



Solaris Zones Guide

Version 8.2



SBAAdmin Solaris Zones Guide

The purpose of this guide is to describe the use of SBAAdmin in a Solaris environment where **non-global zones** are present, and to define what aspects of system recovery are and are not supported when using SBAAdmin with non-global zones. This guide does not include all the features and options available with SBAAdmin, but is intended to be used as an addendum to the *SBAAdmin User Guide* and *SBAAdmin Solaris System Recovery Guide*.

Full System Backup of Non-Global Zones

Non-global zones may only be backed up as part of the global zone backup. It is not possible to install SBAAdmin directly into a non-global zone. When performing a *full system* type backup of the global zone, all non-global zones will be backed up unless otherwise excluded. The non-global zones included in the backup may be running or shutdown.

Performing a system backup and recovery of the global zone is done in the same manner as other Solaris system backups. Refer to the *SBAAdmin Solaris System Recovery Guide* for additional details.

Excluding Non-global Zones in the Global Zone Backup

To exclude non-global zone data from a full-system backup of the global zone, it is necessary to determine all mount points and/or ZFS datasets that are used by the non-global zone. These mount points can then be added to an exclude list and applied to a backup job. Refer to the *SBAAdmin Solaris System Recovery Guide* for additional information on creating exclude lists. To determine which mount points are used by non-global zones, use the “*df -Z*” command on the global zone.

System Recovery of Global Zone with Non-global Zones

Solaris 10 Global Zone

After performing system recovery of a Solaris 10 global zone, all non-global zones will start with no additional steps. If the non-global zone is set to *auto-boot*, it will boot automatically after the global zone is restored. If the non-global zone is not set to *auto-boot* it may be started with the *zoneadm* command.

Solaris 11 Global Zone

After restoring a Solaris 11 global zone, it is necessary to **detach** and **reattach** non-global zones before they will boot. The non-global zones configured to *auto-boot* may require **updating the /var ZFS dataset mountpoint property**. Using the *zfs* command, check the mountpoint property of all zoned filesystems. The following is example output for a non-global zone named “z1”.

```
root@solaris:~# zfs get -r mountpoint rpool/zones/z1
NAME                                     VALUE
rpool/zones/z1                          /zones/z1
rpool/zones/z1/rpool                     /zones/z1/root/rpool
rpool/zones/z1/rpool/ROOT                 legacy
rpool/zones/z1/rpool/ROOT/solaris        /zones/z1/root
rpool/zones/z1/rpool/ROOT/solaris/var
```

```

/zones/z1/root/zones/z1/root/var
rpool/zones/z1/rpool/export          /zones/z1/root/export
rpool/zones/z1/rpool/export/home     /zones/z1/root/export/home

```

The highlighted lines above illustrate that the /var filesystem will be mounted out of place. The reason the mounpoint property is as seen above is because this is the mountpoint property for the ZFS dataset at the time of the backup. This property was applied to the dataset at the time of the recovery. It is necessary for you to change the property using the following commands:

```

root@solaris:~# zfs set zoned=off
rpool/zones/z1/rpool/ROOT/solaris/var

root@solaris:~# zfs set mountpoint=/zones/z1/root/var
rpool/zones/z1/rpool/ROOT/solaris/var

```

Now that the mountpoint property is set properly, the zone must be detached. To detach a zone, the following command should be used (replacing “z1” with the actual zone name).

```

root@solaris:~# zoneadm -z z1 detach

```

Once the zone has been detached, it may be necessary to **remove empty directories from the zone’s root directory** prior to re-attaching it. The directories are present based on the mountpoint properties applied at the time of recovery. The following command shows the “z1” **zonepath**.

```

root@solaris:~# zoneadm list -cv
ID NAME          STATUS    PATH                BRAND    IP
0  global         running   /                   solaris  shared
-  z1             installed /zones/z1         solaris  excl

```

The non-global zone’s root directory will be the next subdirectory in the path. Ensure there are no files or directories in the directory and remove any subdirectories with the following command. Empty directories will prevent the ZFS dataset from mounting during the attach.

```

root@solaris:~# rm -rf /zones/z1/root/*

```

If the zones root directory is not empty, attaching the zone will fail with the following error:

```
root@solaris:~# zoneadm -z z1 attach
Progress being logged to
/var/log/zones/zoneadm.20111213T172709Z.z1.attach
ERROR: Error: Command <zfs mount -o
nodevices,mountpoint=/zones/z1/root
rpool/zones/z1/rpool/ROOT/solaris> exited with status 1
ERROR: ZFS temporary mount of rpool/zones/testzone/rpool/ROOT/solaris
on /zones/z1/root failed.
```

Result: Attach Failed.

Now, the zone may be re-attached. During the attach process, the **zoned property** of the ZFS dataset will become enabled again (set to "on").

```
root@solaris:~# zoneadm -z z1 attach
Progress being logged to
var/log/zones/zoneadm.20111213T140136Z.z1.attach
Converting detached zone boot environment `solaris`.
Attaching...
  Installing: Using existing zone boot environment
  Zone BE root dataset: rpool/zones/z1/rpool/ROOT/solaris
  Cache: Using /var/pkg/publisher.
Updating Image format
  Updating non-global zone: Linking to image /.
  Updating non-global zone: Auditing packages.
No updates necessary for this image.

  Updating non-global zone: Zone updated.
  Result: Attach Succeeded.
Log saved in non-global zone as /zones/z1/root/var/log/zones/
zoneadm.20111213T140136Z.z1.attach
```

The zone may now be booted using the *zoneadm* command.