

Migrating from Solaris 10 UFS boot to Solaris 10 ZFS boot

By Haim Tzadok and Cyril Plisko, Grigale LTD.

About the Authors

Haim Tzadok and Cyril Plisko are the owners of Grigale LTD. A Software house specializing on [Open]Solaris Administration, development and training. They are leaders and participating on many OpenSolaris communities. They can be reached at: info at grigale dot com.

Abstract

In this paper we will explain about migrating from Solaris 10 previous updates(Sparc Editions) to Solaris 10 update 6 or later Sparc edition. Using live upgrade migration process, while manually creating zfs root pool and copying the ufs root filesystem into it using live-upgrade process.

Introduction

ZFS boot is a unique feature introduced in Solaris 10 update 6(10/08).

While performing a fresh install to Solaris 10 update 6 seems to be easier to the Novice System Administrator. Migration from Solaris 10 ufs root filesystem to zfs boot involves upgrading the Solaris

System Requirements

In order to perform this migration you will need the following:

1. Solaris 10 update 6 sparc DVD or iso file.
2. 2 system disks, one occupying the current Solaris 10 ufs root filesystem, the second – free for use.

Prerequisites:

- Prior to starting this procedure you should install latest 10_Recommended.zip patches. Otherwise, the whole process might encounter problems.
- Prior to starting this procedure it is recommended to consolidate all your /, /var, /opt filesystems to one single filesystem.
- For cases that you use Solaris volume manager for mirroring it is recommended to remove all SVM configuration.

You are encouraged to perform a full backup of your root filesystem for safety reasons.

List of the steps

Step 1: Upgrading Live Upgrade(LU) packages on the Solaris 10 update 5 boot environment.

Step 2: Creating Boot Environemnt called sol10u6_stage1 (still on ufs)

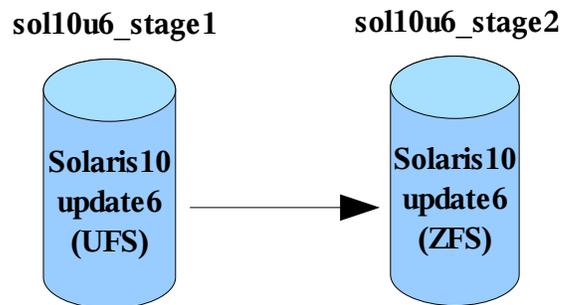
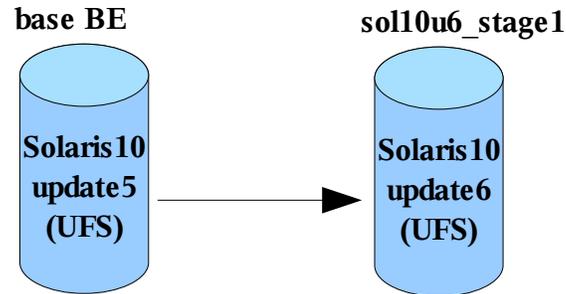
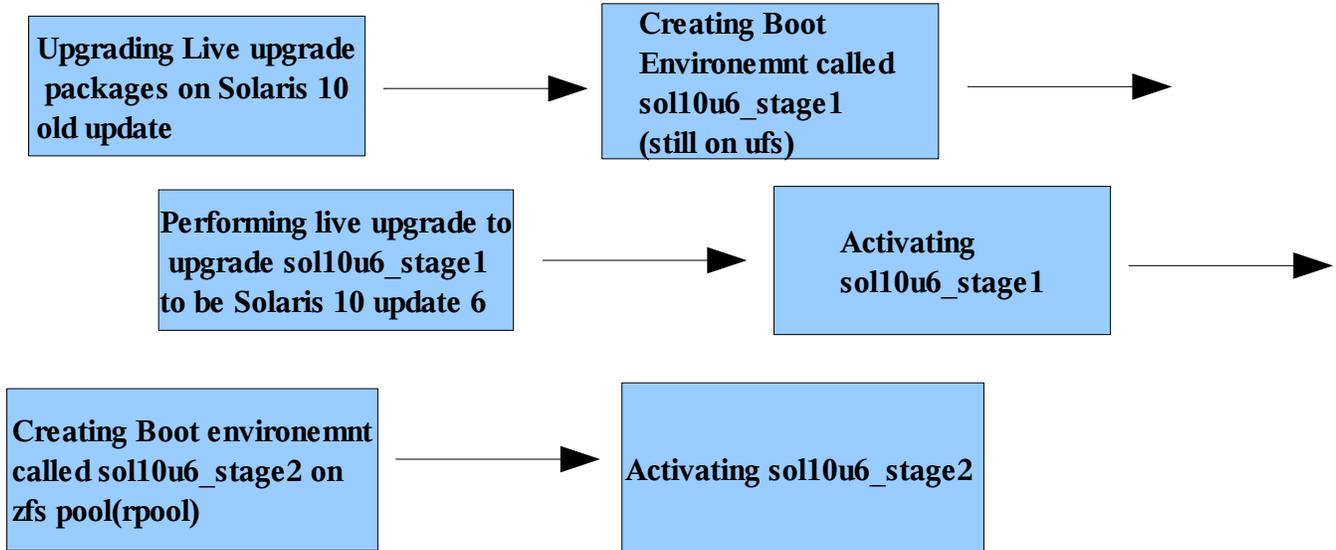
Step 3: Performing live upgrade to upgrade sol10u6_stage1 to be Solaris 10 update 6

