

# Breaking the Chains of Trusting Trust

Vagrant Cascadian <[vagrant@reproducible-builds.org](mailto:vagrant@reproducible-builds.org)>

FOSSY 2023-07-16

# Who am I



	Vagrant
debian user	2001
debian developer	2010
reproducible builds	2015

Different levels of trust:

- `curl http://example.net/hackme | sudo sh`

Different levels of trust:

- `curl http://example.net/hackme | sudo sh`
- `curl -proto '=https' -tlsv1.2 -sSf https://sh.rustup.rs | sh`

Different levels of trust:

- `curl http://example.net/hackme | sudo sh`
- `curl -proto '=https' -tlsv1.2 -sSf https://sh.rustup.rs | sh`
- download file, verify signature ... run code

Different levels of trust:

- `curl http://example.net/hackme | sudo sh`
- `curl -proto '=https' -tlsv1.2 -sSf https://sh.rustup.rs | sh`
- download file, verify signature ... run code
- download source, verify signature, compile from source

Different levels of trust:

- `curl http://example.net/hackme | sudo sh`
- `curl -proto '=https' -tlsv1.2 -sSf https://sh.rustup.rs | sh`
- download file, verify signature ... run code
- download source, verify signature, compile from source
- `emerge -emptytree @world`

Different levels of trust:

- `curl http://example.net/hackme | sudo sh`
- `curl -proto '=https' -tlsv1.2 -sSf https://sh.rustup.rs | sh`
- download file, verify signature ... run code
- download source, verify signature, compile from source
- `emerge --emptytree @world`
- rewrite everything in assembly



Different levels of trust:

- `curl http://example.net/hackme | sudo sh`
- `curl -proto '=https' -tlsv1.2 -sSf https://sh.rustup.rs | sh`
- download file, verify signature ... run code
- download source, verify signature, compile from source
- `emerge --emptytree @world`
- rewrite everything in assembly
- build it up from transitors

Different levels of trust:

- `curl http://example.net/hackme | sudo sh`
- `curl -proto '=https' -tlsv1.2 -sSf https://sh.rustup.rs | sh`
- download file, verify signature ... run code
- download source, verify signature, compile from source
- `emerge --emptytree @world`
- rewrite everything in assembly
- build it up from transitors
- I have a beach, some wood, abundant sunshine, and a lot of time

Ken Thompson

Reflections on trusting trust, 1984

<https://archive.org/details/reflections-on-trusting-trust>

# The Moral of Trusting Trust

"You can't trust code that you did not totally create yourself.  
(Especially code from companies that employ people like me.)  
No amount of source-level verification or scrutiny will protect you  
from using untrusted code." - Ken Thompson

# Did I say 1984, I meant 1974

Karger, 1974

"... insert a trap door into the... compiler...  
the trap door can maintain itself,  
even when the compiler is recompiled"

Since 1974

- 1984: Reflections on trusting trust

Since 1974

- 1984: Reflections on trusting trust
- 1980s: some papers about compiling multiple times

Since 1974

- 1984: Reflections on trusting trust
- 1980s: some papers about compiling multiple times
- 1990s . . . usenet post mumbling about multiple compilers



Since 1974

- 1984: Reflections on trusting trust
- 1980s: some papers about compiling multiple times
- 1990s . . . usenet post mumbling about multiple compilers
- 2000s: some more papers about compiling multiple times

Since 1974

- 1984: Reflections on trusting trust
- 1980s: some papers about compiling multiple times
- 1990s . . . usenet post mumbling about multiple compilers
- 2000s: some more papers about compiling multiple times
- 2005: Countering Trusting Trust through Diverse Double-Compiling

Since 1974

- 1984: Reflections on trusting trust
- 1980s: some papers about compiling multiple times
- 1990s . . . usenet post mumbling about multiple compilers
- 2000s: some more papers about compiling multiple times
- 2005: Countering Trusting Trust through Diverse Double-Compiling
- 2009: Fully Countering Trusting Trust through Diverse Double-Compiling

Since 1974

- 1984: Reflections on trusting trust
- 1980s: some papers about compiling multiple times
- 1990s ... usenet post mumbling about multiple compilers
- 2000s: some more papers about compiling multiple times
- 2005: Countering Trusting Trust through Diverse Double-Compiling
- 2009: Fully Countering Trusting Trust through Diverse Double-Compiling
- ... and some high profile compromises!

# XcodeGhost or should we say Strawhorse?

XcodeGhost, 2015

- Modified version of Apple's Xcode

# XcodeGhost or should we say Strawhorse?

## XcodeGhost, 2015

- Modified version of Apple's Xcode
- Over 4000 compromised apps

# SolarWhat?

SolarWinds, 2020

- Compromised build server...

