Reproducible builds everywhere
eg. in Debian and Fedora and everywhere

Bit by bit identical binaries
from a given source

Dennis Gilmore
Holger 'h01ger' Levsen

DevConf.cz in Brno, Czech Republic
2017-01-27

about Dennis

- `28CA D001 51E6 21DA 1F2D C13B 7EE5 B4E3 663C 50D1`
- Fedora user since Fedora Core 1 (2003)
- Fedora contributor since fedora.us
- Plattform lead at Red Hat
- Day job for the last 8 years is Fedora Release Engineering

about h01ger

- `B8BF 5413 7B09 D35C F026 FE9D 091A B856 069A AA1C`
- Debian user since 1995, contributor since 2001, official developer status since 2007
- DebConf organizer, founded the DebConf video team
  - `http://video.debian.net`
- Debian-Edu (Debian for education)
- Debian QA (quality assurance)
  - `https://piuparts.debian.org`
  - `https://jenkins.debian.net` (1200 jobs continously testing Debian)
- Debian Reproducible builds team member
  - since April 2015 funded by the Linux Foundation

about h01ger

- B8BF 5413 7B09 D35C F026 FE9D 091A B856 069A AA1C
- Debian user since 1995, contributor since 2001, official developer status since 2007
- DebConf organizer, founded the DebConf video team
  - http://video.debian.net
- Debian-Edu (Debian for education)
- Debian QA (quality assurance)
  - https://piuparts.debian.org
  - https://jenkins.debian.net (1200 jobs continously testing Debian)
- Debian Reproducible builds team member
  - since April 2015 funded by the Linux Foundation
- the Debian branding on these slides is obviously my fault...

Debian reproducible builds contributors

akira
Alexis Bienvenüe
Andrew Ayer
Asheesh Laroia
Boyuan Yang
Ceridwen
Chris Lamb
Chris West
Christoph Berg
Clint Adams
Dafydd Harries
Daniel Kahn Gillmor
Daniel Shahaf
Daniel Stender
David Suarez
Dhole
Drew Fisher
Emmanuel Bourg

Emanuel Bronshtein
Esa Peuha
Fabian Wolff
Guillem Jover
Hans-Christoph Steiner
Harlan Lieberman-Berg
Helmut Grohne
Holger Levsen
HW42
Intrigeri
Jelmer Vernooij
josch
Juan Picca
Lunar
Maria Glukhova
Mathieu Bridon
Mattia Rizzolo
Nicolas Boulenguez

Niko Tyni
Paul Wise
Peter De Wachter
Philip Rinn
Reiner Herrmann
Robbie Harwood
Santiago Vila
Sascha Steinbiss
Satyam Zode
Scarlett Clark
Stefano Rivera
Stéphane Glondu
Steven Chamberlain
Tom Fitzhenry
Valerie Young
Valentin Lorentz
Wookey
Ximin Luo

Debian reproducible builds contributors

akira
Alexis Bienvenüe
Andrew Ayer
Asheesh Laroia
Boyuan Yang
Ceridwen
Chris Lamb
Chris West
Christoph Berg
Clint Adams
Dafydd Harries
Daniel Kahn Gillmor
Daniel Shahaf
Daniel Stender
David Suarez
Dhole
Drew Fisher
Emmanuel Bourg

Emanuel Bronshtein
Esa Peuha
Fabian Wolff
Guillem Jover
Hans-Christoph Steiner
Harlan Lieberman-Berg
Helmut Grohne
Holger Levsen
HW42
Intrigeri
Jelmer Vernooij
josch
Juan Picca
Lunar
Maria Glukhova
Mathieu Bridon
Mattia Rizzolo
Nicolas Boulenguez

Niko Tyni
Paul Wise
Peter De Wachter
Philip Rinn
Reiner Herrmann
Robbie Harwood
Santiago Vila
Sascha Steinbiss
Satyam Zode
Scarlett Clark
Stefano Rivera
Stéphane Glondu
Steven Chamberlain
Tom Fitzhenry
Valerie Young
Valentin Lorentz
Wookey
Ximin Luo

Who are you?

Who are you?

- Seen a talk about reproducible builds?

Who are you?

- Seen a talk about reproducible builds?
- Contributed to the effort?

Who are you?

- Seen a talk about reproducible builds?
- Contributed to the effort?
- Uses Debian or a Debian based system?

Who are you?

- Seen a talk about reproducible builds?
- Contributed to the effort?
- Uses Debian or a Debian based system?
- Uses Fedora, RHEL, CentOS or a Fedora derivative based system?

1. Motivation
2. Common ressources
3. Status Debian
4. Status Non-Debian World
5. Status RPM world: Fedora and openSUSE
6. Future work
7. Getting involved
8. Questions, comments, ideas?

The problem: we need to believe

- Free Software is great: one can study, modify, share and use it!

The problem: we need to believe

- Free Software is great: one can study, modify, share and use it!
- We study, modify and share source code.
- We use binaries.

