# Kernel lock-down series

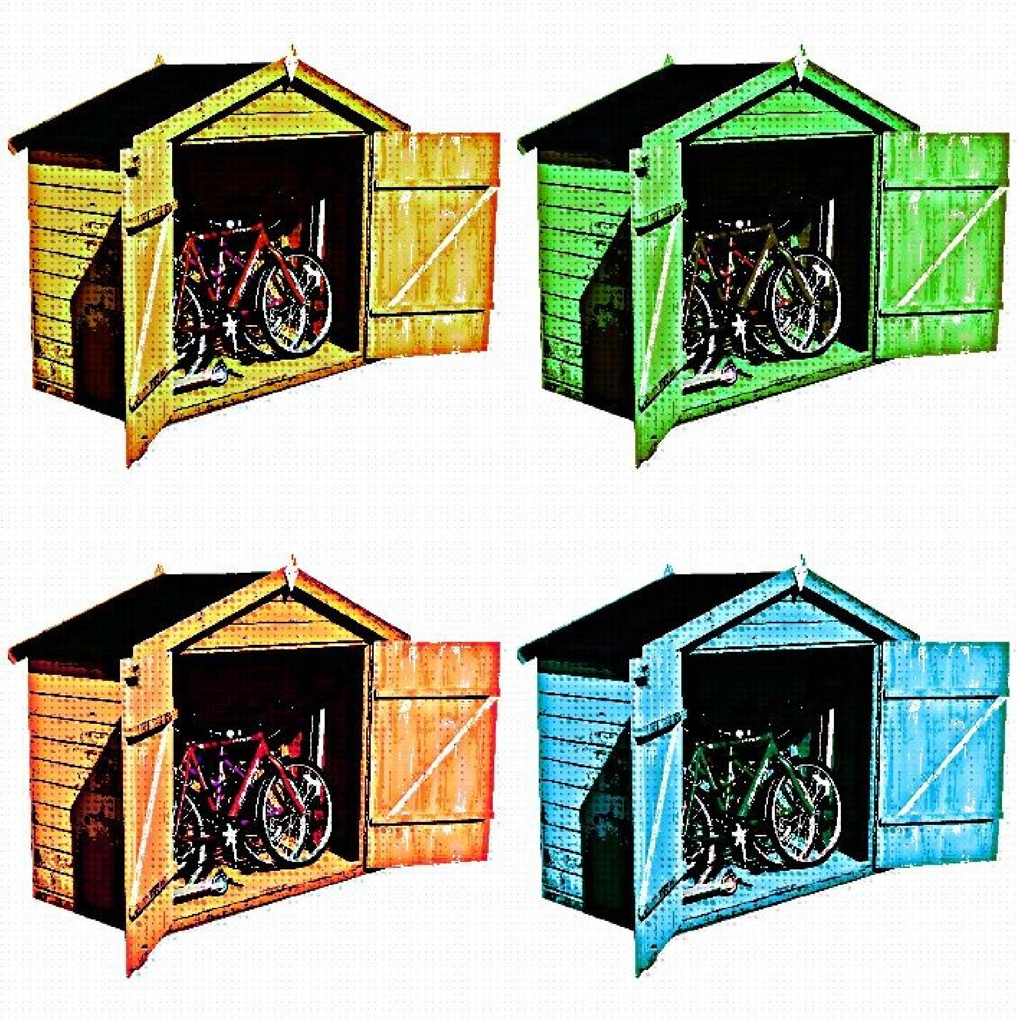

http://outflux.net/slides/2014/lss/lockdown.pdf

Linux Security Summit, Chicago 2014

Kees Cook <keescook@chromium.org>

(pronounced "Case")

chrome

# Overview

- What and why
- Objections/Rebuttals
- Name
- Discuss!

# What, why?

- Verified boot flow wants to keep kernel trusted and userspace untrusted: bright line between kernel memory and userspace memory

lkml thread:
https://lkml.org/lkml/2014/2/26/554

git:
https://git.kernel.org/cgit/linux/kernel/git/kees/linux.git/commit/?h=lockdown

# Objections/Rebuttals

- Should be new capabilities flag
  - Totally orthogonal to capabilites, breaks userspace, not all protections are process-based
- It's not perfect, so it shouldn't happen at all
  - How else can we evolve the protection over time?
- CAP_SYS_RAWIO should be revoked too
  - Needed for things that don't violate ring0/uid0
- Not useful/wouldn't be used
  - Fedora has been carrying it for a while
  - One-off Identical limitations have been added to hibernation and kexec

# Name

- “securelevel”
  - Linus said “No”
- “trusted_kernel”
  - Boot firmware trusts the kernel (via whatever mechanism, including measurement)
- “measured_kernel”
  - Not all cases are measured
- “lockdown_kernel”
  - It's the request being made by whatever wants to enforce the kernel/userspace separation

# Talk amongst yourselves

I'll give you a topic ...

http://outflux.net/slides/2014/lss/firmware.pdf

keescook@chromium.org