# Creating and Administering Oracle® Solaris 11.4 Boot Environments

**Part No: E60980**
November 2019

ORACLE

Creating and Administering Oracle Solaris 11.4 Boot Environments

**Part No: E60980**

**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.

**Référence: E60980**



**Restrictions de licence/Avis d'exclusion de responsabilité en cas de dommage indirect et/ou consécutif**

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

**Exonération de garantie**

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

**Avis sur la limitation des droits**

Si ce logiciel, ou la documentation qui l'accompagne, est livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

**Avis sur les applications dangereuses**

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer un risque de dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour des applications dangereuses.

**Marques**

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Inside sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Epyc, et le logo AMD sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée de The Open Group.

**Avis d'exclusion de responsabilité concernant les services, produits et contenu tiers**

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers, sauf mention contraire stipulée dans un contrat entre vous et Oracle. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

**Date de publication et mention de la version préliminaire de Disponibilité Générale ("Pre-GA")**

Version préliminaire de Disponibilité Générale ("Pre-GA") : 15.01.2020

**Avis sur la version préliminaire de Disponibilité Générale ("Pre-GA") de la documentation**

Si ce document est fourni dans la Version préliminaire de Disponibilité Générale ("Pre-GA") à caractère public ou privé :

Cette documentation est fournie dans la Version préliminaire de Disponibilité Générale ("Pre-GA") et uniquement à des fins de démonstration et d'usage à titre préliminaire de la version finale. Celle-ci n'est pas toujours spécifique du matériel informatique sur lequel vous utilisez ce logiciel. Oracle Corporation et ses affiliés déclinent expressément toute responsabilité ou garantie expresse quant au contenu de cette documentation. Oracle Corporation et ses affiliés ne sauraient en aucun cas être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'utilisation de cette documentation.

**Mention sur les informations confidentielles Oracle**

INFORMATIONS CONFIDENTIELLES ORACLE. Destinées uniquement à un usage autorisé. Ne pas distribuer à des tiers.

**Avis sur la reconnaissance du revenu**

Si ce document est fourni dans la Version préliminaire de Disponibilité Générale ("Pre-GA") à caractère privé :

Les informations contenues dans ce document sont fournies à titre informatif uniquement et doivent être prises en compte en votre qualité de membre du customer advisory board ou conformément à votre contrat d'essai de Version préliminaire de Disponibilité Générale ("Pre-GA") uniquement. Ce document ne constitue en aucun cas un engagement à fournir des composants, du code ou des fonctionnalités et ne doit pas être retenu comme base d'une quelconque décision d'achat. Le développement, la commercialisation et la mise à disposition des fonctions ou fonctionnalités décrites restent à la seule discrétion d'Oracle.

**Accessibilité de la documentation**

Pour plus d'informations sur l'engagement d'Oracle pour l'accessibilité de la documentation, visitez le site Web Oracle Accessibility Program, à l'adresse : http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

**Accès aux services de support Oracle**

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info ou le site http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs si vous êtes malentendant.

# Contents

# Using This Documentation

- **Overview** – Describes how to manage and create boot environments
- **Audience** – Technicians, system administrators, and authorized service providers
- **Required knowledge** – Experience administrating a Oracle Solaris system

## Product Documentation Library

Documentation and resources for this product and related products are available at http://www.oracle.com/pls/topic/lookup?ctx=E37838-01.

## Feedback

Provide feedback about this documentation at http://www.oracle.com/goto/docfeedback.

♦ ♦ ♦ CHAPTER 1

# Introduction to Managing Boot Environments

This guide provides information about Oracle Solaris boot environments (BEs) and how to manage them by using the `beadm` command.

This chapter covers the following topics:



## About Boot Environments

A *boot environment* is a bootable instance of the Oracle Solaris image and its component software packages. When you install Oracle Solaris for the first time, or upgrade the OS, a BE is automatically created on the system. The BE is also activated and becomes the default BE that is used when the system is rebooted.

You can also manually create new BEs by using the `beadm` command. Thus, a system can contain multiple BEs. However, only one boot BE can be active at a time.

### Maintaining Multiple Boot Environments

Multiple BEs reduce risk when updating software, as shown in the following examples.

- You can revert to a backup BE in case of errors occurring when updating software.
- You can install, test, and update different software packages on a cloned BE and thus preserve the original BE.

Although only one BE can be active at a time, you can mount an inactive one on which you can install or update specific packages. For more information, see "Installing a Package Into a New Boot Environment" in *Updating Systems and Adding Software in Oracle Solaris 11.4*.

- While modifying a BE, you can take a snapshot at any stage. For example, if you are doing monthly upgrades to your boot environment, you can capture monthly snapshots.

  Although a snapshot is not bootable, you can create a BE based on that snapshot.
- You can specify any BE to be the default boot BE at the next system reboot.
- Packages are installed or updated on a clone of the active BE. The clone is created automatically and becomes the active and default BE when the system is rebooted. Thus, the original BE remains intact.

# About Boot Environments and Datasets

A BE consists of a *root dataset* and, optionally, additional nested datasets.

A *dataset* is a generic name for ZFS entities such as clones, file systems, or snapshots. In the context of boot environment administration, the dataset more specifically refers to the file system specifications for a particular boot environment.

