Oracle® Exalogic Elastic Cloud

Backup and Recovery Guide Using ExaBR

Release 1.2

E36329-16

April 2016

This document describes how to back up and recover Exalogic components by using ExaBR, a tool that automates the process of backing up and recovering Exalogic data.

ORACLE

Oracle Exalogic Elastic Cloud Backup and Recovery Guide Using ExaBR, Release 1.2

E36329-16

Primary Author: Ashish Thomas

Contributing Authors: Scott Balfour, Neeraj Gupta, Rachid Benkreira, Jeremy Hoyland, Ivan Primorac, Ari Shapiro, Anoop Madhavan, Jeremy Bar, Ajit Kamble

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 that have purchased support 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

ExaBR is an Exalogic tool that allows you to automate the backup process for the Oracle Exalogic system configuration and data.

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 and disaster
- 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 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 using the ExaBR tool to back up the Exalogic system to the storage appliance.

1.2 Overview of ExaBR

You can use ExaBR to automate the process of backup and recovery of Exalogic data.

You can use ExaBR to back up and recover the following components:

- Compute nodes
- InfiniBand switches
- Ethernet management switches
- Storage appliance heads
- Exalogic Control stack
- User vServers

1.3 Supported Platforms

For information about the supported platforms, see the My Oracle Support document ID 1912063.1.

Note: Before running ExaBR for an Exalogic rack that was upgraded to EECS 2.0.6.0.0, you must first synchronize the ECU configuration files with the current configuration of the machine, by running an ECU converter. The ECU converter is a tool that is included with ExaPatch. For more information about the ECU converter, see the ExaPatch User's Guide.

1.4 Known Issues

For information about known issues for ExaBR, see the My Oracle Support document ID 1912063.1.

1.5 Scope

This document covers using the ExaBR tool to back up and restore the configuration of the Exalogic infrastructure components, and backing up and restoring the management components (the Exalogic Control Stack) of the cloud infrastructure and user vServers when the Exalogic rack is deployed in a virtual configuration.

The Exalogic Control Stack in the Exalogic Elastic Cloud Software 2.0.6 release consists of Oracle VM Manager, Oracle Exalogic Control and their repositories deployed to an Oracle Database.

1.6 Backup and Recovery Recommendations

This section provides general backup and recovery recommendations for Exalogic.

- It is strongly recommended that regular backups of your data be scheduled. The backup schedule should be based on the nature of your data.
- These backups should also be copied to an external ZFS Storage appliance or tape on a schedule, based on the nature of your data.
- 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

Using ExaBR to Back Up and Restore Exalogic Data

This chapter describes how to download, configure, and use ExaBR to back up and recover Exalogic.

This chapter contains the following sections:

2.1 Backup Locations

This section describes the backup locations for ExaBR.

2.1.1 Shares Created for the Exalogic Lifecycle Toolkit

When you install the Exalogic Lifecycle Toolkit, the installer creates the following shares, in the `common` project, on the storage appliance:

- `exalogic-lcdata`
- `exalogic-lctools`

The installer mounts the shares in the root directory (`/`) of the compute node from which the installer is run. By default, ExaBR stores backups in the `/exalogic-lcdata/backups` directory. You can copy these backups to an external storage device or tape.

> **Note:** Backups of virtual components, such as the Exalogic Control stack and user vServers created using an EECS 2.0.4 guest base template or earlier, are not stored in the directory you specify. They are stored separately as snapshots, on the storage appliance.

2.1.2 Backup Directories Created by ExaBR

ExaBR creates the following directories in the backup location you specify to back up and recover Exalogic data:

Directory Name	Contents
compute_nodes	OS backups of compute nodes.
ib_gw_switches	Configuration backups of the InfiniBand gateway switches.
ib_spine_switches	Configuration backups of the InfiniBand spine switches.
management_switches	Configuration backups of the management switch.
iloms	Configuration backups of the component ILOMs.
storage_nodes	Configuration backups of the storage nodes.

2.2 Preparing to Use ExaBR

You must complete the following tasks before using ExaBR:

- Task 1, "Installing ExaBR"
- Task 2, "Preparing the ExaBR Configuration File"

Task 1 Installing ExaBR

ExaBR is packaged as a part of the Exalogic Lifecycle Toolkit. To obtain the latest release of the Exalogic Lifecycle Toolkit, you can download the toolkit installer and tar bundle from the My Oracle Support document ID 1912063.1. The My Oracle Support document also contains instructions for installing the toolkit.

The toolkit installer creates the `exalogic-lcdata` and `exalogic-lctools` shares in the `common` project. The toolkit installer mounts these shares in the root directory (`/`) of the compute node from which the installer script is running.

Note: Always run ExaBR on a compute node. In environments that are not STIG-hardened, ExaBR can take backups of all components when run from a compute node. For STIG-hardened compute nodes, ExaBR must be run locally on each compute node to back it up.

Task 2 Preparing the ExaBR Configuration File

Before running ExaBR for the first time, you must generate a configuration file called `exabr.config`. This configuration file contains the host names of the components you wish to back up or recover.

Prepare the ExaBR configuration file by doing the following:

1. Navigate to the `/exalogic-lctools/bin` directory:

   ```
   cd /exalogic-lctools/bin
   ```

2. Generate the ExaBR configuration file by running the `init` command:

   ```
   exabr init address_of_the_first_compute_node
   ```

   In this command, `address_of_the_first_compute_node` is the host name or IP address of the **first** compute node in your Exalogic rack or the compute node on which Exalogic Configuration Utility was run.

Note: In STIG-hardened environments, run ExaBR as an administrator user with `sudoer` permissions as follows:

```
sudo exabr init local_address
```

The `init` command discovers the components of the Exalogic rack and creates the `exabr.config` file in the `/exalogic-lcdata/backups` directory.

Note: If required, you can manually create the `exabr.config` file.

3. Depending on the environment of your Exalogic rack, manually add the host names or IP addresses of the following components:

   - For a **physical environment**, manually add the host names or IP addresses of the following components, by editing the `exabr.config` file.

Table 2–1 Physical Components to Manually Add

Component	Parameter in exabr.config
Management switch	`management_switches`
NM2 36P switch (spine switch)	`ib_spine_switches`
Storage ILOMs	`storage_nodes_iloms`
Solaris only: Solaris zones	`compute_nodes`

   - For a **virtual environment**, manually add to the `user_vservers` parameter the host names or IP addresses of guest vServers created using EECS 2.0.6 Guest Base Template, by editing the `exabr.config` file.

   For a sample configuration file, see Example 2–1.

Example 2–1 Sample ExaBR Configuration File

```
#
# Exabr configuration file.
# Created on: 2013/12/17 14:24
# Please edit the values if needed to match your environment
#

compute_nodes = cn1.example.com, cn2.example.com, cn3.example.com,
cn4.example.com, cn5.example.com, cn6.example.com, cn7.example.com

compute_nodes_iloms = cn1ilom.example.com, cn2ilom.example.com,
cn3ilom.example.com, cn4ilom.example.com, cn5ilom.example.com,
cn6ilom.example.com, cn7ilom.example.com

ib_gw_switches = ib01.example.com, ib02.example.com
ib_spine_switches = ibsp01.example.com

storage_nodes = sn01.example.com, sn02.example.com
storage_nodes_iloms = sn01ilom.example.com, sn02ilom.example.com

# Cisco switch
# For ssh, specify it in the form: user@hostname
management_switches = mgmt.example.com
# connection type=telnet or ssh
```

```
management_switches_connection_type = telnet

# Control stack
emoc = cn1-eoib1-vm011.example.oracle.com
ovmm = cn1-eoib1-vm011.example.oracle.com

# The following 2 entries are only used for 2.0.4.x control stack
proxy_controllers = pc1.example.com,pc2.example.com
db = db.example.com
# End entries for 2.0.4.x
# ######################################
# Add user VMs for backup in this section

user_vservers = 192.168.1.2


# ######################################

#
# how many backups to keep (per component)
# '0' indicates that no backups should be removed automatically
retention_count = 5
```

4. Review the `exabr.config` file to ensure that all components were discovered correctly.

Note: If the compute nodes cannot access Exalogic Control and Oracle VM Manager on the EoIB network, update the control stack addresses in the `exabr.config` file with the `IPoIB-admin` address:

1. Log in to the Exalogic Control BUI as the `root` user.
2. Expand **Servers**.
3. Expand the compute node running the Exalogic Control vServer.
4. Click the Exalogic Control vServer.
5. Click the **Network** tab.
6. Note the `IPoIB-admin` address of the Exalogic Control vServer.
7. Edit the `exabr.config` file.
8. Set the `emoc` and `ovmm` parameters to the `IPoIB-admin` address you noted.

2.3 Configuring ExaBR

This section describes the configuration settings for ExaBR. It contains the following topics:

- Section 2.3.1, "Enabling Key-Based Authentication for ExaBR"
- Section 2.3.2, "Configuring the Connection Protocol to the Management Switch"
- Section 2.3.3, "Configuring the Backup Retention Policy of ExaBR"
- Section 2.3.4, "Scheduling ExaBR Backups"

2.3.1 Enabling Key-Based Authentication for ExaBR

Before enabling key-based authentication for ExaBR, you must generate a private and public key pair using standard commands, such as `ssh-keygen`. You can set up

key-based authentication for SSH by using the `init-ssh` command in the following ways:

> **Note:** Running the `init-ssh` command is necessary only if you want to take non-interactive backups. If you do not run the `init-ssh` command, ExaBR will prompt for passwords when required.

- To set up key-based authentication for the specified components, use the `init-ssh` command followed by the host names of the components:

  ```
  exabr init-ssh hostname1[,hostname2,...]
  ```

- To set up key-based authentication for a set of components, use the `init-ssh` command as follows:

  ```
  exabr init-ssh all-component
  ```

  For a list of `all-component` targets, see Section 2.4.3, "ExaBR Targets."

You can remove key-based authentication for SSH by using the `remove-ssh` command.

> **Note:** If the private SSH key is password protected, use a standard command, such as `ssh-agent`, to load the key on the compute node running ExaBR.

Copy SSH Keys to the Exalogic Control Stack vServers (EECS 2.0.4 Racks Only)

To use key-based authentication when backing up the Exalogic Control stack on EECS 2.0.4 racks, you must manually copy the public key to all the Exalogic Control stack vServers:

1. Log in to the compute node on which you are running ExaBR.
2. Copy the public keys to the Exalogic Control stack vServers by running the `ssh-copy-id` command as follows:

   ```
   ssh-copy-id -i path_to_public_key control_vServer_IP_Address
   ```

   Examples:

   ```
   ssh-copy-id -i ~/.ssh/id_dsa.pub pc01.example.com
   ssh-copy-id -i ~/.ssh/id_dsa.pub pc02.example.com
   ssh-copy-id -i ~/.ssh/id_dsa.pub db.example.com
   ssh-copy-id -i ~/.ssh/id_dsa.pub ec.example.com
   ssh-copy-id -i ~/.ssh/id_dsa.pub ovmm.example.com
   ```

   In these examples, the public key is copied from the compute node to all of the Exalogic Control stack vServers.

2.3.2 Configuring the Connection Protocol to the Management Switch

By default, ExaBR connects to the management switch using the `telnet` protocol. You can change the connection protocol to `SSH` by editing the `management_switches_connection_type` parameter in the `exabr.config` file.

Example:

```
management_switches_connection_type = ssh
```

If the connection protocol is set to SSH, you can also specify the user name in the `management_switches` parameter as follows:

```
management_switches = user@switch_hostname
```

Example:

```
management_switches = cisco@mgmt.example.com
```

Note: The firmware version of the switch that ships with Exalogic does not support SSH. To enable SSH, the switch firmware must be updated. For more information, see the My Oracle Support document ID 1912063.1.

2.3.3 Configuring the Backup Retention Policy of ExaBR

By default, ExaBR retains the last five successful backups for each component. You can change this setting by editing the `retention_count` parameter in the `exabr.config` file that you prepared in Task 2, "Preparing the ExaBR Configuration File". The retention policy applies to each component and is invoked after every successful and failed back up.

The retention policy is separately applied to successful and failed backups. A failed backup is retained, but only counts towards the retention count for failed backups. This applies to successful backups too. Oracle recommends manually deleting failed backups to save space on the storage appliance.

Example:

```
retention_count = 3
```

Note: The retention policy applies to infrastructure component backups such as compute nodes, switches, and so on, as well as to snapshot-based backups of the Exalogic Control stack and user vServers.

2.3.4 Scheduling ExaBR Backups

You can schedule ExaBR backups using `cron` or any other scheduler. Before you schedule ExaBR backups, note the following:

- If you use `cron` to schedule backups, ensure that ExaBR does not output to the console by adding `&>/dev/null` to the end of each ExaBR command.
- If you use `cron` to schedule backups of the Exalogic Control Stack, the `cron` job should run a script that contains all the necessary commands. Do not use `cron` to directly run each command because stopping and starting the Exalogic Control Stack can take a few minutes.
- Use the `init-ssh` command to enable public key SSH authentication for all targets. If public key SSH is not enabled for a target, it is not backed up. For more information on `init-ssh`, see Section 2.4.1, "ExaBR Commands."
- To ensure that users are not prompted for passwords during a scheduled backup, run ExaBR using the `--noprompt` option. For more information on `--noprompt`, see

Section 2.4.2, "ExaBR Options."

- The management switch cannot be backed up in a scheduled backup, because the management switch does not support non-interactive back ups. You must back up the management switch interactively, as described in Section 3.4.1, "Backing Up the Management Switch."

Note: To schedule backups of STIG-hardened components, you must additionally perform the following prerequisites:

- Run the scheduler you use (such as `cron`) as a privileged user.
- Since you must run the scheduler as a privileged user, the `sudo` command is not required when running ExaBR commands non-interactively as follows:

  ```
  ./exabr backup local_address [options]
  ```

2.4 Using ExaBR

Ensure that you have installed and prepared ExaBR by following all the steps in Section 2.2, "Preparing to Use ExaBR." You can run ExaBR either on specified components or on all components of a specified type.

