

Poudrière

Efficient package building

Baptiste Daroussin
bapt@FreeBSD.org



EuroBSDCon 2015
Stockholm
October 4th, 2015

- ▶ Package building system
- ▶ Port tester
- ▶ Quality insurance on packages
- ▶ Package repository generator
- ▶ System stress tool

- ▶ 2010-07: Initial work
- ▶ 2011: Start to be known and used in the french community
- ▶ 2012-01-31: 1.0 - enter the ports tree
- ▶ 2012-04-08: 1.2 - limit network on fetch phase
- ▶ 2012-05-15: 1.3 - pbi support, attract interest of bdrewery@
- ▶ 2012-08-28: 2.0 - parallel build, ugly html UI (bapt as a designer)
- ▶ 2012-10-15: 2.2 - Removal of pbi support, support for "sets"
- ▶ 2013-05-20: 3.0 - ZFS optional, full tmpfs support, nice and reactive web UI (bdrewery designer)
- ▶ 2013-07: Used in the FreeBSD cluster
- ▶ 2013-09-22: 3.0.3 - support staging, initial qemu support
- ▶ 2014-12-04: 3.1.0 - Yet a better web UI



- ▶ Simple
 - ▶ Easy to setup:
 - ▶ only depend on base (by default)
 - ▶ one simple configuration file
 - ▶ few command to prepare the resources
 - ▶ Easy to use
 - ▶ One single command
 - ▶ Simple subcommands
- ▶ Resource efficient
 - ▶ parallel build: by default 1 core == 1 package building
 - ▶ low overhead (resources should be dedicated to build sources not for poudriere itself)
- ▶ Safe and contained
 - ▶ all builds in clean jail(8)
 - ▶ only access network during fetch phase
 - ▶ build as regular user

Subcommands:

- ▶ bulk: Generate packages for given ports
- ▶ jail: Manage the jails used by poudriere
- ▶ ports: Create, update or delete the portstrees

- ▶ Fetch release/snapshot/old releases sets

- ▶ Fetch release/snapshot/old releases sets
- ▶ Build from sources: git, svn, file, support for branches

- ▶ Fetch release/snapshot/old releases sets
- ▶ Build from sources: git, svn, file, support for branches
- ▶ Full support for src.conf

- ▶ Fetch release/snapshot/old releases sets
- ▶ Build from sources: git, svn, file, support for branches
- ▶ Full support for src.conf
- ▶ Support for multiple arches (via qemu user emulation)

- ▶ Fetch release/snapshot/old releases sets
- ▶ Build from sources: git, svn, file, support for branches
- ▶ Full support for src.conf
- ▶ Support for multiple arches (via qemu user emulation)
- ▶ Can have kernel

- ▶ Fetch release/snapshot/old releases sets
- ▶ Build from sources: git, svn, file, support for branches
- ▶ Full support for src.conf
- ▶ Support for multiple arches (via qemu user emulation)
- ▶ Can have kernel
- ▶ Updatable (via sources or freebsd-update)

- ▶ Fetch release/snapshot/old releases sets
- ▶ Build from sources: git, svn, file, support for branches
- ▶ Full support for src.conf
- ▶ Support for multiple arches (via qemu user emulation)
- ▶ Can have kernel
- ▶ Updatable (via sources or freebsd-update)

Creating a jail

```
| poudriere jail -c -j 102 -v 10.2-RELEASE
```

- ▶ Fetch release/snapshot/old releases sets
- ▶ Build from sources: git, svn, file, support for branches
- ▶ Full support for src.conf
- ▶ Support for multiple arches (via qemu user emulation)
- ▶ Can have kernel
- ▶ Updatable (via sources or freebsd-update)

Creating a jail

```
| poudriere jail -c -j 102 -v 10.2-RELEASE
```

Updating a jail

```
| poudriere jail -u -j 102
```


- ▶ Fetch from portsnap, git, svn

- ▶ Fetch from portsnap, git, svn
- ▶ Notion of "default" ports tree

- ▶ Fetch from portsnap, git, svn
- ▶ Notion of "default" ports tree

Creating a ports tree

```
| poudriere ports -c -p portstree
```

- ▶ Fetch from portsnap, git, svn
- ▶ Notion of "default" ports tree

Creating a ports tree

```
| poudriere ports -c -p portstree
```

Updating a ports tree

```
| poudriere ports -u -p portstree
```


- ▶ Associate a ports tree, a jail and a list of packages to build

- ▶ Associate a ports tree, a jail and a list of packages to build
- ▶ Massively parallelized (1 port per core, fine grain tuning possible)

- ▶ Associate a ports tree, a jail and a list of packages to build
- ▶ Massively parallelized (1 port per core, fine grain tuning possible)
- ▶ Support ccache