In the following example, the boot environment `BE1`'s root dataset is `rpool/ROOT/BE1`. `rpool` is the name of the root storage pool. `ROOT` is a special dataset that was created by the installer. `rpool/ROOT` is reserved exclusively for use by boot environment root datasets.

---

**Note -** Storage pools are sometimes called `zpools` because they are administered through the `zpool` command.

---

```
$ zfs list
NAME                   USED  AVAIL  REFER  MOUNTPOINT
rpool                 42.5G  24.4G  4.65M  /rpool
rpool/ROOT            25.6G  24.4G    31K  legacy
rpool/ROOT/BE1        8.91M  24.4G  4.17G  /
rpool/ROOT/BE1/var    3.96M  24.4G   276M  /var
```

*Shared datasets* are located outside a BE's root dataset. Shared datasets are user-defined directories, such as a directory that stores user accounts. These directories are accessible regardless of which BE is currently active. In this example, `tank` and `tank/home` are shared datasets.

```
$ zfs list
NAME                     USED  AVAIL  REFER  MOUNTPOINT
```

```
rpool                   42.5G  24.4G  4.65M  /rpool
rpool/ROOT              25.6G  24.4G    31K  legacy
rpool/ROOT/BE1          8.91M  24.4G  4.17G  /
rpool/ROOT/BE1/var      3.96M  24.4G   276M  /var
tank                     450K   457G    18K  /export
tank/home                315K   457G    21K  /export/home
```

For further reference, see the `zpool`(8) and the `zfs`(8) man pages as well as "Querying ZFS Storage Pool Status" in *Managing ZFS File Systems in Oracle Solaris 11.4*.

# About Boot Environment FMRIs

The boot environment management software uniquely identifies BEs in both the global zone and non-global zones through the use of boot environment FMRIs (BeFMRIs). It interprets BeFMRIs based on the environment in which the FMRI is used.

Depending on whether you are working in the global zone or a non-global zone, the BE name is interpreted differently, based on two schemes that represent BeFMRIs:

- `be`
- `zbe`

If you are in the global zone, the BE name specified in the `be` scheme is interpreted as the name of a global zone BE. If you are in the non-global zone, the BE name specified is interpreted as the name of a non-global zone BE. Additionally, while in the global zone, a user can specify the name of the storage pool on which the BE resides.

The string format for the `be` scheme is `be://[zpool-name]/be-name[@snapshot-name]`, as shown in the following examples:

- `be://rpool/be1` – the boot environment `be1` in the storage pool called `rpool`.
- `be://rpool/be1@backup` – the snapshot `backup` for the boot environment `be1` in the currently running storage pool called `rpool`.
- `be://rpool2/be1` – the boot environment `be1` in the storage pool called `rpool2`.

The `zbe` scheme is supported in the global zone only. This scheme manages non-global zone BEs from the global zone. The zone name specified in the `zbe` scheme must be defined in the running environment. The string format of the `zbe` scheme is `zbe://[zpool-name]/zbe-name [@snapshot-name]`, as shown in the following examples:

- `zbe://zone2/ngz_be2` – a non-global zone boot environment called `ngz_be2` in a zone called `zone2`.

- zbe://zone1/ngz_be2@backup – a snapshot called backup of the boot environment ngz_be2 in zone1.

When you display a list of boot environments on the system, the information includes the BE's corresponding FMRI designations. See examples in "Listing Existing Boot Environments and Snapshots" on page 29.

For more information, see the beadm(8) man page. See also Chapter 2, "beadm Zones Support" for examples of the use of the beadm command to manage boot environments.

# About the beadm Command

Different Oracle Solaris utilities can create BEs. However, for regular administration of BEs, always use the beadm command. Even though the command is based on ZFS technology, it has additional functionalities specific to managing a BE's root dataset and its contents.

The beadm command can also perform the following operations:

- It aggregates all datasets in a BE and performs actions on the entire BE at once. Do not use ZFS commands to modify each dataset individually.
- It manages the dataset structures within BEs. For example, when the beadm command clones a BE that has shared datasets, the command automatically recognizes and manages those shared datasets for the new BE.
- It enables you to perform administrative tasks on BEs in a global zone or in a non-global zone.
- It automatically manages and updates the GRUB menu for x86 systems or the boot menu for SPARC systems.

# Using Rights Profiles to Administer Boot Environments

Oracle Solaris implements role-based access control (RBAC) to control system access. To manage boot environments, you must be assigned at a minimum the Software Installation profile. Other profiles are required if you need to perform additional tasks indirectly related to your current one, such as creating and configuring zones.

An administrator that has the solaris.delegate.* authorization can assign the required profiles to users.

For example, an administrator assigns the Software Installation profile to user jdoe. Before jdoe executes a privileged command, jdoe must be in a profile shell. The shell can be created

by issuing the `pfbash` command. Or, `jdoe` can combine `pfexec` with every privileged command that is issued, for example, `pfexec beadm`.

As an alternative, instead assigning profiles directly to users, a system administrator can create a role that would contain a combination of required profiles to perform a range of tasks.

Suppose that a role `beadmin` is created with the profiles for software installation, unified archive administration, and zone configuration. As an authorized user, `jdoe` uses the `su` command to assume that role. All roles automatically get `pfbash` as the default shell.

