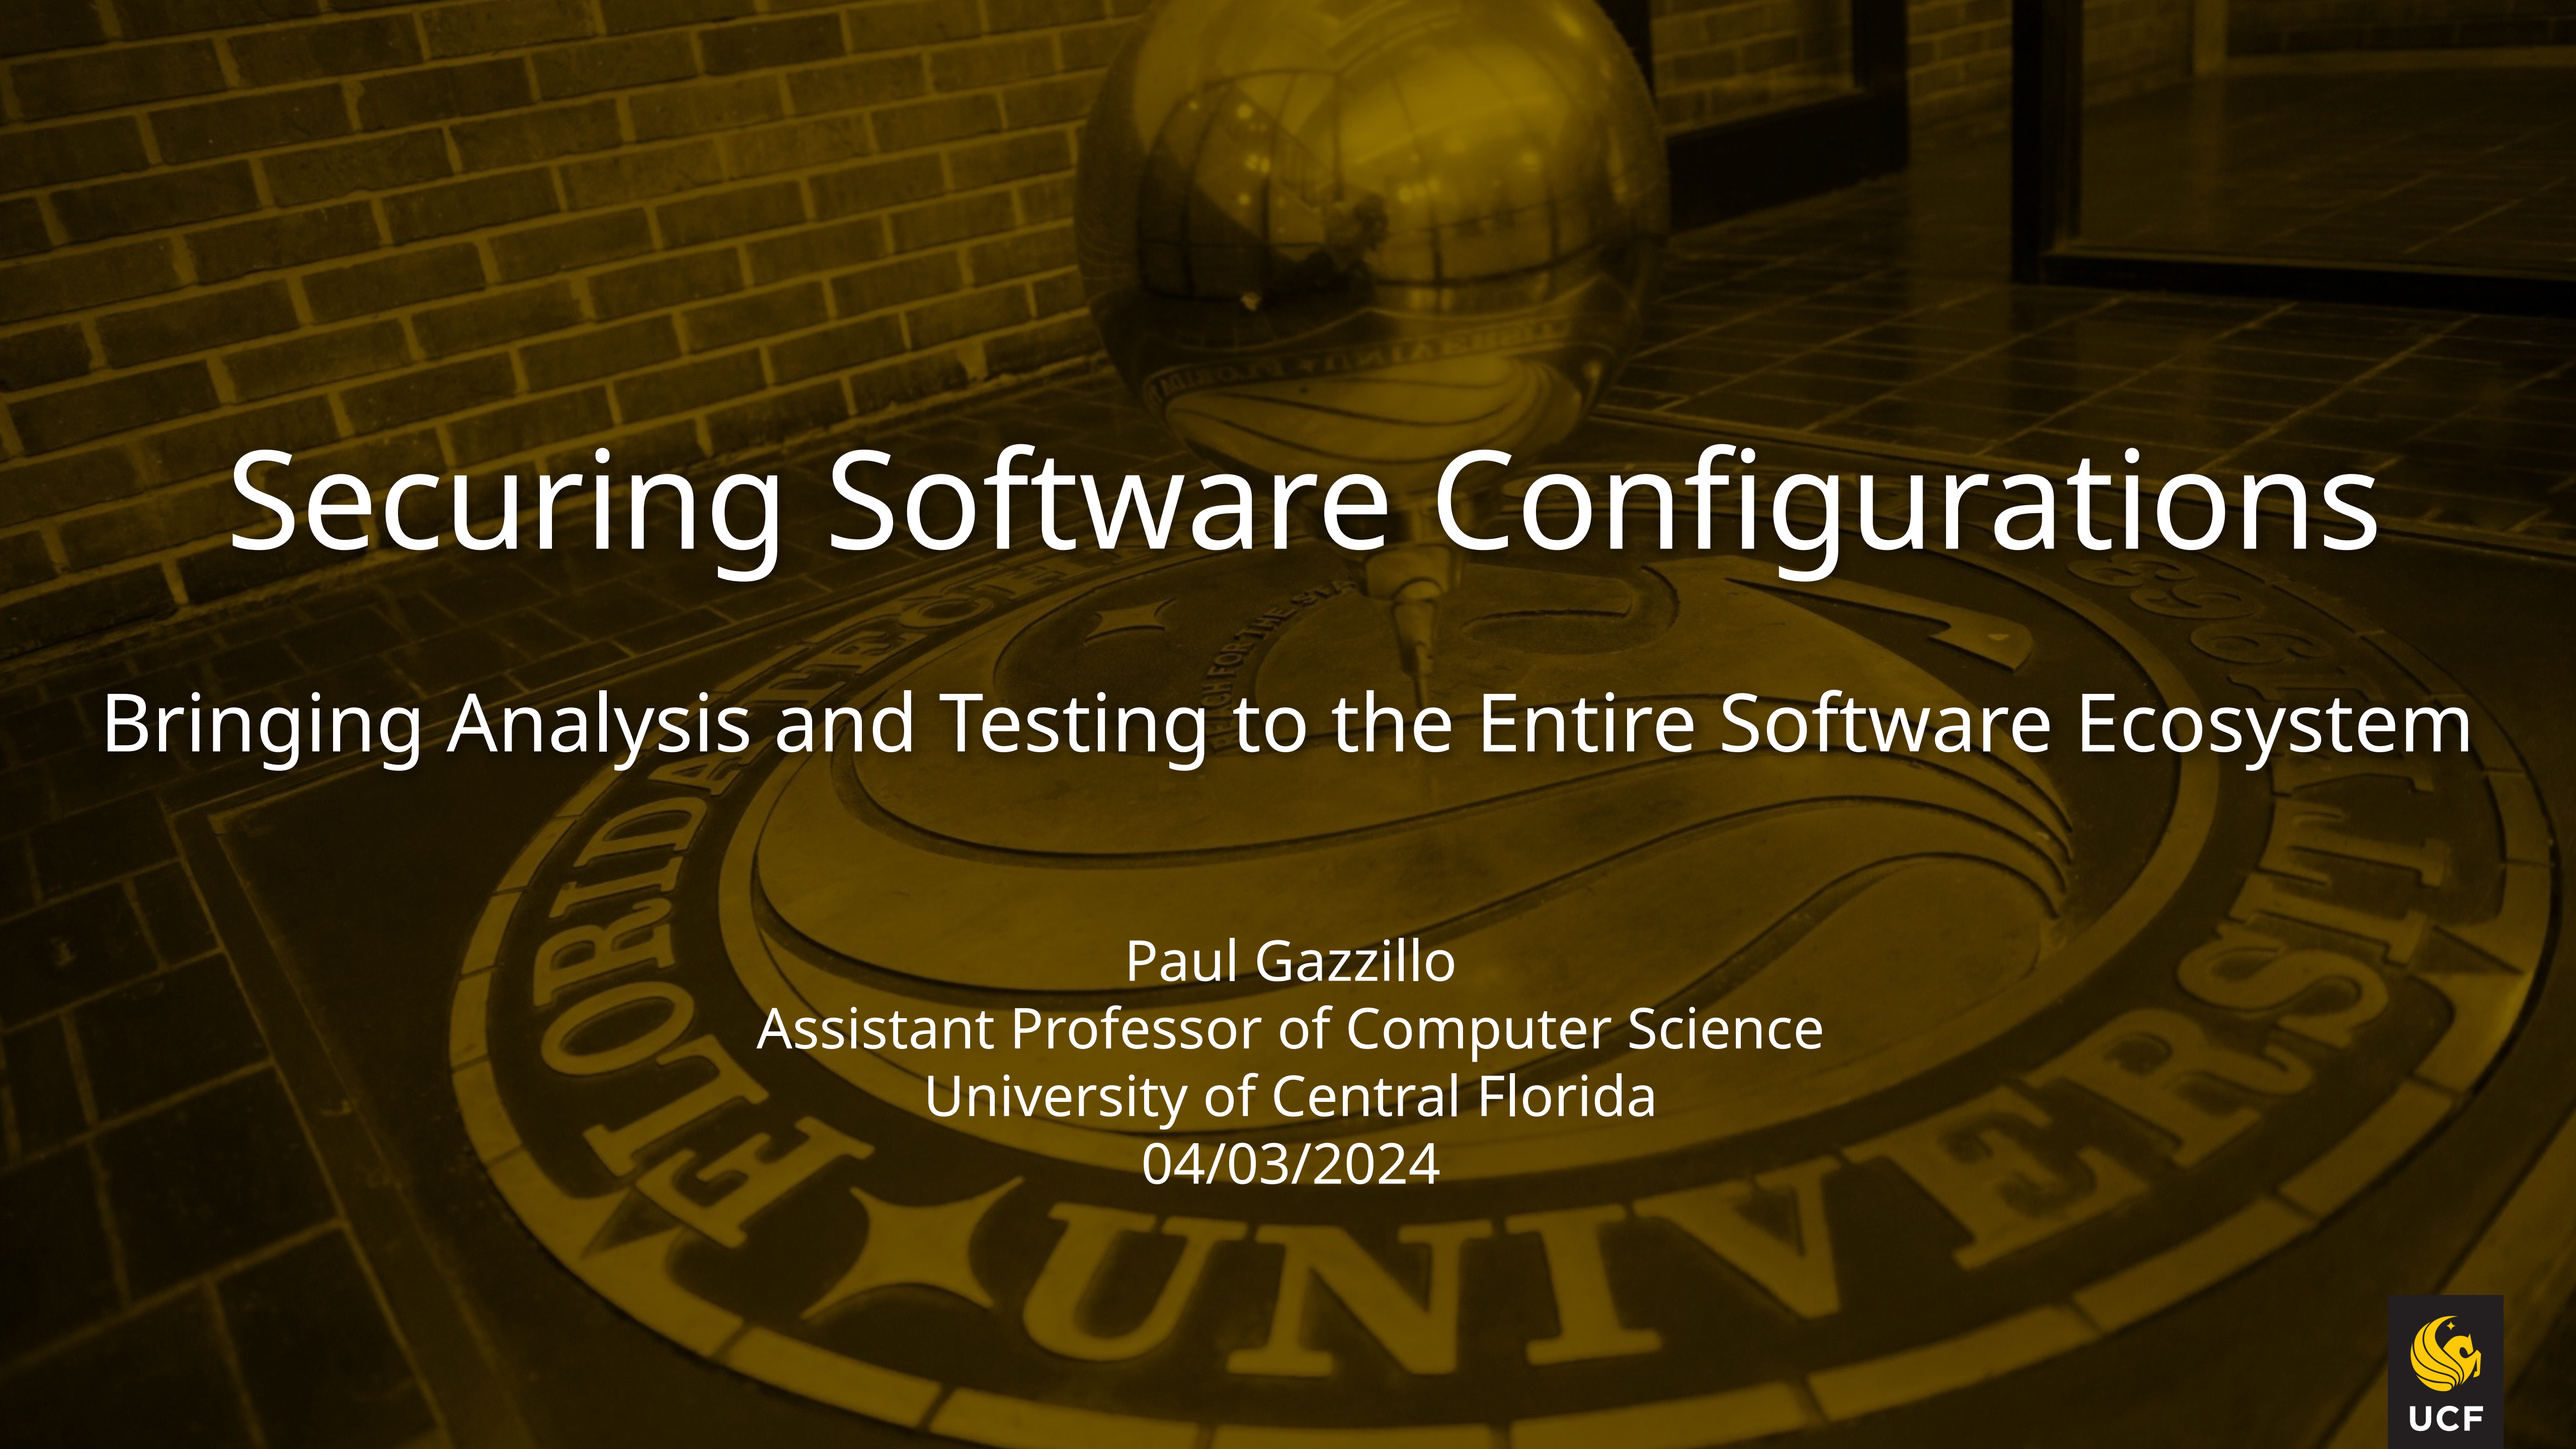

Securing Software Configurations

Bringing Analysis and Testing to the Entire Software Ecosystem

Paul Gazzillo
Assistant Professor of Computer Science
University of Central Florida
04/03/2024

Vision

Expand the scope of software analysis beyond the programming language to the entire software ecosystem to further strengthen and secure software.

Memory Safety Dominates Exploits

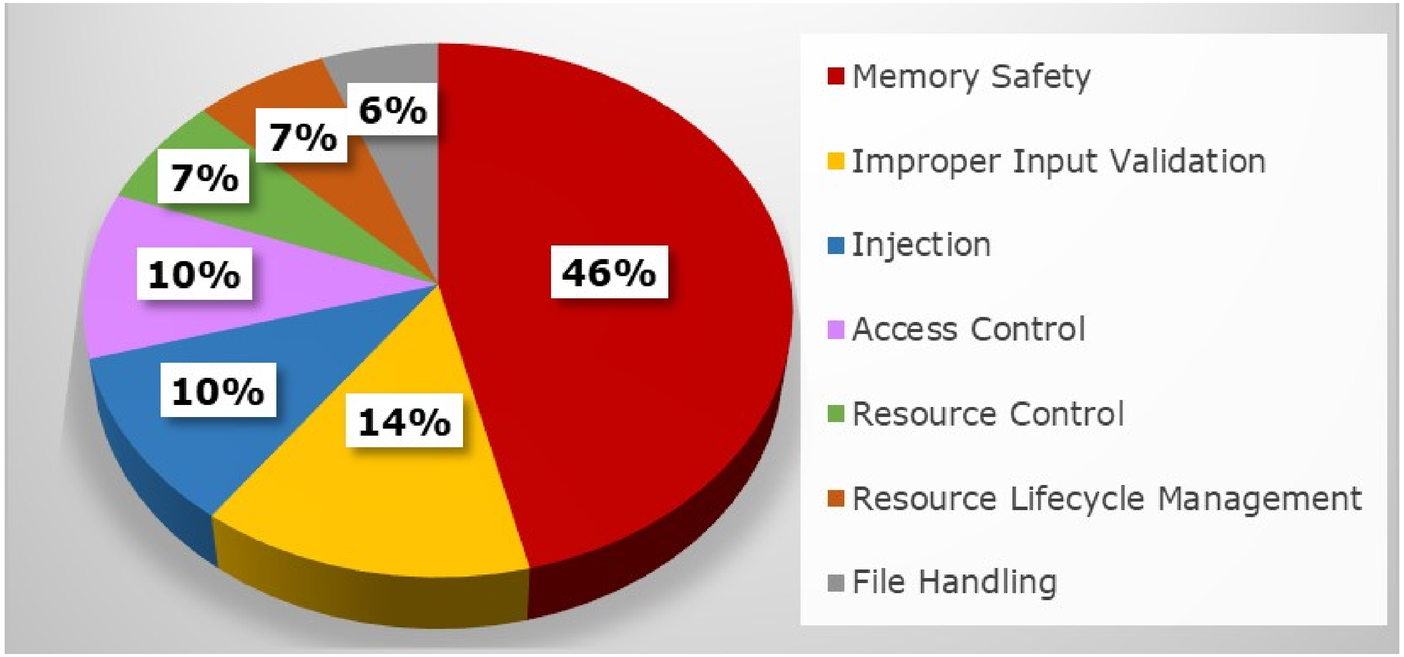