- Using ExaBR on **specified components**: To run ExaBR on the specified components, use the following command:

  ```
  ./exabr command hostname1[,hostname2,...] [options]
  ```

- Using ExaBR on **particular types of components**: To run ExaBR on all components of a particular type, for example on all compute nodes, run ExaBR using a target as follows:

  ```
  ./exabr command target [options]
  ```

- Using ExaBR on **STIG-hardened components**: In a STIG-hardened environment, note the following:
  - ExaBR supports only STIG-hardened Linux compute nodes. ExaBR does not support STIG-hardened user vServers
  - Run ExaBR locally on each component you want to back up or restore.
  - Mount the ELLC shares on each component you want to back up or restore. You can copy the ELLC installer to the node and use the `-m` option to mount the shares as follows:

    ```
    # ./exalogic-lctools-version_number-installer.sh ZFS_Address -m
    ```

  - Run ExaBR as an administrator user with `sudoer` permissions.

  To run ExaBR in a STIG-hardened environment, run ExaBR as an administrator user by using the `sudo` command as follows:

  ```
  sudo exabr command hostname1[,hostname2,...] [options]
  ```

  or

  ```
  sudo exabr command target [options]
  ```

For a list of ExaBR commands, see Section 2.4.1, "ExaBR Commands."

For a list of ExaBR options, see Section 2.4.2, "ExaBR Options."

For a list of ExaBR targets, see Section 2.4.3, "ExaBR Targets."

2.4.1 ExaBR Commands

You can use the following commands when you run ExaBR:

Table 2–2 ExaBR Commands

Command	Description
`backup`	Performs a backup. For more information on performing a backup, see the relevant back up section.
`configure`	Runs the necessary Exalogic Configuration Utility steps to configure a compute node after the compute node has been replaced and re-imaged.
`control-register`	Rediscover the specified component and adds it to the list of assets in Exalogic Control.
`control-unregister`	Removes the specified component from the list of assets in Exalogic Control.
`ib-register`	Registers a component with the InfiniBand fabric after the component has been physically replaced. You must use this command when you replace a compute node or InfiniBand switch as described in Chapter 3, "Backup and Recovery of Infrastructure Components."
`init`	Creates a configuration file for an Exalogic rack. For more information, see Task 2, "Preparing the ExaBR Configuration File".
`init-ssh`	Sets up key-based authentication with a component. For more information, see Section 2.3.1, "Enabling Key-Based Authentication for ExaBR."
`remove-ssh`	Removes key-based authentication with a component.
`restore`	Restores from a backup. For more information on performing a restore, see the relevant restore chapter.
`start control-stack`	Starts the Exalogic Control stack.
`stop control-stack`	Stops the Exalogic Control stack.
`list`	Lists backups in the backup directory. For more information, see Section 2.5.3, "Listing Backups."

2.4.2 ExaBR Options

You can use the following options when you run ExaBR:

Table 2–3 ExaBR Options

Option	Description	Applicable to Command
`-b` or `--backup`	Directs ExaBR to restore from a specified backup directory. By default, ExaBR uses the most recent backup for restoration. **Example:** `./exabr restore cn1.example.com -b` *`201303201409`* `./exabr restore cn1.example.com --backup` *`201303201409`* `./exabr restore control-stack -b` *`exabr_201309111440_control`*	`restore`

Table 2–3 (Cont.) ExaBR Options

Option	Description	Applicable to Command

`--exclude-paths`	Specifies the directories to exclude when backing up or restoring compute nodes and vServers created using EECS 2.0.6 Guest Base Template.
	This option overrides the default exclusion list, but ExaBR always excludes NFS,HSFS, and OCFS2 mounts.
	`./exabr backup all-cn --exclude-paths directory1[,directory2,...]`
	When you use this option, you can use those wild cards that are supported by the `tar --exclude` command. You can use wild cards such as `?,*,[,],\`. In the following example, ExaBR excludes all directories whose names start with wls.
	`./exabr backup all-cn --exclude-paths /wls*`
		`backup` and `restore`

`--dry-run`	The command displays all the operations the `ib-register` command will run, but does not save the changes. Oracle recommends that you first run the `ib-register` command with `--dry-run` to ensure that there are no issues with registering the component on the InfiniBand fabric.	`ib-register`
`-h` or `--help`	Displays commands you can use when running ExaBR.	None

`--include-paths`	Specifies the directories to back up or restore on compute nodes or vServers created using EECS 2.0.6 Guest Base Template. Only the directories specified are backed up or restored. All other directories are excluded.
	`./exabr backup all-cn --include-paths directory1[,directory2,...]`
	When you use this option, you can use those wild cards that are supported by the `tar --exclude` command. You can use wild cards such as `?,*,[,],\`. In the following example, ExaBR includes all files with the extension `.log` in the `log` directory.
	`./exabr backup all-cn --include-paths /log/*.log`
		`backup` and `restore`

`-n` or `--noprompt`	Disables prompts for passwords. This option can be used only when key-based authentication has been enabled. For more information, see Section 2.3.1, "Enabling Key-Based Authentication for ExaBR." Use this option if you want to schedule backups using `cron` or any other scheduler. If this option is used for a component for which key-based authentication is not available, ExaBR does not back up the component.	`backup` for all targets except the management switch

Table 2–3 (Cont.) ExaBR Options

Option	Description	Applicable to Command

`-r` or `--repository`	Specifies the backup location. By default, ExaBR backs up to and restores from the `/exalogic-lcdata/backups` directory.
	■ To specify the backup location for your **current session**, set the `EXABR_REPO` environment variable by using the `export` command:
	`export EXABR_REPO=`*`path_to_mounted_backup_directory`*
	You can verify if the `EXABR_REPO` environment variable has been set, by running the following command:
	`echo $EXABR_REPO`
	■ To specify the backup location **each time** you run ExaBR, use one of the following options:
	The `-r` option:
	`exabr` *`command`* `-r` *`path_to_mounted_backup_directory`*
	Example:
	`exabr backup all-hw -r /custom_dir`
	The `--repository` option:
	`exabr` *`command`* `--repository=`*`path_to_mounted_backup_directory`*
	Example:
	`exabr backup all-hw --repository=/custom_dir`
		All

`--timeout`	Specifies the amount of time, in seconds, after which a command automatically fails if it has not completed. The default value is 2700 seconds.	`start` and `stop` of the Control Stack. `backup` and `restore` of compute nodes and EECS 2.0.6 Guest Base Template vServers.
`-v` or `--verbose`	Displays detailed results from the `list` command.	`list`
`--version`	Displays the version of ExaBR.	None

2.4.3 ExaBR Targets

You can use the following targets to run ExaBR on a set of components:

Note: Oracle recommends that the `all-`*`component`* targets be used only for backup operations. When restoring an Exalogic rack, restore each component individually.

Table 2–4 ExaBR Targets

Target	Components
`all-cn`	All compute nodes and their ILOMs
`all-ib`	All InfiniBand gateway and spine switches and their ILOMs
`all-ilom`	All ILOMs
`all-mgmt`	Ethernet management switch
`all-sn`	All storage appliance heads and their ILOMs
`all-hw`	All hardware components and their ILOMs
`all-itemized-vms`	All user vServers listed in the `exabr.config` file
`control-stack`	The Exalogic Control stack
`user-vm`	All user vServers

2.5 Managing ExaBR Backups

After you run ExaBR, the backups are stored in the `/exalogic-lcdata/backups` directory or directory you specify using the `-r` option.

2.5.1 ExaBR Log File

The ExaBR log file is called `exabr.log`. This file contains the information that is displayed when ExaBR runs in the interactive mode. By default, it is stored in the home directory of the user running ExaBR. You can override the location of this file using the `EXABR_LOG` environment variable.

2.5.2 Status of Backups

After each backup is taken, ExaBR stores the status of the backup in the `backup.info` file located in the backup directory. This file includes errors, if any, along with a description of any errors. If the file does not exist, the backup has failed.

The results of the backup are stored in the `backup.info` file in the backup directory. The following is an example of the `backup.info` log file:

```
[info]
status = OK
description =
command = backup component_host_name
details =
date = 2013-04-02 00:23:24
target = component_host_name
```

2.5.3 Listing Backups

You can view all backups in your backup directory by using the `list` command.

- To list the backups of the specified components, run ExaBR as follows:

  ```
  ./exabr list hostname1[,hostname2,...] [options]
  ```

- To list the backups of a set of components, run ExaBR using an `all-`*component* target as follows:

  ```
  ./exabr list all-component [options]
  ```

Example:

```
./exabr list all-cn -v
```

In this example, ExaBR lists all backups of compute nodes in detail, as the verbose (`-v`) option is used. When the `-v` option is used, ExaBR also validates the backup using `md5` checksums.

3

Backup and Recovery of Infrastructure Components

This chapter provides the steps for backing up and recovering the components of an Exalogic rack.

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

It contains the following sections:

3.1 Exalogic Configuration Utility

The Exalogic Configuration Utility (ECU) is used to configure an Exalogic rack during initial deployment. After the initial deployment is complete, Oracle recommends that you back up the configuration and the runtime files generated by the ECU.

> **Note:** ExaBR does not automate the backup of the ECU files.

3.1.1 Backing Up the ECU Files

You must perform the following steps manually to back up the configuration and runtime files generated by the ECU:

1. Log in to the node on which you installed the Exalogic Lifecycle Toolkit.
2. Create a directory called `ecu` in the `exalogic-lcdata/backups` directory. You should use this directory to manually store backups of the ECU files.
3. If the `ExalogicControl` share is not mounted in the `mnt/ExalogicControl` directory, mount it.
4. In the `ExalogicControl` directory, navigate to the `ECU_ARCHIVE` directory.

5. Copy the ECU file called `ecu_log-`*`date&time_stamp`*`.tgz` to the `ecu` directory you created in step 2. This file contains the following:
   - `ecu_run_time.tgz`: Contains tarball of the ECU configuration files
   - `ecu_home.tgz`: Contains tarball of the ECU scripts
   - `ecu_archive.tgz`: Contains tarball of the ECU log files

3.1.2 Recovering the ECU Files

Recover configuration and runtime files generated by the ECU, by doing the following:

1. Extract the tarball containing the configuration files to the `/opt/exalogic` directory.
2. Extract the tarball containing the runtime files to the `/var/tmp/exalogic` directory.
3. Extract the tarball containing the log files to the `/var/log/exalogic` directory.

3.2 Exalogic Compute Nodes

This section contains the following subsections:

3.2.1 Backing Up Compute Nodes

You can back up compute nodes using ExaBR as follows:

3.2.1.1 Backing Up Linux Compute Nodes

ExaBR can back up compute nodes running Linux using Logical Volume Manager (LVM)-based snapshots.

Note: To run ExaBR on STIG-hardened compute nodes, you must complete the following prerequisites:

- Run ExaBR locally on each compute node you want to back up.
- Run ExaBR with the `sudo` command. Alternatively, you can log in to the compute node and run `su -` to gain root access and run ExaBR commands without the `sudo` command.
- Mount the ELLC shares on each component you want to back up. You can copy the ELLC installer to the node and use the `-m` option to mount the shares as follows:

  ```
  # ./exalogic-lctools-version_number-installer.sh ZFS_Address -m
  ```

- Run ExaBR as a user with `sudoer` permissions. Alternatively, you can log in to the compute node and run `su -` to gain root access and run ExaBR commands without the `sudo` command.

To **automate** the backups of STIG-hardened Linux compute nodes, you must additionally perform the following prerequisites:

- Run the scheduler you use (such as `cron`) as a privileged user.
- Since you must run the scheduler as a privileged user, the `sudo` command is not required when running ExaBR commands non-interactively as follows:

  ```
  ./exabr backup local_address [options]
  ```

Overview

ExaBR backs up a Linux compute node using LVM-based snapshots by doing the following tasks:

1. Creates an LVM-based snapshot of the compute node file system.
2. Mounts the snapshot.
3. Creates a `tar` backup of the mounted file system from the snapshot.
4. Unmounts and deletes the snapshot.

When restoring LVM-based backups, ExaBR replaces files that were backed up. However, any files that were created after the backup was made are not deleted during a restore.

By default, ExaBR excludes the following directories when backing up the compute node:

- `/dev`
- `/proc`
- `/sys`
- `/tmp`
- `/var/tmp`
- `/var/run`
- `/var/lib/nfs`
- NFS and OCFS2 mounted file systems

You can use the `--exclude-paths` option to define your own exclusion list. The `--exclude-paths` option overrides the default exclusion list. The NFS and OCFS2 mounts, however, are automatically added to the exclusion list.

The `--include-paths` option can be used to back up only the specified directories. The `--include-paths` option is useful for backing up customizations. These options are described in Section 2.4.2, "ExaBR Options."

Prerequisites

The following are the prerequisites for using LVM-based snapshots to back up Linux compute nodes with ExaBR:

- When ExaBR takes the first backup, swap space should not be in use. During the first backup, ExaBR allocates 1 GB of the swap space for all subsequent LVM-based snapshots.

  If there is free memory on the compute node, you can free the swap space by logging in to the compute node as the `root` user and running the following commands:

  ```
  swapoff -av
  swapon -a
  ```

- The logical volume group that contains the root (`/`) volume should have at least 1 GB of unused space or more than 512 MB of swap space. Use the `vgs` command to display the amount of available free space (`VFree` column).
- The share you back up to must have enough free space for the extraction of the LVM-based snapshot.

Procedure

To back up compute nodes running Linux or the configuration of their ILOMs by doing the following:

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. Use ExaBR to back up compute nodes or the configuration of their ILOMs in one of the following ways:

   - To back up specific compute nodes or ILOMs, run ExaBR as follows:

     ```
     ./exabr backup hostname1[,hostname2,...] [options]
     ```

   - To back up all compute nodes and their ILOMs, run ExaBR using the `all-cn` target as follows:

     ```
     ./exabr backup all-cn [options]
     ```

   For a complete list of options, see Section 2.4.2, "ExaBR Options."