The problem: we need to believe

- Free Software is great: one can study, modify, share and use it!
- We study, modify and share source code.
- We use binaries.
- We need to believe our binaries come from the source code they are said to made from.

The problem: we need to believe

- Free Software is great: one can study, modify, share and use it!
- We study, modify and share source code.
- We use binaries.
- We need to believe our binaries come from the source code they are said to made from.
- **I don't want to believe.**

The problem in greater detail

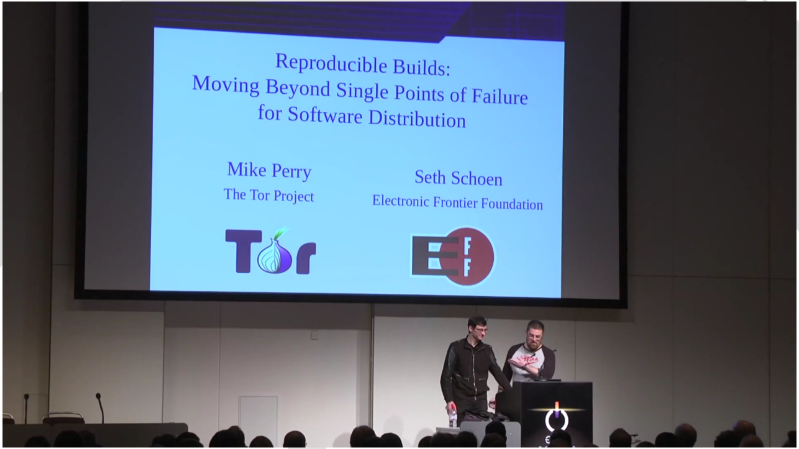

Reproducible Builds:
Moving Beyond Single Points of Failure
for Software Distribution

Mike Perry
The Tor Project

Seth Schoen
Electronic Frontier Foundation

Available on `media.ccc.de`, 31c3

A few examples from that 31c3 talk

- CVE-2002-0083: remote root exploit in `sshd`, a single bit difference in the binary

A few examples from that 31c3 talk

- CVE-2002-0083: remote root exploit in `sshd`, a single bit difference in the binary
- 31c3 talk had a live demo with a kernel module modifying source code in memory only

A few examples from that 31c3 talk

- CVE-2002-0083: remote root exploit in `sshd`, a single bit difference in the binary
- 31c3 talk had a live demo with a kernel module modifying source code in memory only
- How can you be sure what's running on your machine or on a build daemon network connected to the net? Do you ever leave your computers physically alone?

A few examples from that 31c3 talk

- CVE-2002-0083: remote root exploit in `sshd`, a single bit difference in the binary
- 31c3 talk had a live demo with a kernel module modifying source code in memory only
- How can you be sure what's running on your machine or on a build daemon network connected to the net? Do you ever leave your computers physically alone?
- How much do you pay your admins? Enough to withstand a multi million dollar attack?

A few examples from that 31c3 talk

- CVE-2002-0083: remote root exploit in `sshd`, a single bit difference in the binary
- 31c3 talk had a live demo with a kernel module modifying source code in memory only
- How can you be sure what's running on your machine or on a build daemon network connected to the net? Do you ever leave your computers physically alone?
- How much do you pay your admins? Enough to withstand a multi million dollar attack?
- Legal challanges. Could you be forced to backdoor (some of) your software (for some customers)?

Another example from real life

At a CIA conference in 2012:

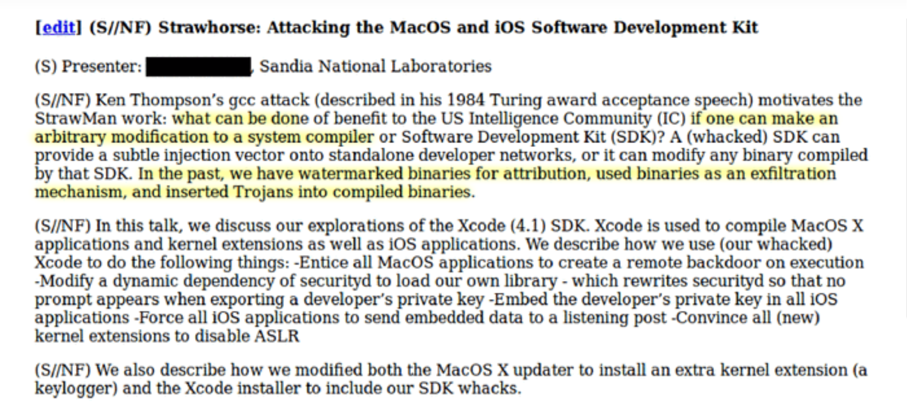

[edit] (S//NF) Strawhorse: Attacking the MacOS and iOS Software Development Kit

(S) Presenter: [redacted], Sandia National Laboratories

(S//NF) Ken Thompson's gcc attack (described in his 1984 Turing award acceptance speech) motivates the StrawMan work: what can be done of benefit to the US Intelligence Community (IC) if one can make an arbitrary modification to a system compiler or Software Development Kit (SDK)? A (whacked) SDK can provide a subtle injection vector onto standalone developer networks, or it can modify any binary compiled by that SDK. In the past, we have watermarked binaries for attribution, used binaries as an exfiltration mechanism, and inserted Trojans into compiled binaries.