Source: [2023 CWE Top 10 KEV Weaknesses List Insights]

C/C++ Is the Origin

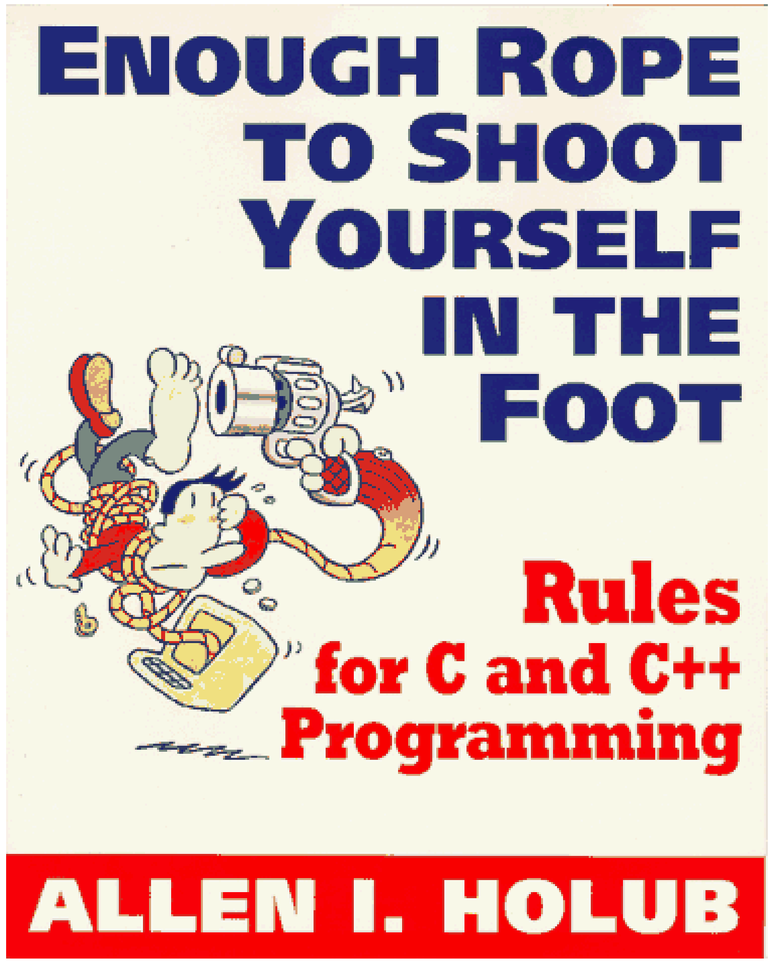

Memory Safe Programming is Solved

CYCLONE

Checked C

UCF

Just a Matter of Time

THE WHITE HOUSE

PRESS RELEASE: Future Software Should Be Memory Safe

FEBRUARY 26, 2024

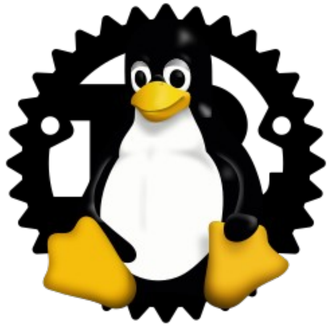

Rust for Linux

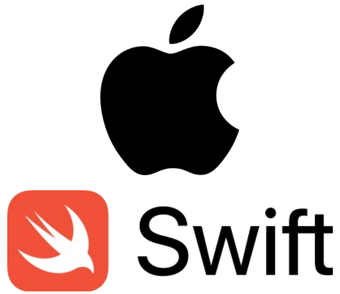

Swift for Apple

What’s Left After Memory-Based Exploits?

