

Trust and the Linux kernel

Greg Kroah-Hartman

gregkh@linuxfoundation.org

git.sr.ht/~gregkh/presentation-linux-trust

Disclaimer

Nothing in here reflects the opinion of the Linux Foundation or any other Linux kernel developer. It's all my personal opinion.

Open source software is more trustworthy than closed source software.

Open source software is more trustworthy than closed source software.

Because it can be audited by anyone at anytime.

Open source software is more trustworthy than closed source software.

Because it can be audited by anyone at anytime and fixed by anyone.

University of Minnesota “episode”

or

How to NOT do research on an
open source community

umn.edu timeline

- 2018-2020: Various small "fixes" from umn.edu were submitted, many were accepted.

umn.edu timeline

- 2020 Aug 9..21: “Hypocrite Commits” patches sent from UMN researchers
 - Attempted to introduce vulnerabilities to see if they would be detected
 - Sent to kernel developers from false identities; without consent, notice, or ethics review
 - 4 were submitted, 1 was accepted

umn.edu timeline

- 2020 Nov: Draft “Hypocrite Commits” paper is published
- 2020 Nov 22: Sarah Jamie Lewis calls attention to paper’s questionable ethics
- 2020 Dec 1: Lewis & others send letter to IEEE S&P, questioning ethics
- 2020 Dec ??: UMN IRB appears to give after-the-fact exemption to research on the basis that it believes the research is not human research
- 2020 Dec 15: UMN issues clarification to IEEE

umn.edu timeline

- 2021 Apr 6: Poor quality patches sent by UMN
after ~7 months of silence
 - Raises spectre of continued attacks

umn.edu timeline

- 2021 Apr 20: I ask submitters to stop sending poor quality patches under the guise of “research on maintainers”

umn.edu timeline

- 2021 Apr 20: I ask submitters to stop sending poor quality patches under the guise of “research on maintainers”
 - Researcher claimed new set of patches was not part of previous research

umn.edu timeline

- 2021 Apr 20: I ask submitters to stop sending poor quality patches under the guise of “research on maintainers”
 - Researcher claimed new set of patches was not part of previous research
 - I reply, umn.edu submissions should be rejected until all of this is figured out.

umn.edu timeline

- 2021 Apr 21: start review of all @umn.edu commits

umn.edu timeline

- 2021 Apr 23: Linux Foundation sends letter to UMN requesting:
 - Id all proposals of known-vulnerable code from any U of MN experiment
 - Withdraw, from formal publication, research where subjects didn't give prior consent
 - Ensure all future U of MN experiments on people first have review and approval
 - Ensure all future reviews of proposed experiments on people will normally ensure the consent of those being experimented on

umn.edu timeline

- 2021 Apr 24: UMN publishes “An open letter to the Linux community”
- 2021 Apr 26: UMN researchers retract "Hypocrite Commits" paper from formal publication
 - Hours before IEEE was about to revoke it

umn.edu timeline

- 2021 Apr 27: UMN published details on commits & replies to LF
 - Paper withdrawn. UMN believes it's not “human subjects research”
 - Will do faculty ethics training in 2021-2022, explore added processes, to prevent similar situations

umn.edu timeline

- 2021 May 3: I post a final set of reverts, along with correct fixes

umn.edu timeline

- 2021 May 5: Linux TAB publishes detailed report, with due diligence audit results

TAB Report summary

- 435 UMN commits were re-reviewed, thanks to 95 Linux kernel developers

TAB Report summary

- Confirmed that all intentionally-vulnerable patches with vulnerabilities were rejected

TAB Report summary

- One (“patch 1”) was intended to be vulnerable, but due to lack of understanding by the submitter, it was valid & was accepted
- “patch 1” was later removed because submission was made under a false name

TAB Report summary

- Huge majority of the reviewed commits (349) were found to be correct
- UMN overall patch quality very poor
 - 25 were already fixed by later commits
 - 39 needed new fixes

TAB Report summary

- Huge majority of the reviewed commits (349) were found to be correct
- UMN overall patch quality very poor
 - 25 were already fixed by later commits
 - 39 needed new fixes
- ~20% committed patches were incorrect.

TAB Report summary

- Almost all UMN changes were for obscure drivers in error handling “cleanup paths”

TAB Report summary

- Almost all UMN changes were for obscure drivers in error handling “cleanup paths”
- ~20% of commits to fix a problem, were incorrect.

“Never attribute to malice that which is adequately explained by stupidity.”

– Hanlon’s razor

“Hypocrite Commits”

- UMN allowed researchers to submit using fake identities, while agreeing to the DCO legal document for submission with the fake identity

“Hypocrite Commits” – patch 1

- Was a valid change and was accepted.
- Researchers claimed it was invalid in their paper.
- Later reverted as it came from a fake identity

9fcddaf2e28d crypto: cavium/nitrox - add an error message to explain the failure of pci_request_mem_regions

“Hypocrite Commits” – patch 2

- Duplicate attempt at a “syzbot fix” that spawned my 2019 Kernel Recipes “CVEs suck” presentation.
- Instantly rejected by the maintainer.

“Hypocrite Commits” – patch 3

- Quickly recognized was incorrect
- Maintainer offered possible solutions
- Maintainer was ignored

“Hypocrite Commits” – patch 4

- Maintainer recognized change was incorrect
- Maintainer offered possible solutions
- Developer apologized for the incorrect submission

“Hypocrite Commits” – patch 5

- Bonus patch!
- Was supposed to be a real fix
- Was sent from a machine set up to send “hypocrite changes”
- Invalid author name caused it to be ignored.

“Hypocrite Commits” – patch 5

- Bonus patch!
- Was supposed to be a real fix
- Was sent from a machine set up to send “hypocrite changes”
- Invalid author name caused it to be ignored.

“James Bond <jameslouisebond@gmail.com>”

“Hypocrite Commits”

- All were caught by maintainers
- This fact was ignored in submitted paper
- Our development model works!

