Why I prefer thick jails over thin jails

Dan Langille
EuroBSDCon 2019
Lillehammer

@dlangille
https://dan.langille.org/

Disclaimer

- Don’t do what I’m doing just because I’m doing it
- It’s right for me - now
- Needs change over time
- Use this talk as catalyst for thoughts about your systems

Terminology

- host - a FreeBSD install
- jail - a VM on a host

What are jails?

- FreeBSD 4.x (2000)
- Poul-Henning Kamp
- simple explanation: simple `chroot`
- security by isolating stuff
- jails can’t see into host/other jails
- good for trying new stuff/isolating stuff

thick jails

- traditional jail
- complete OS installed
- manage it, more or less, like a host
- not a clone of another system
- zfs copy - OK
- zfs clone - no

thin jails

- a jail which is not thick
- ezjail - basejail - the base OS is supplied via a nullfs mount
- created via **`zfs clone`**
- jails designed to run exactly one application

jail managers

- There are a few, but I have used only two

ezjail

- first released 2005-10-14
- I used ezjail since at least 2008
- thin jail
- basejail shared by all jails
- update basejail: all your jails are now updated
- but not quite….

iocage

- Added to FreeBSD ports tree in 2014
- I’ve used it since at least 2015
- thick jails
- can use clones… sort of thin
- originally written as a shell script
- now in Python

Why did I convert?

- outdated jails (my fault) - no mergemaster
- upgrade the basejail, upgrade ALL the jails
- can easily mix jail versions (e.g. 11.3 and 12.0)
- disk space

Why should you convert?

- clones are good for short-lived jails
- Don’t upgrade clones to next release - space penalty
- can run `freebsd-update` from inside jail (not recommended if using a jail manager)
- disk space

The script

- converted from `ezjail` to `iocage` in mid-2019
- https://github.com/dlangille/thin_to_thick
- replaces basejail with jail (proper)
- “This tool is designed to allow you to copy an existing thin jail into a thick jail, ignoring the bits provided by the basejail.”
- specific to ezjail-base jails, but easily modified

Typical ezjail basejail

```
$ ls -l /usr/jails/newjail
total 203
basejail
bin     -> /basejail/bin
boot    -> /basejail/boot
lib     -> /basejail/lib
libexec -> /basejail/libexec
rescue  -> /basejail/rescue
sbin    -> /basejail/sbin
sys     -> usr/src/sys
```

The steps

- `iocage create --thickjail -r 12.0-RELEASE -n myjail`
- set config for new iocage jail: hostname, IP address, etc
- `zfs snapshot -r system/iocage/jails/myjail@clean`
- `ezjail-admin stop myjail`
- `thin_to_thick.sh /usr/jails/newjail \`
  `/usr/jails/myjail/ \`
  `/iocage/jails/myjail/root`
- `iocage start myjail`

Post conversion

- `ezjail-admin config -r norun myjail`
- `iocage set boot=on myjail`

Thick is for you!

- You pick and choose when some jails are upgraded
- You **want** to run jails which are on different versions
- Friends don’t let friends clone jails

Thin is for you!

- Saves space!
- Easy one-step upgrade of all jails
- `rc.d`? Who needs that!
- I'll mergemaster them later!

Template jails

- all my jails have a common sub-set of packages
- e.g. sudo, anvil, bash, joe, xtail
- the same `/etc/resolv.conf`
- the same pkg.conf files
- but I do not use template jails
- Ansible scripts will install what I need, post jail-setup

monitoring tips

- `/usr/local/etc/periodic/security/405.pkg-base-audit`
    - installed by security/base-audit
- `/usr/local/etc/periodic/security/410.pkg-audit`
    - installed by ports-mgmt/pkg
    - code at https://github.com/dlangille/freebsd-nagios-jail

/etc/periodic.conf

```
# for security/405.pkg-base-audit
security_status_baseaudit_enable="YES"
security_status_baseaudit_jails="*"

#for 410.pkg-audit
security_status_pkgaudit_expiry=1

# for many scripts including 405 & 410
pkg_jails='*'
```

Just say no to jail managers!

- Sometimes the jail manager breaks
- when it does, your jails can go offline
- I first used jails without a jail manager, I can do it again

Always use a jail manager!

- The tasks around managing a jail are tedious and boring
- jail managers will do the **`zfs create`** (& more) for you!
- You'll start writing scripts for managing jails
- Let someone else do that!

updating the errant jail

- jail configuration files (inside the jail) are out of date?
- run mergemaster
- but instead of mergemaster….

etcmerge / etcupdate

- run instead of mergemaster
- sysutils/etcmerge
- etcupdate is in base since FreeBSD 10.0
- They both do automatic 3-way merges

How I update my jails

- locage update
- soon to be replaced by freebsd-update

Blame Peter Wemm

- Peter tweeted at me to use plain jails
- I ignored him
- … for a while

one last thought

- very scary

The dark side: jail.conf

- very tempting
- wonderful use of default values
- some tasks are tedious
- I might write a script
- or two
- and package them
- creating a new jail manager….

