

What has ZFS ever done for us?

ZFS boot in the Mission Critical Environment

Philip Scarlett

<http://www.linkedin.com/in/philipscarlett>

What's the problem?

- We have to patch
 - We will need to make config changes
 - We will need to upgrade
-
- Downtime
 - Cost
 - Risk

“It's working – why change it?”

Patching - Traditional vs. ZFS

SVM / VxVM ZFS

| | | |
|--------------|------------------------------|----------------|
| Day 1 - 1h | Backup (u) Backup (zfs send) | Day 1 - 5 mins |
| Day 1 - 1.5h | Split m zfs snap/clone | Day 1 - 5 mins |
| Day 1 - ??h | Pat Patch | Day 1 - ??h |
| Day 2 - 6 | Cooling d "No worries, mate" | Day 2 - 6 |
| Day 7 - 1.5h | Re-mirror Day Off? | Day 7 |

The Money Slide


Savings of ZFS over traditional patching

- 250 servers
- Patch twice a year
- Saving of 4 hours per server
- £50/hour

$$\text{Total} = 250 * 2 * 4 * 50 =$$

£100,000

What could we do with £100,000?



ZFS pre-change process

```
# zfs snapshot rpool/ROOT/std_s10u8@CR54321
# zfs send <snapshot> | ssh <backup_server> zfs recv <backup_pool>/<target>
# zfs clone rpool/ROOT/std_s10u8@CR54321 rpool/ROOT/std_s10u8_CR54321
# zfs set canmount=noauto rpool/ROOT/std_s10u8_CR54321
# zfs set mountpoint=/ rpool/ROOT/std_s10u8_CR54321

# echo ``
#----- Pre-Patch CR54321 -----
#
title Solaris 10 - Pre-Patch CR54321
bootfs rpool/ROOT/std_s10u8_CR54321
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS
module /platform/i86pc/boot_archive
#
# ----- Pre-Patch CR54321 Failsafe -----
#
title Solaris 10 - Pre-Patch CR54321 - Failsafe
bootfs rpool/ROOT/std_s10u8_CR54321
kernel /boot/multiboot -s -B console=ttya
module /boot/amd64/x86.miniroot/safe
#
#----- END Pre-Patch CR54321 -----
" >> /rpool/boot/grub/menu.lst
```

ZFS fail-back process


- Boot failsafe (either original or cloned version)
- Don't elect to mount any discovered root filesystem

```
# zpool import -f rpool
# zfs promote rpool/ROOT/std_s10u8_CR54321
# zfs destroy -r rpool/ROOT/std_s10u8
# zfs rename rpool/ROOT/std_s10u8_CR54321 rpool/ROOT/std_s10u8
# reboot
```

- Once system is up, update the `/rpool/boot/grub/menu.lst` appropriately

Note:- This maintains any snapshots prior to cloning

Traditional pre/post change process



ZFS Boot

Flexible Hardware Choices

- Supports all H/W solutions
 - Single disk
 - S/W mirroring
 - H/W mirroring
 - SAN Boot
 - SSD

More ZFS Boot Features

- Compression
- Swap & Dump management
- ZFS Flash install in U8 (10/09)
- ZFS Supported in Jumpstart
- ZFS Backup Solutions

- *De-dupe?*
- *zpool split/join?*

Summary

- UFS has served us well
- Other VMs are “heavyweight” and/or costly
- UNIX skill sets are changing
- Simplified System Administration

Handy Links

Big Admin

<http://www.sun.com/bigadmin/topics/zfs/>

Boot Disk Management – John Howard, David Deeths

<http://www.sun.com/books/catalog/howard.xml>

OpenSolaris ZFS Resources

<http://www.opensolaris.org/os/community/zfs/>

Solaris 10 – 10/09

<http://docs.sun.com/app/docs/coll/1236.11?l=en>

Lori Alt – Staff Engineer at Sun

http://blogs.sun.com/storage/entry/what_we_re_watching_lori

OpenBackups using ZFS – LOSUG Presentation - Sally Houghton

<http://opensolaris.org/os/project/losug/files/June2009/>

Solaris Internals Wiki – ZFS Root Pool Info

http://www.solarisinternals.com/wiki/index.php/ZFS_Troubleshooting_Guide