(S//NF) In this talk, we discuss our explorations of the Xcode (4.1) SDK. Xcode is used to compile MacOS X applications and kernel extensions as well as iOS applications. We describe how we use (our whacked) Xcode to do the following things: -Entice all MacOS applications to create a remote backdoor on execution -Modify a dynamic dependency of securityd to load our own library - which rewrites securityd so that no prompt appears when exporting a developer's private key -Embed the developer's private key in all iOS applications -Force all iOS applications to send embedded data to a listening post -Convince all (new) kernel extensions to disable ASLR

(S//NF) We also describe how we modified both the MacOS X updater to install an extra kernel extension (a keylogger) and the Xcode installer to include our SDK whacks.

firstlook.org/theintercept/2015/03/10/ispy-cia-campaign-steal-apples-secrets/

The solution

Promise that anyone can always and independently generate identical binary packages from a given source

The solution

We call this:

“Reproducible builds”

Debian demo (skipped)

- Build a package 5 times, get 5 .debs with different checksums
- Build a package 5 times, get 5 .debs with the same checksum

Debian demo (skipped)

- Build a package 5 times, get 5 .debs with different checksums
- Build a package 5 times, get 5 .debs with the same checksum
- Yes, it's really this simple.

Debian demo (skipped)

- Build a package 5 times, get 5 .debs with different checksums
- Build a package 5 times, get 5 .debs with the same checksum
- Yes, it's really this simple.
- And works the same with RPMs.

Debian demo (skipped)

- Build a package 5 times, get 5 .debs with different checksums
- Build a package 5 times, get 5 .debs with the same checksum
- Yes, it's really this simple.
- And works the same with RPMs.
- Signed RPMs are a bit more complicated but the principle stays the same.

This should become the **norm**.

This should become the **norm**.

We want to change the meaning of "free software":

it's only free software if it's reproducible!

More benefits than "just" security...

- Lots and lots of QA benefits - we've found so many subtile bugs.

More benefits than "just" security...

- Lots and lots of QA benefits - we've found so many subtile bugs.
- Google does reproducible builds, to save time and money.

More benefits than "just" security...

- Lots and lots of QA benefits - we've found so many subtile bugs.
- Google does reproducible builds, to save time and money.
- Smaller deltas, thus faster updates possible (for packages and images).

More benefits than "just" security...

- Lots and lots of QA benefits - we've found so many subtile bugs.
- Google does reproducible builds, to save time and money.
- Smaller deltas, thus faster updates possible (for packages and images).
- Side effect: meaningful binary diff between two versions.

More benefits than "just" security...

- Lots and lots of QA benefits - we've found so many subtile bugs.
- Google does reproducible builds, to save time and money.
- Smaller deltas, thus faster updates possible (for packages and images).
- Side effect: meaningful binary diff between two versions.
- ...

1. Motivation
2. Common ressources
3. Status Debian
4. Status Non-Debian World
5. Status RPM world: Fedora and openSUSE
6. Future work
7. Getting involved
8. Questions, comments, ideas?

reproducible-builds.org

- `https://reproducible-builds.org`
- git repositories, IRC channels, mailinglists, webspace

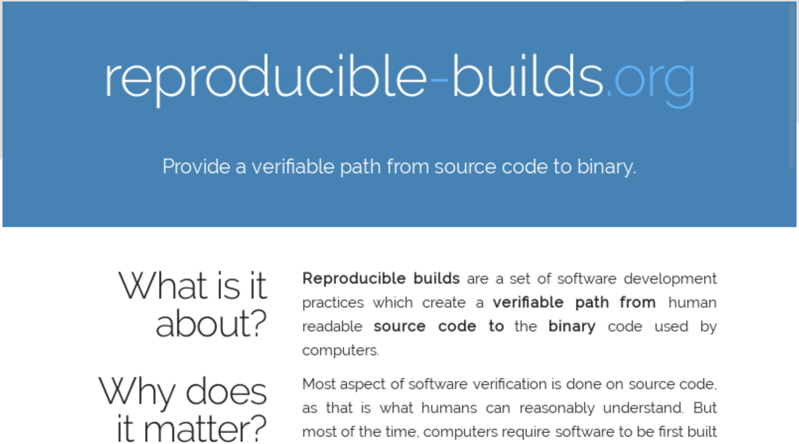

reproducible-builds.org

Provide a verifiable path from source code to binary.

What is it about?

Reproducible builds are a set of software development practices which create a **verifiable path from** human readable **source code to** the **binary** code used by computers.

Why does it matter?