## SolarWinds, 2020

- Compromised build server...
- ...via weak and/or leaked passphrases



## SolarWinds, 2020

- Compromised build server...
- ...via weak and/or leaked passphrases
- signing certificates compromised

## SolarWinds, 2020

- Compromised build server...
- ...via weak and/or leaked passphrases
- signing certificates compromised
- possibly 18000 affected installations

What is the Price. . .  
of Trusting Trust?

## Free and Open Source Software

- Use

## Free and Open Source Software

- Use
- Study

## Free and Open Source Software

- Use
- Study
- Change

## Free and Open Source Software

- Use
- Study
- Change
- Share

## Free and Open Source Software

- Use
- Study
- Change
- Share
- Community



# Share what exactly

## Sharing FOSS...

- source

# Share what exactly

## Sharing FOSS...

- source
- binaries

# Share what exactly

## Sharing FOSS...

- source
- binaries
- files packaged for distribution

# Where do binaries come from

- Source code...

# Where do binaries come from

- Source code...
- Transformed by a toolchain...

# Where do binaries come from

- Source code...
- Transformed by a toolchain...
- Into machine code

## A taste of source

from bash 5.0 assoc.c:

```
assoc_insert (hash, key, value)
```

```
    HASH_TABLE *hash;
```

```
    char *key;
```

```
    char *value;
```

```
{
```

```
    BUCKET_CONTENTS *b;
```

```
    b = hash_search (key, hash, HASH_CREATE);
```

```
    if (b == 0)
```

```
        return -1;
```

```
    /* If we are overwriting an existing element's value, we're not going to  
     use the key.  Nothing in the array assignment code path frees the key  
     string, so we can free it here to avoid a memory leak. */
```

```
    if (b->key != key)
```

```
./configure  
make  
make install
```



# A resulting binary might look like

```
$ head /bin/bash
ELF&@@8 @@8TTTDDPtdDDQtdRtd0<0</lib/ld-linux-aarch64.so.1GNUy;0gUQGNU 04
  #!JzdAPDDB D  @AJ!Ih@i"r
NL@@@AB
OIq(h  @(
  H &RD!D
    $DP‘
  @A4@ABf LO dPCDDBE % 32BX@TD$
  @A%

!O‘O@@bBh
  HBH
Xq@ Y      ‘1B
BdH(O"BB1@
```

<https://reproducible-builds.org/docs/definition/>

A build is reproducible if given the same source code, build environment and build instructions, any party can recreate bit-by-bit identical copies of all specified artifacts.



Reproducible Builds provides...

- strong confidence...



Reproducible Builds provides...

- strong confidence...
- that a binary was produced from a given source...



Reproducible Builds provides...

- strong confidence...
- that a binary was produced from a given source...
- ...probably!



- Historically software was reproducible! Every bit counted.

# Once upon a time

- Historically software was reproducible! Every bit counted.
- Things eventually got more complicated. . .

# Once upon a time

- Historically software was reproducible! Every bit counted.
- Things eventually got more complicated. . .
- Bit for bit reproducible GNU toolchain in the early 90s on 10(?) architectures.



# Once upon a time

- Historically software was reproducible! Every bit counted.
- Things eventually got more complicated. . .
- Bit for bit reproducible GNU toolchain in the early 90s on 10(?) architectures.
- **And we all forgot.**

# Once upon a time

- Historically software was reproducible! Every bit counted.
- Things eventually got more complicated. . .
- Bit for bit reproducible GNU toolchain in the early 90s on 10(?) architectures.
- **And we all forgot.**
- In 2011 and 2012, Bitcoin and Torbrowser were made reproducible. . .

In 2013 folks explore reproducibility for all of Debian

## Debian

- ~34000 source packages

## Debian

- ~34000 source packages
- ~95% reproducible

## Debian

- ~34000 source packages
- ~95% reproducible
- in theory...

## Debian

- ~34000 source packages
- ~95% reproducible
- in theory...
- many submitted patches

gcc and binutils

- test suite logs



## gcc and binutils

- test suite logs
- PGO (Profile Guided Optimization)

## gcc and binutils

- test suite logs
- PGO (Profile Guided Optimization)
- LTO (Link Time Optimization)

linux

- documentation randomness

## linux

- documentation randomness
- other unidentified issues

## linux

- documentation randomness
- other unidentified issues
- fixes available <https://bugs.debian.org/1033663>  
[https://salsa.debian.org/kernel-team/linux/-/merge\\_requests/741](https://salsa.debian.org/kernel-team/linux/-/merge_requests/741)

## linux

- documentation randomness
- other unidentified issues
- fixes available <https://bugs.debian.org/1033663>  
[https://salsa.debian.org/kernel-team/linux/-/merge\\_requests/741](https://salsa.debian.org/kernel-team/linux/-/merge_requests/741)
- well, partial fixes, anyways...