Memory bugs happen here

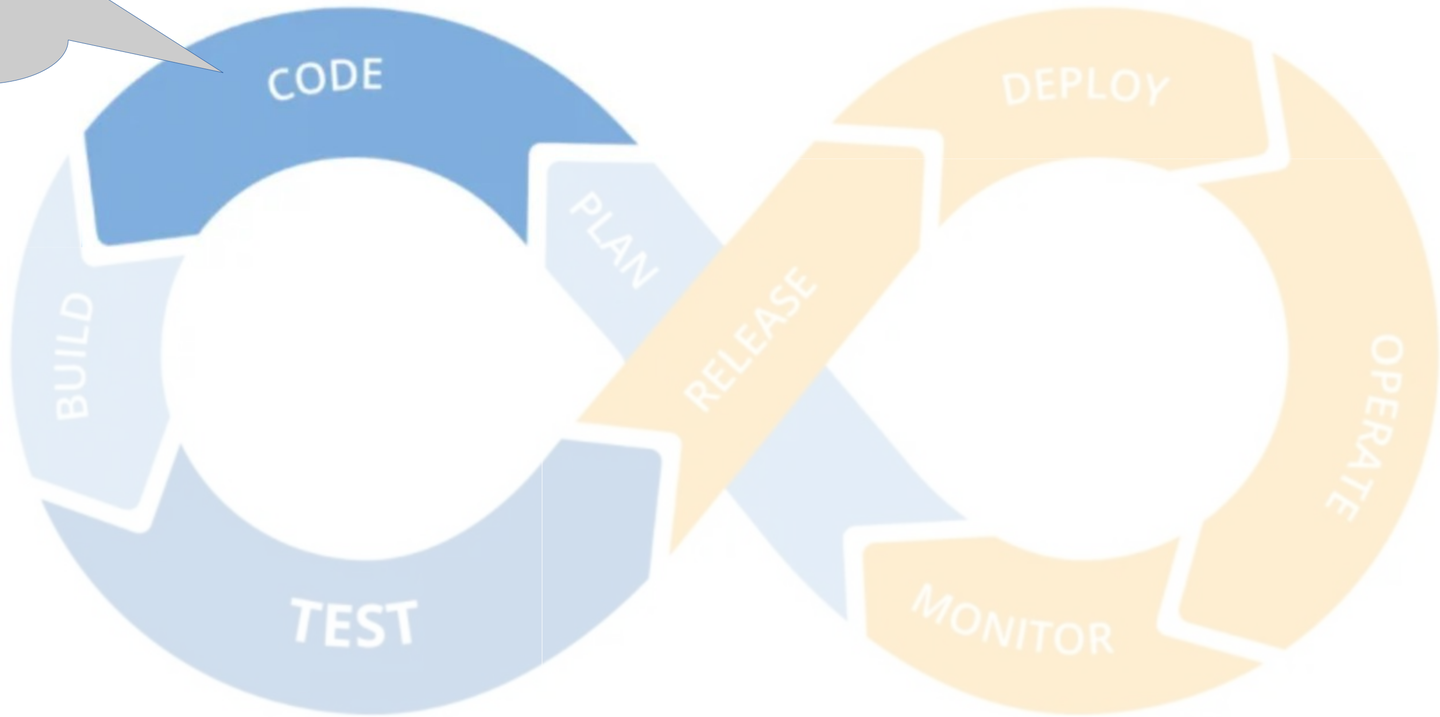

DevOps Phases

Other Phases of Development and Operations

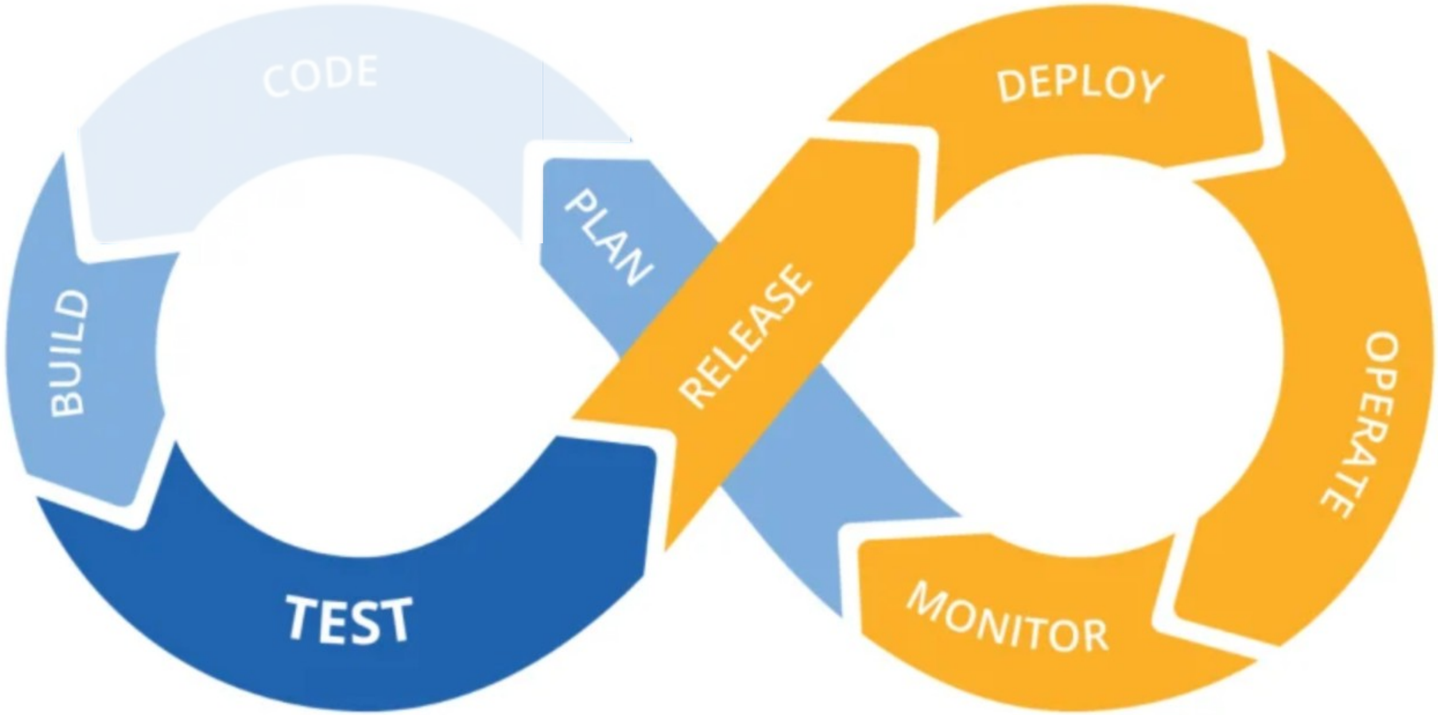

DevOps Phases

High Profile Attacks

solarwinds

- Hacked build system
- Malware in signed code
- “More than 200 victims”

- Feature, not bug
- Disable with configuration setting
- “Most critical vulnerability”

https://www.bloomberg.com/news/articles/2020-12-19/at-least-200-victims-identified-in-suspected-russian-hacking

https://www.theguardian.com/technology/2021/dec/10/software-flaw-most-critical-vulnerability-log-4-shell

UCF

Reflections on Trusting Trust

To what extent should one trust a statement that a program is free of Trojan horses? Perhaps it is more important to trust the people who wrote the software.

KEN THOMPSON

Why Bother Breaking In?

Inferring and Securing Software Configurations

Misconfiguration Vulnerabilities Are Prevalent

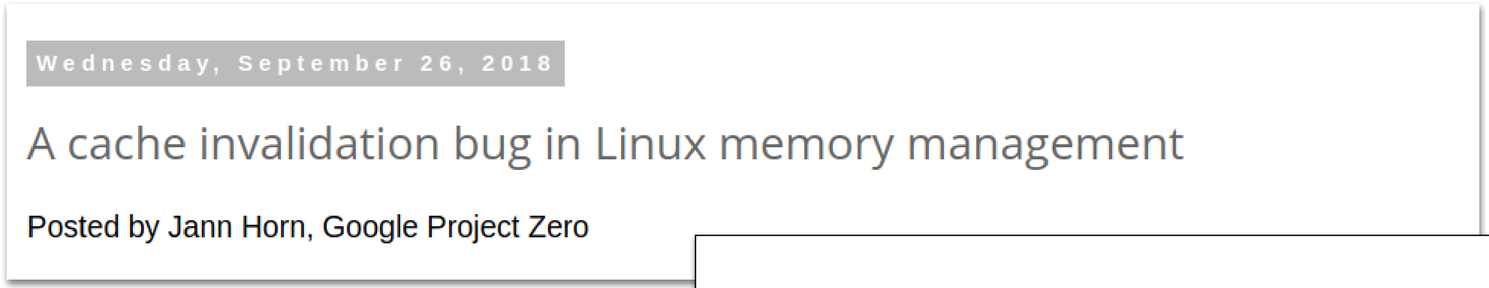

Wednesday, September 26, 2018

A cache invalidation bug in Linux memory management

Posted by Jann Horn, Google Project Zero

"This exploit shows how much impact the kernel configuration can have on how easy it is to write an exploit for a kernel bug."