Most aspect of software verification is done on source code, as that is what humans can reasonably understand. But most of the time, computers require software to be first built

Debugging problems:
`https://try.diffoscope.org`

- Examines differences **in depth**.
- Recursively unpacks archives, uncompresses PDFs, disassembles binaries, unpacks Gettext files, ...
- Easy to extend to new file formats.
- Falls back to binary comparison.
- Outputs HTML or plain text with human readable differences.
- Available from `git`, PyPI, Debian,
  Arch Linux, Guix, Homebrew, Fedora. Works on BSD.
- Maintainers in other distros wanted.
- `https://diffoscope.org/`

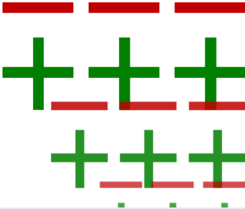

diffoscope example (HTML output)

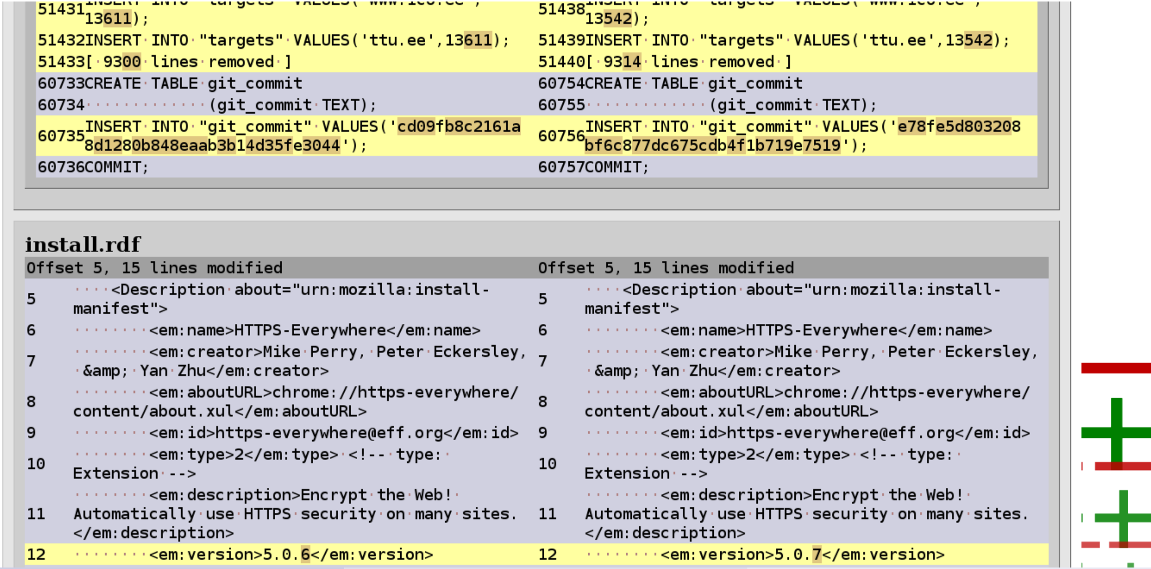

```
51431 13611);                                          51438 13542);
51432INSERT·INTO·"targets"·VALUES('ttu.ee',13611);     51439INSERT·INTO·"targets"·VALUES('ttu.ee',13542);
51433[·9300·lines·removed·]                            51440[·9314·lines·removed·]
60733CREATE·TABLE·git_commit                           60754CREATE·TABLE·git_commit
60734············(git_commit·TEXT);                    60755············(git_commit·TEXT);
60735INSERT·INTO·"git_commit"·VALUES('cd09fb8c2161a     60756INSERT·INTO·"git_commit"·VALUES('e78fe5d803208
     8d1280b848eaab3b14d35fe3044');                         bf6c877dc675cdb4f1b719e7519');
60736COMMIT;                                            60757COMMIT;
```

install.rdf

```
Offset 5, 15 lines modified                            Offset 5, 15 lines modified
5   ····<Description·about="urn:mozilla:install-       5   ····<Description·about="urn:mozilla:install-
    manifest">                                             manifest">
6   ········<em:name>HTTPS-Everywhere</em:name>        6   ········<em:name>HTTPS-Everywhere</em:name>
7   ········<em:creator>Mike·Perry,·Peter·Eckersley,   7   ········<em:creator>Mike·Perry,·Peter·Eckersley,
    ·&amp;·Yan·Zhu</em:creator>                            ·&amp;·Yan·Zhu</em:creator>
8   ········<em:aboutURL>chrome://https-everywhere/    8   ········<em:aboutURL>chrome://https-everywhere/
    content/about.xul</em:aboutURL>                        content/about.xul</em:aboutURL>
9   ········<em:id>https-everywhere@eff.org</em:id>    9   ········<em:id>https-everywhere@eff.org</em:id>
10  ········<em:type>2</em:type>·<!--·type:·           10  ········<em:type>2</em:type>·<!--·type:·
    Extension·-->                                          Extension·-->
    ········<em:description>Encrypt·the·Web!·              ········<em:description>Encrypt·the·Web!·
11  Automatically·use·HTTPS·security·on·many·sites.    11  Automatically·use·HTTPS·security·on·many·sites.
    </em:description>                                      </em:description>
12  ········<em:version>5.0.6</em:version>             12  ········<em:version>5.0.7</em:version>
```

