

Isolated BSD Build Environments Using Virtual and Hardware Machines

@MichaelDexter

Gainframe CTO
iXsystems Senior Analyst
SNIA.org DPCO Member

Gainframe CTO – Support! Hypervisors!

iXsystems Senior Analyst – OpenZFS! FreeNAS!

Storage Industry Networking Association
Standards! Policy! Awareness!

Hypervisors like bhyve are nice and all but...

Are you adequately safeguarding your data?

Bit rot? Ransomware? eDiscovery?

The most pressing issues in storage
are best solved with...

Hypervisors!

Feature Augmentation

Backup Validation

Regression Hunting

Regression Hunting

A murder mystery

svn commit murder?

THE FACTS

Samba smbld on FreeBSD
performance slaughtered with
1000+ files in a directory

Some time between 9.3 and 10.0

THE VICTIMS

Those with electron microscopes
inspecting CPU wafers...

Those counting lumber
at high speed...

```
svn info ~/9.3/src/
```

```
Last Changed Rev: 260789
```

```
svn info ~/10.0/src/
```

```
Last Changed Rev: 268523
```

Only 7734 places to look!

Whodunnit?

“This could take all afternoon”
– *Actual Quotation*

Any responsible sysadmin
begins with...

Blaming Samba!

+/- 80 individual utilities!
Cryptic de facto standard!
Legendary complexity!

Step One

Spin up 9.3 & 10.0 boxes

Perform ad hoc tests

Develop consistent tests

Tweak variables...

```
# touch /mnt/<1,000,000 files>
# \time -h ls /tmp/smbtest/ > /dev/null
# 10.93s real 10.51s user 0.40s sys
```

```
# mount_smbfs -W WORKGROUP -I 0.0.0.0 \
//tmp/smbtest /mnt
```

```
# \time -h ls /mnt > /dev/null
# 49.21s real 11.01s user 3.65s sys
```

45 ~ 60 minutes over the LAN (!)

Step Two

*Contact Developers &
Support Engineers with
my tests & findings, plus beautiful
graphs from clients*

“I once had that many files & NEVER a problem!”
– Actual Quotation

“You didn’t read the (nonexistent) fine manual!”
– Actual Quotation

“So what’s the fix?”
– Actual Quotation

“You’re holding it wrong!”
– Actual Quotation, close enough context

Step Three

Blame `iconv(1)`!

That UTF-8 GSoC project *must*
have broken it... it crashed `vi(1)`

Step Four

Run `truss(1)` on the `smbd(8)` process while benchmarking...

Convince *one* developer, who
broke out `dtrace(1)`

Thank you Ash!

Nine month process. Really?

Is there a better way?

Shortly thereafter...

“So, OpenJDK isn’t building...”

– Actual Quotation

Poodlerare? Pootlepair? Poudrierr?

Without a doubt...

*Spinning up random systems to build
random versions of software to chase
needles in haystacks is insane*

“It builds on my box Bro!”

– Actual Quotation

Define a Scope

1972 AT&T UNIX in C?

1993 FreeBSD 1.0?

1994 FreeBSD 2.0?

1995 4.4BSD-Lite2?

1998 FreeBSD 3.0?

2000 FreeBSD 4.0?

2003 FreeBSD 5.0?

Define a Scope

*Your regression is somewhere between 1972
and whatever FreshBSD.org says*

*“My Docker Container keeps
restarting again and again”*

– Actual Quotation

“Don’t look back Bro!”

*Works great if you never ship a
single bug or regression*

*The 9.3~10.0 regression occurred
in EOL software*

It's Open Source

The history is available
commit by commit

Let's build some infrastructure...

A Local SVN Mirror

`freebsd.org/pub/FreeBSD/development/subversion/`

`svnmirror-base-r290116.tar.xz` 1761997 KB 10/28/15

A Local SVN Mirror

```
svnsync sync file:///pub/svnmirror/base
```

```
svnlite co file:///pub/svnmirror/base/release@5.0.0 ...
```

Please update these snapshots!