   Example:

   ```
   ./exabr backup cn1.example.com,cn1ilom.example.com --exclude-paths
   /var,/tmp,/sys,/proc,/dev,/test
   ```

   In this example, ExaBR backs up the first compute node and the configuration of its ILOM. It excludes the `/var`, `/tmp`, `/sys`, `/proc`, `/dev`, and `/test` directories, as specified by the `--exclude-paths` option, as well as the NFS and OCFS2 mounted file systems.

ExaBR stores the backup in the following files:

File	Description
cnode_backup.tgz	Contains the backup of the compute node.
ilom.backup	Contains the configuration backup of the ILOM.
guids.backup	Contains the InfiniBand GUIDs for the compute node.
version.backup	Contains the firmware version number of the compute node and ILOM.
backup.info	Contains the metadata of the backup of the compute node and ILOM.
checksums.md5	Contains the md5 checksums of the backup.

The files for compute nodes and ILOMs are stored in different directories. The files for compute nodes are stored in the compute_nodes directory and the files for ILOMs in the iloms directory, as described in Section 2.1.2, "Backup Directories Created by ExaBR."

By default, the backups are stored in the /exalogic-lcdata/backups directory.

3.2.1.2 Backing Up Solaris Compute Nodes

The operating system of an Exalogic rack is installed on the local disk of each compute node. ExaBR creates a backup of the root file system and the customizations, if any.

> **Note:** To back up Solaris zones using ExaBR, you must install gtar on the zone. For information on backing up Solaris zones, see Backing Up Solaris Zones.

By default, ExaBR excludes the following folders while backing up a Solaris compute node:

- /proc
- /system
- /tmp
- /var/tmp
- /var/run
- /dev
- /devices
- NFS and HSFS mounted file systems

You can use --exclude-paths option to define your own exclusion list. The --exclude-paths option overrides the default exclusion list. The NFS and HSFS mounts are automatically added to the exclusion list.

Note: For **Solaris** compute nodes and zones, apart from the directories listed in the default exclusion list, some files and directories cannot be backed up. In addition, some files and directory can be backed up by ExaBR, but cannot be restored to live Solaris file systems. The following are examples of such files and directories:

- `/home`
- `/etc/mnttab`
- `/etc/dfs/sharetab`
- `/etc/dev`
- `/etc/sysevent/devfsadm_event_channel/1`
- `/etc/sysevent/devfsadm_event_channel/reg_door`
- `/etc/sysevent/piclevent_door`
- `/lib/libc.so.1`
- `/net`
- `/nfs4`

For example, to back up the file system of a Solaris compute node, run the following command:

```
./exabr backup cn1.example.com --exclude-paths
/var,/tmp,/system,/proc,/dev,/devices,/home,/etc/mnttab,/etc/dfs/sh
aretab,/lib/libc.so.1,/net,/nfs4,/etc/sysevent/devfsadm_event_
channel/1,/etc/sysevent,/piclevent_door
```

This command backs up everything except the exclusion list and the additional files and directories listed in this note.

The `--include-paths` option can be used to only back up certain directories. The `--include-paths` option is useful for backing up customizations. These options are described in Section 2.4.2, "ExaBR Options."

To back up compute nodes running Solaris or the configuration of their ILOMs, do the following:

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. Use ExaBR to back up compute nodes or the configuration of their ILOMs in one of the following ways:
   - To back up specific compute nodes or ILOMs, run ExaBR as follows:

     ```
     ./exabr backup hostname1[,hostname2,...] [options]
     ```

   - To back up all compute nodes and their ILOMs, run ExaBR using the `all-cn` target as follows:

     ```
     ./exabr backup all-cn [options]
     ```

   For a complete list of options, see Section 2.4.2, "ExaBR Options."

   Example:

   ```
   ./exabr backup cn1.example.com, cn1ilom.example.com --exclude-paths
   /var,/tmp,/system,/proc,/dev,/devices,/home,/etc/mnttab,/etc/dfs/sharetab,/lib/
   libc.so.1,/net,/nfs4,/etc/sysevent/devfsadm_event_
   ```

```
channel/1,/etc/sysevent/piclevent_door
```

In this example, ExaBR backs up the first compute node and configuration of its ILOM. As the `--exclude-paths` option is used, it also excludes the directories specified in the previous note and NFS mounted file systems on the compute nodes. The `--exclude-paths` option overrides the default exclusion list.

ExaBR stores the backup in the following files:

File	Description
`cnode_backup.tgz`	Contains the backup of the compute node.
`ilom.backup`	Contains the configuration backup of the ILOM.
`guids.backup`	Contains the InfiniBand GUIDs for the compute node.
`version.backup`	Contains the firmware version number of the compute node and ILOM.
`backup.info`	Contains the metadata of the backup of the compute node and ILOM.
`checksums.md5`	Contains the `md5` checksums of the backup.

The files for compute nodes and ILOMs are stored in different directories. The files for compute nodes are stored in the `compute_nodes` directory and the files for ILOMs in the `iloms` directory, as described in Section 2.1.2, "Backup Directories Created by ExaBR."

By default, the backups are stored in the `/exalogic-lcdata/backups` directory.

Backing Up Solaris Zones

To back up Solaris zones, do the following:

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. Ensure that `gtar` is installed in the zone you want to back up.
3. Ensure that the IP address of the zone you want to back up has been added to the `compute_nodes` parameter of the `exabr.config` file.
4. Use ExaBR to back up zones in one of the following ways:
   - To back up specific zones, run ExaBR as follows:

     ```
     ./exabr backup zone1[,zone2,...] [options]
     ```

   - To back up only the global zone, run ExaBR while excluding the non-global zones with the `--exclude-paths` option as follows:

     ```
     ./exabr backup globalzone --exclude-paths path_to_nonglobalzone1,[path_to_
     nonglobalzone2,...][other_directories_to_exclude] [options]
     ```

   - To back up all the zones on a compute node, run ExaBR as follows:

     ```
     ./exabr backup HostnameOfComputeNode[options]
     ```

   For a complete list of options, see Section 2.4.2, "ExaBR Options."

   Example:

   ```
   ./exabr backup cn1.example.com --exclude-paths
   /var,/tmp,/system,/proc,/dev,/devices,/home,/etc/mnttab,/etc/dfs/sharetab,/lib/
   libc.so.1,/net,/nfs4,/etc/sysevent/devfsadm_event_
   channel/1,/etc/sysevent/piclevent_door
   ```

In this example, ExaBR backs up the compute node `cn1.example.com`. As the `--exclude-paths` option is used, it also excludes the directories specified in the previous note and NFS mounted file systems on the compute nodes. The `--exclude-paths` option overrides the default exclusion list.

3.2.1.3 Backing Up Customizations on Oracle VM Server Nodes

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

Note: Oracle VM Server nodes are stateless, so you do not need to back them up. You can use ExaBR to back up customizations on your Oracle VM Server node, by using the `--include-paths` option. You can reimage and recover an Oracle VM Server node using ExaBR, by following the steps described in Section 3.2.2.4, "Reimaging and Recovering Oracle VM Server Nodes in a Virtual Environment."

By default, ExaBR excludes the following folders while backing up the compute node:

- `/dev`
- `/proc`
- `/sys`
- `/tmp`
- `/var/tmp`
- `/var/run`
- `/var/lib/nfs`
- `poolfs`, `ExalogicPool`, `ExalogicRepo`, and NFS and OCFS2 mounted file systems. The `ExalogicPool` and the `ExalogicRepo` file systems are mounted over NFS.

You can use the `--exclude-paths` option to define your own exclusion list. The `--exclude-paths` option overrides the default exclusion list. The NFS and OCFS2 mounts, however, are automatically added to the exclusion list.

The `--include-paths` option can be used to back up only the specified directories. The `--include-paths` option is useful for backing up customizations. These options are described in Section 2.4.2, "ExaBR Options."

To back up customizations on Oracle VM Server nodes or the configuration of their ILOMs, do the following:

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. Use ExaBR to back up customizations on Oracle VM Server nodes or back up the configuration of their ILOMs in one of the following ways:
   - To back up customizations on specific compute nodes, run ExaBR as follows:

     ```
     ./exabr backup hostname1[,hostname2,...] --include-paths path_to_
     customization1[,path_to_customization2...] [options]
     ```

   - To back up customizations on all compute nodes and back up their ILOMs, run ExaBR using the `all-cn` target as follows:

     ```
     ./exabr backup all-cn --include-paths path_to_customization1[,path_to_
     customization2...] [options]
     ```

For a complete list of options, see Section 2.4.2, "ExaBR Options."

Example:

```
./exabr backup cn1.example.com,cn1ilom.example.com --include-paths /custom
--exclude-paths /var,/tmp,/sys,/proc,/dev,/test
```

In this example, ExaBR backs up the `/custom` directory on the first compute node and also backs up the configuration of its ILOM. It excludes the `/var`, `/tmp`, `/sys`, `/proc`, `/dev`, `/test` directories, NFS and OCFS2 mounted file systems on the compute nodes as the `--exclude-paths` option is used. The `--exclude-paths` option overrides the default exclusion list.

ExaBR stores the backup in the following files:

File	Description
`cnode_backup.tgz`	Contains the backup of the compute node.
`ilom.backup`	Contains the configuration backup of the ILOM.
`guids.backup`	Contains the InfiniBand GUIDs for the compute node.
`version.backup`	Contains the firmware version number of the compute node and ILOM.
`backup.info`	Contains the metadata of the backup of the compute node and ILOM.
`checksums.md5`	Contains the `md5` checksums of the backup.

The files for compute nodes and ILOMs are stored in different directories. The files for compute nodes are stored in the `compute_nodes` directory and the files for ILOMs in the `iloms` directory, as described in Section 2.1.2, "Backup Directories Created by ExaBR."

By default, the backups are stored in the `/exalogic-lcdata/backups` directory.

3.2.2 Recovering Compute Nodes

You can recover a compute node using ExaBR as follows:

- Section 3.2.2.1, "Recovering Compute Nodes in a Physical Environment"
- Section 3.2.2.2, "Recovering Compute Nodes on Failure of the InfiniBand HCA"
- Section 3.2.2.3, "Reimaging and Recovering Compute Nodes in a Physical Environment"
- Section 3.2.2.4, "Reimaging and Recovering Oracle VM Server Nodes in a Virtual Environment"

Note: To run ExaBR on STIG-hardened compute nodes, you must complete the following prerequisites:

- Run ExaBR locally on each compute node you want to restore.
- Mount the ELLC shares on each compute node you want to restore. You can copy the ELLC installer to the node and use the `-m` option to mount the shares as follows:

  ```
  # ./exalogic-lctools-version_number-installer.sh ZFS_Address -m
  ```

- Run ExaBR as an admin user with `sudoer` permissions.

3.2.2.1 Recovering Compute Nodes in a Physical Environment

ExaBR restores the backed up directories to the compute node. By default, ExaBR excludes the following directories when restoring a backup:

Table 3–1 Default Excluded Directories by Platform

Platform	Directories Excluded From Restore
Linux	■ /boot/grub/grub.conf ■ /boot/grub/menu.lst ■ /boot/grub/stage2 ■ /etc/grub.conf
Solaris	■ /home ■ /etc/mnttab ■ /etc/dfs/sharetab ■ /etc/dev ■ /etc/sysevent/devfsadm_event_channel/1 ■ /etc/sysevent/devfsadm_event_channel/reg_door ■ /etc/sysevent/piclevent_door ■ /lib/libc.so.1 ■ /net ■ /nfs4

You can modify the list of directories that are excluded during a restore by using the `--exclude-paths` option as described in Table 2–3, " ExaBR Options".

Restore Exalogic compute nodes, by doing the following:

Note: If a compute node you want to restore has corrupted data in one of the following ways:

- Not accessible through SSH.
- Essential binaries such as `bash` and `tar` do not exist.

Reimage and recover the compute node as described in steps 2 and 3 of Section 3.2.2.3, "Reimaging and Recovering Compute Nodes in a Physical Environment."

---

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. View a list of backups by running ExaBR as follows:

   ```
   ./exabr list hostname [options]
   ```

   Example:

   ```
   ./exabr list cn2.example.com -v
   ```

   In this example, ExaBR lists the backups made for `cn2.example.com` in detail, because the `-v` option is used.
3. Restore a compute node or the configuration of an ILOM, by running ExaBR as follows:

   ```
   ./exabr restore hostname [options]
   ```

For a complete list of options, see Section 2.4.2, "ExaBR Options."

Example:

```
./exabr restore cn2.example.com -b 201308230428 --exclude-paths /temp
```

In this example, ExaBR restores the second compute node from the backup directory `201308230428`, while excluding the `/temp` directory from the restore.

Note: When restoring a Linux compute node, if you see the following error:

`"ERROR: directory_name Cannot change ownership to ..."`

When running the restore command, use the `--exclude-paths` option to exclude the directory displayed in the error.

For an error such as:

```
ERROR tar: /u01/common/general: Cannot change ownership to uid
1000, gid 54321:
```

Run the restore command as in the following example:

```
./exabr restore cn2.example.com -b 201308230428 --exclude-paths
/u01/common/general
```

You can use the `--include-paths` option to restore specified files or directories from a backup.

Note: If the compute node that is being recovered is the master node of the Exalogic rack, restore the Exalogic Configuration Utility (ECU) configuration files that you backed up, as described in Section 3.1, "Exalogic Configuration Utility.". The master node in an Exalogic rack is the node from which the ECU was run.

3.2.2.2 Recovering Compute Nodes on Failure of the InfiniBand HCA

To recover compute nodes after a failure of their InfiniBand Host Channel Adapters (HCA), do the following:

1. This step is **required only for a virtual environment**. Remove the compute node from the list of assets in Exalogic Control by running ExaBR as follows:

   ```
   ./exabr control-unregister hostname_of_node_with_failed_HCA
   ```

2. Replace the failed HCA by following the standard replacement procedure.

3. Register the new GUIDs of the compute node HCA with the InfiniBand partitions that the compute node was previously registered to by running ExaBR as follows:

   ```
   ./exabr ib-register hostname_of_node_with_replaced_HCA [options]
   ```

   `hostname_of_node_with_replaced_HCA` is the IP address or hostname of the compute node whose failed HCA you replaced in the previous step. InfiniBand port GUIDs are saved during the compute node back up. The `ib-register` command searches the ExaBR backup directory for old GUID values. If no GUIDs are found, you are prompted to enter the GUIDs.