#6 in OWASP top ten most critical security risks

most common risk reported

Highly-Configurable Software is Widespread

Linux kernel

- 70% of mobile devices
- 70% of IoT developers
- 40% of servers

billions of devices

Apache web server

- 40% of servers

Misconfiguration Vulnerabilities are Rooted in Software Configuration Management

Manages *change* to a software system

Allows customizing software without reprogramming

Falls outside of classic program analysis

Goal: a world without misconfigurations

Solution approach: formal methods to validate and generate software configurations

Challenges: a lack of existing specifications, an enormous state space

Research Goals

Create a rigorous definition of configuration specifications

Mechanize the generation of valid configurations

Automatically discover secure configurations

Motivating Example: OptionsBleed

A Limit Directive Restricts Access to HTTP Methods in an Apache Webserver

```
<Limit PUT DELTE BIND>
</Limit>
```

OptionsBleed Leaks Arbitrary Memory Contents of an Apache Webserver

invalid http method exposes
a use-after-free bug

Subtle Interactions Between Configuration Mechanisms Influence OptionsBleed's Occurrence

```
<Limit PUT DELTE BIND>
</Limit>
```

BIND is only valid with the WebDAV HTTP extension

Subtle Interactions Between Configuration Mechanisms Influence OptionsBleed's Occurrence

WebDAV is enabled only with a compile-time flag and run-time module loader

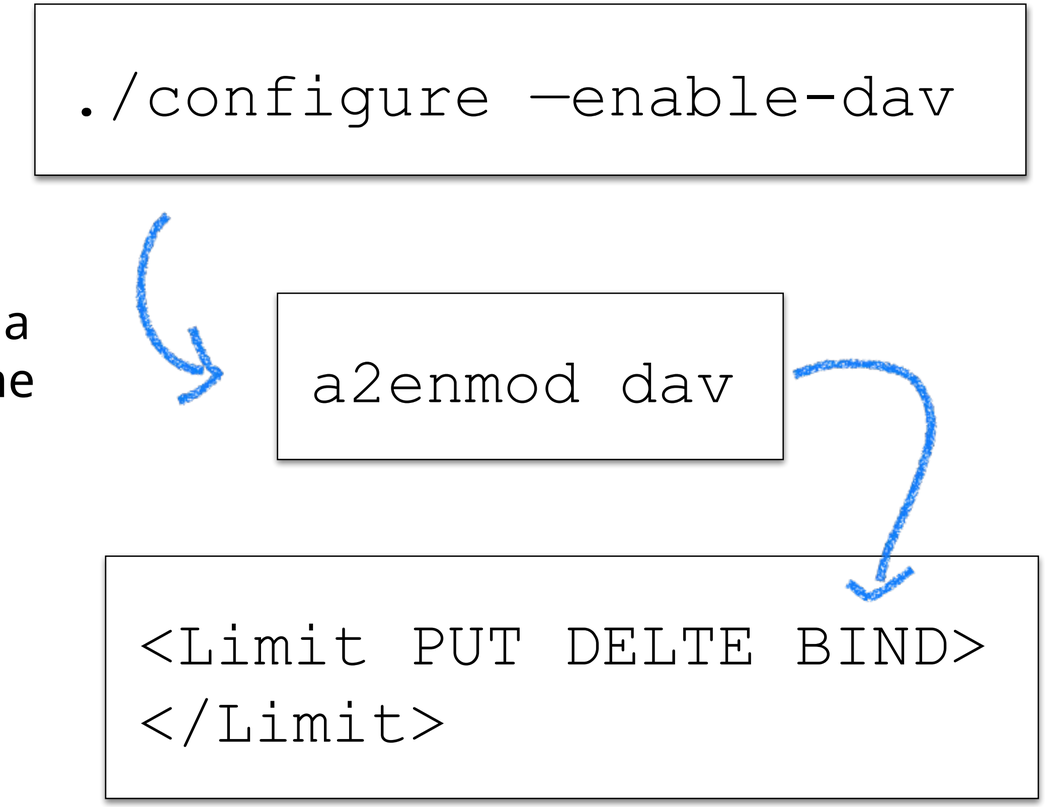

Solution approach: automatically validate and generate software configurations

Automation needs a unified global view of configuration specifications

Configuration options are long-lived values, global to an entire software system

Formalize Valid Configurations as Constraints Among All Configuration Options

```
    limit.method = PUT
or  limit.method = DELETE
or (limit.method = BIND
    and build.enable-dav =
True
    and module.dav = True)
```

configuration validity is satisfiability

build

```
./configure –enable-dav
```

module

```
a2enmod dav
```

limit

```
<Limit PUT DELTE BIND>
</Limit>
```

UCF

Formalizing the Linux Build and Configuration System

The Linux Kernel Build System

Example: Linux Kernel

Linux Kernel

70% of mobile devices

70% of IoT developers

40% of servers

UCF

The Kernel is Ultra-Configurable

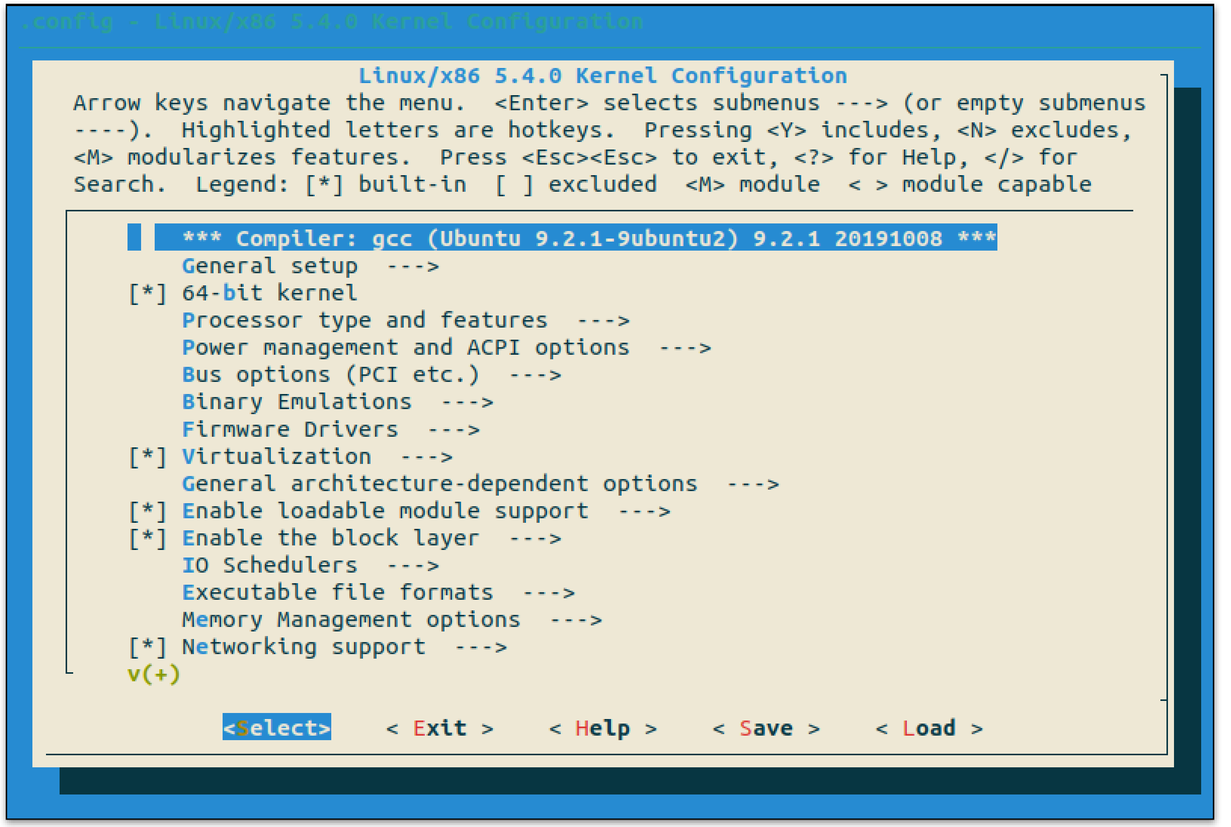

Configurability Makes Maintenance Harder

given a patch, what configurations does it affect? (jmake, lawall et al)

given a bug, what configurations does it appear in? (config-bisect)

what’s a minimal configuration that includes specific source? (config-bisect)

what code is no longer configurable in the kernel? (undertaker, tarlet et al)