For more information about rights profiles, see "Using Your Assigned Administrative Rights" in *Securing Users and Processes in Oracle Solaris 11.4*.

♦ ♦ ♦ CHAPTER 2

# beadm Zones Support

This chapter describes how to administer boot environments in relation to zones.

For more information about zones, refer to the guides under the *Creating and Using Oracle Solaris Virtual Environments* shelf at https://docs.oracle.com/cd/E37838_01/.

This chapter covers the following topics:

- "Using the beadm Command in Non-Global Zones"
- "Unbootable Boot Environments"
- "Zones and Shared Datasets"
- "Working With Non-Global Zone BEs From the Global Zone"

**Note -** To manage boot environments of zones and issue commands described in this chapter, you must have the appropriate rights profiles. See "Using Rights Profiles to Administer Boot Environments" on page 14.

## Using the beadm Command in Non-Global Zones

The beadm command is the tool for managing BEs in both global and non-global zones. However, the supported command options inside non-global zones are restricted compared to those available in the global zone. Refer to the beadm(8) man page.

Further, the beadm command does not provide zone information. For this purpose, use appropriate zoneadm subcommands instead. See the zoneadm(8) man page for additional zoneadm options.

# Unbootable Boot Environments

Both global zones and non-global zones contain BEs. Each BE in a non-global zone is associated with a parent BE in the global zone. If a global zone BE is inactive, the related non-global zone BE is unbootable. The unbootable BE is flagged with `!` when you display a list of BEs. Booting into the parent BE in the global zone causes the related non-global BE to become bootable.

Note that if a BE becomes unbootable, then no other action can be performed on it by any `beadm` subcommands.

# Zones and Shared Datasets

The `beadm` command automatically handles all zone-naming tasks that are related to BEs.

A zone's root dataset uses the format `rpool/zones/`*zonename*, such as `rpool/zones1/zone1`.

The root dataset for the BE is in the zone's root dataset, such as `rpool/ROOT/BE1` in the following example:

```
rpool/zones/zone1/rpool/ROOT/BE1
```

You create a shared dataset during zone configuration. The shared dataset is placed in the zone's root dataset but outside the root dataset for the BE. Thus a shared dataset such as `/export` would be listed as follows:

```
rpool/zones/zone1/rpool/export
```

When a zone is copied from one BE to another, only the datasets under the zone's root dataset are copied. Shared datasets are not cloned when the zone root dataset is cloned. See the examples in Chapter 3, “Creating Boot Environments and Snapshots”.

# Working With Non-Global Zone BEs From the Global Zone

You can manage non-global zone BEs from the global zone, provided that you use BeFMRIs when performing management tasks. See “About Boot Environment FMRIs” on page 13.

The following example shows how you can destroy as well as rename non-global zone BEs directly from the global zone.

```
global$ beadm list -z zone1
Non-global Zone BE Names   Parent BE       Flags Mountpoint               Space
------------------------   --------------   ---- ------------------------ -------
zbe://zone1/BE10           -               RO    -                        705.10M
zbe://zone1/BE3            be://rpool/BE3  R     -                        803.73M
zbe://zone1/solaris        be://rpool/BE2  N     /system/zones/zone1/root 1.04M
zbe://zone1/zbe1           be://rpool/BE2  R     -                        17.10M
zbe://zone1/zbe2           be://rpool/BE2  -     -                        2.00K

global$ beadm destroy zbe://zone1/BE10

global$ beadm rename zbe://zone1/BE3 zbe://zone1/sol-11be

global$ beadm list -z zone1
Non-global Zone BE Names   Parent BE       Flags Mountpoint               Space
------------------------   --------------   ---- ------------------------ -------
zbe://zone1/sol-11be       be://rpool/BE3  R     -                        803.73M
zbe://zone1/solaris        be://rpool/BE2  N     /system/zones/zone1/root 1.04M
zbe://zone1/zbe1           be://rpool/BE2  R     -                        17.10M
zbe://zone1/zbe2           be://rpool/BE2  -     -                        2.00K
```

♦ ♦ ♦ CHAPTER 3

# Creating Boot Environments and Snapshots

This chapter describes creating BEs and BE snapshots. It covers the following topics:

- "Creating a Boot Environment"
- "Creating and Copying Snapshots"

---

**Note -** To create boot environments and snapshots and issue commands described in this chapter, you must have the appropriate rights profiles. See "Using Rights Profiles to Administer Boot Environments" on page 14.

---

## Creating a Boot Environment

To modify an existing BE, a best practice approach is to directly work on the BE's backup in order to preserve the original BE. The BE backup or clone would be listed as an alternate boot environment in the GRUB menu on x86 systems or in the boot menu on SPARC systems.

When you clone a BE, all supported zones in that BE are copied into the clone.

### `beadm create` Command Options

To clone BEs, use the `beadm create` subcommand as follows:

```
$ beadm create BeFMRI
```

Ensure that the name you provide for the new BE is not an existing name. The new BE is created based on the currently active BE, but the new BE is inactive.

The command does not create partial BEs. The command either successfully creates a full BE, or fails.

The following options are available:

`-a` – Activate the new BE upon creation. By default, new BEs are not activated.