[base] Index of /release × +

https://svnweb.freebsd.org/base/release/

Search



/[base]/release

Index of /release



Files shown: 0

Directory **revision:** [321354](#) (of [322831](#))

Sticky Revision:

Set

File	Rev.	Age	Author	Last log entry
Parent Directory				
10.0.0/	260789	3 years	gjb	Copy releng/10.0 to release/10.0.0 as part of the 10.0-RELEASE cycle. Approved ...
10.1.0/	274417	2 years	gjb	Tag releng/10.1@274401 as release/10.1.0 now that the 10.1-RELEASE builds have c...
10.2.0/	286717	2 years	gjb	Tag releng/10.2@286666 as release/10.2.0 (10.2-RELEASE). Approved by: re (impli...
10.3.0/	297553	16 months	maris	Tag releng/10.3@297264 as release/10.3.0 (10.3-RELEASE).
11.0.0/	306280	10 months	gjb	Tag releng/11.0@ r306211 as release/11.0.0. Approved by: re (implicit) Sponsored...
11.0.1/	306421	10 months	gjb	Tag releng/11.0@ r306420 as release/11.0.1 for 11.0-RELEASE-p1. Approved by: re ...
11.1.0/	321354	4 weeks	gjb	Tag releng/11.1@ r321309 as 11.1-RELEASE. Approved by: re (implicit) Sponsored b...
2.0/	4780	22 years	jkh	Tell stupid users to write the bootcode.
2.0.5/	9201	22 years	cvs2svn	This commit was manufactured by cvs2svn to create tag 'RELENG_2_0_5_RELEASE'.
2.1.0/	13149	21 years	cvs2svn	This commit was manufactured by cvs2svn to create tag 'RELENG_2_1_0_RELEASE'.
2.1.5/	18741	20 years	cvs2svn	This commit was manufactured by cvs2svn to create tag 'RELENG_2_1_5_RELEASE'.

Interrogate Your SVN Mirror

```
# svnlite co file:///pub/svnmirror/base/release/11.1.0 \  
~/11.1/src/usr/src > /dev/null && \  
tar -cf ~/11.1/src.txz -C ~/11.1/src . && \  
svn info ~/11.1/src/usr/src > ~/11.1/svninfo.txt
```

Repeat...

Interrogate Your SVN Mirror

```
# find . -name dialog
# find . -name sshd
# find . -name jail | sort
# find . -name zpool | sort
# find . -name ahci.c | sort
# find . -name xen | sort
# find . -name xz | sort
# find . -name pkg | sort
# find . -name virtio | sort
# find . -name clang | sort
# find . -name bhyve | sort
# find . -name uefi.8 | sort
```

Interrogate Your SVN Mirror

dialog	2.0.5
ssh/sshd	4.0
jail	4.0
AMD64	5.2
jexec	5.4
zpool	7.0
AHCI	8.0
Xen DomU	8.0
xz	8.1
pkg	8.4
VirtIO	8.4
Clang	9.0
bhyve	10.0



Michael Dexter

@michaeldexter

Follow


Sure would be helpful to use #bhyve to run #FreeBSD versions...



49 votes · Final results

1:26 AM - 27 Aug 2017

5 Retweets 6 Likes



5 5 6



Michael Dexter @michaeldexter · Aug 27

Replying to @michaeldexter
@sevanjaniyan Give others a chance to vote!

5 5 6



Michael Dexter @michaeldexter · Aug 28

FreeBSD 5.0 looks like a good place to start!

(Stability Aside)

Jail Management Tools
compat_4~1
/boot/kernel/kernel
AMD64 come 5.2
bhyveload(8) likes it

Remarkably familiar!

We have src.txz...
/boot/kernel/...

Woah there! base.aa base.ab...

Hello ftp-archive.freebsd.org

Fortunately, history is stationary
(when available)

Interrogate ftp-archive

```
# fetch <url>/5.0/5.0-RELEASE-i386-disc1.iso -o \
FreeBSD-5.0-RELEASE-i386-disc1.iso

# mdconfig -af FreeBSD-5.0-RELEASE-i386-disc1.iso
md0

# mount_cd9660 /dev/md0 /mnt

# cp /mnt/* ~/5.0/iso

# cat ~/5.0/iso/base/base.?? > ~/5.0/base.tgz

...
```

*Learn about the history of
FreeBSD Release Engineering*

*Avoid violating the
Prime Directive...*

Repack, rinse, repeat...

(All via a script)

Result: base.txz src.txz doc.txz*

Sound familiar?

ZFS with lz4 eh? libarchive(3) don' care!