There’s About 15,000 Configuration Options

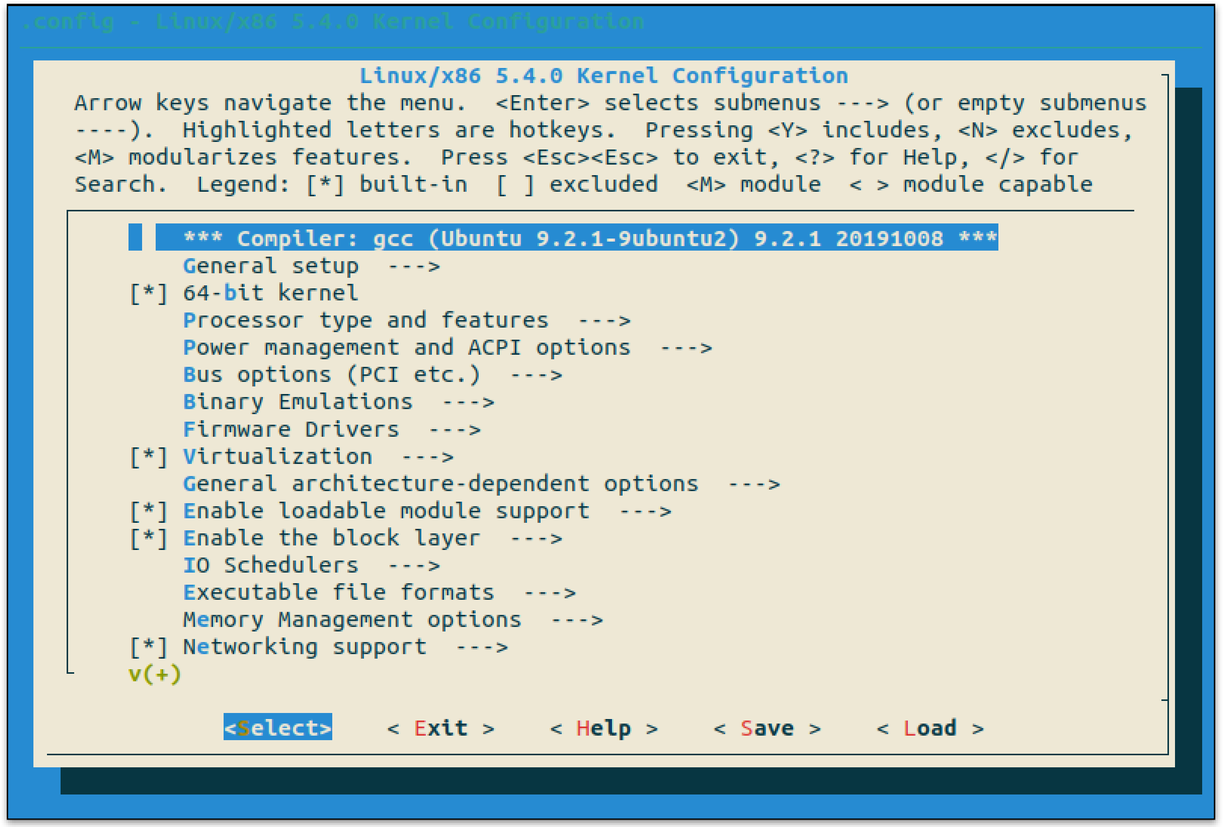

Written in about 150,000 Lines of Kconfig

there’s around 1,500 Kconfig files

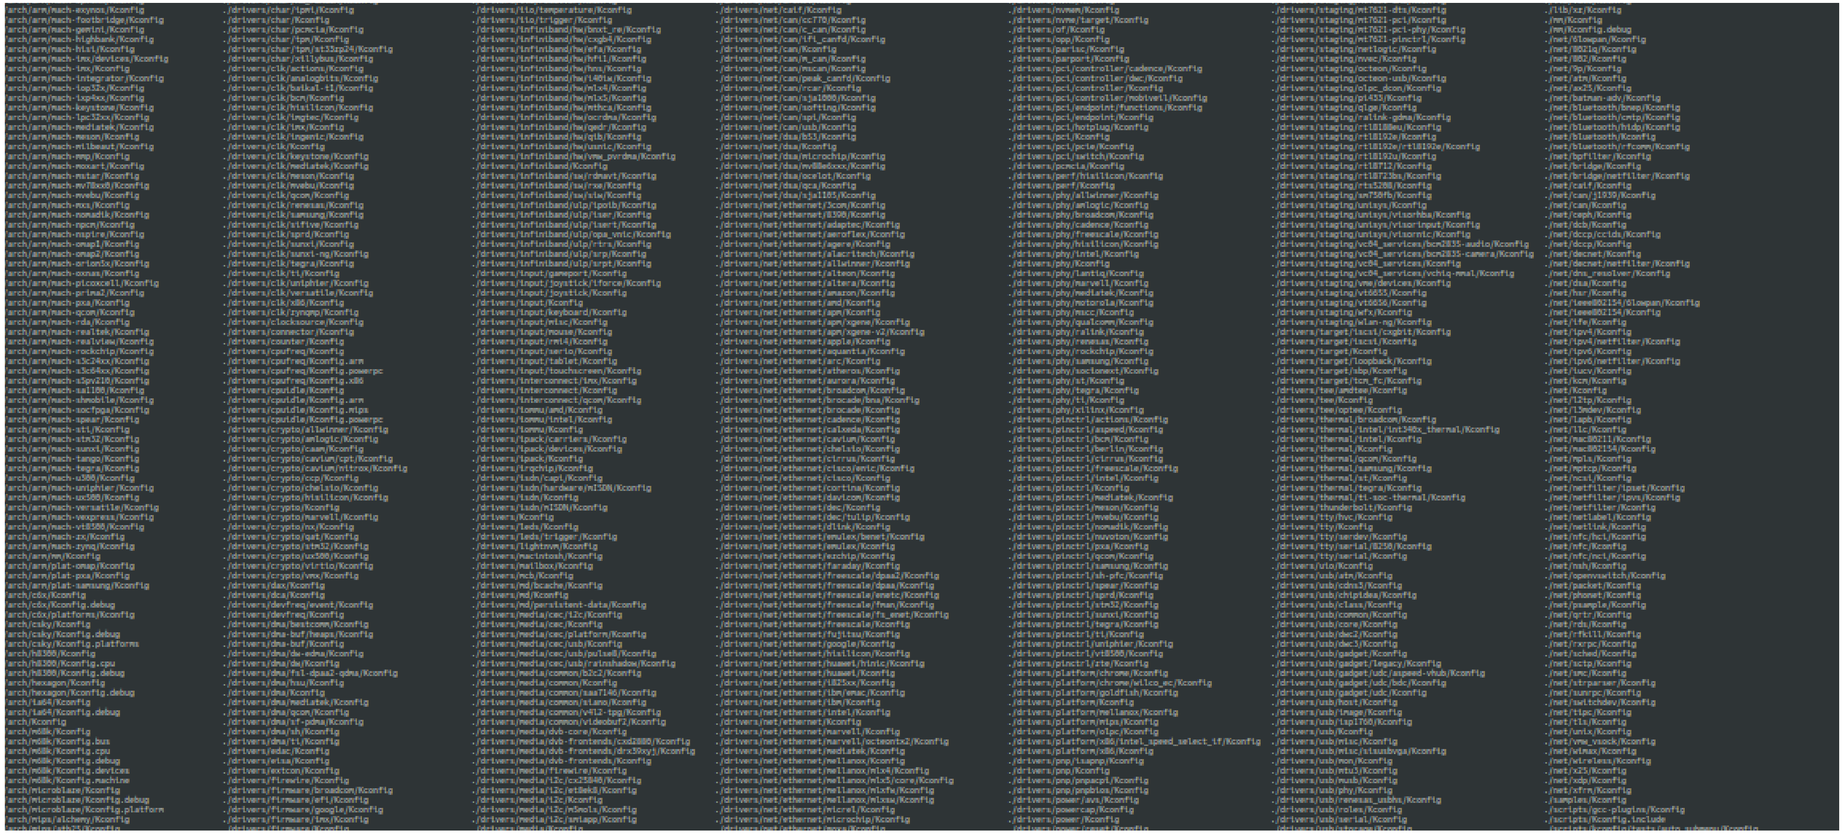

Can Have Trillions of Program in One Codebase

- Allows system builders to reuse existing software

Configuration options enable/disable features

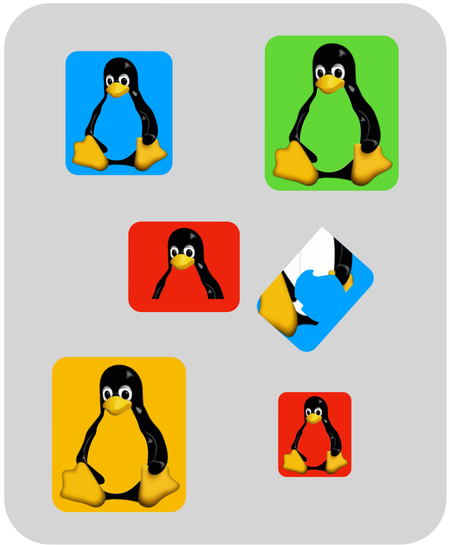

Linux build system generates many variations

Build customized software without reprogramming

Configurability Complicates Maintenance

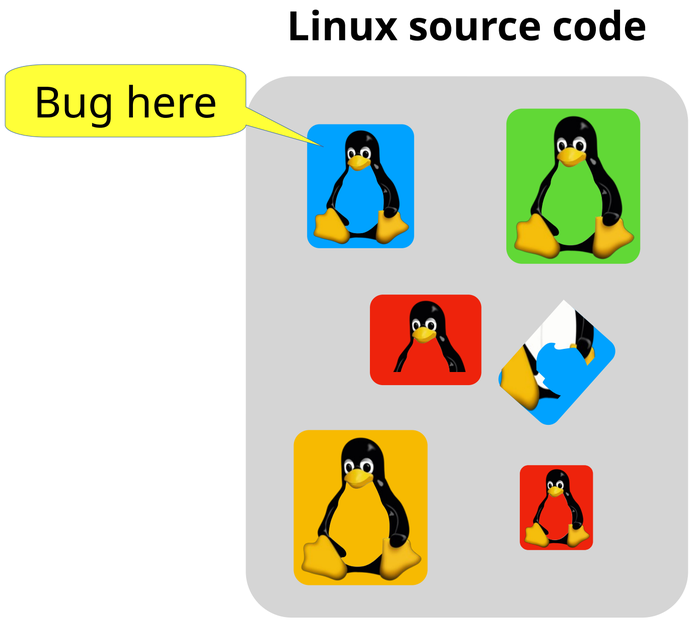

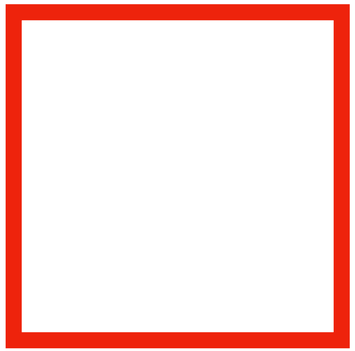

- Even if one variant program is correct, another might be broken
- Have to test all variations that might be used
- Automated testing typically works on one variant at a time