- ▶ Associate a ports tree, a jail and a list of packages to build
- ▶ Massively parallelized (1 port per core, fine grain tuning possible)
- ▶ Support ccache
- ▶ Tuneable via: make.conf (fine grained
[<jailname>-[<setname>-[<portstree>-]]]make.conf)

- ▶ Associate a ports tree, a jail and a list of packages to build
- ▶ Massively parallelized (1 port per core, fine grain tuning possible)
- ▶ Support ccache
- ▶ Tuneable via: make.conf (fine grained
[<jailname>-[<setname>-[<portstree>-]]]make.conf)
- ▶ Nice WebUI (static files made dynamic via js)

- ▶ Associate a ports tree, a jail and a list of packages to build
- ▶ Massively parallelized (1 port per core, fine grain tuning possible)
- ▶ Support ccache
- ▶ Tuneable via: make.conf (fine grained
[<jailname>-[<setname>-[<portstree>-]]]make.conf)
- ▶ Nice WebUI (static files made dynamic via js)
- ▶ Nice cli (with colors and SIGINFO support)

- ▶ Associate a ports tree, a jail and a list of packages to build
- ▶ Massively parallelized (1 port per core, fine grain tuning possible)
- ▶ Support ccache
- ▶ Tuneable via: make.conf (fine grained
[<jailname>-[<setname>-[<portstree>-]]]make.conf)
- ▶ Nice WebUI (static files made dynamic via js)
- ▶ Nice cli (with colors and SIGINFO support)
- ▶ Hooks support

- ▶ Associate a ports tree, a jail and a list of packages to build
- ▶ Massively parallelized (1 port per core, fine grain tuning possible)
- ▶ Support ccache
- ▶ Tuneable via: make.conf (fine grained
[<jailname>-[<setname>-[<portstree>-]]]make.conf)
- ▶ Nice WebUI (static files made dynamic via js)
- ▶ Nice cli (with colors and SIGINFO support)
- ▶ Hooks support
- ▶ Repository generation support (including signature)

- ▶ Associate a ports tree, a jail and a list of packages to build
- ▶ Massively parallelized (1 port per core, fine grain tuning possible)
- ▶ Support ccache
- ▶ Tuneable via: make.conf (fine grained
[<jailname>-[<setname>-[<portstree>-]]]make.conf)
- ▶ Nice WebUI (static files made dynamic via js)
- ▶ Nice cli (with colors and SIGINFO support)
- ▶ Hooks support
- ▶ Repository generation support (including signature)
- ▶ Default ports tree support

- ▶ Associate a ports tree, a jail and a list of packages to build
- ▶ Massively parallelized (1 port per core, fine grain tuning possible)
- ▶ Support ccache
- ▶ Tuneable via: make.conf (fine grained
[<jailname>-[<setname>-[<portstree>-]]]make.conf)
- ▶ Nice WebUI (static files made dynamic via js)
- ▶ Nice cli (with colors and SIGINFO support)
- ▶ Hooks support
- ▶ Repository generation support (including signature)
- ▶ Default ports tree support
- ▶ Incremental support (aggressive)

- ▶ Associate a ports tree, a jail and a list of packages to build
- ▶ Massively parallelized (1 port per core, fine grain tuning possible)
- ▶ Support ccache
- ▶ Tuneable via: make.conf (fine grained
[<jailname>-[<setname>-[<portstree>-]]]make.conf)
- ▶ Nice WebUI (static files made dynamic via js)
- ▶ Nice cli (with colors and SIGINFO support)
- ▶ Hooks support
- ▶ Repository generation support (including signature)
- ▶ Default ports tree support
- ▶ Incremental support (aggressive)
- ▶ Restricted support

- ▶ Associate a ports tree, a jail and a list of packages to build
- ▶ Massively parallelized (1 port per core, fine grain tuning possible)
- ▶ Support ccache
- ▶ Tuneable via: make.conf (fine grained
[<jailname>-[<setname>-[<portstree>-]]]make.conf)
- ▶ Nice WebUI (static files made dynamic via js)
- ▶ Nice cli (with colors and SIGINFO support)
- ▶ Hooks support
- ▶ Repository generation support (including signature)
- ▶ Default ports tree support
- ▶ Incremental support (aggressive)
- ▶ Restricted support
- ▶ Saving workdir after failure

- ▶ Associate a ports tree, a jail and a list of packages to build
- ▶ Massively parallelized (1 port per core, fine grain tuning possible)
- ▶ Support ccache
- ▶ Tuneable via: make.conf (fine grained
[<jailname>-[<setname>-[<portstree>-]]]make.conf)
- ▶ Nice WebUI (static files made dynamic via js)
- ▶ Nice cli (with colors and SIGINFO support)
- ▶ Hooks support
- ▶ Repository generation support (including signature)
- ▶ Default ports tree support
- ▶ Incremental support (aggressive)
- ▶ Restricted support
- ▶ Saving workdir after failure
- ▶ Autodetection of rebuild

