /backup
   ```

9. Go to the `/backup/VirtualMachines` directory.
10. Find the virtual machine configuration files for the customer vServer that you want to restore by using the following command:

    ```
    grep -i vmName /backup/VirtualMachines/*/vm.cfg
    ```

    Example:

    ```
    grep -i wls_vm /backup/VirtualMachines/*/vm.cfg
    0004fb000006000038a23e7e6d307e12/vm.cfg:OVM_simple_name = wls_vm
    ```

11. Find the location and the name of the virtual-disk image file for the customer vServer that you want to restore.

    The location and name of the virtual-disk image file are indicated by the `disk` parameter in the `vm.cfg` file, as shown in the following example:

    ```
    cat 0004fb000006000038a23e7e6d307e12/vm.cfg | grep file
    disk =
    ['file:/OVS/Repositories/0004fb00000300007d5117500d92ae54/VirtualDisks/0004fb00
    001200003d3b8b4058ee7682.img,hda,w'
    ]
    ```

12. Copy the `vm.cfg` files identified earlier to the `/OVS/Repositories/*/VirtualMachines/vmGUID` directory.

    Note that `vmGUID` is name of the directory in which the `vm.cfg` file is located. For example, the following was the previously identified location of the `wls-vm` vServer:

    ```
    0004fb000006000038a23e7e6d307e12/vm.cfg:OVM_simple_name = wls-vm
    ```

    In this example, `0004fb000006000038a23e7e6d307e12` is the `vmGUID`.

13. Copy the virtual-disk image file identified earlier, to the location indicated by the `disk` parameter in the `vm.cfg` file.
14. Log in to the Exalogic Control BUI, and start the customer vServer.

6.2.2 Restoring Customer vServers from a Full Backup

To restore customer vServers from a full backup, do the following:

Note: In some cases, after the vServer is recovered, it may be necessary to run the `fsck` file check utility to fix any file system inconsistencies.

1. Mount the backup location as defined in Section 5.2.4, "Creating a Full Backup of the Exalogic Control Artifacts."
2. Identify the backup file from which you want to restore the customer vServer.
3. Copy the backup files to a directory of your choice.
4. Go to the directory in which you extracted the backup files.
5. Copy the `vm.cfg` file to the `/OVS/Repositories/*/VirtualMachines/`*`vmGUID`* directory.

   Note that `vmGUID` is name of the directory in which the `vm.cfg` file is located. For example, the following was the previously identified location of the `wls-vm` vServer:

   ```
   0004fb000006000038a23e7e6d307e12/vm.cfg:OVM_simple_name = 'wls-vm'
   ```

   In this example, `0004fb000006000038a23e7e6d307e12` is the `vmGUID`.
6. Copy the virtual-disk image file to the location indicated by the `disk` parameter in the `vm.cfg` file
7. Log in to the Exalogic Control BUI, and start the customer vServers.

6.3 Re-creating vServers

To recover a vServer, do the following:

1. Log in to the Exalogic Control BUI as a Cloud User.
2. Identify the vServer that needs to be recovered.

   The vServer should already be shut down. If it is still running, stop it.
3. Copy the required configuration files of the vServer to a location of your choice.
4. Delete the vServer from the account.
5. If the vServer should be restored with the same IP address, log in to the Exalogic Control BUI as the `Cloud Admin` user, allocate the entire network, and select the address you want.

   For more information about allocating IP addresses, see the *Oracle Exalogic Cloud Administrator's Guide* at `http://docs.oracle.com/cd/E18476_01/doc.220/e25258/toc.htm`.
6. Create a vServer as described in the *Oracle Exalogic Cloud Administrator's Guide*.

   Note: Before creating a vServer, make sure that the required volumes, networks, and so on exist in the environment.

7. After the vServer is created, mount the NFS directory containing the customer vServer backups and copy the configurations files you backed up in step 3 to the vServer.