UCF

The Linux Kernel has a Very Active Codebase

Linux-next commit history

index : kernel/git/next/linux-next.git

The linux-next integration testing tree

Linux Next Group

about summary refs log tree commit diff stats

Age	Commit message (Expand)	Author	Files	Lines
8 hours	Add linux-next specific files for 20240213 HEAD next-20240213 master	Stephen Rothwell	4	-0/+9642
8 hours	fixup for "drm/amd: Stop evicting resources on APUs in suspend"	Stephen Rothwell	1	-1/+1
9 hours	Merge branch 'for-next/kspp' of git://git.kernel.org/pub/scm/linux/kernel/git...	Stephen Rothwell	45	-185/+336
9 hours	Merge branch 'bitmap-for-next' of https://github.com/norov/linux.git	Stephen Rothwell	49	-417/+634
9 hours	Merge branch 'for-next/execve' of git://git.kernel.org/pub/scm/linux/kernel/g...	Stephen Rothwell	1	-1/+1
9 hours	Merge branch 'rust-next' of https://github.com/Rust-for-Linux/linux.git	Stephen Rothwell	14	-49/+177
9 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/mic/linu...	Stephen Rothwell	16	-91/+1183
9 hours	Merge branch 'slab/for-next' of git://git.kernel.org/pub/scm/linux/kernel/git...	Stephen Rothwell	7	-126/+115
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/kris...	Stephen Rothwell	3	-20/+42
9 hours	Merge branch 'zstd-next' of https://github.com/terrelln/linux.git	Stephen Rothwell	58	-2594/+4789
9 hours	Merge branch 'mhi-next' of git://git.kernel.org/pub/scm/linux/kernel/git/mani...	Stephen Rothwell	9	-88/+423
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/srin...	Stephen Rothwell	6	-69/+83
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/srin...	Stephen Rothwell	1	-4/+4
9 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/joel/fsi...	Stephen Rothwell	2	-2/+18
9 hours	Merge branch 'for-next/seccomp' of git://git.kernel.org/pub/scm/linux/kernel/...	Stephen Rothwell	3	-14/+73
9 hours	Merge branch 'ntb-next' of https://github.com/jonmason/ntb.git	Stephen Rothwell	2	-2/+2
9 hours	Merge branch 'libnvdimm-for-next' of git://git.kernel.org/pub/scm/linux/kerne...	Stephen Rothwell	3	-3/+3
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/live...	Stephen Rothwell	0	-0/+0
9 hours	Merge branch 'kunit' of git://git.kernel.org/pub/scm/linux/kernel/git/shuah/l...	Stephen Rothwell	2	-3/+4
9 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/shuah/li...	Stephen Rothwell	32	-121/+340
9 hours	Merge branch 'pwm/for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/...	Stephen Rothwell	12	-384/+373
9 hours	Merge branch 'renesas-pinctrl' of git://git.kernel.org/pub/scm/linux/kernel/g...	Stephen Rothwell	4	-53/+276
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/linu...	Stephen Rothwell	24	-94/+130
9 hours	Merge branch 'gpio/for-next' of git://git.kernel.org/pub/scm/linux/kernel/git...	Stephen Rothwell	46	-626/+2510
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/remo...	Stephen Rothwell	12	-193/+271
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/mkp/...	Stephen Rothwell	24	-883/+727
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/jejb...	Stephen Rothwell	43	-471/+1385
9 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/tj/c...	Stephen Rothwell	1	-8/+12
9 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/vkoul/dm...	Stephen Rothwell	15	-242/+615
10 hours	Merge branch 'counter-next' of git://git.kernel.org/pub/scm/linux/kernel/git/...	Stephen Rothwell	2	-2/+1
10 hours	Merge branch 'staging-next' of git://git.kernel.org/pub/scm/linux/kernel/git/...	Stephen Rothwell	38	-4359/+168
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/krzk...	Stephen Rothwell	3	-2/+61
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/vkoul/so...	Stephen Rothwell	2	-4/+2
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/phy/linu...	Stephen Rothwell	40	-1257/+3441
10 hours	Merge branch 'togreg' of git://git.kernel.org/pub/scm/linux/kernel/git/jic23/...	Stephen Rothwell	67	-455/+2724
10 hours	Merge branch 'icc-next' of git://git.kernel.org/pub/scm/linux/kernel/git/djak...	Stephen Rothwell	12	-1116/+1599
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/fpga...	Stephen Rothwell	4	-54/+104
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/coresigh...	Stephen Rothwell	30	-803/+1377
10 hours	Merge branch 'habanalabs-next' of git://git.kernel.org/pub/scm/linux/kernel/g...	Stephen Rothwell	13	-396/+899
10 hours	Merge branch 'char-misc-next' of git://git.kernel.org/pub/scm/linux/kernel/gi...	Stephen Rothwell	6	-10/+35
10 hours	Merge branch 'tty-next' of git://git.kernel.org/pub/scm/linux/kernel/git/greg...	Stephen Rothwell	67	-1615/+2168
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/westeri/...	Stephen Rothwell	12	-50/+278
10 hours	Merge branch 'usb-next' of git://git.kernel.org/pub/scm/linux/kernel/git/greg...	Stephen Rothwell	77	-1038/+5655
10 hours	Merge branch 'driver-core-next' of git://git.kernel.org/pub/scm/linux/kernel/...	Stephen Rothwell	7	-22/+37
10 hours	Merge branch 'for-leds-next' of git://git.kernel.org/pub/scm/linux/kernel/git...	Stephen Rothwell	22	-181/+613
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/sre/...	Stephen Rothwell	1	-1/+1
10 hours	Merge branch 'for-firmware-next' of git://git.kernel.org/pub/scm/linux/kernel...	Stephen Rothwell	1	-1/+1
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/pdx8...	Stephen Rothwell	19	-244/+669
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/tj/w...	Stephen Rothwell	9	-302/+1284
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/denn...	Stephen Rothwell	0	-0/+0
10 hours	Merge branch 'next' of https://github.com/kvm-x86/linux.git	Stephen Rothwell	85	-745/+1631
10 hours	Merge branch 'riscv_kvm_next' of https://github.com/kvm-riscv/linux.git	Stephen Rothwell	2	-11/+15
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/kvmarm/k...	Stephen Rothwell	28	-84/+247
10 hours	scsi: core: Make scsi_bus_type const	Ricardo B. Marliere	2	-2/+2
10 hours	Merge branch kvm-arm64/misc into kvmarm/next	Oliver Upton	1	-1/+1
10 hours	Merge branch 'rcu/next' of git://git.kernel.org/pub/scm/linux/kernel/git/paul...	Stephen Rothwell	32	-564/+841
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/trac...	Stephen Rothwell	0	-0/+0
10 hours	Merge branch 'edac-for-next' of git://git.kernel.org/pub/scm/linux/kernel/git...	Stephen Rothwell	27	-295/+3949
10 hours	Merge branch 'timers/drivers/next' of git://git.linaro.org/people/daniel.lezc...	Stephen Rothwell	5	-7/+22
10 hours	Merge branch 'master' of git://git.kernel.org/pub/scm/linux/kernel/git/tip/ti...	Stephen Rothwell	204	-1114/+4829
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/broo...	Stephen Rothwell	65	-586/+1037
10 hours	KVM: selftests: Print timer ctl register in ISTATUS assertion	Oliver Upton	1	-1/+1
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/krzk...	Stephen Rothwell	4	-9/+9
10 hours	Merge branch 'for-next' of git://git.kernel.org/pub/scm/linux/kernel/git/robh...	Stephen Rothwell	21	-652/+661
10 hours	scsi: core: Really include kunit tests with SCSI_LIB_KUNIT_TEST	Lukas Bulwahn	1	-1/+1
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/pcmoore/...	Stephen Rothwell	2	-4/+2
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/joro/iom...	Stephen Rothwell	9	-483/+457
10 hours	Merge branch 'next' of git://git.kernel.org/pub/scm/linux/kernel/git/jarkko/l...	Stephen Rothwell	1	-3/+3
10 hours	Merge branch 'next' of git://github.com/cschaufler/smack-next	Stephen Rothwell	2	-41/+86

~30k mailing list messages per month

~6k commits per month, 100s per day

e.g., ~13k commits between v5.12 and v5.13

All These Code Changes Need Testing

Most active 5.12 bug reporters

kernel test robot	184	16.1%
Syzbot	111	9.7%
Abaci Robot	107	9.4%
Dan Carpenter	44	3.9%
Hulk Robot	41	3.6%
Stephen Rothwell	28	2.5%
Randy Dunlap	19	1.7%
Kent Overstreet	12	1.1%
Guenter Roeck	11	1.0%
TOTE Robot	11	1.0%
Colin Ian King	9	0.8%
Andrii Nakryiko	8	0.7%
Juan Vazquez	7	0.6%
Arnd Bergmann	6	0.5%

https://lwn.net/Articles/853039/

Intel 0-day kernel test robot

- Suite of static and dynamic testing tools
  - compile, boot, performance, etc.
- continuously runs on new commits in linux-next

Google syzbot

- syzkaller system call fuzz tester
- continuously tests the kernel
- runs on linux-next, other versions