   Example:

```
./exabr ib-register cn2.example.com --dry-run
./exabr ib-register cn2.example.com
```

In the first example command, the `ib-register` command is run with the `--dry-run` option, so ExaBR displays the operations that are run without saving the changes.

In the second example command, ExaBR registers the GUIDs of the HCA card to the InfiniBand partitions.

Note: The `ib-register` command performs operations such as partition registration and vNIC registration on the InfiniBand switches. For **Solaris** compute nodes with a replaced HCA, you should perform additional operations on the compute node, such as a restart of the network interface, configuration of link aggregation, and so on.

---

4. This step is **required only for a virtual environment**. Rediscover the compute node and add it to the list of assets in Exalogic Control using ExaBR by running the `control-register` command as follows:

   ```
   ./exabr control-register hostname_of_replaced_compute_node
   ```

3.2.2.3 Reimaging and Recovering Compute Nodes in a Physical Environment

A compute node should be reimaged when it has been irretrievably damaged, or when multiple disk failures cause local disk failure. A replaced compute node should also be reimaged. During the reimaging procedure, the other compute nodes in the Exalogic rack are available. After reimaging the compute node, you can use ExaBR to recover your compute node using a backup. You should manually restore any scripting, cron jobs, maintenance actions, and other customizations performed on top of the Exalogic base image.

To reimage and recover a compute node in a physical environment, do the following:

1. This step is only required if you are **replacing** the entire compute node.

   a. 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 from a running 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.

   b. Replace the failed compute node, by following the standard replacement procedure.

2. Download the Oracle Exalogic base image from Oracle Software Delivery Cloud and patch-set updates (PSUs) from My Oracle Support.

3. Reimage the compute node.

   The compute node can be reimaged either 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 reimage the compute node, log in to the ILOM of the compute node through SSH, set the `boot_device` to `pxe` and restart 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.example.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 imaging process is complete. 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

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

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

g. Select CD-ROM from the drop-down list and then click **Save**.

h. In the **Remote Control** tab, navigate to the **Remote Power Control** tab.

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

j. 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 prompts you to provide the network configuration.

4. Recover the compute node by running the steps described in Section 3.2.2.1, "Recovering Compute Nodes in a Physical Environment."

5. This step is only required if you **replaced** the entire compute node.

Note: Do not run this step if you are restoring a repaired or existing compute node.

---

a. Restore the configuration of the ILOM of the compute node by running ExaBR as follows:

Note: ExaBR cannot **restore** the ILOM of the compute node from which it is running. To restore the ILOM of that compute node, run ExaBR on a different compute node.

```
./exabr restore hostname_of_ILOM [options]
```

Example:

```
./exabr restore cn2ilom.example.com -b 201308230428
```

In this example, ExaBR restores the configuration of the ILOM of the second compute node from the backup directory `201308230428`.

b. Register the new GUIDs of the replacement compute node with the same InfiniBand partitions that the compute node was previously registered to, using ExaBR as follows:

```
./exabr ib-register hostname_of_replaced_compute_node [options]
```

InfiniBand port GUIDs are saved during the compute node back up. The `ib-register` command searches the ExaBR backup directory for the GUIDs of the replaced compute node. If no GUIDs are found, the user is prompted to enter the GUIDs. This command also creates new vNICs with the GUIDs of the replacement compute node.

Example:

```
./exabr ib-register cn2.example.com --dry-run
./exabr ib-register cn2.example.com
```

In the first example command, the `ib-register` command is run with the `--dry-run` option, so ExaBR displays the operations that are run without saving the changes.

In the second example, the GUIDs of the replacement compute node are registered to the InfiniBand partitions that were registered to the failed compute node. This command also creates vNICs with the GUIDs of the replacement compute node.

Note: The `ib-register` command performs operations, such as partition registration and vNIC registration on the InfiniBand switches. For replaced **Solaris** compute nodes, you should perform additional operations on the compute node, such as a restart of the network interface, configuration of link aggregation, and other tasks required by your system administrator.

3.2.2.4 Reimaging and Recovering Oracle VM Server Nodes in a Virtual Environment

An Oracle VM Server node (compute node in a virtual environment) should be reimaged when it has been irretrievably damaged, or when multiple disk failures cause local disk failure. A replaced Oracle VM Server node should also be reimaged. During the reimaging procedure, the other Oracle VM Server nodes in the Exalogic rack are available. You should restore any scripting, cron jobs, maintenance actions, and other customizations performed on top of the Exalogic base image.

To reimage and recover a Oracle VM Server node in a virtual environment, do the following:

1. If an Oracle VM Server node from which Exalogic Control stack vServers were running is down or powered off, you must migrate the vServers:

   a. If the Oracle VM Server node from which the Exalogic Control vServer was running is down or powered off, migrate the vServer by performing the steps

in Section A.1 on a running Oracle VM Server node.

b. If the Oracle VM Server node from which either of the Proxy Controllers were running is down or powered off, migrate the vServer by performing the steps in Section A.2 on a running Oracle VM Server node.

c. Stop the components of the Exalogic Control stack by running ExaBR as follows:

```
./exabr stop control-stack
```

This command stops the Proxy Controller 2, Proxy Controller 1, and Exalogic Control vServers, in that order.

d. Restart the components of the Exalogic Control stack by running ExaBR as follows:

```
./exabr start control-stack
```

This command starts the Exalogic Control, Proxy Controller 1, and Proxy Controller 2 vServers, in that order.

2. Remove the Oracle VM Server node from the list of assets in Exalogic Control using ExaBR as follows:

```
./exabr control-unregister hostname_of_node
```

3. Verify that the Oracle VM Server node and ILOM are not displayed in the **Assets** accordion by logging into Exalogic Control as a user with the Exalogic System Administrator role.

4. Verify that the Oracle VM Server node is not displayed in Oracle VM Manager, by doing the following:

a. Log in to Oracle VM Manager as the `root` user.

b. Click the **Servers and VMs** tab.

c. Expand **Server Pools**.

d. Under **Server Pools**, verify that the Oracle VM Server node you removed in step 2 is not displayed.

5. If required, you can now restart all non-HA vServers using Exalogic Control. HA-enabled vServers automatically restart, if an Oracle VM Server node fails.

6. Follow steps 1 to 3 from Section 3.2.2.3, "Reimaging and Recovering Compute Nodes in a Physical Environment."

Note: Oracle recommends that you use the same credentials for the replacement Oracle VM Server node as those on the replaced Oracle VM Server node.

Note: Oracle VM Server nodes are stateless, so you do not need to back up or restore them. If required, you can use ExaBR to back up customizations on your Oracle VM Server node, by using the `--include-paths` option.

---

7. In the `/var/tmp/exalogic/ecu/cnodes_current.json` file, update the following for the replaced compute node:

Note: To find the correct values to update `cnodes_current.json`, use the information in `/opt/exalogic/ecu/config/cnodes_target.json`.

- The host name of the node.
- The host name of the ILOM.
- The IP address of the node on the `IPoIB-default` network.
- The IP address of the node on the `eth-admin` network.

8. Configure the Oracle VM Server node by using ExaBR to run the required ECU steps as follows:

   ```
   ./exabr configure Eth_IP_Address_of_CN
   ```

   The Oracle VM Server node restarts.

9. After the Oracle VM Server node has restarted, run the `configure` command again to synchronize the clock on the Oracle VM Server node:

   ```
   ./exabr configure address_of_node
   ```

   Note: For `address_of_node`, use the address specified in the `exabr.config` file.

10. You can update the credentials for the discovery profile of the Oracle VM Server node and ILOM in Exalogic Control by doing the following:

    Note: Do this step only if:

    - You replaced the Oracle VM Server node.
    - You configured different credentials for the replacement Oracle VM Server node than the credentials of the replaced Oracle VM Server node.

    a. Log in to the Exalogic Control BUI.

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

    c. Enter the host name of the Oracle VM Server node in the search box and click **Search**.

       The IPMI and SSH credential entries for the Oracle VM Server node are displayed.

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

       i. Select the entry for the credentials and click **Edit Credentials**.

          The Update Credentials dialog box is displayed.

       ii. Enter your password in the **Password** and **Confirm Password** fields.

       iii. Click **Update**.

       iv. Repeat steps i to iii for each entry.

11. Patch the compute node to the version that the other Oracle VM Server nodes are at using either the PSU or the stand-alone patch that you used previously.

12. Rediscover the Oracle VM Server node and add it to the list of assets in Exalogic Control using ExaBR as follows:

    ```
    ./exabr control-register hostname_of_node
    ```

13. Register the new GUIDs of the Oracle VM Server node HCA with the InfiniBand partitions, that the Oracle VM Server node was previously registered to, using ExaBR as follows:

    Note: This step is only required if the Oracle VM Server node was replaced.

    ```
    ./exabr ib-register hostname_of_node [options]
    ```

    InfiniBand port GUIDs are saved during the Oracle VM Server node back up. The `ib-register` command searches the ExaBR backup directory for old GUID values. If no GUIDs are found, the user is prompted to enter the GUIDs.

    Example:

    ```
    ./exabr ib-register cn2.example.com --dry-run
    ./exabr ib-register cn2.example.com
    ```

    In the first example, because the `ib-register` command is run with the `--dry-run` option, ExaBR displays what operations will be run without saving the changes.

    In the second example, the GUIDs of the replacement Oracle VM Server node are registered to InfiniBand partitions that were registered with the failed Oracle VM Server node.

14. This step is required only for an Exalogic rack that was upgraded from EECS 2.0.4 to EECS 2.0.6. Remove the warning icon for the Oracle VM Server node you restored by doing the following:

    a. Log in to the Oracle VM Server node you restored.

    b. Edit the `/etc/sysconfig/o2cb` file.

    c. Set the `O2CB_ENABLED` parameter to `true`.

    ```
    O2CB_ENABLED=true
    ```

    d. Log in to Oracle VM Manager as the `admin` user.

    e. Click the **Servers and VMs** tab.

    f. Expand **Server Pools**.

    g. Expand the server pool which contains the Oracle VM Server node you restored.

    h. Under the server pool, click the Oracle VM Server node you restored.

    i. From the Perspective drop down box, select **Events**.

    j. Click **Acknowledge All**.

3.3 NM2 Gateway and 36p Switches

The InfiniBand switches are a core part of an Exalogic rack and the configurations of all the Infiniband switches must be backed up regularly.

This section contains the following subsections:

- Section 3.3.1, "Backing Up InfiniBand Switches"
- Section 3.3.2, "Recovering InfiniBand Switches"
- Section 3.3.3, "Replacing InfiniBand Switches in a Virtual Environment"

3.3.1 Backing Up InfiniBand Switches

When backing up the InfiniBand switches, you must back up all the switches in the InfiniBand fabric.

> **Note:** In **virtual environments**, do not back up the InfiniBand switches separately. When backing up the Exalogic Control stack, ExaBR also backs up the InfiniBand switches. To back up the InfiniBand switches in virtual environments, Oracle recommends backing up the Exalogic Control Stack as described in Section 4.1, "Backing Up the Exalogic Control Stack."

ExaBR backs up the InfiniBand data in the following files:

File	Description
`switch.backup`	Contains the configuration backup of the InfiniBand switch made by the built-in ILOM back up.
`partitions.current`	Backup of InfiniBand partitions.
`smnodes`	Backup of the IP addresses of the SM nodes.
`version.backup`	Contains the firmware version number of the switch.
`backup.info`	Contains metadata of the backup of the switch.
`checksums.md5`	Contains the `md5` checksums of the backup.
`bx.conf`	Backup of bridge configuration file. This file is backed up for only the NM2 gateway switches.
`bxm.conf`	Backup of additional gateway configuration file. This file is backed up for only the NM2 gateway switches.
`opensm.conf`	Backup of OpenSM configuration file. This file is backed up for only the NM2 gateway switches.

For the NM2 gateway switches, the following files are backed up only for verification:

- `/conf/bx.conf`
- `/conf/bxm.conf`
- `/etc/opensm/opensm.conf`

You can compare these files with those in the backup directory to ensure that the switch content, such as VNICs, VLANs, and so on, were backed up correctly. ExaBR restores the data in these files from the `switch.backup` file.

You can back up the InfiniBand switches, by doing the following:

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. Use ExaBR to back up the InfiniBand switches in one of the following ways:
   - To back up a specific InfiniBand switch, run ExaBR as follows:

```
./exabr backup hostname1[,hostname2,...] [options]
```

- To back up all the InfiniBand switches, run ExaBR using the `all-ib` target as follows:

  ```
  ./exabr backup all-ib [options]
  ```

  For a complete list of options, see Section 2.4.2, "ExaBR Options."

  Example:

  ```
  ./exabr backup all-ib --noprompt
  ```

  In this example, ExaBR backs up all InfiniBand switches without prompting for passwords because the `--noprompt` option is used. This option can only be used when key-based authentication has been enabled as described in Section 2.3.1, "Enabling Key-Based Authentication for ExaBR." This option can be used to schedule backups, as described in Section 2.3.4, "Scheduling ExaBR Backups."

  Backups must be created for all the switches in the fabric.

By default, the backups are stored in the `/exalogic-lcdata/backups` directory.

For security reasons, encrypted entries are removed from the backed up files. These entries include the following:

- `/SP/services/servicetag/password`
- `/SP/users/ilom-admin/password`
- `/SP/users/ilom-operator/password`

3.3.2 Recovering InfiniBand Switches

ExaBR restores the configuration of the switch using the built-in SP restore of the switch. Ensure that no configuration changes are made while performing the restore. Configuration changes include vServer creation and vNet creation.