## libzstd

- recent regression

## GNU Guix

- ~87% reproducible



## GNU Guix

- ~87% reproducible
- in practice!

## GNU Guix

- ~87% reproducible
- in practice!
- 21594 Reproducible!

## GNU Guix

- ~87% reproducible
- in practice!
- 21594 Reproducible!
- 1559 Unreproducible...

## GNU Guix

- ~87% reproducible
- in practice!
- 21594 Reproducible!
- 1559 Unreproducible...
- 1692 Unknown...

## GNU Guix

- reproducible by design

## GNU Guix

- reproducible by design
- normalized build environment

## GNU Guix

- reproducible by design
- normalized build environment
- guix challenge

## GNU Guix

- reproducible by design
- normalized build environment
- guix challenge
- two build farms to compare against



Arch Linux

<https://reproducible.archlinux.org/>

- ~14000 packages

Arch Linux

<https://reproducible.archlinux.org/>

- ~14000 packages
- ~86% reproducible

Arch Linux

<https://reproducible.archlinux.org/>

- ~14000 packages
- ~86% reproducible
- in practice!

# But wait, there is more!

- NetBSD 84%

# But wait, there is more!

- NetBSD 84%
- OpenWRT 96%-100%

# But wait, there is more!

- NetBSD 84%
- OpenWRT 96%-100%
- Coreboot 100%

# But wait, there is more!

- NetBSD 84%
- OpenWRT 96%-100%
- Coreboot 100%
- NixOS 95%-99.7%

# But wait, there is more!

- NetBSD 84%
- OpenWRT 96%-100%
- Coreboot 100%
- NixOS 95%-99.7%
- Yocto 99.98%



# But wait, there is more!

- NetBSD 84%
- OpenWRT 96%-100%
- Coreboot 100%
- NixOS 95%-99.7%
- Yocto 99.98%
- openEuler 96%

## But wait, there is more!

- NetBSD 84%
- OpenWRT 96%-100%
- Coreboot 100%
- NixOS 95%-99.7%
- Yocto 99.98%
- openEuler 96%
- openSUSE mostly reproducible (caveats apply)

# We miss you!

We once had testing for...

- Alpine
- Fedora

# I am not picky about the color of your hat

Wishlist based on current events...

- AlmaLinux
- Rocky Linux

Reproducible Builds is totally possible. . .  
. . . But it only provides one strong link in a chain

# Building on a solid foundation of turtles

<https://bootstrappable.org>

Compiling your C compiler with a C compiler

And a C compiler to compile the other C compiler

...Ad infinitum

## Java bootstrapping

- openjdk17 needs...

## Java bootstrapping

- openjdk17 needs...
- openjdk16 which needs...



## Java bootstrapping

- openjdk17 needs...
- openjdk16 which needs...
- ...

## Java bootstrapping

- openjdk17 needs...
- openjdk16 which needs...
- ...
- openjdk9 ... etc.

## Rust bootstrapping

- rust 1.64 needs...

## Rust bootstrapping

- rust 1.64 needs...
- rust 1.63 which needs...

## Rust bootstrapping

- rust 1.64 needs...
- rust 1.63 which needs...
- ...

## Rust bootstrapping

- rust 1.64 needs...
- rust 1.63 which needs...
- ...
- rust 1.54 can be built with mrustc

## Rust bootstrapping

- rust 1.64 needs...
- rust 1.63 which needs...
- ...
- rust 1.54 can be built with mrustc
- mrustc is written in C++

David A. Wheeler

Fully Countering Trusting Trust through Diverse Double-Compiling, 2009

<https://dwheeler.com/trusting-trust/dissertation/html/wheeler-trusting-trust-ddc.html>



# A beautiful Mes

GNU Mes is a Scheme interpreter and C compiler for bootstrapping the GNU System.  
<https://www.gnu.org/software/mes/>

# We made the same Mes

Bit-for-bit identical Mes built on three different distributions  
<https://reproducible-builds.org/news/2019/12/21/reproducible-bootstrap-of-mes-c-compiler/>

GNU Guix: The Reduced Binary Seed Bootstrap

https:

[//guix.gnu.org/en/manual/devel/en/guix.html#Reduced-Binary-Seed-Bootstrap](https://guix.gnu.org/en/manual/devel/en/guix.html#Reduced-Binary-Seed-Bootstrap)

- ...