The Build System Causes Blindspots in Testing

Code Hidden by the Build System

Configuration options determine what's compiled

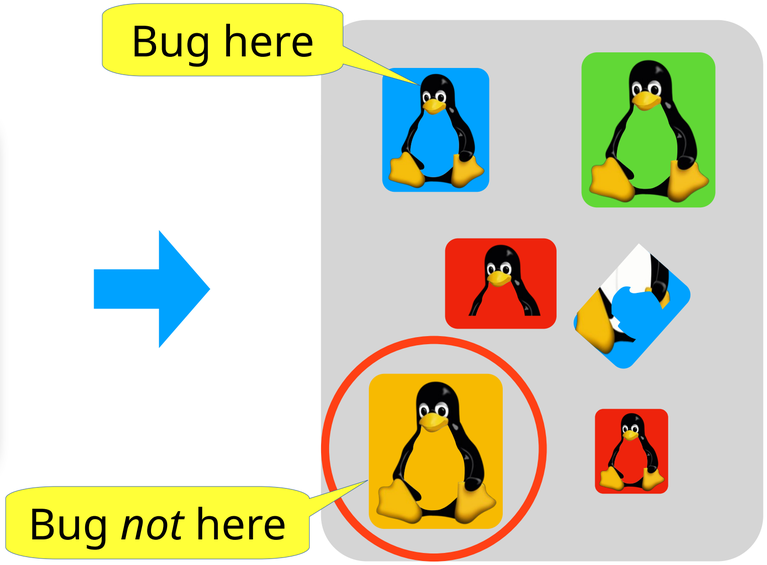

A new bug may only appear in some configurations

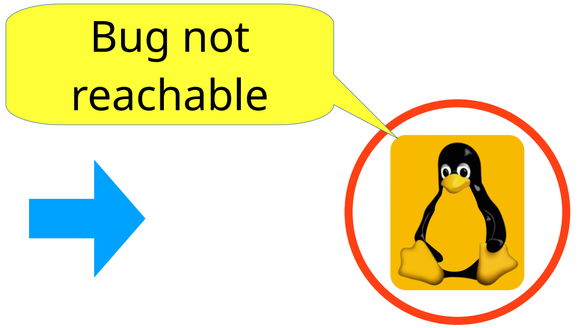

Configuration-dependent bugs not always reachable

Test Robots Miss Most Code Changes

Configuration	Patch Coverage
Default	22%
Random	30%
Maximal	89%

10x longer build time

No runtime testing

No variation

Test Robots Miss Most Code Changes

Configuration	Patch Coverage
Default	22%
Random	30%
Maximal	89%

Test robots mostly use these configurations

Maximal Testing is Limited

Most active 5.12 bug reporters

kernel test robot	184	16.1%
Syzbot	111	9.7%
Abaci Robot	107	9.4%
Dan Carpenter	44	3.9%
Hulk Robot	41	3.6%
Stephen Rothwell	28	2.5%
Randy Dunlap	19	1.7%
Kent Overstreet	12	1.1%
Guenter Roeck	11	1.0%
TOTE Robot	11	1.0%
Colin Ian King	9	0.8%
Andrii Nakryiko	8	0.7%
Juan Vazquez	7	0.6%
Arnd Bergmann	6	0.5%