“Hypocrite Commits”

- All were caught by maintainers
- This fact was ignored in submitted paper
- Our development model works!
- We got lucky

umn.edu timeline

- 2021 May 6: UMN met with me, Kees and LF to discuss productive ways to move forward and improve.
- 2021 May 6: IEEE publishes statement about how the paper violated ethical guidelines and what would be put into place to prevent it happening again.

umn.edu timeline

- 2021 May 7: UMN responds to TAB report, verifying it is correct
 - Identifies one further set of patches from their team, using a private email address in February 2021. All were rejected by the community as they were invalid changes.
 - Stated that they had only done this for the Linux kernel, not for any other open source project:

Furthermore, we want to state unequivocally that no other Linux components or any other open software systems were affected by the 'hypocrite commits' case study or by any of our other research projects. Our “hypocrite commit” work was limited to the Linux Kernel only and consisted of only the four patches (one is valid) submitted between August 9, 2020 and August 21, 2020

umn.edu timeline

- 2021 May 20: All broken UMN commits are reverted and fixed properly in the main kernel tree (5.13-rc3)
 - 2021 May 26: reverted and fixed in 5.12.y, 5.10.y, and 5.4.y
 - 2021 June 3: reverted and fixed in 4.19.y, 4.14.y, 4.9.y, and 4.4.y

umn.edu timeline

- 2021 Nov: UMN professor asks Kees and I if they can start sending patches.

umn.edu timeline

- 2021 Nov: UMN professor asks Kees and I if they can start sending patches.
- We say no.

umn.edu timeline

- 2021 Dec: Patches start coming from umn.edu

umn.edu timeline

- 2021 Dec: Patches start coming from umn.edu
- 2022 Jan: Developers notice umn.edu patches are incorrect

umn.edu timeline

- 2021 Dec: Patches start coming from umn.edu
- 2022 Jan: Developers notice umn.edu patches are incorrect
- 2022 Jan: UMN is notified that their researchers are not abiding by the rules they agreed with.

umn.edu timeline

- 2022 Feb: University claims ignorance, they never got around to doing the training at all, will really do it this time.

Researcher Guidelines

- [Documentation/process/researcher-guidelines.rst](#)
- We expect contributors are acting in good faith.
- Passive research on public data is allowed.
- Active research on developer behavior must be open only.

Researcher Submissions

Must describe in changelog:

- Specific problem found
- How could be reached on running system
- How it was found
- What version of the kernel it was found
- What was changed to solve this
- Why is this change correct
- How the change was build and run-time tested
- What commit this fixed

umn.edu timeline

- 2022 April: University brings in kernel developer to help fix their program.

umn.edu “episode”

Proof that you can go back in time
and audit code based on new
information.