Note: Restore only the switch that is not the SM master. If you want to restore the SM master, relocate the SM master as follows:

1. Log in to the switch that you want to restore.
2. Relocate the SM master by running the `disablesm` command on the SM master.
3. Verify that the SM master has relocated to a different gateway switch by running the `getmaster` command.
4. Once a different gateway switch has become master, re-enable the SM on the switch you want to restore by running the `enablesm` command.
5. Confirm that the switch you logged in to is not the SM master by running the `getmaster` command.

You can recover the InfiniBand switches by doing the following:

Note: To restore the InfiniBand switch, you must run ExaBR as the `root` user.

---

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."

2. Verify the smnodes list on the switch, by doing the following:

   a. Log in to the InfiniBand switch you are recovering as the `root` user.

   b. List the Subnet Manager nodes by running the following command:

   ```
   # smnodes list
   ```

   Note: If the `smnodes` commands do not work, you may need to run the following commands:

   ```
   # disablesm
   # enablesm
   ```

   c. If you want the switch to run Subnet Manager, ensure that the IP address of the switch is listed in the output of the `smnodes list` command. The list of nodes can be modified by using the following commands:

   ```
   # smnodes add IP_Address_of_Switch
   # smnodes delete IP_Address_of_Switch
   ```

   Where `IP_Address_of_Switch` is the IP address of the switch you want to run Subnet Manager.

   Note: On all the gateway switches running the Subnet Manager, ensure that the output of the `smnodes list` command is identical.

3. View a list of backups by running ExaBR as follows:

   ```
   ./exabr list hostname [options]
   ```

   Example:

   ```
   ./exabr list ib01.example.com -v
   ```

   In this example, ExaBR lists the backups for `ib01.example.com` in detail, because the `-v` option is used.

4. Use ExaBR to restore an InfiniBand switch as follows:

   ```
   ./exabr restore hostname [options]
   ```

   For a complete list of options, see Section 2.4.2, "ExaBR Options."

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

   Example:

   ```
   ./exabr restore ib01.example.com -b 201308230428
   ```

   In this example, ExaBR restores an InfiniBand switch. The data is restored from the backup directory `201308230428`, because the `-b` option is used.

   Note: ExaBR binds to port 80 to send the configuration backup to the InfiniBand ILOM over HTTP. If the machine from which you run ExaBR has a firewall, you must open HTTP port 80.

5. The InfiniBand partitions are not restored automatically. If you want to restore the partitions, do the following:

   Note: Oracle recommends that the partitions always be restored when the Exalogic rack is in a virtual environment.

   a. Review the `partitions.current` file that was backed up by ExaBR, to ensure that your partitions were backed up correctly.

   b. Copy `partitions.current` to the master switch by running the following command:

      ```
      scp mounted_exabr_backup/ib_gw_switches/switch_hostname/backup_
      directory/partitions.current ilom-admin@master_switch_hostname:/tmp
      ```

      In this command, `mounted_exabr_backup` is the path to the share you created for ExaBR, `switch_hostname` is the host name of the switch you backed up, `backup_directory` is the directory from which you are restoring the switch, and `master_switch_hostname` is the host name of your master switch.

   c. Log in to the master switch and copy `partitions.current` from the `/tmp` directory to the `/conf` directory.

   d. Run the following command on the master switch to propagate the partitions and commit the change:

      ```
      smpartition start && smpartition commit
      ```

3.3.3 Replacing InfiniBand Switches in a Virtual Environment

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

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. Remove the InfiniBand switch from the list of assets in Exalogic Control using ExaBR by running the `control-unregister` command as follows:

   ```
   ./exabr control-unregister hostname_of_IB_switch
   ```

3. Replace the failed switch by following the standard replacement procedure.

   Note: Before connecting the switch to the InfiniBand fabric, disable Subnet Manager by running the `disablesm` command on the switch.

4. Restore the switch from a backup you created, by running **all** steps in Section 3.3.2, "Recovering InfiniBand Switches."
5. Register the gateway port GUIDs with the EoIB partitions using ExaBR by running the `ib-register` command as follows:

   ```
   ./exabr ib-register hostname_of_IB_switch [options]
   ```

   Example:

   ```
   ./exabr ib-register ib02.example.com --dry-run
   ./exabr ib-register ib02.example.com
   ```

In the first example, ExaBR displays what operations will be run without saving the changes because the `ib-register` command is run with the `--dry-run` option.

In the second example, the gateway port GUIDs of the InfiniBand switch are registered to the EoIB partitions.

6. You can update the credentials of the discovery profile for the InfiniBand switch and ILOM in Exalogic Control by doing the following:

   Note: Only do this step if:

   - You replaced the InfiniBand switch.
   - You configured different credentials for the replacement switch node than the credentials of the replaced switch.

   a. Log in to the Exalogic Control BUI.

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

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

      The IPMI and SSH credential entries for the InfiniBand 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**.

7. Rediscover the switch and add it to the list of assets in Exalogic Control by running ExaBR with the `control-register` command as follows:

   ```
   ./exabr control-register hostname_of_IB_switch
   ```

3.4 Management Switch

The 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"
- Section 3.4.3, "Replacing the Management Switch in a Virtual Environment"

3.4.1 Backing Up the Management Switch

ExaBR uses the built-in mechanism of the management switch to back up the configuration data. You can back up the management switch, by doing the following:

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. Use ExaBR to back up the management switch as follows:

   ```
   ./exabr backup hostname [options]
   ```

Example:

```
./exabr backup mgmt.example.com
```

In this example, ExaBR backs up the management switch.

> **Note:** The management switch supports only interactive back ups. You will be prompted to enter both the login password and the password to enter privileged mode.

By default, the backups are stored in the `/exalogic-lcdata/backups` directory.

By default, ExaBR connects to the management switch using the `telnet` protocol. You can change the connection protocol to SSH, as described in Section 2.3.2, "Configuring the Connection Protocol to the Management Switch."

ExaBR stores the backup in the following files:

File	Description
`switch.backup`	Contains the backup of the management switch.
`backup.info`	Contains metadata of the backup.
`checksums.md5`	Contains the `md5` checksums of the backup.

3.4.2 Recovering the Management Switch

You can recover the management switch by doing the following:

> **Note:** Ensure that configuration changes are not made while performing the restore.

1. Perform the tasks steps described in Section 2.2, "Preparing to Use ExaBR."
2. Use ExaBR to restore the management switch as follows:

   ```
   ./exabr restore hostname [options]
   ```

   For a complete list of options, see Section 2.4.2, "ExaBR Options."

   Example:

   ```
   ./exabr restore mgmt.example.com
   ```

   In this example, ExaBR restores the management switch.
3. Verify that the switch was restored successfully.

   > **Note:** If the management switch does not run as expected, you can restart the switch to revert the restore.

4. If the switch is running correctly, do the following to make the restoration persist:
   a. Log in to the CLI of the management switch.
   b. Run the following command:

      ```
      copy running-config startup-config
      ```

By default, ExaBR connects to the management switch using the `telnet` protocol. You can change the connection protocol to SSH, as described in Section 2.3.2, "Configuring the Connection Protocol to the Management Switch."

3.4.3 Replacing the Management Switch in a Virtual Environment

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

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. Remove the management switch from the list of assets in Exalogic Control using ExaBR, by running the `control-unregister` command as follows:

   ```
   ./exabr control-unregister hostname_of_management_switch
   ```

3. Replace the failed management switch, by following the standard replacement process.
4. Perform the steps in Section 3.4.2, "Recovering the Management Switch" to restore the switch from the latest backup.
5. Rediscover the switch and add it to the list of assets in Exalogic Control using ExaBR by running `control-register` command as follows:

   ```
   ./exabr control-register hostname_of_management_switch
   ```

3.5 Storage Appliance Heads

The storage appliance in an Exalogic rack 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 when clustering is performed.

3.5.1 Backing Up the Configuration of the Storage Appliance Heads

To back up the configuration of the storage appliance heads or back up the configuration of their ILOMs, do the following:

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. Use ExaBR to back up the configuration of a storage appliance head or back up the configuration of their ILOMs in one of the following ways:

   - To back up the configuration of a specific storage appliance head or an ILOM, run ExaBR as follows:

     ```
     ./exabr backup hostname1[,hostname2,...] [options]
     ```

   - To back up the configuration of all storage appliance heads and back up their ILOMs, run ExaBR using the `all-sn` target as follows:

     ```
     ./exabr backup all-sn [options]
     ```

   For a complete list of options, see Section 2.4.2, "ExaBR Options."

   Example:

   ```
   ./exabr backup sn01.example.com --noprompt
   ./exabr backup sn01ilom.example.com --noprompt
   ```

In these examples, ExaBR backs up the configuration of the first storage appliance head and backs up the configuration of its ILOM without prompting for passwords. This option can only be used when key-based authentication has been enabled as described in Section 2.3.1, "Enabling Key-Based Authentication for ExaBR." This option can be used to schedule backups, as described in Section 2.3.4, "Scheduling ExaBR Backups."

For more information on the contents of the configuration backup taken by ExaBR for the storage appliance head, see the following link:

http://docs.oracle.com/cd/E22471_01/html/821-1792/maintenance__system__configurationbackup.html#maintenance__system__configurationbackup__backup_contents

ExaBR stores the backup in the following files:

File	Description
zfssa.backup	Contains the backup of the storage appliance head.
ilom.backup	Contains the configuration backup of the ILOM, if it was backed up.
checksums.md5	Contains the md5 checksums of the backup.
backup.info	Contains metadata of the backup of the storage appliance head and ILOM.

The files for storage appliance head and ILOMs are stored in different directories. The files for switches are stored in the storage_nodes directory and the files for ILOMs in the iloms directory as described in Section 2.1.2, "Backup Directories Created by ExaBR."

3.5.2 Recovering the Configuration of Storage Appliance Heads

To recover the configuration of the storage appliance head, do the following:

> **Note:** The restore process takes several minutes to complete and will impact service to clients, while the active networking configuration and data protocols are reconfigured. A configuration restore should be done only on a development system, or during a scheduled downtime.

1. Run the steps described in Section 2.2, "Preparing to Use ExaBR."
2. Identify the active and passive nodes by doing the following:
   a. Log in to the storage appliance BUI.
   b. Click the **Configuration** tab.

      The Configuration page is displayed.
   c. Click **Cluster**.

      The Cluster page is displayed as shown in Figure 3–1.

Figure 3–1 Cluster Page

d. The node described as `Active (takeover completed)` is the active node and the node described as `Ready (waiting for failback)` is the passive node. In this example, `slce23sn01` is the active node and `slce23sn02` is the passive node.

3. Perform a factory reset the **passive** node (`slce23sn02`) by doing the following:

   a. Log in to the storage appliance BUI of the passive node (`slce23sn02`) as the `root` user.

   b. Click the **Maintenance** tab.

   The Hardware page is displayed.

   c. Click **System**.

   The System page is displayed.

   d. Click the **FACTORY RESET** button.

   The storage node restarts.

4. Before the **passive** node (`slce23sn02`) restarts completely, you must uncluster the **active** node (`slce23sn01`) by doing the following:

   a. Log in to the storage appliance BUI of the active node (`slce23sn01`) as the `root` user.

   b. Click the **Configuration** tab.

   The Services page is displayed.

   c. Click **Cluster**.

   The Cluster page is displayed.

   d. Click the **UNCONFIG** button.

   e. Wait for the passive node (`slce23sn02`) to restart. To check if the passive node has restarted, log in to the ILOM of the passive node.

5. Restore the **active** node (`slce23sn01`) as it owns all the resources, by running ExaBR as follows:

   ```
   ./exabr restore hostname1 [options]
   ```

   For a complete list of options, see Section 2.4.2, "ExaBR Options."

   Example:

   ```
   ./exabr restore slce23sn01.example.com -b 201308230428
   ```

   In this example, ExaBR restores the active storage node, in our example `slce23sn01.example.com` and from the backup directory `201308230428`.

6. After the restoration is successful, reconfigure the clustering by doing the following:

   - For X2-2 and X3-2 racks:

a. Log in as `root` to the storage appliance BUI of the active node (`slce23sn01`).

b. Click the **Configuration** tab.

The Services page is displayed.

c. Click **Cluster**.

The Cluster page is displayed.

d. Click the **SETUP** button.

f. In the **host name** field, enter the host name of the passive node (`slce23sn02`).

g. In the **password** field, enter the password of the passive node (`slce23sn02`).

i. Click the **COMMIT** button.

The passive node joins the cluster and copies the configuration from the restored active node.

- For X4-2 and later racks:

a. Log in as `root` to the storage appliance BUI of the active node (`slce23sn01`).

b. Click the **Configuration** tab.

The Services page is displayed.

c. Click **Cluster**.

The Cluster page is displayed.

d. Click the **SETUP** button.

e. Click the **COMMIT** button.

f. In the **Appliance Name** field, enter the host name of the passive node (`slce23sn02`).

g. In the **Root Password** field, enter the password of the passive node (`slce23sn02`).

h. In the **Confirm Password** field, enter the password of the passive node (`slce23sn02`).

i. Click the **COMMIT** button.

The passive node should join the cluster and copy the configuration from the restored active node.