`diffoscope` is "just" for debugging

- Reminder: `diffoscope` is for **debugging**
- "reproducible" according to our definition means: **bit by bit identical**. So the tools for testing whether something is reproducible are either `diff` or `sha256sum`!

diffoscope is "just" for debugging

- Reminder: diffoscope is for **debugging**
- "reproducible" according to our definition means: **bit by bit identical**. So the tools for testing whether something is reproducible are either diff or sha256sum!
- https://try.diffoscope.org

tests.reproducible-builds.org

- Continuously testing Debian `testing`, `unstable` and `experimental`
- Also testing: coreboot, OpenWrt, LEDE, NetBSD, FreeBSD, Arch Linux, Fedora and soon F-Droid too
- 44 nodes (amd64/i386/arm64/armhf), 200 cores and 1 TB RAM
- 486 jenkins jobs running on jenkins.debian.net
- 43 scripts in Python and Bash, 283 lines of code in average
- 37 contributors for `jenkins.debian.net.git`

Variations (when testing Debian)

variation	first build	second build
hostname	jenkins	i-capture-the-hostname
domainname	debian.net	i-capture-the-domainname
env TZ	GMT+12	GMT-14
env LANG	C	fr_CH.UTF-8
env LC_ALL	not set	fr_CH.UTF-8
env USER	pbuilder1	pbuilder2
uid	1111	2222
gid	1111	2222
UTS namespace	shared with the host	*modified using /usr/bin/unshare --uts*
kernel version	Linux 3.16 or 4.X	on amd64 always varied, on armhf sometimes
umask	0022	0002
CPU type	varied on i386	
	on armhf varied a bit, not on amd64	
filesystem	same for both builds on amd64: (tmpfs), on armhf ext3/4	
		(and we have disorderfs*, but the code is disabled)*
year, month, date	on amd64: 398 days variation, on armhf not yet	
hour, minute	hour is usually the same... usually, the minute differs...	
everything else	*is likely the same...*	

Common problems

- time stamps
- timezones
- locales
- build paths
- everything else (seperated into known issues and the blurry rest)

Documentation about common problems

- `https://reproducible-builds.org/docs`
- Lunar's talk from CCCamp 2015 also on `https://media.ccc.de`

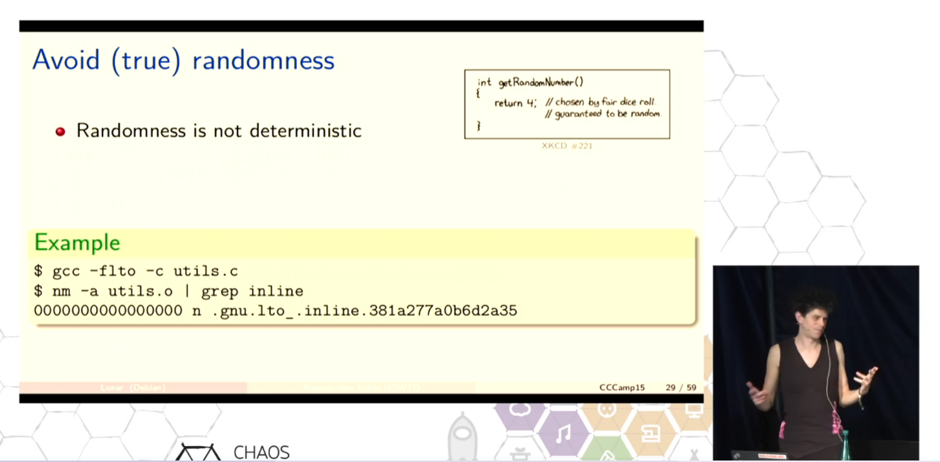

Avoid (true) randomness

- Randomness is not deterministic

```
int getRandomNumber()
{
    return 4;  // chosen by fair dice roll.
               // guaranteed to be random.
}
```

XKCD #221

Example

```
$ gcc -flto -c utils.c
$ nm -a utils.o | grep inline
0000000000000000 n .gnu.lto_.inline.381a277a0b6d2a35
```

Lunar (Debian) Reproducible builds HOWTO CCCamp15 29 / 59

SOURCE_DATE_EPOCH

- Build date (timestamps) usually not useful for the user
- `SOURCE_DATE_EPOCH` is defined as the last modification of the source, since the epoch (1970-01-01)
- can be used instead of current date
- can also be used for random seeds etc.
- in Debian, set from the latest `debian/changelog` entry
- can be set to the latest git commit too or the latest file modification date

SOURCE_DATE_EPOCH

- SOURCE_DATE_EPOCH spec available:
- https://reproducible-builds.org/specs/
- many upstreams support it already
- has been adopted by other distributions (openSUSE, OpenWrt, LEDE, NetBSD, FreeBSD, Arch Linux, coreboot, Guix, ...) and many many upstreams (GCC, dpkg, rpm, mkisofs, ghostscript, libxslt, sphinx, texlive-bin, ...)

two more tools

- `strip-nondeterminism`

two more tools

- `strip-nondeterminism`
- `reprotest`

1. Motivation
2. Common ressources
3. Status Debian
4. Status Non-Debian World
5. Status RPM world: Fedora and openSUSE
6. Future work
7. Getting involved
8. Questions, comments, ideas?