`-d` *description* – Provide a custom description to be used as the title in the x86 GRUB menu or the SPARC boot menu to describe the new BE. If this option is not used, *BeFMRI* is used for the title.

`-e` *non-activeBeFMRI* – Create a new BE from an existing but inactive BE. By default, clones are created from the active BE.

`-e` *BeFMRI@snapshot* – Create a BE from an existing BE snapshot.

`-o` *property=value* – Create datasets for a new BE with specific ZFS properties. Multiple `-o` options can be specified. See the `zfs`(8) man page for more information about the `-o` option.

## How to Clone a Boot Environment

**Before You Begin** Ensure that your role has the appropriate rights profiles to perform this procedure. See "Using Rights Profiles to Administer Boot Environments" on page 14.

1. **Clone the boot environment.**

   ```
   $ beadm create BeFMRI
   ```

   To clone an inactive BE, use the `-e` *inactive-BeFMRI* option.

2. **(Optional) Use the `beadm mount` command to mount the new BR.**

   ```
   $ beadm mount BeFMRI mount-point
   ```

   You might perform this step, for example, if you want to modify some configuration files in the new BE before rebooting into it. Note that mounting the new BE does **not** cause the BE to become active.

   ---

   **Note -** If the directory for the mount point does not exist, the `beadm` command creates the directory, then mounts the BE on that directory.

   If the BE is already mounted, the `beadm mount` command fails and does not remount the BE at the newly specified location.

   ---

3. **(Optional) Activate the boot environment.**

   ```
   $ beadm activate BeFMRI
   ```

On reboot, the newly active BE is displayed as the default selection in the x86 GRUB menu or the SPARC boot menu.

**Note -** The GRUB menu or boot menu always displays the most recently activated BE as the default.

# Examples of Cloning Boot Environments

The following examples illustrate how to clone BEs, and how the cloning process varies depending on the zone and dataset structure:


- Example 1, “Cloning a Boot Environment in a Global Zone That Contains Non-Global Zones,” on page 23
- Example 2, “Cloning a New Boot Environment with Datasets,” on page 24
- Example 3, “Creating a New Boot Environment With Existing Shared Datasets,” on page 26


**EXAMPLE 1** Cloning a Boot Environment in a Global Zone That Contains Non-Global Zones

If the BE being cloned has an associated zone BE in a non-global zone, that associated BE is also cloned.

This example makes the following assumptions:

- The original BE on the global zone is `solaris`. Its root dataset at `rpool/ROOT/solaris`.
- The non-global zone `z1` exists with its dataset at `rpool/zones/z1`.
- The global zone BE `solaris` has an associated BE in `z1`. The zone BE is also named `solaris`, whose root dataset is `rpool/zones/z1/rpool/ROOT/solaris`.

```
$ zfs list -r rpool
NAME                                 USED  AVAIL  REFER MOUNTPOINT
rpool                               11.5G  3.89G  4.46M  /rpool
rpool/ROOT                          8.47G  3.89G    31K  legacy
rpool/ROOT/solaris                  2.98M  3.89G  2.49G  /
rpool/dump                          1.03G  3.92G  1.00G  -
rpool/export                         120K  3.89G    32K  /export
rpool/export/home                   88.5K  3.89G    32K  /export/home
rpool/export/home/user1             56.5K  3.89G  56.5K  /export/home/user1
rpool/swap                          1.03G  3.92G  1.00G  -
rpool/zones                          672M  3.89G    32K  /zones
rpool/zones/z1                       672M  3.89G    32K  /zones/z1
```

```
rpool/zones/z1/rpool                    671M  3.89G    31K  /rpool
rpool/zones/z1/rpool/ROOT               671M  3.89G    31K  legacy
rpool/zones/z1/rpool/ROOT/solaris       671M  3.89G   591M  /zones/z1/root
rpool/zones/z1/rpool/export              62K  3.89G    31K  /export
rpool/zones/z1/rpool/export/home         31K  3.89G    31K  /export/home
```

To clone the global BE and create `test`, you would proceed as follows. The results are highlighted in the output.

```
$ beadm create test
$ zfs list -r rpool
NAME                                     USED  AVAIL  REFER MOUNTPOINT
rpool                                   11.5G  3.89G  4.46M  /rpool
rpool/ROOT                              8.47G  3.89G    31K  legacy
rpool/ROOT/solaris                      2.98M  3.89G  2.49G  /
rpool/ROOT/test                           71K  3.89G  2.50G  /
rpool/dump                              1.03G  3.92G  1.00G  -
rpool/export                             120K  3.89G    32K  /export
rpool/export/home                       88.5K  3.89G    32K  /export/home
rpool/export/home/user1                 56.5K  3.89G  56.5K  /export/home/user1
rpool/swap                              1.03G  3.92G  1.00G  -
rpool/zones                              672M  3.89G    32K  /zones
rpool/zones/z1                           672M  3.89G    32K  /zones/z1
rpool/zones/z1/rpool                     671M  3.89G    31K  /rpool
rpool/zones/z1/rpool/ROOT                671M  3.89G    31K  legacy
rpool/zones/z1/rpool/ROOT/solaris        671M  3.89G   591M  /zones/z1/root
rpool/zones/z1/rpool/ROOT/test             2K  3.89G   591M  /
rpool/zones/z1/rpool/export               62K  3.89G    31K  /export
rpool/zones/z1/rpool/export/home          31K  3.89G    31K  /export/home
```

The clone is named `test`, with a root dataset at `rpool/ROOT/test`.