Trust

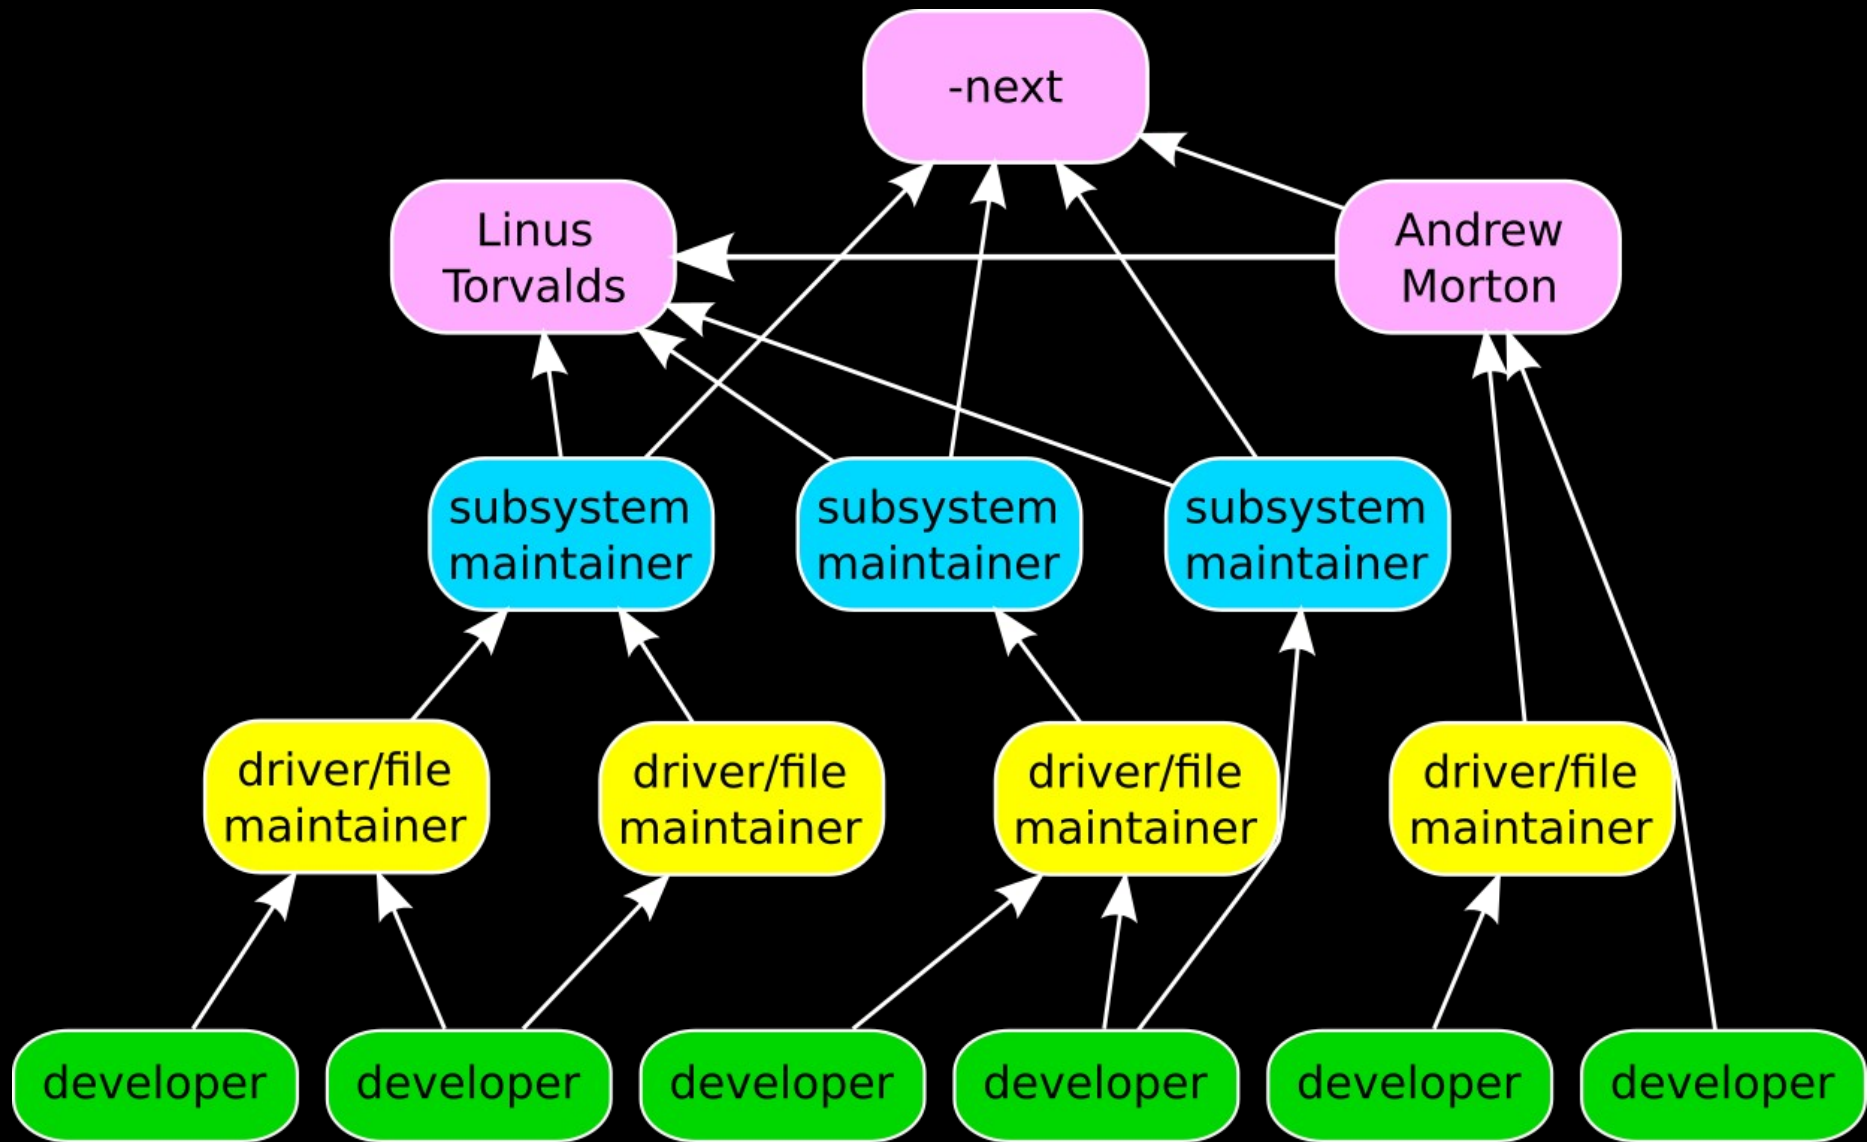
Trust

NO WARRANTY

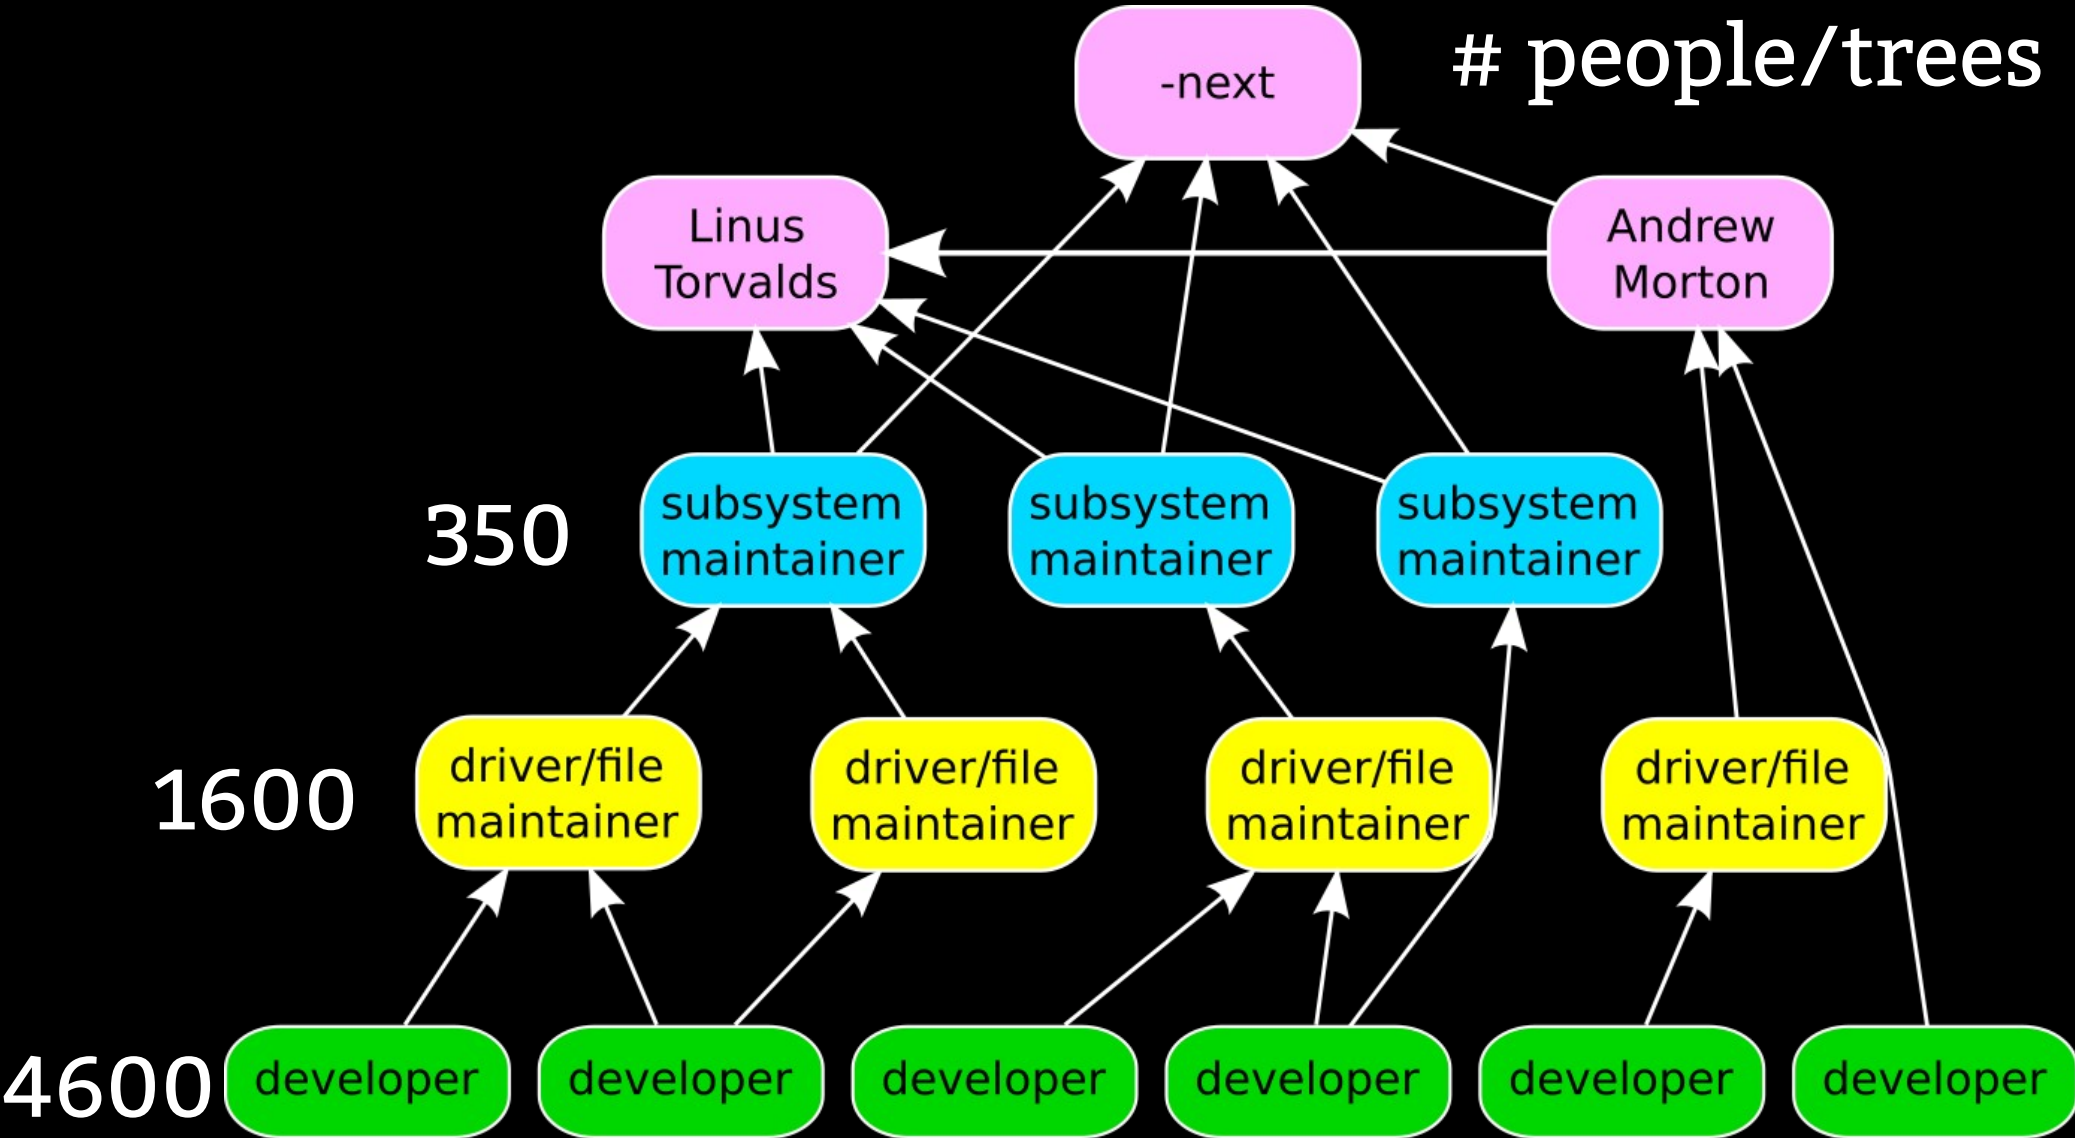
11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

Trust

“You need to verify all developers to ensure you know who they are.”



people/trees



Development stats for 2021

79.662 total commits

Fixes for 2021

79.662 total commits

13.587 commits marked with Fixes: tag

17% are fixes

Fixes for 2021

79.662 total commits

13.587 commits marked with Fixes: tag

17% are fixes

Found after commits hit subsystem trees

Fixes for 2021

79.662 total commits

13.587 commits marked with Fixes: tag

17% are fixes

Found after commits hit subsystem trees

26% of the fixes were for issues before -final

2021 changes

~ 12% of all commits were fixes for problems in older releases.