Recommended reading

- Michael W Lucas
- FreeBSD Mastery: Jails

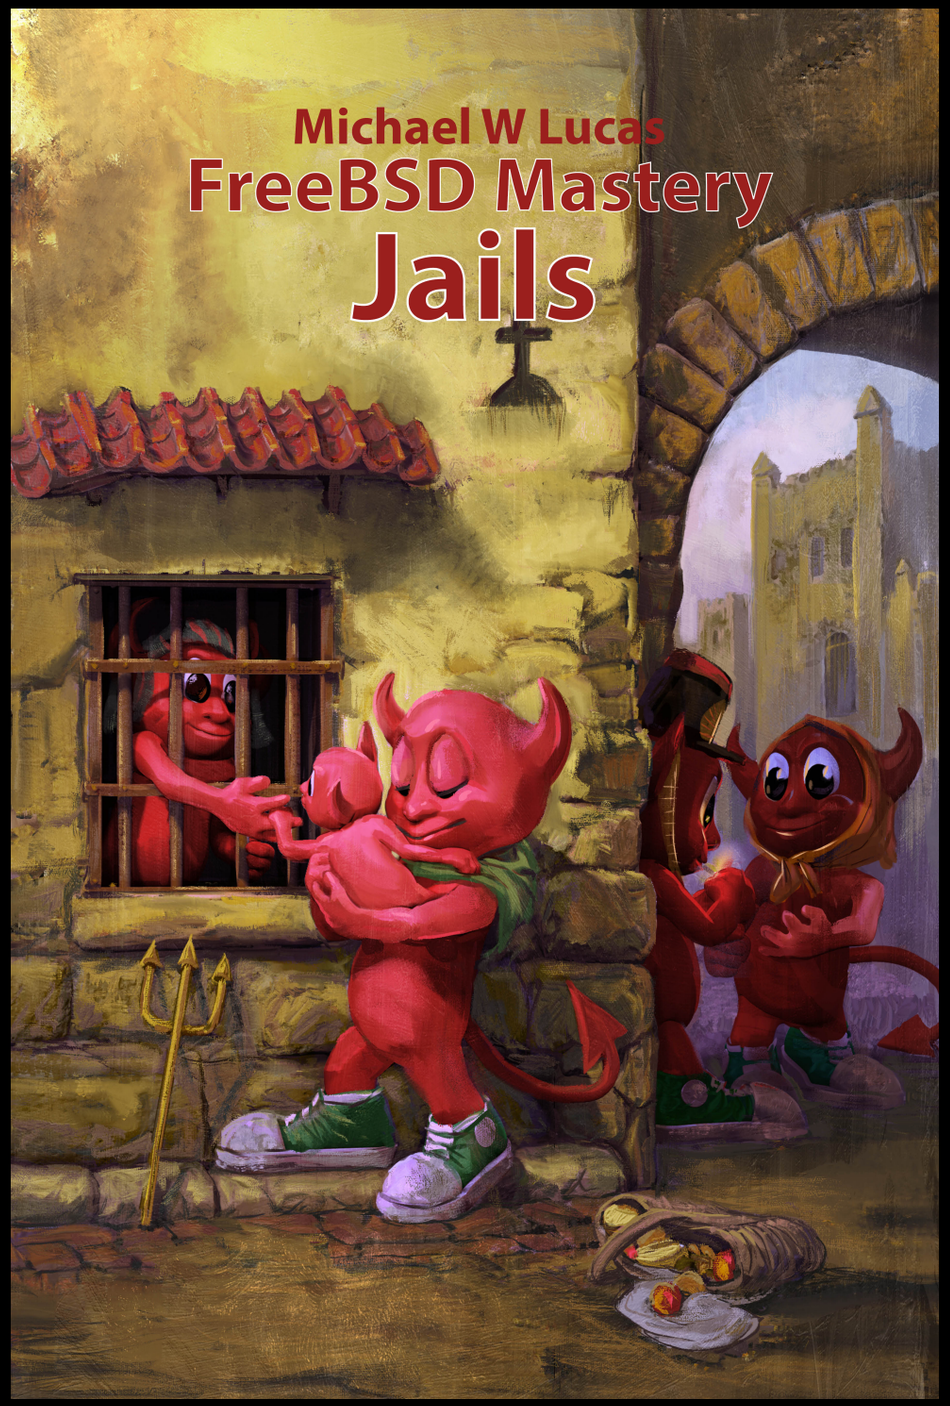

jail all the things!

Dan Langille
EuroBSDCon 2019
Lillehammer

@dlangille
https://dan.langille.org/