Likewise, the associated BE `solaris` in `z1` is also cloned as `test`, whose dataset is `zones/z1/rpool/ROOT/test`.

**EXAMPLE 2** Cloning a New Boot Environment with Datasets

This example illustrates how datasets are set up in a newly created BE. This example does not involve multiple zones.

Suppose that the original BE called `OracleSolaris` had a root dataset `rpool/ROOT/OracleSolaris` with a nested dataset `var`.

If you create the clone `BE2`, the nested dataset of the original BE is also cloned.

```
$ zfs list
```

```
NAME                          USED  AVAIL  REFER  MOUNTPOINT
rpool                        42.5G  24.4G  4.65M  /rpool
rpool/ROOT                   25.6G  24.4G    31K  legacy
rpool/ROOT/OracleSolaris     8.91M  24.4G  4.17G  /
rpool/ROOT/OracleSolaris/var 3.96M  24.4G   276M  /var

$ beadm create BE2
$ zfs list

NAME                          USED  AVAIL  REFER  MOUNTPOINT
rpool                        42.5G  24.4G  4.65M  /rpool
rpool/ROOT                   25.6G  24.4G    31K  legacy
rpool/ROOT/OracleSolaris     8.91M  24.4G  4.17G  /
rpool/ROOT/OracleSolaris/var 3.96M  24.4G   276M  /var
rpool/ROOT/BE2               8.91M  24.4G  4.17G  /
rpool/ROOT/BE2/var           3.96M  24.4G   276M  /var
```

In the case of a system that has non-global zones, the cloning of nested datasets also applies if the global zone BE has associated BEs in the non-global zone. Specifically, if a global zone BE with nested datasets is cloned, the nested datasets of the associated BE in the non-global zone are also cloned.

Suppose that the original BE information is as follows:

```
$ zfs list -r rpool
NAME                                  USED  AVAIL  REFER MOUNTPOINT
rpool                                11.5G  3.89G  4.46M  /rpool
rpool/ROOT                           8.47G  3.89G    31K  legacy
rpool/ROOT/solaris                   2.98M  3.89G  2.49G  /
rpool/ROOT/solaris/var                428K  3.89G   298M  /var
...
rpool/zones                           672M  3.89G    32K  /zones
rpool/zones/z1                        672M  3.89G    32K  /zones/z1
rpool/zones/z1/rpool                  671M  3.89G    31K  /rpool
rpool/zones/z1/rpool/ROOT             671M  3.89G    31K  legacy
rpool/zones/z1/rpool/ROOT/solaris     671M  3.89G   591M  /zones/z1/root
rpool/zones/z1/rpool/ROOT/solaris/var  79.3M  3.89G 78.4M  /zones/z1/root/var
```

After cloning `solaris` to `newBE`, the new information would be as follows:

```
$ zfs list -r rpool
NAME                                  USED  AVAIL  REFER MOUNTPOINT
rpool                                11.5G  3.89G  4.46M  /rpool
rpool/ROOT                           8.47G  3.89G    31K  legacy
rpool/ROOT/solaris                   2.98M  3.89G  2.49G  /
rpool/ROOT/solaris/var                428K  3.89G   298M  /var
rpool/ROOT/newBE                     2.98M  3.89G  2.49G  /
rpool/ROOT/newBE/var                  428K  3.89G   298M  /var
...
```

```
rpool/zones                             672M  3.89G    32K  /zones
rpool/zones/z1                          672M  3.89G    32K  /zones/z1
rpool/zones/z1/rpool                    671M  3.89G    31K  /rpool
rpool/zones/z1/rpool/ROOT               671M  3.89G    31K  legacy
rpool/zones/z1/rpool/ROOT/solaris       671M  3.89G   591M  /zones/z1/root
rpool/zones/z1/rpool/ROOT/solaris/var  79.3M  3.89G 78.4M   /zones/z1/root/var
rpool/zones/z1/rpool/ROOT/newBE           2K  3.89G   591M  /zones/z1/root
rpool/zones/z1/rpool/ROOT/newBE/var       1K  3.89G  78.4M  /zones/z1/root/var
```

**EXAMPLE 3** Creating a New Boot Environment With Existing Shared Datasets

In this example, `BE1` is the original BE, and the shared datasets are `rpool/export` and `rpool/export/home`. This example does not involve multiple zones.