2021 - Top developers

Christoph Hellwig	960 (1.2%)
Lee Jones	737 (0.9%)
Andy Shevchenko	704 (0.9%)
Mauro Carvalho Chehab	642 (0.8%)
Pavel Begunkov	624 (0.8%)
Vladimir Oltean	600 (0.8%)
Sean Christopherson	597 (0.7%)
Colin Ian King	573 (0.7%)
Arnd Bergmann	535 (0.7%)
Geert Uytterhoeven	487 (0.6%)

2021 - Top fixers

Dan Carpenter	340 (2.5%)
Arnd Bergmann	227 (1.7%)
Colin Ian King	165 (1.2%)
Sean Christopherson	160 (1.2%)
Vladimir Oltean	143 (1.1%)
Christophe JAILLET	142 (1.0%)
Randy Dunlap	140 (1.0%)
Geert Uytterhoeven	132 (1.0%)
Johan Hovold	125 (0.9%)
Eric Dumazet	119 (0.9%)

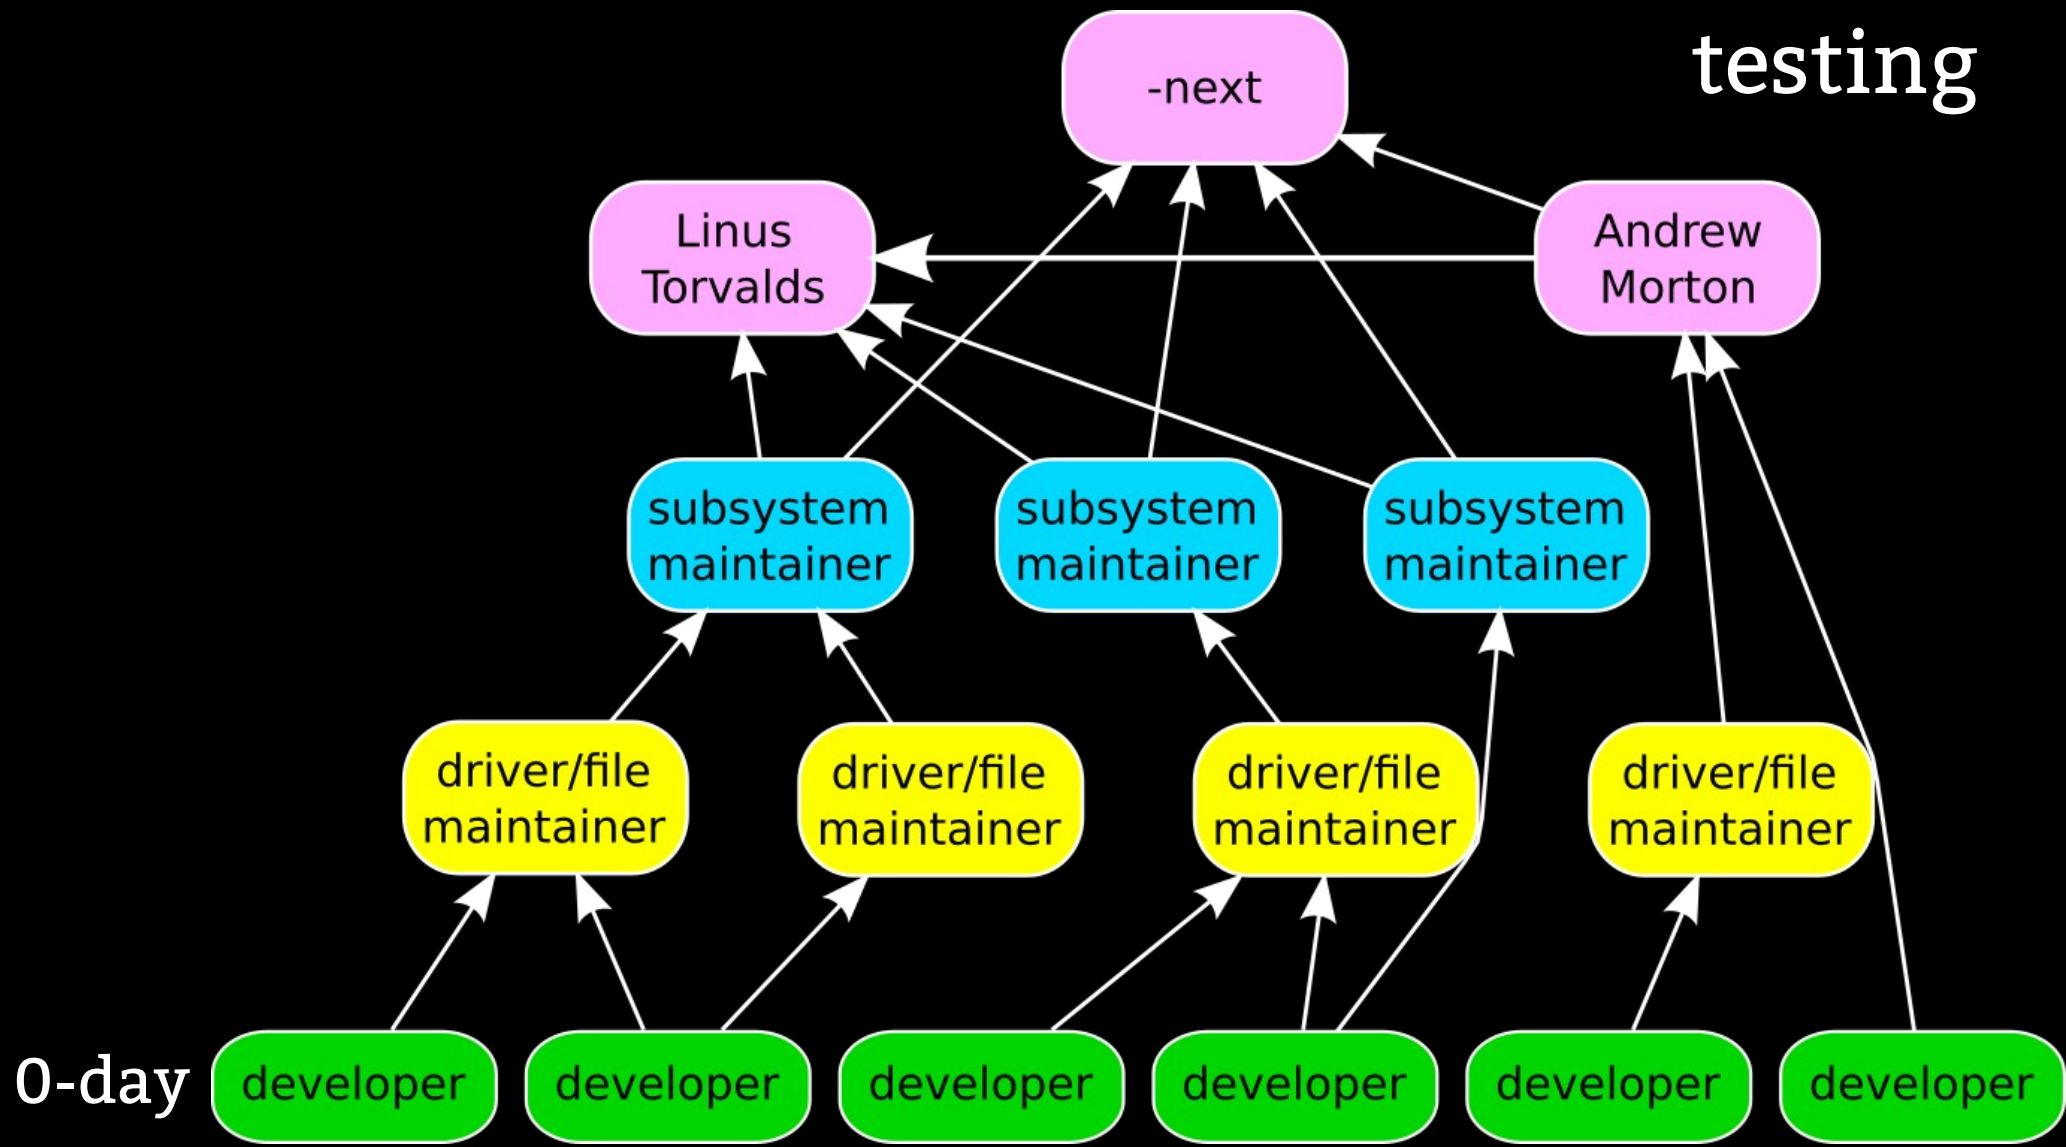
2021 - Authors of commits fixed

[masqué]	207 (1.5%)
[masqué]	161 (1.1%)
[masqué]	121 (0.8%)
[masqué]	109 (0.8%)
[masqué]	93 (0.7%)
[masqué]	89 (0.6%)
[masqué]	77 (0.5%)
[masqué]	75 (0.5%)
[masqué]	69 (0.5%)
[masqué]	67 (0.5%)

Over time, the most prolific developers
will write the most bugs.

Over time, the most prolific developers will write the most bugs.