7. Configure the access to the ethernet management network for the storage node you restored (the active node `slce23sn01`) by doing the following:

- For X2-2 and X3-2 racks:

a. SSH to the active node as the `root` user.

b. Destroy the interface of the storage node you restored. In this example, destroy the `igb1` interface because you reset the second storage node (`slce23sn02`). If you reset the `slce23sn01` storage node, you must destroy the `igb0` interface.

```
configuration net interfaces destroy igb1
```

c. Re-create the interface you destroyed in the previous step and the route of the node you restored (the active node `slce23sn01`) by running the following commands:

```
configuration net interfaces ip
set v4addrs=ip address of passive node (slce23sn02) in CIDR notation
set links=igb1
set label=igb1
commit
configuration net routing create
set family=IPv4
set destination=0.0.0.0
set mask=0
set interface=igb1
set gateway=<gateway ip>
commit
```

d. Change the ownership of the `igb1` resource, by running the following commands:

```
configuration cluster resources select net/igb1
set owner=<inactive head host name>
commit
commit
The changes have been committed. Would you also like to fail back? (Y/N)
N
```

- For X4-2 and later racks:

  a. SSH to the active node as the `root` user.

  b. Destroy the interface of the storage node you restored. In this example, destroy the `ixgbe1` interface because you reset the second storage node (`slce23sn02`). If you reset the `slce23sn01` storage node, you must destroy the `ixgbe0` interface.

  ```
  configuration net interfaces destroy ixgbe1
  ```

  c. Re-create the interface you destroyed in the previous step and the route of the node you restored (the active node `slce23sn01`) by running the following commands:

  ```
  configuration net interfaces ip
  set v4addrs=ip address of passive node (slce23sn02) in CIDR notation
  set links=ixgbe1
  set label=ixgbe1
  commit
  configuration net routing create
  set family=IPv4
  set destination=0.0.0.0
  set mask=0
  set interface=ixgbe1
  set gateway=<gateway ip>
  commit
  ```

  d. Change the ownership of the `ixgbe1` resource, by running the following commands:

  ```
  configuration cluster resources select net/ixgbe1
  set owner=<inactive head host name>
  commit
  commit
  The changes have been committed. Would you also like to fail back? (Y/N)
  N
  ```

8. Perform a takeover to make the passive node (`slce23sn02`) the active node by doing the following:
   a. Log in to the ILOM of the passive node (`slce23sn02`) as the `root` user.
   b. Run the following command to start the console:

      ```
      start /SP/console
      ```

   c. Log in to the passive node (`slce23sn02`) as the `root` user.
   d. Perform a takeover to make the node active:

      ```
      configuration cluster takeover
      ```

      The active node (`slce23sn01`) restarts and becomes the passive node.
   e. Wait for the formerly active node (`slce23sn01`) to start up.
9. Set the resource you re-created to private by doing the following:
   - For X2-2 and X3-2 racks:

     a. Log in to the new active node (`slce23sn02`) as the `root` user.

     b. Run the following commands:

     ```
     configuration cluster resources select net/igb1
     set type=private
     commit
     commit
     The changes have been committed. Would you also like to fail back? (Y/N)
     N
     ```

   - For X4-2 and later racks:

     a. Log in to the new active node (`slce23sn02`) as the `root` user.

     b. Run the following commands:

     ```
     configuration cluster resources select net/ixgbe1
     set type=private
     commit
     commit
     The changes have been committed. Would you also like to fail back? (Y/N)
     N
     ```

     Note: If you want to make original head the active node, log in to the node as the `root` user and run step 8 d.

4

Backup and Recovery of the Exalogic Control Stack

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

4.1 Backing Up the Exalogic Control Stack

The Exalogic Control stack repository contains vServer disk images, templates, and vServer configuration files for all the vServers running on the Exalogic rack. The Exalogic Control repository resides in the `ExalogicRepo` share in the ZFS storage appliance. ExaBR takes a ZFS snapshot of this share to back up the Exalogic Control stack.

Note: To take non-interactive backups of the Exalogic Control Stack, enable key-based authentication for the compute nodes and storage heads using the `init-ssh` command. The usage of the `init-ssh` command is described in Section 2.3.1, "Enabling Key-Based Authentication for ExaBR." For EECS 2.0.4 racks, you must manually copy the keys to the Exalogic Control stack vServers as described in Copy SSH Keys to the Exalogic Control Stack vServers (EECS 2.0.4 Racks Only).

Oracle recommends that the components of the Exalogic Control stack and their repositories be backed up after all lifecycle operations, such as adding and removing accounts, users, and vServers.

Note: When you back up the Exalogic Control stack, ExaBR also backs up the InfiniBand switches. When you restore the Exalogic Control stack, you must also restore the corresponding InfiniBand switch backup.

To ensure data consistency, the vServers running the Exalogic Control stack must be shut down before a backup is taken.

You can back up the Exalogic Control stack, by doing the following:

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."

2. Stop the components of the Exalogic Control stack by running ExaBR as follows:

```
./exabr stop control-stack
```

> **Note:** Ensure that the Exalogic Control stack is not being used when you stop it.

This command stops the Proxy Controller 2, Proxy Controller 1, and Exalogic Control vServers, in that order.

3. Back up the Exalogic Control stack and InfiniBand switches, by running ExaBR as follows:

```
./exabr backup control-stack [options]
```

For a complete list of options, see Section 2.4.2, "ExaBR Options."

Example:

```
./exabr backup control-stack --noprompt
```

In this example, ExaBR backs up the Exalogic Control stack and InfiniBand switches without prompting for passwords, using the `--noprompt` option.

This command takes a snapshot of the `ExalogicRepo` share as a backup for the Exalogic Control stack. This command also backs up the InfiniBand switches as described in Section 3.3.1, "Backing Up InfiniBand Switches."

4. Restart the components of the Exalogic Control stack, by running ExaBR as follows:

```
./exabr start control-stack
```

This command starts the Exalogic Control, Proxy Controller 1, and Proxy Controller 2 vServers, in that order.

4.2 Restoring the Exalogic Control Stack

The Exalogic Control stack can be recovered to the point in time when the last backup was made by ExaBR. To ensure data consistency, all the Exalogic Control vServers are restored together.

To restore the Exalogic Control stack, do the following:

> **Caution:** Restoring the Exalogic Control stack reverts its data to the point in time when the backup was created. After you restore the Exalogic Control stack:
>
> - Any vServers that were created after the backup will not be visible in Exalogic Control.
> - Any vServers that were deleted after the backup will be visible in Exalogic Control, even though they do not exist.
>
> You can clean up such vServers by following the steps described in Appendix B.

You can restore the Exalogic Control stack by doing the following:

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. Stop the components of the Exalogic Control stack, by running ExaBR as follows:

   ```
   ./exabr stop control-stack
   ```

   Note: Ensure that the Exalogic Control stack is not being used when you stop it.

   This command stops the Proxy Controller 2, Proxy Controller 1, and Exalogic Control vServers, in that order.
3. View a list of backups by running ExaBR as follows:

   ```
   ./exabr list control-stack
   ```

   Example:

   ```
   ./exabr list control-stack
   ```

   In this example, ExaBR lists the backups made for the Exalogic Control stack. Note the name of the backup you want to use to restore the Exalogic Control stack and InfiniBand switch.
4. Restore the Exalogic Control stack by running ExaBR as follows:

   ```
   ./exabr restore control-stack [options]
   ```

   For a complete list of options, see Section 2.4.2, "ExaBR Options."

   Example:

   ```
   ./exabr restore control-stack -b exabr_201309041714_control
   ```

   In this example, ExaBR restores the control stack from the backup called `exabr_201309041714_control` using the `-b` option.

   Running this command does the following:

   a. Clones the snapshot of the `ExalogicRepo` share taken by ExaBR during the control stack back up.
   b. Identifies the Exalogic Control vServers in the snapshot.
   c. Restores the Exalogic Control vServers data from the clone to the Oracle VM server repository.
   d. Deletes the clone.

   Note: Restoration of the Exalogic Control stack can take at least one hour.
5. Wait until the restoration of the Exalogic Control stack is complete.
6. Restore each InfiniBand switch as described in Section 3.3.2, "Recovering InfiniBand Switches." Ensure that you restore the InfiniBand switch from the same backup you used to restore the Exalogic Control stack. In this example, the InfiniBand switch is restored from the backup directory `201309041714`.

Note: To ensure data consistency, you must restore the InfiniBand switches when restoring the Exalogic Control stack. You must restore the switch from the backup of the switch that was taken when backing up the Exalogic Control stack.

7. Restoring the Exalogic Control stack reverts its data to the point in time when the backup was created. After you restore the Exalogic Control stack:

   - Any vServers that were created after the backup will not be visible in Exalogic Control.
   - Any vServers that were deleted after the backup will be visible in Exalogic Control, even though they do not exist.

   You can clean up such vServers, by following the steps described in Appendix B, "Removing Orphan and Ghost vServers After Restoring the Exalogic Control Stack."

8. Restart the components of the Exalogic Control stack by running ExaBR as follows:

   ```
   ./exabr start control-stack
   ```

   This command starts the Exalogic Control, Proxy Controller 1, and Proxy Controller 2 vServers, in that order.

5

Backup and Recovery of User vServers

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

- Backing Up User vServers
- Restoring User vServers

5.1 Backing Up User vServers

You can back up user vServers using ExaBR as follows:

- Section 5.1.1, "Backing Up User vServers Created Using EECS 2.0.6 Guest Base Template"
- Section 5.1.2, "Backing Up User vServers Created Using EECS 2.0.4 Guest Base Template or Earlier"

5.1.1 Backing Up User vServers Created Using EECS 2.0.6 Guest Base Template

User vServers created using EECS 2.0.6 Guest Base Template support Logical Volume Manager (LVM)-based partitioning. ExaBR can back up these vServers using LVM-based snapshots. By default, these LVM snapshots stored in the `/exalogic-lcdata/backups` directory.

Note: ExaBR does not support STIG-hardened user vServers.

Overview

ExaBR backs up a user vServer using LVM-based snapshots, by doing the following tasks:

1. Creates an LVM-based snapshot of the vServer file system.
2. Mounts the snapshot.
3. Creates an archive of the mounted file system from the snapshot.
4. Unmounts and deletes the snapshot.

When restoring LVM-based backups, ExaBR replaces files that were backed up. However, any files that were created after the backup was made are not deleted during a restore

By default, ExaBR excludes the following directories when backing up user vServers:

- `/tmp`

- /var/tmp
- /var/run
- /var/lib/nfs

ExaBR backs up only files from the root (/) and /boot directories. It does not back up the NFS mounts or mounted volumes on a vServer, if any.

You can use the --exclude-paths option to define your own exclusion list. You can use the --include-paths option to back up only the specified directories. The --include-paths option is useful for backing up customizations. These options are described in Section 2.4.2, "ExaBR Options."

Prerequisites

The following are the prerequisites for using LVM-based snapshots to back up Linux compute nodes with ExaBR:

- When ExaBR takes the first backup, swap space should not be in use. During the first backup, ExaBR allocates 1 GB of the swap space for all subsequent LVM-based snapshots.

  If there is free memory on the user vServer, you can free the swap space by logging in to the user vServer as the root user and running the following commands:

  ```
  swapoff -av
  swapon -a
  ```

- The logical volume group that contains the root (/) volume should have at least 1 GB of unused space or more than 512 MB of swap space. Use the vgs command to display the amount of available free space (VFree column).
- The share you back up to must have enough free space for the extraction of the LVM-based snapshot.

Procedure

You can use ExaBR to take LVM-based snapshots as backups for these vServers, by doing the following:

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. Add the addresses of the vServers you want to back up to the exabr.config file under the user-vm parameter. You can use the virt-admin address or, if applicable, the public IP address of the vServer.
3. Use ExaBR to back up user vServers, in one of the following ways:
   - To back up specific user vServers, run ExaBR as follows:

     ```
     ./exabr backup address_of_vServer [options]
     ```

   - To back up all user vServers listed in the exabr.config file, run ExaBR as follows:

     ```
     ./exabr backup all-itemized-vms [options]
     ```

   For a complete list of options, see Section 2.4.2, "ExaBR Options."

   Example:

   ```
   ./exabr backup all-itemized-vms --exclude-paths
   /var/tmp,/var/lib/nfs,/tmp,/dev,/custom
   ```

In this example, ExaBR backs up all user vServers listed in the `exabr.config` to the backup location you specified. It also excludes the `/var/tmp`, `/var/lib/nfs`, `/tmp`, `/dev`, and `/custom` directories as the `--exclude-paths` option is used.

The `--exclude-paths` option overrides the default exclusion list.

By default, the backups are stored in the `/exalogic-lcdata/backups` directory.

To restore backups made using this procedure, see Section 5.2.1, "Restoring User vServers Created Using EECS 2.0.6 Guest Base Template."

5.1.2 Backing Up User vServers Created Using EECS 2.0.4 Guest Base Template or Earlier

All the artifacts for the user vServers are stored in the `ExalogicRepo` share on the ZFS storage appliance. These artifacts can be backed up either by using ExaBR to create ZFS snapshots or by performing a full back up to an external storage device using your existing back up strategy (for example: agent-based back up, NDMP, or ZFS replication). ExaBR backs up users vServers created using EECS 2.0.4 guest base template or earlier by taking a snapshot and storing it on the storage appliance. ExaBR does not store these backups in the backup directories as described in Section 2.1, "Backup Locations."

Note: To take non-interactive backups of user vServers, enable key-based authentication for the first compute node and storage heads using the `init-ssh` command. The usage of the `init-ssh` command is described in Section 2.3.1, "Enabling Key-Based Authentication for ExaBR."

Note: When you back up the Exalogic Control stack using ExaBR, it also backs up user vServers. If you have backed up the control stack, you do not need to run the following steps.

You can back up user vServers using ExaBR, by doing the following:

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. You can back up user vServers by running ExaBR as follows:

   ```
   ./exabr backup user-vm [options]
   ```

   For a complete list of options, see Section 2.4.2, "ExaBR Options."

   ExaBR backs up all user vServers by taking a snapshot. This operation backs up the local virtual disks of a vServer. It does not back up the NFS mounts of a vServer or mounted volumes on a vServer, if any.

   Note: If a vServer you are backing up hosts a critical application, ensure that you back up the vServer when the vServer is shut down.

   Example:

   ```
   ./exabr backup user-vm
   ```

In this example, ExaBR backs up all user vServers by taking a snapshot of the `ExalogicRepo` share.

To restore backups made using this procedure, see Section 5.2.2, "Restoring User vServers Created Using EECS 2.0.4 Guest Base Template or Earlier."

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`

5.2 Restoring User vServers

You can restore user vServers using ExaBR as follows:

- Section 5.2.1, "Restoring User vServers Created Using EECS 2.0.6 Guest Base Template"
- Section 5.2.2, "Restoring User vServers Created Using EECS 2.0.4 Guest Base Template or Earlier"

5.2.1 Restoring User vServers Created Using EECS 2.0.6 Guest Base Template

This section describes restoring vServers that you backed up using LVM snapshots, as described in Section 5.1.1, "Backing Up User vServers Created Using EECS 2.0.6 Guest Base Template."