When you clone `BE1`, its datasets are also cloned but not the shared datasets Instead, both `BE1` and its clone would have access and use of `rpool/export` and `rpool/export/home`.

```
$ zfs list
NAME                      USED  AVAIL  REFER  MOUNTPOINT
rpool                    42.5G  24.4G  4.65M  /rpool
rpool/ROOT               25.6G  24.4G    31K  legacy
rpool/ROOT/BE1           8.91M  24.4G  4.17G  /
rpool/ROOT/BE1/var       3.96M  24.4G   276M  /var
rpool/export             450K   457G    18K  /rpool/export
rpool/export/home        315K   457G    21K  /rpool/export/home

$ beadm create BE2

$ zfs list
NAME                      USED  AVAIL  REFER  MOUNTPOINT
rpool                    42.5G  24.4G  4.65M  /rpool
rpool/ROOT               25.6G  24.4G    31K  legacy
rpool/ROOT/BE1           8.91M  24.4G  4.17G  /
rpool/ROOT/BE1/var       3.96M  24.4G   276M  /var
rpool/ROOT/BE2           8.91M  24.4G  4.17G  /
rpool/ROOT/BE2/var       3.96M  24.4G   276M  /var
rpool/export             450K   457G    18K  /rpool/export
rpool/export/home        315K   457G    21K  /rpool/export/home
```

# Creating and Copying Snapshots

You can manually create a snapshot of a BE for reference. This snapshot is a read-only image of a dataset or boot environment at a given point in time. You can create a custom name for the

snapshot that indicates when the snapshot was created or what it contains. You can then copy that snapshot.

Note that a clone is a complete bootable copy of an image that can be much larger than a mere snapshot. The snapshot merely records what's changed in the datasets instead of copying all the datasets contents.

## Creating a Snapshot of a Boot Environment

The following command creates a BE snapshot.

**`beadm create`** *BeFMRI@snapshotdescription*

*BeFMRI* is the name of an existing BE whose snapshot you want to create. If the existing BE name is not valid, the command fails. The custom description in the name helps to identify the date or purpose of the snapshot.

---

**Note -** If you do not follow the snapshot name format, the command performs a cloning operation instead of creating snapshot.

---

Note the following snapshot sample names and descriptions:

- `BE1@0312200.12:15pm` – The name for a snapshot of the existing `BE1` boot environment. The custom description, `0312200.12:15pm`, records the date and time that the snapshot was taken for future reference.
- `BE2@backup` – The name for a snapshot of `BE2`. The snapshot description merely notes that this is a backup of `BE2`.
- `BE1@march132008` – The name for a snapshot of `BE1`. The snapshot description records the date that the snapshot was taken.

Some other system functions automatically take snapshots of a BE. Names for such snapshots automatically include a timestamp that indicates when the snapshot was taken.

## Cloning a Boot Environment From an Existing Snapshot

A snapshot is not bootable. However, you can clone a new BE from an existing snapshot. Then you can activate and boot that new BE.

## ▼ How to Clone a Boot Environment From a Snapshot

**Before You Begin** Ensure that your role has the appropriate rights profiles to perform this procedure. See "Using Rights Profiles to Administer Boot Environments" on page 14.

### 1. Clone a new BE from a snapshot.

```
$ beadm create -e snapshot-name  NewName
```

As previously indicated, the snapshot name is in the format *BEFMRI@snapshotdescription*. For example:

```
$ beadm create -e BE1@now BE2
```

### 2. (Optional) Activate the new BE.

```
$ beadm activate BE2
```

At the next system reboot, `BE2` will be used and becomes the default BE.

♦ ♦ ♦ CHAPTER 4

# Administering Boot Environments

This chapter describes administration tasks related to boot environments. It covers the following topics:



## Listing Existing Boot Environments and Snapshots

To display information about snapshots, BEs, and datasets, use the `beadm list` command. The command also displays BEs that are created by the `pkg` command.

To view information for a specific BE, specify its name. By default all BEs are listed without additional information.

The following options are available:

`-a` – Lists all available information about the BE. This information includes subordinate datasets and snapshots.

`-d` – Lists information about all subordinate datasets that belong to the BE.

`-s` – Lists information about the snapshots of the BE.

`-H` – Lists information in machine-parsable format. Each field in the output is separated by a semicolon.

`-o` *attributes* – Lists only the attributes that you specify.

-Z and -z – lists all the zone BEs that are associated with a specific global zone BE (-Z), or lists all the BEs that are associated with a specific zone (-z).

# Viewing Information About Boot Environments

The -a option shows full information for one or all BEs, including all dataset and snapshot information. This information includes flags to indicate whether the boot environment is active or not, the mountpoint for the dataset, the amount of space used by the dataset, the policy and the date the dataset was created.

The values for the Flags column are as follows:

R – Active on reboot

N – Active now

NR – Active now and active on reboot

O – Orphaned non-global zone BE, only visible inside a non-global zone

- – Inactive

! – Unbootable BEs in a non-global zone

The following example displays full information for BE5.

```
$ beadm list -a BE5
BE Name Flags Mountpoint Space  Policy Created
------- ----- ---------- ------ ------ ----------------
BE5     -     -          80.76M static 2016-12-19 16:28

    Details:
    --------
    Zpool.......... rpool
    UUID........... 67cda37d-c1a2-4ccf-9415-06ed96b79244
    FMRI........... be://rpool/BE5
    Root Dataset... rpool/ROOT/BE5


    In-BE Datasets     Mountpoint Space  Created
    ------------------ ---------- ------ ----------------
    rpool/ROOT/BE5     -          80.76M 2016-12-19 16:28
    rpool/ROOT/BE5/var -          85.00K 2016-12-19 16:28


    Snapshots             Space  Created
    --------------------- ------ ----------------
    be://rpool/BE2@backup 67.00K 2016-12-19 16:32
```

For more information about orphaned BEs see "About Orphaned Zone Boot Environments and Clones" in *Creating and Using Oracle Solaris Zones*.

