



Zones + ZFS on personal server

Lubos Kocman
UNIX IT Specialist IBM
CZOSUG Leader

Agenda

- What are zones?
- Why to use zones?
- Zones + ZFS simplifies filesystem management
- How to set up new zone in OpenSolaris
- Some tips at the end
- QA

What are zones?

- Virtualization technique
- Can be described as Operating System level virtualization
- Two main purposes
 - Application isolation
 - Simple resource management
- Supports migration between systems
- Zones are using shared or emulated kernel
- Alternatives AIX's wpars or BSD's jails

Types of zones

- 2 main types of zones
 - Supporting system call translation
 - native (default in Solaris 10)
 - ipkg (default since OpenSolaris 2008.11)
 - cluster (solaris cluster zones)
 - labeled (used for zones in Trusted Environment)
 - Not-Supporting system call transl. (BrandZ)
 - lx (supporting linux 2.4 kernel)
 - lx26 (supporting linux 2.6 kernel – only in OSOL)
 - solaris8, solaris9
 - S10brand (available in OSOL)

Some facts about zones



?

Some facts about zones

- Maximum is 8191 non-global zones per system
- You'll basically use only two commands zoneadm and zonecfg.
- You can't mount NFS share from the global zone in a non-global zone (*no autofs*) and you can't run NFS server in non-global zone
- You can use lx/lx26 to run your linux unportable Linux applications in Solaris

Some facts about zones

- Zone can be easily migrated between systems using
 - pax archives / scp on VxFS/UFS and zoneadm dettach/attach.
 - **Or more simply and faster using zfs send / scp / receive and zoneadm attach/dettach on ZFS**


Why to use zones on your server?

- Zones can make your system safer by isolating unstable applications and limiting them resources (memory, max. cpu usage ...)
- You can provide an access to an isolated environment without having a fear that somebody will put down the whole box
- Good as a system for building. Zone can be easily re-installed/rollbacked without bigger downtime. A clean environment anytime.

What advantages does zones/zfs have

- **Easy to rollback / backup the whole zone**
- You can assign zfs dataset to a zone, so privileged user can manage it.
 - This includes snapshots/rollbacks new datasets
 - New dataset can be very useful when you need to set some attributes e.g. compression, mountpoint or quota on specific places
- You can add also new FS (not dataset)
 - But users will not be able to manage it

Setting up zones



Basic setup of the zones

!! Be sure that zones service is online (\$svcs zones)

```
•$pfexec zfs create rpool/zones
$pfexec zonecfg -z myzone
> create -t SUNWipkg
> set zonepath=/rpool/zones/myzone
>add net
  > set physical=bge0
  > set address=192.168.0.10/24
  > set defrouter=192.168.0.254
  >end
> verify
>commit
>^D
```

Resources

- Linux in Solaris - <http://opensolaris.cz>
 - Step by step debian/centos 5 installations
 - For non-czech guys “let the google translate be with you”. Honestly **code doesn't have to be translated**. And see mentioned quote or docs.sun.com for some theory.

- <http://docs.sun.com>

Installing and booting and the zone

```
$ pfexec zoneadm -z myzone install
```

In the case that something went wrong
Just uninstall zone and you can try it again

```
( $ pfexec zoneadm -z myzone uninstall )
```

```
$ pfexec zoneadm -z myzone boot
```

First configuration (interactive)

- The system is after installation unconfigured and you have to go through few forms. Like hostname, default terminal, root password, ip address ...
- Console of the zone can be accessed using
 - `$ pfexec zlogin -C myzone`
 - **You can leave console by pressing ~.**
- You'll be automatically requested to fill up the data after the boot process

First configuration (non-interactive)

```
globalzone# cat > /rpool/zones/myzone/root/etc/sysidcfg
system_locale=C
terminal=dtterm
network_interface=primary {
hostname=myzone
}
timeserver=localhost
security_policy=NONE
name_service=NONE
timezone=US/Eastern
root_password=""
^D
```

Example of installing of lx zone



DEMO

Some useful commands for zones

```
$ for x in zones_cmds; do tell_something(x);done
```

Logging in

○ Using zlogin

```
$ pfexec zlogin [-C ] myzone [ -l username ]  
[command]
```

- -C stands for console – you'll need to enter followin escape sequence “~.” to leave console

○ You can use ssh as well (including X11-forwarding), but you should configure sshd/network first

Booting halting the zone

○ Boot zone

- zoneadm -z myzone boot
- Autoboot can be set during configuration
 - set autoboot=true

○ Halt zone

- zoneadm -z myzone halt

○ Reboot zone

- zoneadm -z myzone reboot

Listing zones

- \$ pfexec zoneadm list -vc

....

Possible states of zone

- Configured: configuration was completed and committed
- Incomplete: Transition state during install or uninstall operation
- Installed: the packages have been successfully installed
- Ready: the virtual platform has been established
- Running: the zone booted successfully and is now running

Possible states of zone

- Shutting down: the zone is in the process of shutting down - this is a temporary state, leading to "Down"
- Down: the zone has completed the shut down process and is down - this is a temporary state, leading to "Installed"

Getting info and reconfiguring zone

- zonecfg -z myzone info

```
$pfexec zonecfg -z myzone
```

```
> help # ;-)
```

- You'll generally use add/remove/commit

Setting up resources (SUNWrcap)

- Simple example
- You need to have SUNWrcap installed

```
$ pfexec zonecfg -z myzone
```

```
add capped-cpu
```

```
set ncpus=0.4 # 1 stands for once cpu ...
```

```
end
```

```
add capped-memory
```

```
set physical=1G
```

```
set swap=2G
```

```
set locked=10M
```

Setting resources

- You can use add and help to get some hints
 - \$ pfexec zonecfg -z myzone
add
... list of what can be added
add fs
help
... usage
- Visit <http://docs.sun.com> form more info about resource management pools and projects

Working with ZFS and zones

- New filesystem to zone can be easily added using only few steps

- Adding zfs dataset (users are able to manage it)

```
$ pfexec zfs create rpool/mydataset
```

```
$ pfexec zonecfg -z myzone
```

```
add dataset
```

```
    set name=rpool/mydataset
```

```
end
```

```
commit
```

```
^D
```

Working with ZFS and zones

- Adding zfs as fs (users are not-able to manage it)

```
$ pfexec zfs create rpool/mydataset
```

```
$ pfexec zonecfg -z myzone
```

```
add fs
```

```
    set type=zfs
```

```
    set dir=/mount/point
```

```
    set special=/rpool/mydataset
```

```
end
```

```
commit
```

```
^D
```

Removing the zone

- First you should remove the data occupied by zone
 - `pfexec zoneadm -z myzone uninstall`
- Then delete info from the system
 - `pfexec zonecfg -z myzone delete`

Some tips on the end

- You can't add zfs dataset to a linux zone
- Listing of processes in global zone will show you even processes from the non-global zones
 - Use **ps -ef -z zonename** instead

Q&A

Any questions?

A question for a price ;-)



What does number 8191 stands on?

Thank you!

THE END