Step 4: Activating sol10u6_stage1

Step 5: Creating Boot environment called sol10u6_stage2 on zfs pool(rpools)

Step 6: Activating sol10u6_stage2 (Optional)

Diagram of the process



Step 1: Upgrading Live Upgrade(LU) packages on the Solaris 10 update 5 boot environment.

LU packages are available on Solaris 10 update 6 sparc DVD or iso.

Mount your DVD or iso file and Install the packages by typing:

`/cdrom/cdrom0/s0/installer`

(if you have access to X-Server it will open a GUI window, otherwise it will continue in CLI mode).

Example CLI installation:

```
Java Accessibility Bridge for GNOME loaded.
Solaris Web Start will assist you in installing software for Solaris 10
Software.
  <Press ENTER to continue> ENTER
IMPORTANT NOTICE
Patches Needed to Run Solaris Live Upgrade
Correct operation of Solaris Live Upgrade requires that a limited set of patch
revisions be installed for a given OS version.
Before installing or running Live Upgrade, you are required to install a
limited set of patch revisions. Make sure you have the most recently updated
patch list by consulting sunsolve.sun.com. Search for the info doc 72099 on the
SunSolve(tm) web site.
Live Upgrade may fail to work properly if the latest limited set of patch
revisions are not installed on this system.
  <Press ENTER to continue> ENTER
Please select the type of install to perform from the following choices:
Typical - Software will be installed with the most common
options. Recommended for most users.
Custom - You may choose the options you want to install.
Recommended for advanced users.
1. Typical
2. Custom
  What would you like to do [1]?
Checking disk space.

The following items will be installed:
Product: Live Upgrade
Location: /
Size: 2.49 MB
-----
  Live Upgrade (root)  1.27 MB
Live Upgrade (usr)   1.21 MB
Live Upgrade (config) 17.26 KB
Ready to Install
1. Install Now
2. Start Over
3. Exit Installation
  What would you like to do [1]? ENTER
```

Step 2: Creating Boot Environment called sol10u6_stage1 (still on ufs)

As a first stage, upgrade to Solaris 10 update 6 on a ufs filesystem.

Your current Solaris 10 update 5 boot environment will be called base. Your new Solaris 10 update 6 on ufs filesystem will be called sol10u6_stage1.

Copy the vtocs of the current root disk to the free disk.

Assuming c1t1d0s2 is the root disk and c1t2d0s2 is the free disk type the following:

```
prtvtoc /dev/rdisk/c1t1d0s2 | fmthard -s - /dev/rdisk/c1t2d0s2
```

Create base and sol10u6_stage1 boot environments by typing:

```
lucreate -c base -m /:/dev/dsk/c1t2d0s0:ufs -m -:/dev/dsk/c1t2d0s1:swap -n sol10u6_stage1
```

(This may take a substantial quite of time depend on the capacity of your system).