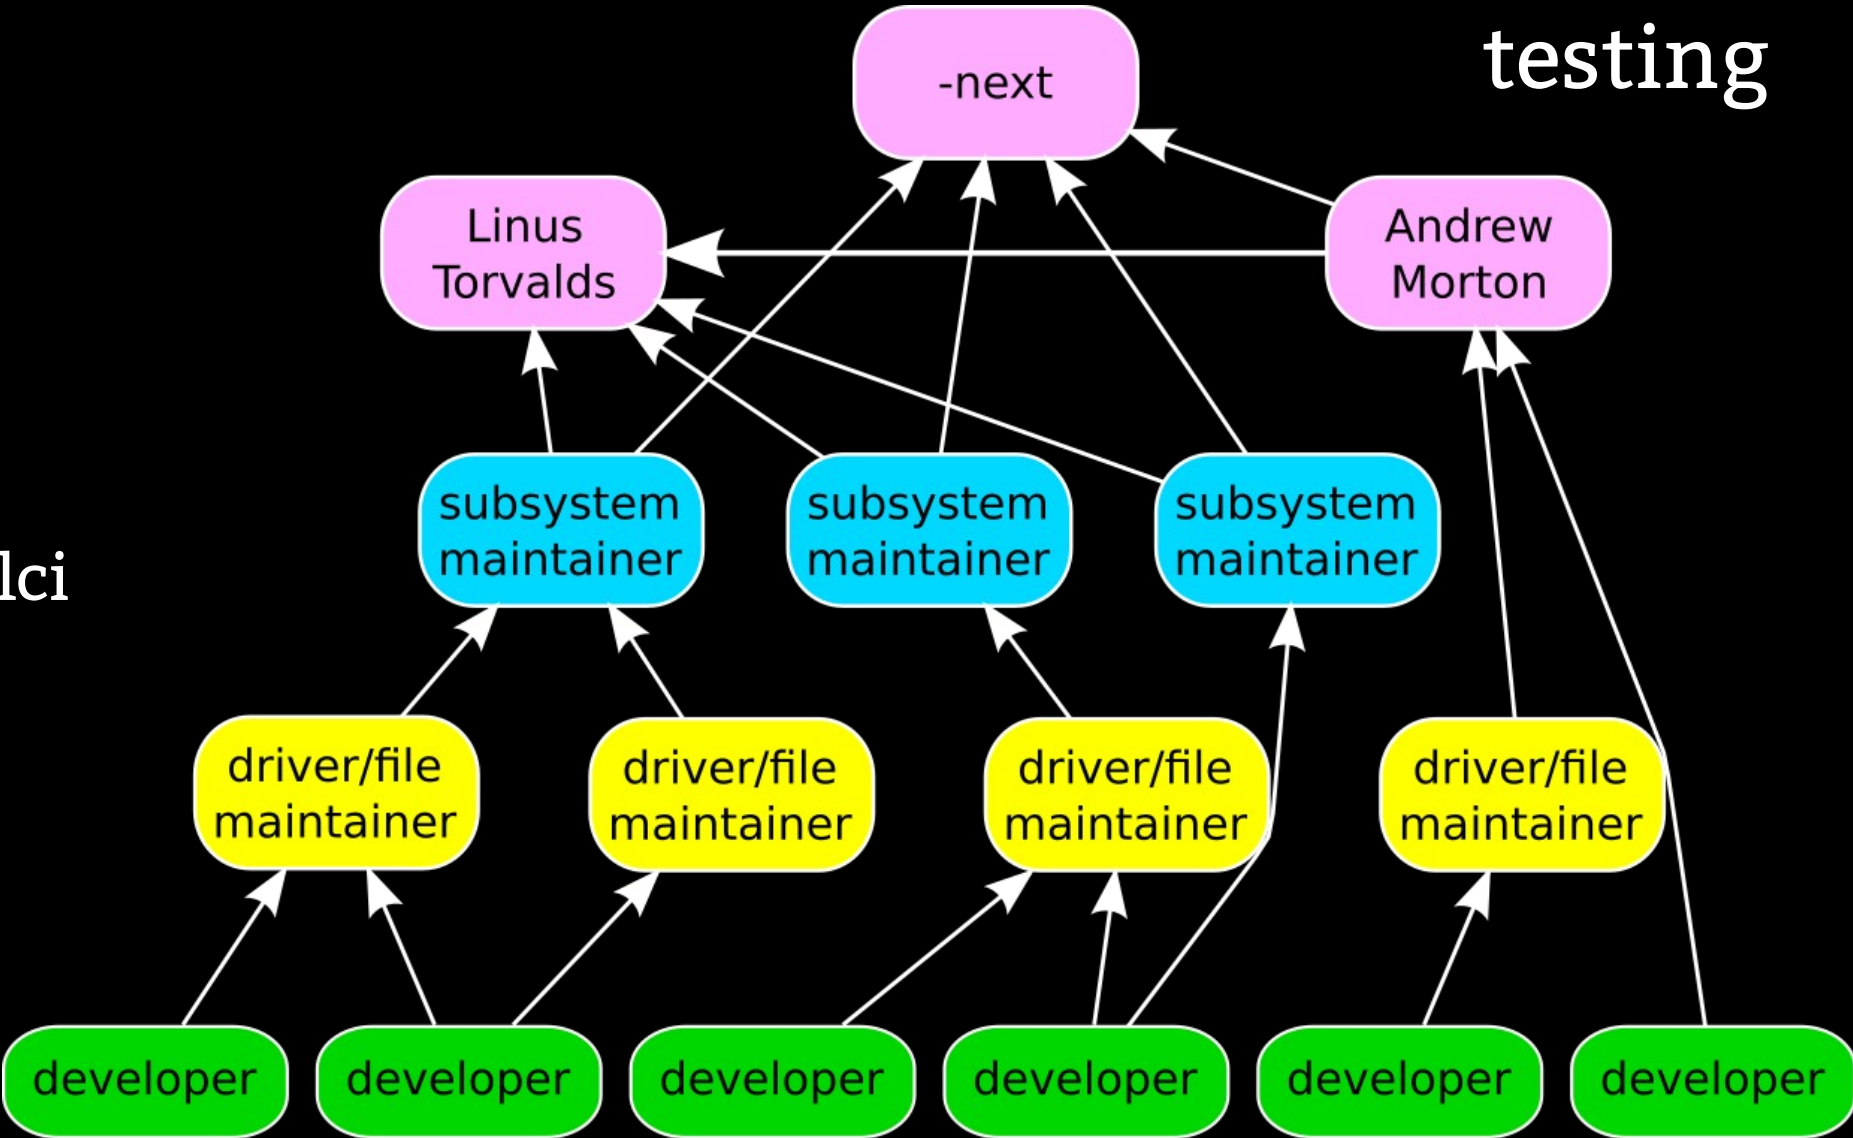
So make it easy to find and fix those bugs.