Take a breath

Suddenly, you can install
FreeBSD 5.0 with `bsdinstall(8)`
12-CURRENT

Cruft be gone!

POLA: Principle Of Least
Astonishment *in action*

Interesting Facts

FreeBSD has supported
root on NFS from day one

FreeBSD 12-CURRENT
supports root on UFS1

File and Block boot options

Interesting Facts

bhyve NE2000 emulation appears to work

bhyve e1000 emulation works quite well

bhyve ATA emulation is incomplete

bhyve AHCI emulation works quite well

Interesting Facts

While bhyve is awesome...

QEMU is famous for compatibility

Xen is on FreeBSD

VirtualBox is still around

While FreeBSD 5.0 on a zvol sounds great,
an NFS root is better for *science*

Nested VT-x would be great but would involve
an insane amount of work

Where is he heading with this?

bhyveload(8) can load FreeBSD 5.0 onward
bhyveload -e can feed in boot parameters

*Extract base.txz and kernel.txz, create an
fstab, adjust the /etc/ttys if necessary,
share over NFS and FreeBSD 5.0 through
12-CURRENT are at your fingertips*

So about those regressions...

Obviously they are hiding between releases

Obviously building world is “slow”

FreeBSD can be significantly modularized with `src.conf` knobs

“Meta mode” builds and `ccache(1)` can speed incremental builds

My `STUDENT` kernel that boots under `bhyve`
builds in seconds to minutes depending
on your hardware

Minimum `world` builds in seconds to
minutes depending on your hardware

Regression Hunting

Define Success with a Test

Ports are a pain but *-dvd1.iso is your friend

With 7734 places to look, start at 3867 and
work up and down to 1933 and 5800

Regression Hunting

FreeBSD 9.3 will *not* build on 11.1

Isolate your build in a VM

Perform your test in a *second VM*

Regression Hunting

root on NFS is half way to PXE

Perform your tests on *real hardware*
as appropriate using the same infrastructure

So about *that* regression...



⚙ Optimize pathologic case of telldir()

 Closed PublicAuthored by **mav** on Apr 16 2017, 8:16 AM.

Details

Reviewers jhb
 jilles

Commits [rS317064: Optimize pathologic case of telldir\(\) for Samba.](#)

☰ SUMMARY

When application reads large directory, calling `telldir()` for each entry, it creates exponential performance drop as number of entries reach tenths to hundreds of thousands. It is caused by full search through the internal list, that never finds matches in that scenario, but creates $O(n^2)$ delays. My patch optimizes that search, limiting it to entries of the same buffer, turning time from $O(n^2)$ closer to $O(n)$ in the case of linear directory scan.

PR: 218622

Edit Revision

Update Diff

Download Raw Diff

Edit Related Revisions... ▶

Edit Related Objects... ▶

Subscribe

Award Token

Flag For Later

Tags

None

Revision 269204



Jump to revision:



Author:

jhb

Date: Tue Jul 29 00:16:33 2014 UTC (3 years, 1 month ago)

Changed paths: 2

Log Message:

If telldir() is called immediately after a call to seekdir(), POSIX requires the return value of telldir() to equal the value passed to seekdir(). The current seekdir code with SINGLEUSE enabled breaks this case as each call to telldir() allocates a new cookie. Instead, remove the SINGLEUSE code and change telldir() to look for an existing cookie for the directory's current location rather than always creating a new cookie.

CR: <https://phabric.freebsd.org/D490>

PR: 121656

Reviewed by: jilles

MFC after: 1 week

Changed paths

Path

Details

 [head/lib/libc/gen/directory.3](#)

[modified](#) , [text changed](#)

 [head/lib/libc/gen/telldir.c](#)

[modified](#) , [text changed](#)

Lessons Learned

Regressions can live *anywhere* in history

History is far easier to *maintain* than *reconstruct*

Don't assume anything in software

Questions to Answer

What past and future RELEASEs
can each RELEASE build?

What make -jN's will each RELEASE support?

Your question here!

TO DO

Plug `src.conf` leaks. You will thank yourself.

Update the SVN mirror TAR balls,
with each release?

Fill in ftp-archive. Thank you gad@freebsd.org!

Thank You!

Thank you Team Verisign!

@MichaelDexter
dexter@freebsd.org