Output example:

```
(root@b1000: /)# lucreate -c base -m /:/dev/dsk/c1t2d0s0:ufs -m -:/dev/dsk/c1t2d0s1:swap -n sol10u6_stage1
Discovering physical storage devices
Discovering logical storage devices
Cross referencing storage devices with boot environment configurations
Determining types of file systems supported
Validating file system requests
Preparing logical storage devices
Preparing physical storage devices
Configuring physical storage devices
Configuring logical storage devices
Analyzing system configuration.
No name for current boot environment.
Current boot environment is named <base>.
Creating initial configuration for primary boot environment <base>.
The device </dev/dsk/c1t1d0s0> is not a root device for any boot environment; cannot get BE ID.
PBE configuration successful: PBE name <base> PBE Boot Device </dev/dsk/c1t1d0s0>.
Comparing source boot environment <base> file systems with the file
system(s) you specified for the new boot environment. Determining which
file systems should be in the new boot environment.
Updating boot environment description database on all BEs.
Searching /dev for possible boot environment filesystem devices
```

Continue output example:

```
Updating system configuration files.
The device </dev/dsk/c1t2d0s0> is not a root device for any boot environment; cannot get BE ID.
Creating configuration for boot environment <sol10u6_stage1>.
Source boot environment is <base>.
Creating boot environment <sol10u6_stage1>.
Creating file systems on boot environment <sol10u6_stage1>.
Creating <ufs> file system for </> in zone <global> on </dev/dsk/c1t2d0s0>.
Mounting file systems for boot environment <sol10u6_stage1>.
Calculating required sizes of file systems for boot environment <sol10u6_stage1>.
Populating file systems on boot environment <sol10u6_stage1>.
Checking selection integrity.
Integrity check OK.
Populating contents of mount point </>.
Copying.
Creating shared file system mount points.
Creating compare databases for boot environment <sol10u6_stage1>.
Creating compare database for file system </>.
Updating compare databases on boot environment <sol10u6_stage1>.
Making boot environment <sol10u6_stage1> bootable.
Population of boot environment <sol10u6_stage1> successful.
Creation of boot environment <sol10u6_stage1> successful.
```

Check that you have 2 boot environments by typing:

lustatus

Output example:

```
(root@b1000: /)# lustatus
Boot Environment      Is      Active Active   Can   Copy
Name                  Complete Now   On Reboot Delete Status
-----
base                  yes    yes  yes    no    -
sol10u6_stage1       yes    no   no     yes   -
```

Step 3: Performing live upgrade to upgrade sol10u6_stage1 to be Solaris 10 update 6

Perform live upgrade to sol10u6_stage1 by typing:

```
luupgrade -u -n sol10u6_stage1 -s /net/pronto/export/install/sparc/s10u6
```

Output example:

```
(root@b1000: /)# luupgrade -u -n sol10u6_stage1 -s /net/pronto/export/install/sparc/s10u6/

42092 blocks
miniroot filesystem is <lofs>
Mounting miniroot at </net/pronto/export/install/sparc/s10u6//Solaris_10/Tools/Boot>
Validating the contents of the media </net/pronto/export/install/sparc/s10u6/>.
The media is a standard Solaris media.
The media contains an operating system upgrade image.
The media contains <Solaris> version <10>.
Constructing upgrade profile to use.
Locating the operating system upgrade program.
Checking for existence of previously scheduled Live Upgrade requests.
Creating upgrade profile for BE <sol10u6_stage1>.
Determining packages to install or upgrade for BE <sol10u6_stage1>.
Performing the operating system upgrade of the BE <sol10u6_stage1>.
CAUTION: Interrupting this process may leave the boot environment unstable
or unbootable.
Upgrading Solaris: 100% completed
Installation of the packages from this media is complete.
Updating package information on boot environment <sol10u6_stage1>.
Package information successfully updated on boot environment <sol10u6_stage1>.
Adding operating system patches to the BE <sol10u6_stage1>.
The operating system patch installation is complete.
INFORMATION: The file </var/sadm/system/logs/upgrade_log> on boot
environment <sol10u6_stage1> contains a log of the upgrade operation.
INFORMATION: The file </var/sadm/system/data/upgrade_cleanup> on boot
environment <sol10u6_stage1> contains a log of cleanup operations required.
WARNING: <4> packages failed to install properly on boot environment <sol10u6_stage1>.
INFORMATION: The file </var/sadm/system/data/upgrade_failed_pkgadds> on
boot environment <sol10u6_stage1> contains a list of packages that failed
to upgrade or install properly.
INFORMATION: Review the files listed above. Remember that all of the files
are located on boot environment <sol10u6_stage1>. Before you activate boot
environment <sol10u6_stage1>, determine if any additional system
maintenance is required or if additional media of the software
distribution must be installed.
The Solaris upgrade of the boot environment <sol10u6_stage1> is partially complete.
```

Step 4: Activating sol10u6_stage1

In order to activate sol10u6_stage1 type:

```
luactivate sol10u6_stage1
```

```
init 6
```

Step 5: Creating Boot environment called sol10u6_stage2 on zfs pool(rpoul)

We need to delete base boot environment by typing:

```
ludelete base
```

Output example:

```
(root@b1000: Product)# ludelete base
Determining the devices to be marked free.
Updating boot environment configuration database.
Updating boot environment description database on all BEs.
Updating all boot environment configuration databases.
Boot environment <base> deleted.
```

Now, we need to create a new vtoc(partition table) to the first disk, so it will only have partition 0 which is a copy of partition 2 – the backup partition. All other partitions should be empty.