Progress in Debian testing ("stretch")

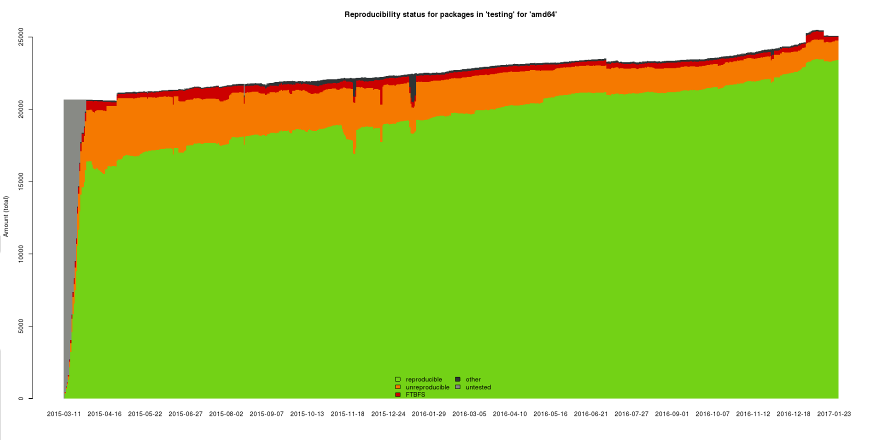

23,405 (93.3%) out of 25,067 source packages are reproducible in our test framework on amd64

Progress in Debian `unstable`

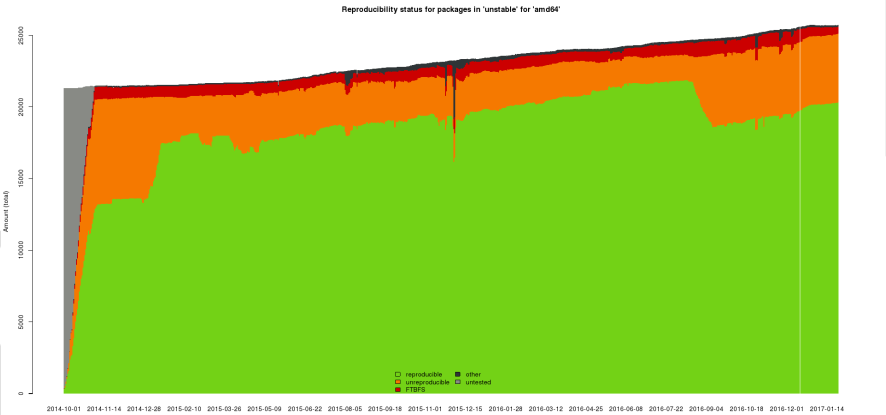

20,309 (78.9%) out of 25,734 source packages are reproducible in our test framework on `amd64` (difference due to build path variations)

Details on tests.reproducible-builds.org

- `https://reproducible.debian.net/$src`
- 48 package sets
- 282 categorised distinct issues
- 7,413 notes
- 1,595 unreproducible packages in `stretch/amd64` (testing), but only 111 without a note (5,288 in `unstable` but also only 154 without a note)
- maintained in `notes.git` by 49 contributors
- currently Debian only, but cross distro notes are planned

Debian `.buildinfo` files

- Aggregates in the same file:
  - Sources (checksums)
  - Generated binaries (checksums)
  - Packages used to build (with specific version, checksums coming soon)
- Can be later used to exactly recreate environment
- For Debian, all versions are available from `snapshot.debian.org`

Progress in the Debian bug tracker

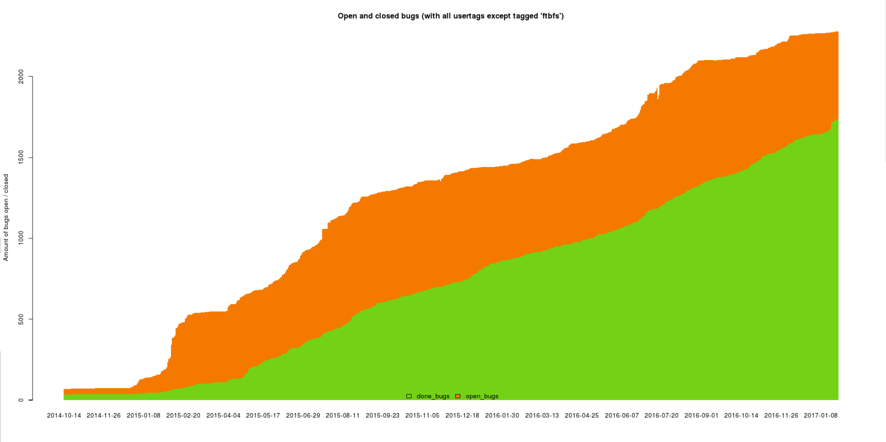

As a rule, we file bugs with patches.
There are very few exceptions.