When restoring LVM-based backups, ExaBR replaces files that were backed up. However, any files that were created after the backup was made are not deleted during a restore.

Restore user vServers that you backed up, by doing the following:

Note: You can restore a vServer only if the vServer exists in Exalogic Control. If you deleted the vServer in Exalogic Control, you cannot restore the vServer.

Note: ExaBR does not support STIG-hardened user vServers.

---

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. Ensure that the addresses of the vServers you want to restore exist in the `exabr.config` file in the `user-vm` parameter.
3. View a list of backups by running ExaBR as follows:

   ```
   ./exabr list all-itemized-vm
   ```

   Example:

   ```
   ./exabr list all-itemized-vm
   ```

   In this example, ExaBR lists the backups made for user vServers. Note the name of the backup you want to use to restore your user vServer.
4. If a user vServer you want to restore has corrupted data in one of the following ways:
   - Not accessible through SSH.

- Essential binaries such as `bash` and `tar` do not exist.

Do the following:

a. Log in to Exalogic Control.

b. Note the name, IP address, and vServer type of the vServer that has corrupted data.

c. Delete the vServer in Exalogic Control.

d. Recreate the vServer with the name, IP address, and vServer type you noted in step b.

5. You can restore user vServers by running ExaBR as follows:

- If the data of the user vServer was not corrupted, run ExaBR as follows:

```
./exabr restore address_of_vServer [options]
```

For a complete list of options, see Section 2.4.2, "ExaBR Options."

Examples:

```
./exabr restore 192.168.1.2 -b 201309121315
```

In the first example, ExaBR restores the user vServer with the address `192.168.1.2` from the backup called `201309121315`.

```
./exabr restore 192.168.1.2 -b 201309121315 --include-paths /home/Oracle
```

In the second example, ExaBR restores the `/home/Oracle` directory of the user vServer with the address `192.168.1.2` from the backup called `201309121315`.

- If the data of the user vServer was corrupted as described in step 4, run ExaBR as follows:

```
./exabr restore address_of_vServer --exclude-paths
/etc/sysconfig/network-scripts [options]
```

For a complete list of options, see Section 2.4.2, "ExaBR Options."

Note: When restoring a user vServer, if you see the following error:

```
"ERROR: directory_name Cannot change ownership to ..."
```

When running the restore command, use the `--exclude-paths` option to exclude the directory displayed in the error.

For an error such as:

```
ERROR tar: /u01/common/general: Cannot change ownership to uid
1000, gid 54321:
```

Run the restore command as in the following example:

```
./exabr restore 192.168.1.2 -b 201308230428 --exclude-paths
/u01/common/general
```

5.2.2 Restoring User vServers Created Using EECS 2.0.4 Guest Base Template or Earlier

Restore user vServers that you backed up in Section 5.1.2, "Backing Up User vServers Created Using EECS 2.0.4 Guest Base Template or Earlier", by doing the following:

Note: You can restore a vServer only if the vServer exists in Exalogic Control. If you deleted the vServer in Exalogic Control, you cannot restore the vServer.

1. Perform the tasks described in Section 2.2, "Preparing to Use ExaBR."
2. Use Exalogic Control to stop the vServer you want to restore.
3. View a list of backups by running ExaBR as follows:

   ```
   ./exabr list user-vm
   ```

   Example:

   ```
   ./exabr list user-vm
   ```

   In this example, ExaBR lists the backups made for user vServers. Note the name of the backup you want to use to restore your user vServer.
4. You can restore user vServers by running ExaBR as follows:

   ```
   ./exabr restore user-vm --vm vm_name [options]
   ```

   For a complete list of options, see Section 2.4.2, "ExaBR Options."

   Example:

   ```
   ./exabr restore user-vm --vm test -b exabr_201309121315
   ```

   In this example, ExaBR restores the user vServer called test from the backup called `exabr_201309121315`. If more than one vServer exists with the same name, use the vServer GUID instead, as in the following example:

   ```
   ./exabr restore user-vm --vm 0004fb0000060000399e5199ecae0041 -b exabr_
   201309121315
   ```

   You can identify the vServer GUID using Exalogic Control in the **Summary** page of the vServer as the Domain Name.
5. After the restoration is complete, log in to all the restored vServers
6. Log in to Exalogic Control as the `root` user.
7. From the navigation pane on the left, click **vDC Management**.
8. Expand your cloud, such as `MyCloud`.
9. Expand **Accounts**.
10. Expand the account to which the vServer you restored belongs, such as `Dept1`.

    All the vServers in the account are displayed.
11. Select the **vServer** you restored.

    The dashboard of the vServer is displayed.
12. From the actions pane on the right, click **Launch Virtual Console** and monitor the vServer startup. If any disk issue is detected, `fdisk` starts automatically.

A

Recovery of the Exalogic Control Stack From Hardware Failures

This appendix describes how to recover components of the Exalogic Control stack if a Oracle VM Server node running an Exalogic Control or Proxy Controller vServer crashes.

> **Note:** Use the procedures in this appendix only when directed to in Section 3.2.2.4, "Reimaging and Recovering Oracle VM Server Nodes in a Virtual Environment."

A.1 Recovering the Exalogic Control vServer From Hardware Failures

Oracle VM Manager, Enterprise Manager Ops Center, and Database are deployed in the Exalogic Control vServer. By default, the Exalogic Control vServer is deployed to the first Oracle VM Server node in the first server pool. If the Exalogic Control vServer crashes, both Oracle VM manager and Exalogic Control are not operational.

Migrate the Exalogic Control vServer, by doing the following:

1. Identify the Oracle VM Server node to which you should migrate the Exalogic Control vServer.

Failed Nodes	Migrate to Node
Node 1	Node 4
Node 4	Node 1
Nodes 1 and 4	Any running node

2. Start the Exalogic Control vServer, by performing the following steps:
   a. Log in as the `root` user on the Oracle VM Server node you identified in the previous step.
   b. Find the absolute path to the virtual machine configuration file for the Exalogic Control vServer.

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

      ```
      # grep "ExalogicControl" /OVS/Repositories/*/*/*/vm.cfg
      ```

Note: On an Exalogic rack that was upgraded to EECS 2.0.6 from EECS 2.0.4, the `ExalogicControl` vServer is called `ExalogicControlOVMM`.

Example:

```
# grep "ExalogicControl" /OVS/Repositories/*/*/*/vm.cfg

/OVS/Repositories/0004fb00000300007d4eef3af41ca807/VirtualMachines/0004fb00
00060000014361b5c6f404/vm.cfg:OVM_simple_name='ExalogicControlOpsCenterPC1'

/OVS/Repositories/0004fb00000300007d4eef3af41ca807/VirtualMachines/0004fb00
0006000040e5af16d3288845/vm.cfg:OVM_simple_name = 'ExalogicControl'

/OVS/Repositories/0004fb00000300007d4eef3af41ca807/VirtualMachines/0004fb00
0060000cbf5bca84ab4b65/vm.cfg:OVM_simple_name='ExalogicControlOpsCenterPC2
```

In this example, the absolute path of the Exalogic Control vServer is as follows:

```
/OVS/Repositories/0004fb00000300007d4eef3af41ca807/VirtualMachines/0004fb00
0006000040e5af16d3288845/vm.cfg
```

c. Start the Exalogic Control vServer, by using the `xm create` command as follows:

```
xm create absolute_path_to_vm.cfg
```

Example:

```
xm create
/OVS/Repositories/0004fb00000300007d4eef3af41ca807/VirtualMachines/0004fb00
0006000040e5af16d3288845/vm.cfg
```

Exalogic Control can take at least five minutes to start.

3. Verify that the Exalogic Control vServer is running, by logging in to the Exalogic Control BUI as the `root` user.
4. If the Proxy Controller vServers are down, follow the procedure in Section A.2, "Recovering the Proxy Controller vServers From Hardware Failures" to start the vServers manually.
5. Identify the IP addresses of the Proxy Controller vServers, by doing the following:
   a. Log in to any Oracle VM Server node on the Exalogic rack as the `root` user.
   b. If the `ExalogicControl` share is not mounted in the `mnt/ExalogicControl` directory, mount it.
   c. Navigate to the `/mnt/ExalogicControl/ECU_ARCHIVE` directory.
   d. Extract the archive of the ECU files called `ecu_log-`*`date&time_stamp`*`.tgz`.

      The following files are extracted `ecu_run_time.tgz`, `ecu_home.tgz`, `ecu_archive.tgz`.
   e. Extract `ecu_run_time.tgz` that contains the ECU configuration files.

      A directory called `ecu` that contains all the configuration files as extracted.
   f. Navigate to the extracted `ecu` directory.

```
cd ecu
```

g. Identify the IP addresses of the Proxy Controller 1 and 2 vServers respectively, by running the following commands:

```
# grep ecu_pc_IPoIB-admin_primary ops_center.properties
# grep ecu_pc_IPoIB-admin_secondary ops_center.properties
```

h. Note the IP addresses of the Proxy Controller vServers.

6. After restarting the Exalogic Control vServer, restart the Proxy Controllers by doing the following:

a. Log in to any Oracle VM Server node on the Exalogic rack as the `root` user.

b. SSH to the Proxy Controller vServer you want to restart, with the IP address you identified in step 5, as the `root` user.

Example:

```
# ssh root@192.168.20.12
```

c. Stop the Proxy Controller services by running the following command:

```
[root@hostname-pc1 ~]#  /opt/sun/xvmoc/bin/proxyadm stop -w

proxyadm: Shutting down Proxy Controller using SMFlite...
application/scn/proxy-available:default... ... stopped.
application/scn/uce-proxy:default... ... stopped.
application/management/common-agent-container:scn-proxy... ... stopped.
application/scn/proxy-enable:default... ... stopped.
proxyadm: Proxy Controller services have stopped
```

d. Start the Proxy Controller services by running the following command:

```
[root@hostname-pc1 ~]#  /opt/sun/xvmoc/bin/proxyadm start -w

proxyadm: Starting Proxy Controller with SMFlite...
application/scn/proxy-enable:default... ...started.
application/scn/uce-proxy:default... ...started.
application/management/common-agent-container:scn-proxy... ...started.
application/scn/proxy-available:default... ...started.
proxyadm: Proxy Controller services have started
```

e. Run the `exit` command to exit the Proxy Controller vServer.

f. Repeat steps a to e for the other Proxy Controller vServer.

A.2 Recovering the Proxy Controller vServers From Hardware Failures

The Proxy Controller is deployed as two vServers. By default, the first Proxy Controller vServer is deployed to the first Oracle VM Server node in the first pool and the second vServer is deployed to the second Oracle VM Server node in the first pool. If either of the Proxy Controller vServers is down, Exalogic Control functionality is affected.

Start the Proxy Controller vServers manually, by doing the following:

1. Verify that the Exalogic Control vServer is running, by logging in to the Exalogic Control BUI as the `root` user.

2. Identify the IP addresses of the Proxy Controller vServers, by doing the following:

a. Log in to any Oracle VM Server node on the Exalogic rack as the `root` user.

b. If the `ExalogicControl` share is not mounted in the `mnt/ExalogicControl` directory, mount it.

c. Navigate to the `/mnt/ExalogicControl/ECU_ARCHIVE` directory.

d. Extract the archive of the ECU files called `ecu_log-`*`date&time_stamp`*`.tgz`.

The following files are extracted `ecu_run_time.tgz`, `ecu_home.tgz`, `ecu_archive.tgz`.

e. Extract `ecu_run_time.tgz` that contains the ECU configuration files.

A directory called `ecu` that contains all the configuration files as extracted.

f. Navigate to the extracted `ecu` directory.

```
cd ecu
```

g. Identify the IP addresses of the Proxy Controller 1 and 2 vServers respectively, by running the following commands:

```
# grep ecu_pc_IPoIB-admin_primary ops_center.properties
# grep ecu_pc_IPoIB-admin_secondary ops_center.properties
```

h. Note the IP addresses of the Proxy Controller vServers.

3. Identify the Oracle VM Server node on which you should start the Proxy Controller vServer.

Proxy Controller	Failed Nodes	Migrate to Node
Proxy Controller 1	Node 1	Node 3
	Node 3	Node 1
	Nodes 1 and 3	Any running node that is not running Proxy Controller 2
Proxy Controller 2	Node 2	Node 4
	Node 4	Node 2
	Nodes 2 and 4	Any running node that is not running Proxy Controller 1

4. Start the Proxy Controller vServer, by doing the following:

a. Log in as the `root` user to the Oracle VM Server node you identified in the previous step.

b. Find the absolute path to the virtual machine configuration file for the Proxy Controller vServer. Run the following `grep` command to identify the configuration file corresponding to the Proxy Controller vServer:

```
# grep "ExalogicControlOpsCenterPC" /OVS/Repositories/*/*/*/vm.cfg
```

Example:

```
# grep "ExalogicControlOpsCenterPC" /OVS/Repositories/*/*/*/vm.cfg

/OVS/Repositories/0004fb00000300007d4eef3af41ca807/VirtualMachines/0004fb00
00060000014361b5c6f404/vm.cfg:OVM_simple_name='ExalogicControlOpsCenterPC1'

/OVS/Repositories/0004fb00000300007d4eef3af41ca807/VirtualMachines/
```

```
0000cbf5bca84ab4b658/vm.cfg:OVM_simple_name = 'ExalogicControlOpsCenterPC2
```

In this example, the absolute paths of the Proxy Controller 1 and Proxy Controller 2 vServers respectively are as follows:

```
/OVS/Repositories/0004fb00000300007d4eef3af41ca807/VirtualMachines/0004fb00
00060000014361b5c6f404/vm.cfg

/OVS/Repositories/0004fb00000300007d4eef3af41ca807/VirtualMachines/0004fb00
00060000cbf5bca84ab4b658/vm.cfg
```

c. Start the Proxy Controller vServer by using the `xm create` command as follows:

```
xm create absolute_path_to_vm.cfg
```

Example:

```
xm create
/OVS/Repositories/0004fb00000300007d4eef3af41ca807/VirtualMachines/0004fb00
00060000014361b5c6f404/vm.cfg
```

In this example, you are starting the Proxy Controller 1 vServer.