testing

0-day kernelci

0-day

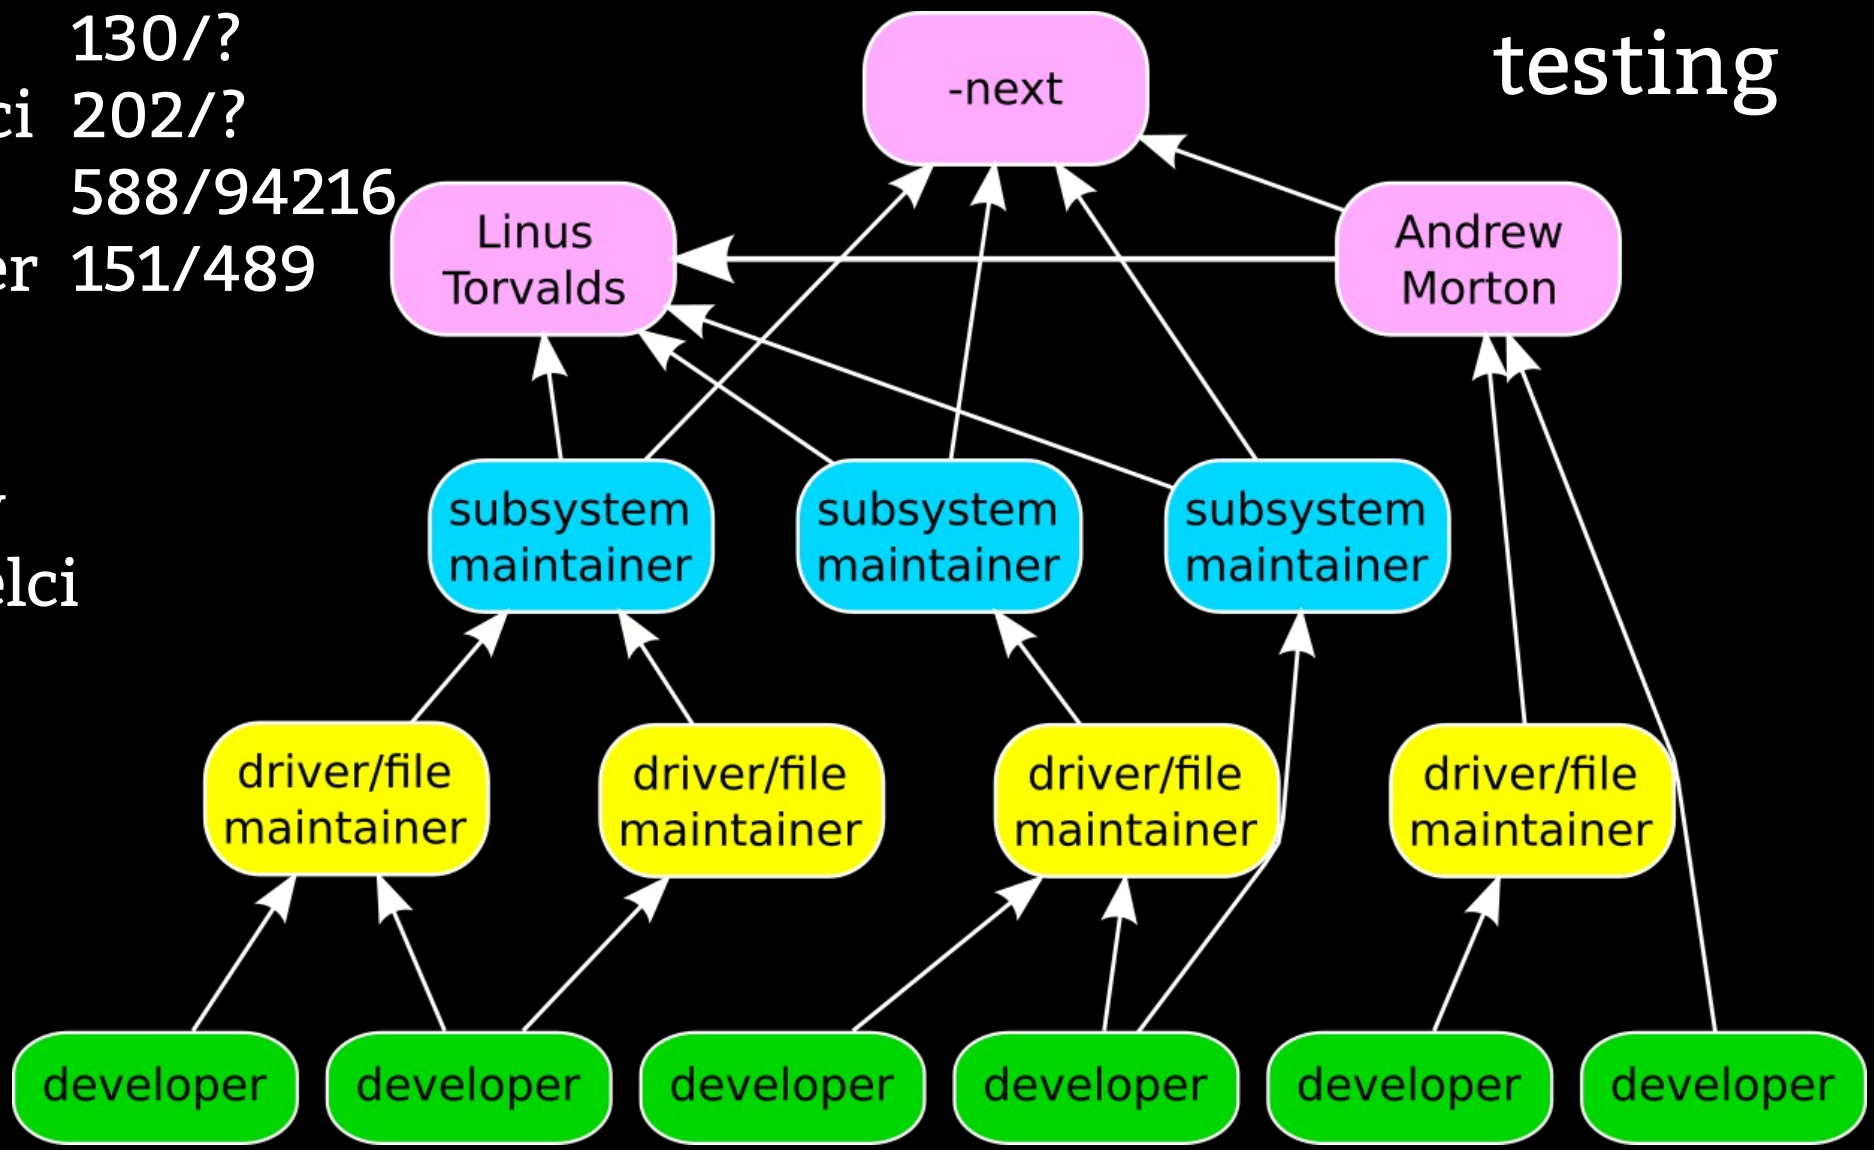



0-day 130/?
kernelci 202/?
lkft 588/94216
Guenter 151/489

testing

0-day
kernelci

0-day





Testing every release

kernelci

lkft

Guenter

Shuah

Android

Huawei

Nvidia

Debian

Fedora

Many others

Trust in Linux kernel development

Trust but verify.

Trust in Linux kernel development

Trust but test.

Trust in Linux kernel development

We trust not that you will always get it right, but that you will be there to fix it when you get it wrong.