Sending progress upstream

- So we filed a lot of bugs... with patches...!
- ... but only in Debian and we rely on Debian maintainers sending them upstream.

Sending progress upstream

- So we filed a lot of bugs... with patches...!
- ... but only in Debian and we rely on Debian maintainers sending them upstream.
- Bernard Wiedemann (from openSUSE) thought that wasn't good enough and created `https://github.com/orgs/distropatches`

Debian summary / What's left to do

- This is/was a proof-of-concept, Debian is neither 93.3% reproducible nor 78.9%. (and 10% $>$ 2,500 sources packages!)

Debian summary / What's left to do

- This is/was a proof-of-concept, Debian is neither 93.3% reproducible nor 78.9%. (and 10% > 2,500 sources packages!)
- All our required changes are finally in Debian now!
- Debian 9, "stretch", will only be partially reproducible.
- Because, Debian does not (yet?) do full rebuilds before releasing... so stuff is in the archive which is not reproducible unless it's rebuild.

Debian summary / What's left to do

- This is/was a proof-of-concept, Debian is neither 93.3% reproducible nor 78.9%. (and 10% > 2,500 sources packages!)
- All our required changes are finally in Debian now!
- Debian 9, "stretch", will only be partially reproducible.
- Because, Debian does not (yet?) do full rebuilds before releasing... so stuff is in the archive which is not reproducible unless it's rebuild.
- And then we don't distribute `.buildinfo` files yet. That (and user tools) still needs more *design and code.*

Debian summary continued

- Debian 9, "stretch", will only be partially reproducible.
- Canonical can take our work now and make Ubuntu 17.04 (partyl) reproducible...

Debian summary continued

- Debian 9, "stretch", will only be partially reproducible.
- Canonical can take our work now and make Ubuntu 17.04 (partyl) reproducible...
- Debian 10, "buster", will be partly reproducible in 2019.

Debian summary continued

- Debian 9, "stretch", will only be partially reproducible.
- Canonical can take our work now and make Ubuntu 17.04 (partyl) reproducible...
- Debian 10, "buster", will be partly reproducible in 2019.
- We hope will have `debian-policy` will mandate 100% reproducible builds for Debian 11, "bullseye", in 2021.

Tell the world & collaborate

- "We don't care about Debian (only), we care about free and open source software."

Tell the world & collaborate

- "We don't care about Debian (only), we care about free and open source software."
- 90 Weekly reports since May 2015

Tell the world & collaborate

- "We don't care about Debian (only), we care about free and open source software."
- 90 Weekly reports since May 2015
- First Reproducible World Summit in December 2015 (Athens, Greece)
  - `reproducible.debian.net` has become `tests.reproducible-builds.org`
- Second Reproducible World Summit in December 2016 in Berlin
- Third summit planned for 2017, probably a hackathon in spring 2017 too

Tell the world & collaborate

- "We don't care about Debian (only), we care about free and open source software."
- 90 Weekly reports since May 2015
- First Reproducible World Summit in December 2015 (Athens, Greece)
  - `reproducible.debian.net` has become `tests.reproducible-builds.org`
- Second Reproducible World Summit in December 2016 in Berlin
- Third summit planned for 2017, probably a hackathon in spring 2017 too
- GSoC and Outreachy

1. Motivation
2. Common ressources
3. Status Debian
4. Status Non-Debian World
5. Status RPM world: Fedora and openSUSE
6. Future work
7. Getting involved
8. Questions, comments, ideas?

Skipping some...

- https://tests.r-b.org/coreboot
- https://tests.r-b.org/netbsd
- https://tests.r-b.org/freebsd
- paused: https://tests.r-b.org/archlinux
- not yet: https://tests.r-b.org/f-droid
- https://tests.r-b.org/openwrt
- https://tests.r-b.org/lede

Skipping some more...

- Cygnus.com (1992)
- Bitcoin (2011)
- Tor (2013)
- NixOS, GNU Guix, ElectroBSD
- openSUSE
- Qubes, Tails, webconverger
- Google Bazil
- ducible (build tool for Windows)
- very few commercial, propietary software

Detour: what, reproducible commercial Software???

- Guess which

Detour: what, reproducible commercial Software???

- Guess which
- windows? (the source is available)
- medical devices in your body?
- arms?
- critical infrastructure like in nuclear powerplants?
- cars?

Detour: what, reproducible commercial Software???

- Guess which
- windows? (the source is available)
- medical devices in your body?
- arms?
- critical infrastructure like in nuclear powerplants?
- cars?
- Gambling machines!