5. Verify whether the Proxy Controller vServer is up by doing the following:

a. Log in to any Oracle VM Server node on the Exalogic rack as the `root` user.

b. From the Oracle VM Server node, log in to the Proxy Controller vServer you started as the `root` user.

Example:

```
# ssh root@192.168.20.12
```

c. Verify the Proxy Controller is online, by running the following command:

```
[root@hostname-pc1 ~]# /opt/sun/xvmoc/bin/proxyadm status
```

d. If the Proxy Controller is offline, run step 6 in Section A.1, "Recovering the Exalogic Control vServer From Hardware Failures."

6. If you want to restart the other Proxy Controller vServer, repeat steps 4 and 5.

B

Removing Orphan and Ghost vServers After Restoring the Exalogic Control Stack

When you restore the Exalogic Control stack from a backup, it reverts to the point in time when that backup was created.

- vServers that were created after that backup will not be displayed in Exalogic Control (orphan vServers).
- vServers that were deleted after that backup will be displayed in Exalogic Control, but they do not exist (ghost vServers).

So after restoring the Exalogic Control stack, you must remove such vServers, by performing the following steps:

1. Log in to Oracle VM Manager as the `root` user.
2. Identify the vServer GUID of orphan vServers, by doing the following:
   a. Click the **Servers and VMs** tab.
   b. Expand **Server Pools**.
   c. Click a server pool.
   d. From the Perspective drop down box, select **Virtual Machines**.
   e. Expand a vServer called `ORPHAN_`*`GUID`*.
   f. From the **Configuration** tab of the vServer, note the vServer GUID (ID) of the vServer.
   g. Repeat steps e and f for all vServers called `ORPHAN_`*`GUID`*.
   h. If you have more than one server pool, repeat steps c to g for each server pool.
3. Identify the names and vNIC GUIDs of the orphan vServers, by doing the following:
   a. Log in to any running Oracle VM Server node as the `root` user.
   b. Note the names of orphan vServers, using the GUIDs you identified in step 2, by running the following command:

      ```
      grep simple_name /OVS/Repositories/Repository_ID/VirtualMachines/vServer_
      GUID/vm.cfg
      ```

      Example:

      ```
      grep simple_name
      /OVS/Repositories/0004fb0000030000041471155f08fb017b/VirtualMachines/0004fb00
      000600008c6cfab8bb065dff/vm.cfg OVM_simple_name = 'vsrvr'
      ```

In this example, the vServer with the GUID `0004fb00000600008c6cfab8bb065dff` is called `vsrvr`.

c. Note the vNIC GUIDs of orphan vServers, using the GUIDs you identified in step 2, by running the following command:

```
grep exalogic_vnic /OVS/Repositories/Repository_ID/VirtualMachines/vServer_
GUID/vm.cfg
```

Example:

```
grep exalogic_vnic
/OVS/Repositories/0004fb0000030000414715 5f08fb017b/VirtualMachines/0004fb00
000600008c6cfab8bb065dff/vm.cfg
exalogic_vnic = [{'pkey': ['0x8006'], 'guid': '0xfdc7b8890128889c', 'port':
'1'}, {'pkey': ['0x8006'], 'guid': '0xfdc7b8890128889d', 'port': '2'}]
```

In this example, the orphan vServer with the GUID `0004fb00000600008c6cfab8bb065dff` has the vNIC GUIDs `fdc7b8890128889c` and `fdc7b8890128889d`. Ensure that you note the GUIDs without the `0x` prefix.

d. Repeat step b and c for all vServer GUIDs you noted in step 2.

4. Refresh the Oracle VM Manager repository, by doing the following:
   a. Click the **Repositories** tab.
   b. Select the repository.
   c. Click the **refresh** icon.

   After the repository refresh, orphaned vServers will display their actual names and ghost vServers will display errors.
5. Identify the vServer GUID (ID) of ghost vServers, by doing the following:
   a. Click the **Servers and VMs** tab.
   b. Expand **Server Pools**.
   c. Click a server pool.
   d. From the Perspective drop down box, select **Virtual Machines**.
   e. Expand the vServer for which the status is stopped and the Event Severity displays an error.
   f. Note the GUID (ID) and name of the ghost vServer.
   g. If you have more than one server pool, repeat steps c to f for each server pool.
6. Acknowledge the Oracle VM Manager events of vServers that cannot start due to errors, by doing the following:
   a. Click the **Servers and VMs** tab.
   b. Expand **Server Pools**.
   c. Click the server pool running the vServer that is in the error state.
   d. From the Perspective drop down box, select **Virtual Machines**.
   e. Right click on a vServer for which its Event Severity displays an error.
   f. Click **Display Events**.
   g. Click **Acknowledge All**.

h. Repeat steps e to g for all vServers that cannot run due to errors.

i. Refresh the Server Pools view, by clicking the **Repository** tab and then the **Servers and VMs** tab.

j. Verify that a green status icon is displayed for all Oracle VM Server nodes, once all vServers in the error state have been acknowledged.

7. Delete all orphan vServers in Oracle VM Manager, by doing the following:

   a. Select an orphan vServer.

      You can identify these vServers by using the names you noted in step 3 and GUIDs you noted in 2.

   b. Click the **Stop** button.

      The confirmation box appears.

   c. Click **OK**.

      A job to stop the vServer is created.

   d. Wait for the job to stop the vServer to complete.

   e. Select the vServer you stopped.

   f. Click the **Delete** button.

      The confirmation box appears.

   g. Under Select virtual disks to delete, select all virtual disks that are displayed.

   h. Click **OK**.

   i. Repeat steps a to h for all orphan vServers.

8. Delete vNICs associated with all orphan vServers, by doing the following:

   a. Log in to an InfiniBand Gateway switch as `root`.

   b. Identify the connector and vNIC ID of the vNIC whose GUID you identified in step 3:

      ```
      showvnics | egrep -i "vnic_guid1|vnic_guid2"
      ```

      Example:

      ```
      showvnics | egrep -i "fdc7b8890128889c|fdc7b8890128889d"
      89 WAIT-IOA  N  FDC7B8890128889C alpha05cn04 EL-C 192.168.54.74 0000
      00:21:F6:CC:BB:AA 20 8006   0A-ETH-1
      ```

      In this example, for the GUID `FDC7B8890128889C`, the vNIC ID is `89` and the connector is `0A-ETH-1`.

      Note: If this command does not display any results, run the command on the other gateway switch as described in step d.

   c. Delete the vNIC:

      ```
      deletevnic connector vnic_id
      ```

      Example:

      ```
      deletevnic  0A-ETH-1  89
      ```

d. On all other InfiniBand Gateway switches, run steps a to c.

e. Repeat steps a to d for all the vNICs whose GUIDs you identified in step 3.

9. Before deleting ghost vServers, you must create `vm.cfg` files for these vServers by doing the following:

a. Log in to any running Oracle VM Server node as the `root` user.

b. Run the following command to navigate to the `VirtualMachines` directory of your repository:

```
cd /OVS/Repositories/Repository_ID/VirtualMachines/
```

`Repository_ID` is the ID of your repository.

c. Verify that no directory exists for the ghost vServer, by running the `ls` command with the vServer GUID of a ghost vServer:

```
ls vServer_GUID
```

Example:

```
ls 0004fb00000600008c6cfab8bb064dbb
ls: 0004fb00000600008c6cfab8bb064dbb: No such file or directory
```

d. Create a directory using the vServer GUID, say `0004fb00000600008c6cfab8bb064dbb`, as follows:

```
mkdir 0004fb00000600008c6cfab8bb064dbb
```

e. Create an empty `vm.cfg` file for the vServer, say with vServer GUID `0004fb00000600008c6cfab8bb064dbb`, as follows:

```
touch 0004fb00000600008c6cfab8bb064dbb/vm.cfg
```

f. Repeat steps c to e for all ghost vServers.

g. Refresh the repository as described in step 4.

vServers for which you created `vm.cfg` files should be shown in the error state again.

10. Delete all ghost vServers that you created `vm.cfg` files for, by doing the following:

a. Log in to the Exalogic Control BUI as a user with the Exalogic Systems Administrator role.

b. From the left navigation pane, click **vDC Accounts**.

c. Under vDC Accounts, expand the name of your Account.

d. Click the ghost vServer you are deleting, using the name you noted in step 5.

The vServer dashboard appears.

e. Verify that the domain ID of the vServer matches the vServer GUID you noted in step 5.

f. From the Actions pane on the right, click the **delete** button.

The Delete vServer confirmation screen appears.

g. Click **Delete** to confirm.

h. Repeat steps d to g for all ghost vServers for which you created `vm.cfg` files.

C

Increasing the Size of the Logical Volume Group of a vServer

When ExaBR backs up a vServer using LVM-based snapshots, it uses the swap space of the vServer to take the snapshot. However, if ExaBR finds enough free space in the default Volume Group (`VolGroup00`) it does not use the swap space of the vServer. This appendix describes how to increase the size of the default Volume Group (`VolGroup00`) of a guest vServer to ensure that ExaBR uses the unused space of the default Volume Group for the snapshot, and not the swap space of the vServer.

> **Note:** The procedure described in this chapter is relevant to vServers created by using the EECS 2.0.6 Guest Base Template only. *Do not* use this procedure to modify the disks of vServers creating using a Guest Base Template that is earlier than EECS version 2.0.6.

You can increase the size of the Volume Group, by performing the following tasks:

1. Create a volume.
2. Attach the volume to the vServer.
3. Format and add the volume to the default Volume Group of the vServer.

C.1 Creating a Volume

To create a volume, complete the following steps:

1. Log in to the Exalogic Control BUI as a Cloud User.
2. From the left navigation pane, click **vDC Management**.
3. Under vDC Accounts, click the name of your account, such as `Dept1`.

   The vDC Account dashboard is displayed.
4. Click **Storage** on the top navigation bar.
5. Click the **Volumes** tab.
6. Under Volumes, click the **+** icon.
7. In the **Volume Name** field, enter a name for the volume. For example, `Volume1`.
8. In the **Description** field, enter a short description.
9. Click **Next**.
10. Do not select the **Shared** option.

11. Set the size of the volume in GB. Oracle recommends setting the size of the volume to at least 2 GB.
12. Click **Next**.

    The Volume Summary screen is displayed.
13. Review the summary, and click **Finish** to create the volume in your account.

C.2 Attaching the Volume to a vServer

You can attach a volume to a vServer as follows:

1. Log in to the Exalogic Control BUI as a Cloud User.
2. From the left navigation pane, click **vDC Management**.
3. Under vDC Accounts, expand the name of your account, such as `Dept1`.
   All the vServers in the account are displayed.
4. Select the vServer (for example, `vserver2`) to which you wish to attach a volume.
   The `vserver2` dashboard is displayed.
5. From the actions pane on the right, click **Stop vServer**. Wait till the job succeeds in the jobs pane.
6. From the actions pane on the right, click **Attach vServer Volumes**.
   The Attach vServer Volumes wizard is displayed.
7. Select the volume you wish to attach to `vserver2`.
8. Click the right arrow icon.
9. Click **Next**.
   The confirmation screen is displayed.
10. Click **Finish**. Wait till the job succeeds in the jobs pane.
11. From the actions pane on the right, click the **Start vServer** button to restart the vServer.

C.3 Formatting the Volume on the vServer

Format the volume on the vServer, by doing the following:

1. Log in to the vServer as the `root` user.
2. Examine the current partitioning by running the following command:

   ```
   # df -h
   ```

   The following is an example of the output of this command:

   ```
   Filesystem                      Size  Used Avail Use% Mounted on
   /dev/mapper/VolGroup00-LogVol00  5.1G  3.3G  1.6G  68% /
   /dev/xvda1                       99M   23M   71M  25% /boot
   tmpfs                           4.0G     0  4.0G   0% /dev/shm
   ```

3. Examine the available physical volumes on the vServer by running the following command:

   ```
   # cat /proc/partitions
   major minor  #blocks  name

    202        0    6145024 xvda
   ```

```
 202        1     104391 xvda1
 202        2    6040440 xvda2
 253        0    5505024 dm-0
 253        1     524288 dm-1
 202       16  104857600 xvdb
```

`/dev/xvdb` is the newly attached volume.

4. Run the `fdisk` command, as shown in the following example:

> **Note:** The user input required at various stages while running the `fdisk` command is indicated by **bold** text.

```
# fdisk /dev/xvdb
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF
disklabel
Building a new DOS disklabel. Changes will remain in memory only,
until you decide to write them. After that, of course, the previous
content won't be recoverable.

The number of cylinders for this disk is set to 13054.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:
1) software that runs at boot time (e.g., old versions of LILO)
2) booting and partitioning software from other OSs
(e.g., DOS FDISK, OS/2 FDISK)
Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)
Command (m for help): p

Disk /dev/xvdb: 107.3 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

    Device Boot      Start         End      Blocks   Id  System

Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-13054, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-13054, default 13054):
Using default value 13054

Command (m for help): t
Selected partition 1
Hex code (type L to list codes): 8e
Changed system type of partition 1 to 8e (Linux LVM)

Command (m for help): p

Disk /dev/xvdb: 107.3 GB, 107374182400 bytes
255 heads, 63 sectors/track, 13054 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

    Device Boot      Start         End      Blocks   Id  System
/dev/xvdb1               1       13054   104856223+  8e  Linux LVM
```

```
Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

5. Create a physical volume:

```
# pvcreate /dev/xvdb1
  Writing physical volume data to disk "/dev/xvdb1"
  Physical volume "/dev/xvdb1" successfully created
```

6. Extend the volume group `VolGroup00` with the physical volume `/dev/xvdb1`:

```
# vgextend VolGroup00 /dev/xvdb1
  Volume group "VolGroup00" successfully extended
```

7. Verify that the volume group was extended successfully, by running the following command:

```
# vgs
VG         #PV #LV #SN     Attr  VSize  VFree
VolGroup00   2   2   0   wz--n-   7.5G   3.6G
```

   The amount of free space appears under the `VFree` column.