We will do this by typing:

```
format
```

Following the instructions in the output example:

```
(root@b1000: /)# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. c1t1d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
     /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w2100002037e3b3af,0
  1. c1t2d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
     /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w2100002037b2d617,0
Specify disk (enter its number): 0 ENTER
p-> ENTER p-> ENTER

partition> modify
Select partitioning base:
  0. Current partition table (original)
  1. All Free Hog
Choose base (enter number) [0]? 1

Part  Tag  Flag  Cylinders    Size      Blocks
  0  root  wm    0          0  (0/0/0)      0
  1  swap  wu    0          0  (0/0/0)      0
  2  backup wu    0 - 24619  33.92GB (24620/0/0) 71127180
  3  unassigned  wm    0          0  (0/0/0)      0
  4  unassigned  wm    0          0  (0/0/0)      0
  5  unassigned  wm    0          0  (0/0/0)      0
  6  usr    wm    0          0  (0/0/0)      0
  7  unassigned  wm    0          0  (0/0/0)      0
```

Continue output example:

```
Do you wish to continue creating a new partition
table based on above table[yes]? ENTER
Do you wish to continue creating a new partition
table based on above table[yes]?
Free Hog partition[6]? 0
Enter size of partition '1' [0b, 0c, 0.00mb, 0.00gb]: 0
Enter size of partition '3' [0b, 0c, 0.00mb, 0.00gb]: 0
Enter size of partition '4' [0b, 0c, 0.00mb, 0.00gb]: 0
Enter size of partition '5' [0b, 0c, 0.00mb, 0.00gb]: 0
Enter size of partition '6' [0b, 0c, 0.00mb, 0.00gb]: 0
Enter size of partition '7' [0b, 0c, 0.00mb, 0.00gb]: 0

Part  Tag  Flag  Cylinders      Size      Blocks
0   root  wm   0 - 24619    33.92GB (24620/0/0) 71127180
1   swap  wu    0          0    (0/0/0)      0
2  backup  wu   0 - 24619    33.92GB (24620/0/0) 71127180
3 unassigned  wm    0          0    (0/0/0)      0
4 unassigned  wm    0          0    (0/0/0)      0
5 unassigned  wm    0          0    (0/0/0)      0
6   usr  wm    0          0    (0/0/0)      0
7 unassigned  wm    0          0    (0/0/0)      0

Okay to make this the current partition table[yes]? ENTER
Enter table name (remember quotes): test

Ready to label disk, continue? y

partition> q-> ENTER q-> ENTER
```

Now we will check that our zpool supports zfs boot property by typing:

zpool upgrade -v

Output example:

```
(root@b1000: /)# zpool upgrade -v
This system is currently running ZFS pool version 10.
The following versions are supported:

VER DESCRIPTION
-----
1  Initial ZFS version
2  Ditto blocks (replicated metadata)
3  Hot spares and double parity RAID-Z
4  zpool history
5  Compression using the gzip algorithm
6  bootfs pool property
7  Separate intent log devices
8  Delegated administration
9  refquota and reservation properties
10 Cache devices

For more information on a particular version, including supported releases, see:
http://www.opensolaris.org/os/community/zfs/version/N
Where 'N' is the version number.
```

If you see that you have bootfs pool property you may proceed to create a root pool by typing:

```
zpool create rpool c1t1d0s0
```

Check that the rpool is created by typing:

```
zpool list
```

Output example:

```
(root@b1000: /)# zpool list
NAME  SIZE  USED  AVAIL  CAP  HEALTH  ALTROOT
rpool 33.8G 111K 33.7G  0%  ONLINE  -
```

proceed to create live upgrade environment on rpool by typing:

```
lucreate -c sol10u6_stage1 -n sol10u6_stage2 -p rpool
```