1. Motivation
2. Common ressources
3. Status Debian
4. Status Non-Debian World
5. **Status RPM world: Fedora and openSUSE**
6. Future work
7. Getting involved
8. Questions, comments, ideas?

reproducible openSUSE

- `https://build.opensuse.org/package/show/home:bmwiedemann:reproducible/rpm?expand=0`
- Bernhard Wiedemann has built openSUSE twice (with some variations):
  - build-succeeded: 3172
  - bit-by-bit-identical: 2117
  - not-bit-by-bit-identical: 1055

tests.r-b.org/fedora

- used to test Fedora 23, could be made working again
- or build elsewhere and machine readable exported

Fedora basics

- `diffoscope` is available in Fedora
- `yum` and `dnf` might create non-identical environments
- `rpm`-4.13 has an option to override hostname via rpmmacros
- signed RPMs -> re-apply signature, will match for identical builds

TODO: design `.buildinfo` files from koji/mock/zypper

- rfc822 format?
- needs to define the environment
- needs to define the sources (input)
- needs to define the binaries (output)

1. Motivation
2. Common ressources
3. Status Debian
4. Status Non-Debian World
5. Status RPM world: Fedora and openSUSE
6. Future work
7. Getting involved
8. Questions, comments, ideas?

Future work

- So far we mostly worked on making reproducible builds possible...

Future work

- So far we mostly worked on making reproducible builds possible...
- We'll need constant tests for future code.

Future work

- So far we mostly worked on making reproducible builds possible...
- We'll need constant tests for future code.
- And then, this still needs tools, infrastructure and policies to become meaningful and to be used in practice.

Rebuilds and sharing signed checksums

- Almost no work has been done here yet. We are just at the first step: being able to rebuild reproducibly...
- Different projects, different solutions?

Rebuilds and sharing signed checksums

- Almost no work has been done here yet. We are just at the first step: being able to rebuild reproducibly...
- Different projects, different solutions?
  - something like `.buildinfo` files (defining the environment, the input and the output(s)) will be needed everywhere:
  - implemented for Debian (both in sbuild and well as buildinfo.debian.net)
  - work has begun for coreboot, LEDE/OpenWrt and Fedora (mock/koji) and maybe openSUSE (OpenBuildService)

Rebuilders and sharing signed checksums, cont.

- Individuelly signed checksums (think web of trust) could work in the Debian case (we have a gpg web of trust), but IMO won't scale.
- Another idea: rebuilders, run by large organisations (ACLU, CCC, Deutsche Bank, Greenpeace, NASA, NSA, US-Army).
- Fedora rebuilds Debian, Debian rebuilds openSUSE, openSUSE rebuilds NetBSD, etc...
- Big customers could just rebuild everything themselves.

Integration in user tools

- "Do you really want to install this unreproducible software (y/N)"

Integration in user tools

- "Do you really want to install this unreproducible software (y/N)"
- "Do you want to build those packages which have unconfirmed checksums, before installing? (Y/n)"

Integration in user tools

- "Do you really want to install this unreproducible software (y/N)"
- "Do you want to build those packages which have unconfirmed checksums, before installing? (Y/n)"
- "How many signed checksums do you require to call a package 'reproducible'?" - and whom do you trust?

1. Motivation
2. Common ressources
3. Status Debian
4. Status Non-Debian World
5. Status RPM world: Fedora and openSUSE
6. Future work
7. Getting involved
8. Questions, comments, ideas?

As a software developer

- Stop using build dates
- Use `SOURCE_DATE_EPOCH` instead
- See `https://reproducible-builds.org/specs/`

Form your reproducible builds team!

- Why?
  - Every distribution should be reproducible!
  - Learn something new everyday
  - Change the (software) world!
  - `https://tests.reproducible-builds.org/fedora` needs **your** help
- How to get started?
  - Build something twice, run diffoscope on the results.
  - Experiment - learning by doing
  - RTFM, there is lots of documentation
  - Talk to Dennis or h01ger here or talk to us on IRC or via mail.

1. Motivation
2. Common ressources
3. Status Debian
4. Status Non-Debian World
5. Status RPM world: Fedora and openSUSE
6. Future work
7. Getting involved
8. Questions, comments, ideas?

Thanks to...! ...and thank **you**, too!

- All "Reproducible Builds" contributors
  (you are just **so** awesome!)
- DevConf.cz

LINUX FOUNDATION

CORE INFRASTRUCTURE INITIATIVE

PROFITBRICKS The IaaS-Company.

dennis@ausil.us	28CA D001 51E6 21DA 1F2D C13B 7EE5 B4E3 663C 50D1
holger@debian.org	B8BF 5413 7B09 D35C F026 FE9D 091A B856 069A AA1C

Questions, comments, ideas?

- https://reproducible-builds.org/
- #reproducible-builds on irc.OFTC.net
- https://lists.reproducible-builds.org
- twitter: @ReproBuild

Questions, comments, ideas?

- `https://reproducible-builds.org/`
- `#reproducible-builds` on `irc.OFTC.net`
- `https://lists.reproducible-builds.org`
- twitter: @ReproBuild
- Mike and Seth's talk from 31c3 about motivations
- Lunar's talk about fixing reproducible issues from CCCamp 15

The source of this document is available from https://anonscm.debian.org/git/reproducible/presentations.git.