# Viewing Information About Boot Environments in Machine-Parsable Output

The `-H` option suppresses header titles and displays results separated by semicolons. The following example shows information for all BEs.

```
$ beadm list -H
BE2;4659d6ee-76a0-c90f-e2e9-a3fcb570ccd5;;;55296;static;1211397974
BE3;ff748564-096c-449a-87e4-8679221d37b5;;;339968;static;1219771706
BE4;1efe3365-02c5-6064-82f5-a530148b3734;;;16541696;static;1220664051
BE5;215b8387-4968-627c-d2d0-f4a011414bab;NR;/;7786206208;static;1221004384
```

Each field in the output is separated by a semicolon. The output fields display as follows.

| | |
|---|---|
| Field 1 | BE name |
| Field 2 | UUID |
| Field 3 | Flags |
| Field 4 | Mountpoint |
| Field 5 | Space; size in bytes |
| Field 6 | Policy |
| Field 7 | Creation time (in seconds since 00:00:00 UTC, Jan 1, 1970) |

In this example, the universally unique identifier (UUID) for the BE is provided in the second field. The UUID for BE5 is `215b8387-4968-627c-d2d0-f4a011414bab`. For a BE in a non-global zone, the UUID field represents the parent ID with which that BE is associated.

# Viewing Snapshot Specifications

The `-s` option displays information about existing snapshots including the time these were created.

```
$ beadm list -s BE2
```

```
BE Name Flags Mountpoint Space  Policy Created
------- ----- ---------- ------ ------ ----------------
BE2     -     -          80.83M static 2016-12-19 16:28


    Snapshots              Space  Created
    ---------------------- ------ ----------------
    be://rpool/BE2@backup  67.00K 2016-12-19 16:32
    be://rpool/BE2@backup2 1.00K  2016-12-19 16:57
```

# Viewing Information About Specific Boot Environment Attributes

Instead of displaying a predefined list of BE information, you can filter specific attributes or properties to be displayed by using the -o in combination with a list of properties.

Filtering is helpful especially in a system with a big number of zones. In a system with multiple zone configurations, listing the BEs would take a long time to process depending on the number of zones on the system. Or, the command would generate information beyond what you need or want. The -o option enables you to directly retrieve only the information you need in a shorter time.

In the following example, the command displays only the information about the name, date and time created, and the root dataset of all BEs:

```
$ beadm list -o name,created,root_dataset
BE Name        Created          Root Dataset
-------------- ---------------- --------------
BE1            2015-10-16 10:45 rpool/ROOT/BE1
BE2            2015-10-19 07:04 rpool/ROOT/BE2
BE3            2015-10-20 05:59 rpool/ROOT/BE3
solaris        2015-10-16 08:34 rpool/ROOT/solaris
```

The following example lists only the BE name, subordinate file systems, and creation date for all BEs that match the BeFmri pattern in the command.

```
$ beadm list -o name,datasets,created be://rpool/BE*
BE Name        Datasets                Created
-------------- ----------------        ----------------
BE1            rpool/ROOT/BE1          2015-10-16 10:45
               rpool/ROOT/BE1/var
               rpool/ROOT/BE1/my_fs
BE2            rpool/ROOT/BE2          2015-10-19 07:04
               rpool/ROOT/BE2/var
BE3            rpool/ROOT/BE3          2015-10-20 05:59
```

```
                rpool/ROOT/BE3/var
```

## Viewing Zone BEs Associated with a Specific Global Zone

Use the -Z to view information about zone BEs that are associated with a specific global zone. This option is available only from the global zone.

```
global$ beadm list -Z BE2
BE Name Flags Mountpoint Space Policy Created
------- ----- ---------- ----- ------ ----------------
BE2     NR    /          4.37G static 2015-10-19 07:04

        Non-global Zone BE Names  Flags Mountpoint               Space
        ------------------------- ----- ------------------------ -------
        zbe://zone1/solaris       N     /system/zones/zone1/root 1.04M
        zbe://zone1/zbe2          -     -                        2.00K
        zbe://zone1/zbe1          R     -                        17.10M
        zbe://zone2/solaris       NR    /system/zones/zone2/root 803.80M
        zbe://zone4/solaris       NR    /system/zones/zone4/root 803.87M
        zbe://zone3/solaris       NR    /system/zones/zone3/root 803.84M
```

## Viewing the BEs of a Specific Zone

Use the -z to view information about a zone's BEs. This option is also available only from the global zone.

```
global$ beadm list -z zone1
Non-global Zone BE Names  Parent BE      Flags Mountpoint               Space
------------------------- -------------- ----- ------------------------ -------
zbe://zone1/BE10          -              RO    -                        705.10M
zbe://zone1/BE3           be://rpool/BE3 R     -                        803.73M
zbe://zone1/solaris       be://rpool/BE2 N     /system/zones/zone1/root 1.04M
zbe://zone1/zbe1          be://rpool/BE2 R     -                        17.10M
zbe://zone1/zbe2          be://rpool/BE2 -     -                        2.00K
```

# Changing the Default Boot Environment

Changing an inactive BE to become active means that the BE will be used when the system is next rebooted. This BE then becomes the default environment.

```
$ beadm activate BeFMRI
```

**Note -** When a BE is created, regardless of whether it is active or inactive, an entry is created for the BE on the x86 GRUB menu or the SPARC boot menu. The default BE is the last BE that was activated.