Output example:

```
(root@b1000: /)# lucreate -c sol10u6_stage1 -n sol10u6_stage2 -p rpool
Analyzing system configuration.
Comparing source boot environment <sol10u6_stage1> file systems with the
file system(s) you specified for the new boot environment. Determining
which file systems should be in the new boot environment.
Updating boot environment description database on all BEs.
Updating system configuration files.
The device </dev/dsk/c1t1d0s0> is not a root device for any boot environment; cannot get BE ID.
Creating configuration for boot environment <sol10u6_stage2>.
Source boot environment is <sol10u6_stage1>.
Creating boot environment <sol10u6_stage2>.
Creating file systems on boot environment <sol10u6_stage2>.
Creating <zfs> file system for </> in zone <global> on <rpool/ROOT/sol10u6_stage2>.
Populating file systems on boot environment <sol10u6_stage2>.
Checking selection integrity.
Populating contents of mount point </>.
Copying.
Creating shared file system mount points.
Creating compare databases for boot environment <sol10u6_stage2>.
Creating compare database for file system </rpool/ROOT>.
Creating compare database for file system </>.
Updating compare databases on boot environment <sol10u6_stage2>.
Making boot environment <sol10u6_stage2> bootable.
Creating boot_archive for /.alt.tmp.b-oL.mnt
updating /.alt.tmp.b-oL.mnt/platform/sun4u/boot_archive
Population of boot environment <sol10u6_stage2> successful.
Creation of boot environment <sol10u6_stage2> successful.
```

Step 6: Activating sol10u6_stage2

Check the status of sol10u6_stage2 by typing:

lustatus

```
(root@b1000: /)# lustatus
Boot Environment      Is   Active Active  Can  Copy
Name                  Complete Now  On Reboot Delete Status
-----
sol10u6_stage1       yes  yes  yes   no  -
sol10u6_stage2       yes  no   no    yes -
```

If everything is ok, Activate the sol10u6_stage2 by typing:

luactivate sol10u6_stage2

init 6

After performing the reboot, check to see that indeed you system is using zfs boot by typing:

zpool list

zfs list

Output example:

```
(root@b1000: /)# zpool list
NAME  SIZE  USED  AVAIL  CAP  HEALTH  ALTROOT
rpool 33.8G 11.1G 22.6G  32%  ONLINE -
(root@b1000: /)# zfs list
NAME                USED  AVAIL  REFER  MOUNTPOINT
rpool               15.0G 18.2G 21.5K  /rpool
rpool/ROOT           9.12G 18.2G  18K  /rpool/ROOT
rpool/ROOT/sol10u6_stage2 9.12G 18.2G 9.12G  /
rpool/dump           2.01G 18.2G 2.01G  -
rpool/swap           3.91G 22.1G  16K  -
```

Now you have your system upgraded to Solaris 10 update 6 with ZFS boot.

(OPTIONAL)Step 6: Activating ZFS mirroring

After checking that indeed everything is running and working. If you want to activate mirroring:

Delete sol10u6_stage1 by typing:

```
ludelete sol10u6_stage1
```

Output example:

```
Determining the devices to be marked free.  
Updating boot environment configuration database.  
Updating boot environment description database on all BEs.  
Updating all boot environment configuration databases.  
Boot environment <sol10u6_stage1> deleted.
```

Now Copy the first disk vtoc to the second disk by typing:

```
prtvtoc /dev/rdisk/c1t1d0s2 |fmthard -s - /dev/rdisk/c1t2d0s2
```

Add c1t2d0s0 to the rpool pool by typing:

```
zpool attach rpool c1t1d0s0 c1t2d0s0
```

Now check that indeed rpool is in mirror state by typing:

```
zpool status
```

Output example:

```
(root@b1000: /)# zpool status  
pool: rpool  
state: ONLINE  
status: One or more devices is currently being resilvered. The pool will  
continue to function, possibly in a degraded state.  
action: Wait for the resilver to complete.  
scrub: resilver in progress for 0h0m, 2.36% done, 0h25m to go  
config:  
  
NAME      STATE  READ WRITE CKSUM  
rpool    ONLINE  0   0   0  
mirror   ONLINE  0   0   0  
  c1t1d0s0 ONLINE  0   0   0  
  c1t2d0s0 ONLINE  0   0   0  
  
errors: No known data errors
```

Wait for resilvering(copy process) to finish.

Now you have a mirror configuration on your system using ZFS.