GNU Guix: The Reduced Binary Seed Bootstrap

https:

[//guix.gnu.org/en/manual/devel/en/guix.html#Reduced-Binary-Seed-Bootstrap](https://guix.gnu.org/en/manual/devel/en/guix.html#Reduced-Binary-Seed-Bootstrap)

- ...
- Reduced to 145MB of bootstrap binaries (from 250MB)

## GNU Guix: The Reduced Binary Seed Bootstrap

https:

[//guix.gnu.org/en/manual/devel/en/guix.html#Reduced-Binary-Seed-Bootstrap](https://guix.gnu.org/en/manual/devel/en/guix.html#Reduced-Binary-Seed-Bootstrap)

- ...
- Reduced to 145MB of bootstrap binaries (from 250MB)
- Using Mes and guile...

## GNU Guix: The Reduced Binary Seed Bootstrap

https:

[//guix.gnu.org/en/manual/devel/en/guix.html#Reduced-Binary-Seed-Bootstrap](https://guix.gnu.org/en/manual/devel/en/guix.html#Reduced-Binary-Seed-Bootstrap)

- ...
- Reduced to 145MB of bootstrap binaries (from 250MB)
- Using Mes and guile...
- Builds from source GCC, binutils, glibc, etc.

## GNU Guix: The Reduced Binary Seed Bootstrap

https:

[//guix.gnu.org/en/manual/devel/en/guix.html#Reduced-Binary-Seed-Bootstrap](https://guix.gnu.org/en/manual/devel/en/guix.html#Reduced-Binary-Seed-Bootstrap)

- ...
- Reduced to 145MB of bootstrap binaries (from 250MB)
- Using Mes and guile...
- Builds from source GCC, binutils, glibc, etc.
- 145MB of binaries is still not really auditable...

# Before The Mes and Beyond

GNU Guix: The Full-Source Bootstrap

<https://guix.gnu.org/en/blog/2023/>

[the-full-source-bootstrap-building-from-source-all-the-way-down/](https://guix.gnu.org/en/blog/2023/the-full-source-bootstrap-building-from-source-all-the-way-down/)

Now available via guix pull!

- hex0 (357-byte binary)



# Before The Mes and Beyond

GNU Guix: The Full-Source Bootstrap

<https://guix.gnu.org/en/blog/2023/>

[the-full-source-bootstrap-building-from-source-all-the-way-down/](https://guix.gnu.org/en/blog/2023/the-full-source-bootstrap-building-from-source-all-the-way-down/)

Now available via guix pull!

- hex0 (357-byte binary)
- hex1

# Before The Mes and Beyond

GNU Guix: The Full-Source Bootstrap

<https://guix.gnu.org/en/blog/2023/>

[the-full-source-bootstrap-building-from-source-all-the-way-down/](https://guix.gnu.org/en/blog/2023/the-full-source-bootstrap-building-from-source-all-the-way-down/)

Now available via guix pull!

- hex0 (357-byte binary)
- hex1
- M0

# Before The Mes and Beyond

GNU Guix: The Full-Source Bootstrap

<https://guix.gnu.org/en/blog/2023/>

[the-full-source-bootstrap-building-from-source-all-the-way-down/](https://guix.gnu.org/en/blog/2023/the-full-source-bootstrap-building-from-source-all-the-way-down/)

Now available via guix pull!

- hex0 (357-byte binary)
- hex1
- M0
- hex2

# Before The Mes and Beyond

GNU Guix: The Full-Source Bootstrap

<https://guix.gnu.org/en/blog/2023/>

[the-full-source-bootstrap-building-from-source-all-the-way-down/](https://guix.gnu.org/en/blog/2023/the-full-source-bootstrap-building-from-source-all-the-way-down/)

Now available via guix pull!

- hex0 (357-byte binary)
- hex1
- M0
- hex2
- M1

# Before The Mes and Beyond

GNU Guix: The Full-Source Bootstrap

<https://guix.gnu.org/en/blog/2023/>

[the-full-source-bootstrap-building-from-source-all-the-way-down/](https://guix.gnu.org/en/blog/2023/the-full-source-bootstrap-building-from-source-all-the-way-down/)

Now available via guix pull!

- hex0 (357-byte binary)
- hex1
- M0
- hex2
- M1
- mescc-tools

# Before The Mes and Beyond

GNU Guix: The Full-Source Bootstrap

<https://guix.gnu.org/en/blog/2023/>

[the-full-source-bootstrap-building-from-source-all-the-way-down/](https://guix.gnu.org/en/blog/2023/the-full-source-bootstrap-building-from-source-all-the-way-down/)

Now available via guix pull!