# Mounting and Updating an Inactive Boot Environment

To update packages on an existing, inactive BE, mount BE first, in order to gain access to it, Then, you can update packages on it.

**Note -** When you mount a BE, the supported zones in that environment are mounted relative to the mount points for the environment.

The `beadm mount` command mounts a BE at a specified mount point. The mount point, if it exists, must be empty. If the directory for the mount point does not exist, the `beadm` command creates the directory, then mounts the BE on that directory. However, it remains inactive.

If the specified BE is already mounted, the `beadm mount` command aborts.

## ▼ How to Mount and Update a Boot Environment

1. **Mount the boot environment.**

   ```
   $ beadm mount BeFMRI mount-point
   ```

2. **(Optional) Update packages on the boot environment.**
   For example, to update specific packages on the BE, you would type the following:

   ```
   $ pkg -R /mnt install package-name
   ```

   */mnt* is the mount point for the boot environment.

# Unmounting Boot Environments

When you unmount a BE, the zones in that environment are also unmounted. All mount points are returned to their states prior to being mounted.

**Note -** You cannot unmount the BE that is currently booted.

The command syntax is as follows.

**beadm unmount [-f]** *BeFMRI*

The `-f` option forcefully unmounts the BE even if it is currently busy.

# Destroying a Boot Environment

Destroy unused BEs to free space on the system. The command has the following syntax:

**beadm destroy [[-fF]** *BeFMRI* | *BeFMRI@snapshot***]|[-fF] -O**

`beadm destroy` destroys the specified BE or snapshot. The command prompts the user to provide confirmation before destroying the BE.

- `-f` – Forces destruction of the BE even if it is mounted.
- `-F` – Forces destruction of the BE without prompting for confirmation.
- `-O` – Destroys all orphan non-global zone BEs. Only works in a non-global zone.

Destroying a BE removes its entry from the x86 GRUB menu or the SPARC boot menu. Any zone BEs that are associated with that BE are also destroyed.

**Note -** You cannot destroy the boot environment that is currently booted.

In the following example, the `/tank` and `/tank/home` are shared datasets, which are unaffected when `BE2` is destroyed.

```
$ zfs list
NAME                      USED  AVAIL  REFER  MOUNTPOINT
rpool                    42.5G  24.4G  4.65M  /rpool
rpool/ROOT               25.6G  24.4G    31K  legacy
rpool/ROOT/BE1           8.91M  24.4G  4.17G  /
rpool/ROOT/BE1/var       3.96M  24.4G   276M  /var
rpool/ROOT/BE2           8.91M  24.4G  4.17G  /
rpool/ROOT/BE2/var       3.96M  24.4G   276M  /var
tank                      450K   457G    18K  /export
tank/home                 315K   457G    21K  /export/home

$ beadm destroy BE2
```

```
$ zfs list
NAME                      USED  AVAIL  REFER  MOUNTPOINT
rpool                    42.5G  24.4G  4.65M  /rpool
rpool/ROOT               25.6G  24.4G    31K  legacy
rpool/ROOT/BE1           8.91M  24.4G  4.17G  /
rpool/ROOT/BE1/var       3.96M  24.4G   276M  /var
tank                      450K   457G    18K  /export
tank/home                 315K   457G    21K  /export/home
```

The command also destroys nonshared datasets of the BE. Shared datasets outside the BE root dataset area remain intact.

For more information about destroying orphaned BEs see "About Orphaned Zone Boot Environments and Clones" in *Creating and Using Oracle Solaris Zones*.

# Creating Custom Names for Boot Environments

The `beadm rename` command enables you to rename an existing BE for a more meaningful identity. For example, you could rename a BE to specify how you customized that environment. The BEs dataset name is also changed to conform to the new BE name.

Renaming a BE does not affect zones or zone dataset names. Likewise, the change does not affect the relationships between the zones and their related boot environments.

An active BE cannot be renamed. Only inactive BEs can be renamed. Further, you cannot rename the BE that is currently booted. Activate a different BE first and boot to that environment. Then can rename the first BE. The command syntax is as follows:

**`beadm rename`** *BeFMRI newBeFMRI*

The command fails if the new name already exists.

# Creating Additional Datasets for Boot Environments

To create additional datasets for a BE, use the `zfs` command. The new dataset is located hierarchically under the root dataset of the BE.

This new dataset will not be automatically cloned to existing, associated zone BEs. Also, because the new dataset is under the root dataset of a BE, the new dataset will be private to that BE. The new dataset will not be a shared dataset.

After the dataset has been created under the root dataset, if you clone that entire BE, the new dataset would be included in the new, cloned BE.

**Note -** When creating additional datasets for BEs, the `canmount` ZFS property of the dataset must be set to `noauto`. The dataset's mountpoint is inherited from the BE's root dataset and hence should be mounted at `/myfs` when this BE is in use.

For example, to create a new, non-shared dataset mounted at `/myfs` for the `BE1`, you would type:

```
$ zfs create -o canmount=noauto rpool/ROOT/BE1/myfs
```

For further information, see the `zfs`(8) man page.

# Index

## U



## V



## Z