Building packages:

```
| poudriere bulk -j 102 -f listofpackages.txt
```

Building packages with Q/A:

```
| poudriere bulk -j 102 -t -f listofpackages.txt
```

Building all ports

```
| poudriere bulk -j 102 -a
```

Building all ports with a special "set"

```
| poudriere bulk -z test1 -j 102 -a
```

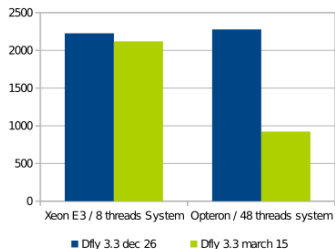
In FreeBSD:

- ▶ ZFS deadlocks
- ▶ tmpfs deadlocks
- ▶ nullfs deadlocks
- ▶ tons of fixes in sh(1) in particular regarding job control
- ▶ highlight contentions

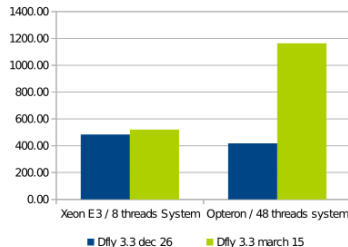
In Dragonfly:

- ▶ Used as a benchmark tool in 2013
- ▶ Lots of performance improvement between December 26, 2012 and March 15, 2013 (released in 3.4)
- ▶ Lots of scalability improvements on large multi-core
- ▶ Lots of panics fixed

Build time in minutes



Number of packages per hours



Poudrière: under the hood

- ▶ Mostly coded in sh(1) (clean and maintainable shell is possible!)

- ▶ Mostly coded in sh(1) (clean and maintainable shell is possible!)
- ▶ Small bits in C

- ▶ Mostly coded in sh(1) (clean and maintainable shell is possible!)
- ▶ Small bits in C
- ▶ Lots of care made on efficiency:
 - ▶ avoid subshells as much as possible
 - ▶ parallelize as many things as possible
 - ▶ reuse resources as much as possible

- ▶ Mostly coded in sh(1) (clean and maintainable shell is possible!)
- ▶ Small bits in C
- ▶ Lots of care made on efficiency:
 - ▶ avoid subshells as much as possible
 - ▶ parallelize as many things as possible
 - ▶ reuse resources as much as possible
- ▶ Use filesystem as a Key/Value DB (on tmpfs for speed)

syscall	seconds	calls	errors
fcntl	0.000012803	1	0
fork	0.000131565	1	0
getegid	0.000011390	1	0
geteuid	0.000023009	2	0
getgid	0.000011620	1	0
getpid	0.000011494	1	0
getppid	0.000011767	1	0
getuid	0.000011717	1	0
[...]			
mmap	0.000370901	22	0
open	0.000182884	6	0
openat	0.000122017	6	0
close	0.000213574	14	0
fstat	0.000233375	11	0
lstat	0.000166041	6	1
write	0.000021875	1	0
access	0.000048576	3	0
sigaction	0.000112313	6	0
sigprocmask	0.000173640	10	0
__getcwd	0.000079150	1	0
pipe	0.000017026	1	0
munmap	0.000115129	8	0
read	0.001368036	12	0
wait4	0.000079314	1	0
sysarch	0.000012172	1	0
	-----	-----	-----
	0.003621732	124	2

```
#!/bin/sh
testfct() {
    echo yes
}

test=$(testfct)
echo $test
```

```
#!/bin/sh
testfct() {
    setvar "$1" "yes"
    # or more posix eval $1="yes"
}

testfct test
echo $test
```

- ▶ abuse xargs!
- ▶ learn awk! stop the "`| grep | sed | grep | cut`" (proper string matching)
- ▶ learn sed! stop the "`| sed | sed | sed`"
- ▶ sh(1) can play with file descriptors (only 10 on POSIX shells)
- ▶ set -e !

- ▶ Associate jails, packages and overlays
- ▶ Able to generate usable images:
 - ▶ Isos: with or without mfsroot
 - ▶ Usb disk: with or without mfsroot
 - ▶ GPT base firmwares (NanoBSD-like)
 - ▶ plain mfsroot
 - ▶ rawdisk (VMs)
- ▶ Reusing code/ideas from NanoBSD/Crochet

- ▶ <https://github.com/freebsd/poudriere/>
- ▶ <https://github.com/freebsd/poudriere/wiki>
- ▶ <https://www.freebsd.org/doc/handbook/ports-poudriere.html>
- ▶ <https://www.freebsd.org/doc/en/books/porters-handbook/testing-poudriere.html>

Questions?



Thanks