- hex0 (357-byte binary)
- hex1
- M0
- hex2
- M1
- mescc-tools
- M2-Planet

# Before The Mes and Beyond

GNU Guix: The Full-Source Bootstrap

<https://guix.gnu.org/en/blog/2023/>

[the-full-source-bootstrap-building-from-source-all-the-way-down/](https://guix.gnu.org/en/blog/2023/the-full-source-bootstrap-building-from-source-all-the-way-down/)

Now available via guix pull!

- hex0 (357-byte binary)
- hex1
- M0
- hex2
- M1
- mescc-tools
- M2-Planet
- Mes

# Before The Mes and Beyond

GNU Guix: The Full-Source Bootstrap

<https://guix.gnu.org/en/blog/2023/>

[the-full-source-bootstrap-building-from-source-all-the-way-down/](https://guix.gnu.org/en/blog/2023/the-full-source-bootstrap-building-from-source-all-the-way-down/)

Now available via guix pull!

- hex0 (357-byte binary)
- hex1
- M0
- hex2
- M1
- mescc-tools
- M2-Planet
- Mes
- TinyCC (patched)



# Before The Mes and Beyond

GNU Guix: The Full-Source Bootstrap

<https://guix.gnu.org/en/blog/2023/>

[the-full-source-bootstrap-building-from-source-all-the-way-down/](https://guix.gnu.org/en/blog/2023/the-full-source-bootstrap-building-from-source-all-the-way-down/)

Now available via guix pull!

- hex0 (357-byte binary)
- hex1
- M0
- hex2
- M1
- mescc-tools
- M2-Planet
- Mes
- TinyCC (patched)
- old versions of GCC, binutils, glibc, gzip, tar ...

# Before The Mes and Beyond

GNU Guix: The Full-Source Bootstrap

<https://guix.gnu.org/en/blog/2023/>

[the-full-source-bootstrap-building-from-source-all-the-way-down/](https://guix.gnu.org/en/blog/2023/the-full-source-bootstrap-building-from-source-all-the-way-down/)

Now available via guix pull!

- hex0 (357-byte binary)
- hex1
- M0
- hex2
- M1
- mescc-tools
- M2-Planet
- Mes
- TinyCC (patched)
- old versions of GCC, binutils, glibc, gzip, tar ...
- modern GCC and everything

`https://github.com/fossilinux/live-bootstrap`

- A live environment

`https://github.com/fossilinux/live-bootstrap`

- A live environment
- From kernel and a bit of source code

`https://github.com/fossilinux/live-bootstrap`

- A live environment
- From kernel and a bit of source code
- To a reproducibly bootstrapped toolchain

`https://github.com/fossilinux/live-bootstrap`

- A live environment
- From kernel and a bit of source code
- To a reproducibly bootstrapped toolchain
- no pregenerated "source" code shortcuts

# UEFI based bootstrap

Work-in-progress UEFI bootstrap

<https://git.stikonas.eu/andrius/stage0-uefi>

Only stage0...

Stage0 on Bare Metal?

<https://git.savannah.nongnu.org/cgit/stage0.git/tree/>



Full bootstrap only available on x86

...x86 toolchain can then cross-compile to x86\_64

- arm
- riscv64
- powerpc64le or powerpc64el

Free/Libre and Open Source Software  
Allows arbitrary third-party verification

No need to Trust, All you need is:

- Free/Libre and Open Source Software

No need to Trust, All you need is:

- Free/Libre and Open Source Software
- Reproducible Builds

No need to Trust, All you need is:

- Free/Libre and Open Source Software
- Reproducible Builds
- Bootstrapping

No need to Trust, All you need is:

- Free/Libre and Open Source Software
- Reproducible Builds
- Bootstrapping
- Diverse compilation

No need to Trust, All you need is:

- Free/Libre and Open Source Software
- Reproducible Builds
- Bootstrapping
- Diverse compilation
- ... and lots of compile cycles



# Copyright and attributions

Copyright 2019-2023 Vagrant Cascadian <[vagrant@reproducible-builds.org](mailto:vagrant@reproducible-builds.org)> Portions by contributors to the reproducible-builds.org website.

Copyright 2019 Holger Levsen <[holger@layer-acht.org](mailto:holger@layer-acht.org)>

This work is licensed under the Creative Commons Attribution-ShareAlike 4.0 International License.

To view a copy of this license, visit <https://creativecommons.org/licenses/by-sa/4.0/>

snippet from bash assoc.c:

Copyright (C) 2008,2009,2011 Free Software Foundation, Inc.

Bash is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version.

<http://www.gnu.org/licenses/>