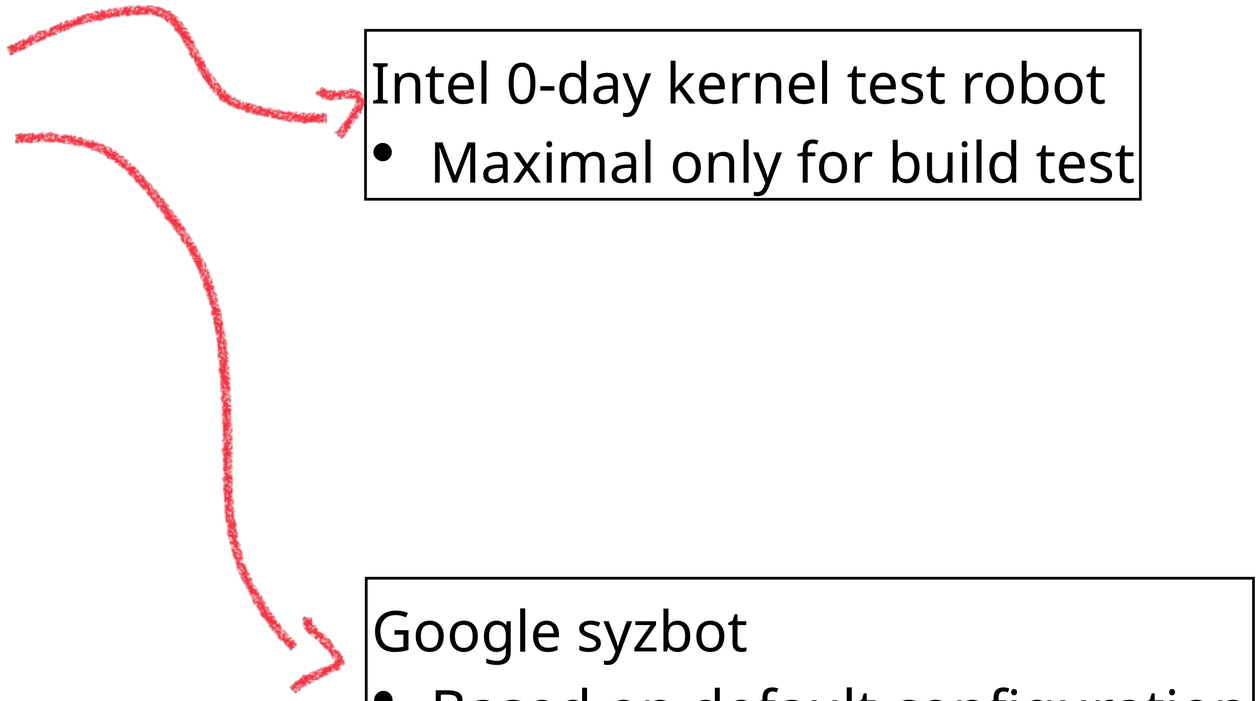

Intel 0-day kernel test robot
- Maximal only for build test

Google syzbot
- Based on default configuration

https://lwn.net/Articles/853039/

Build System Turns Configurations into Binaries

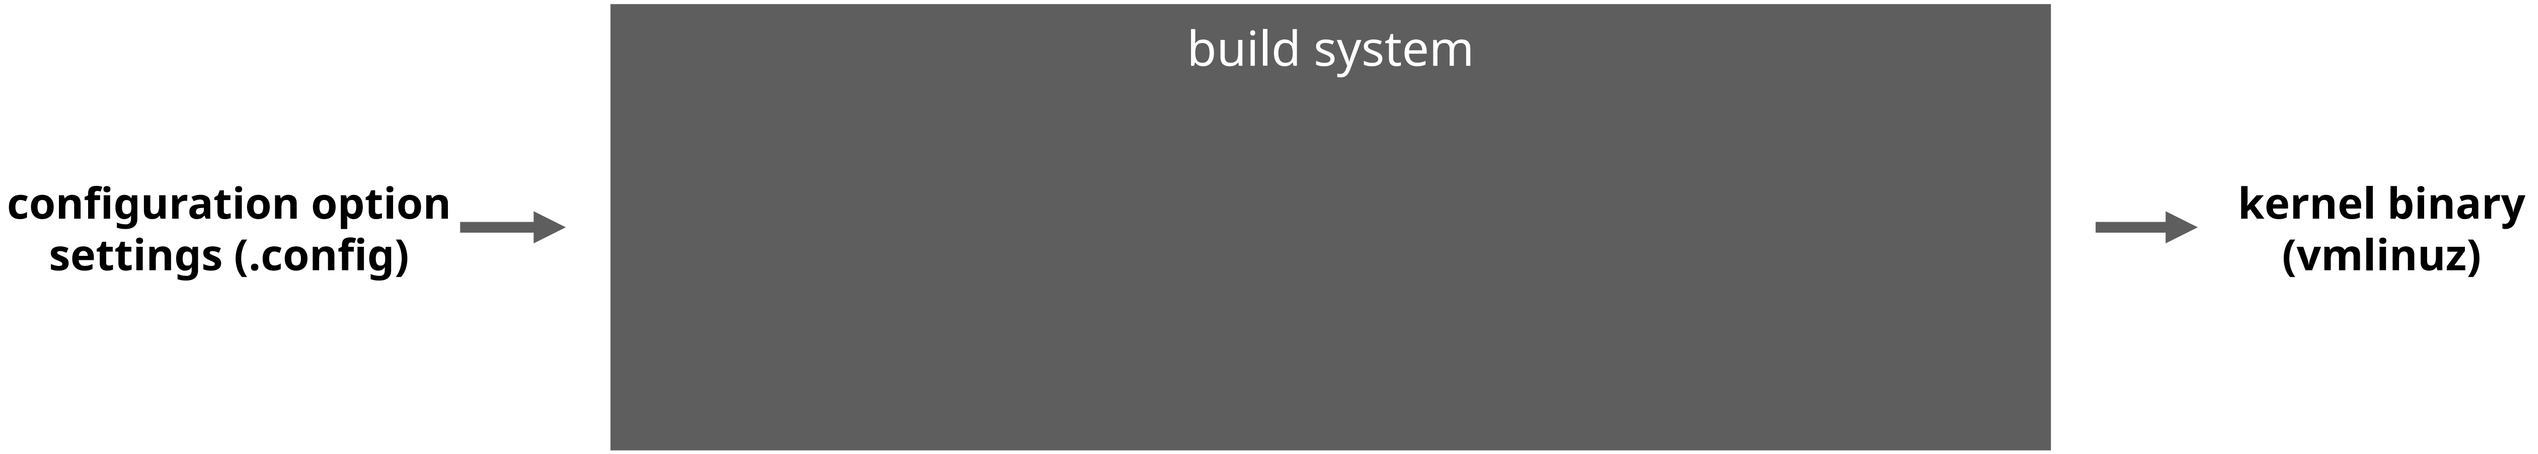

Build System Comprises Several Tools

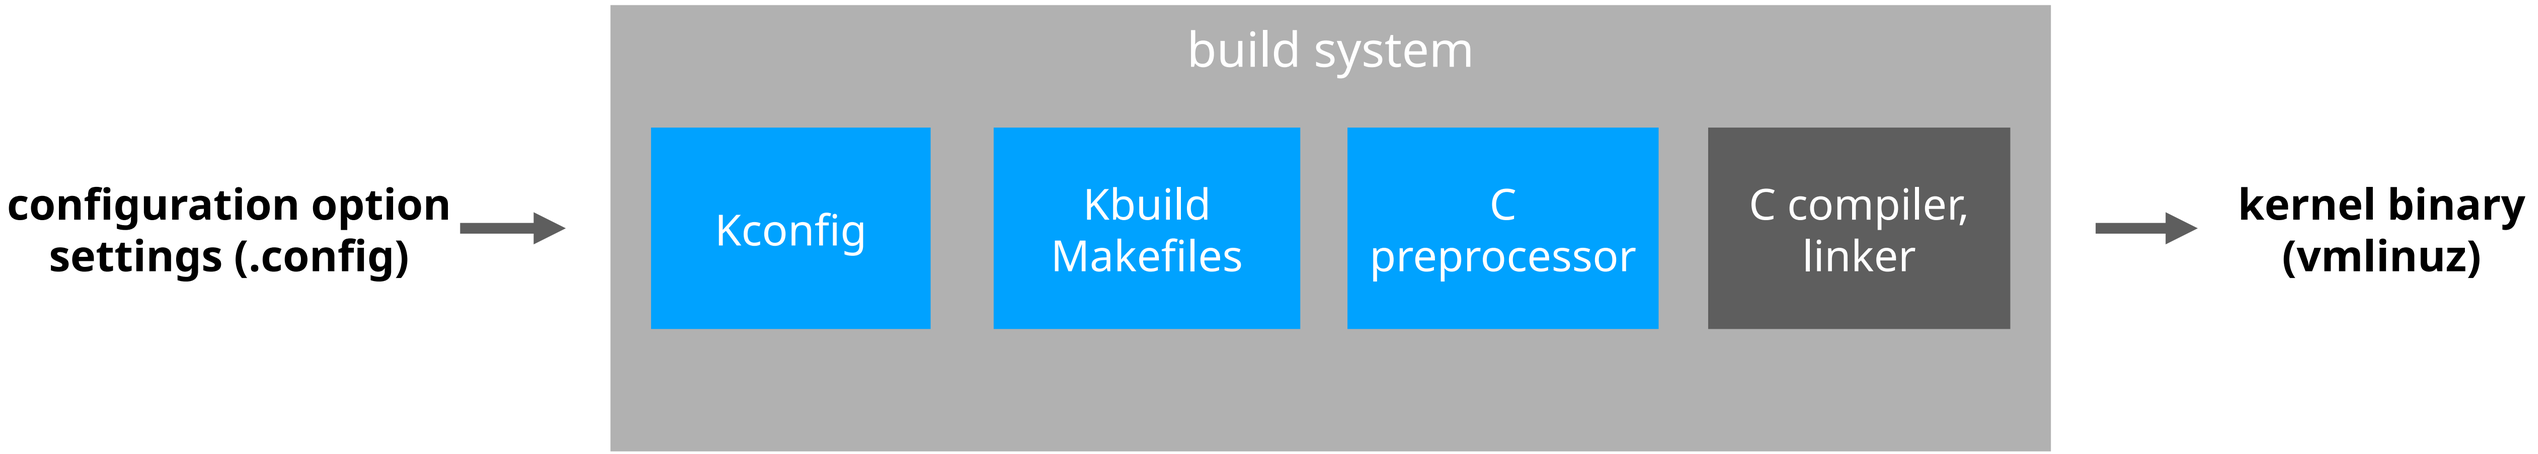

Build System Comprises Several Tools

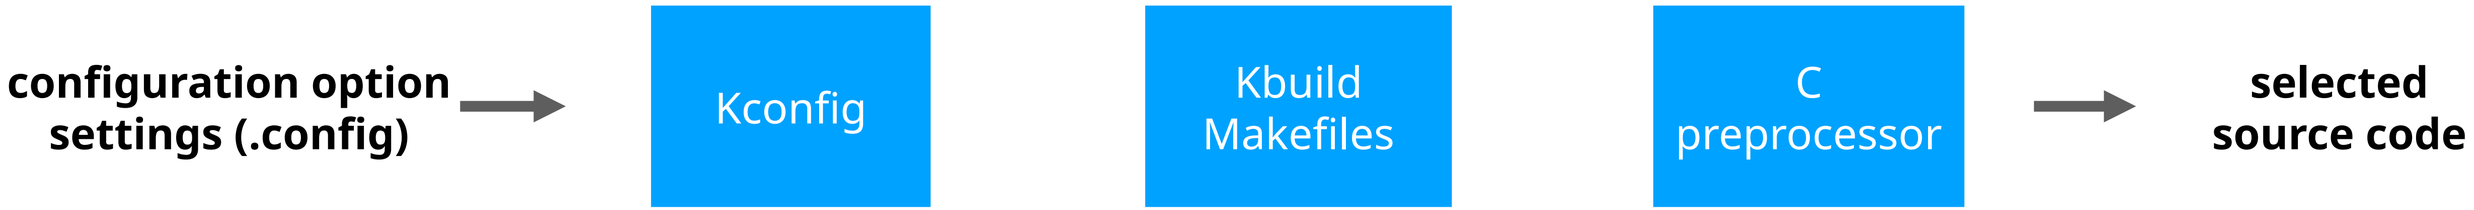

Use Program Analysis on Build Tools

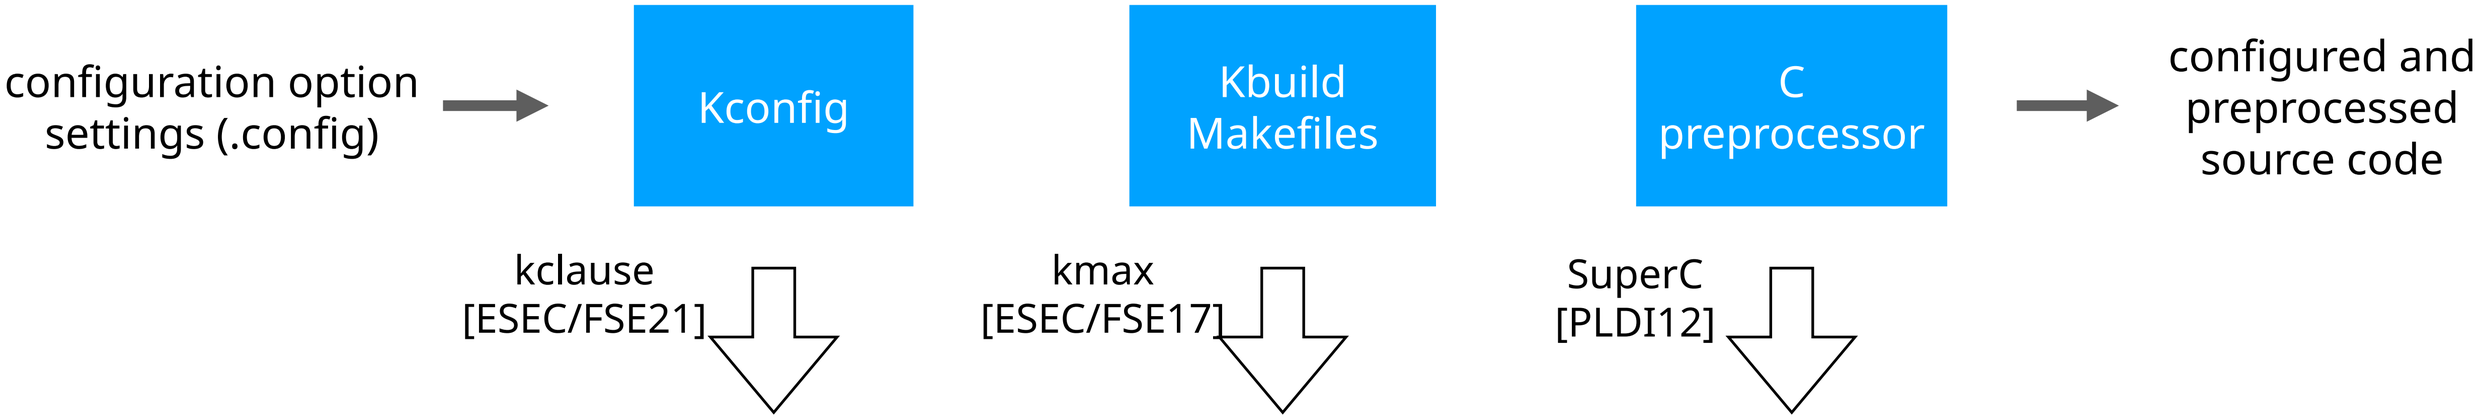

logical constraints on .config files for each file:line

Formally Model Build System Behavior

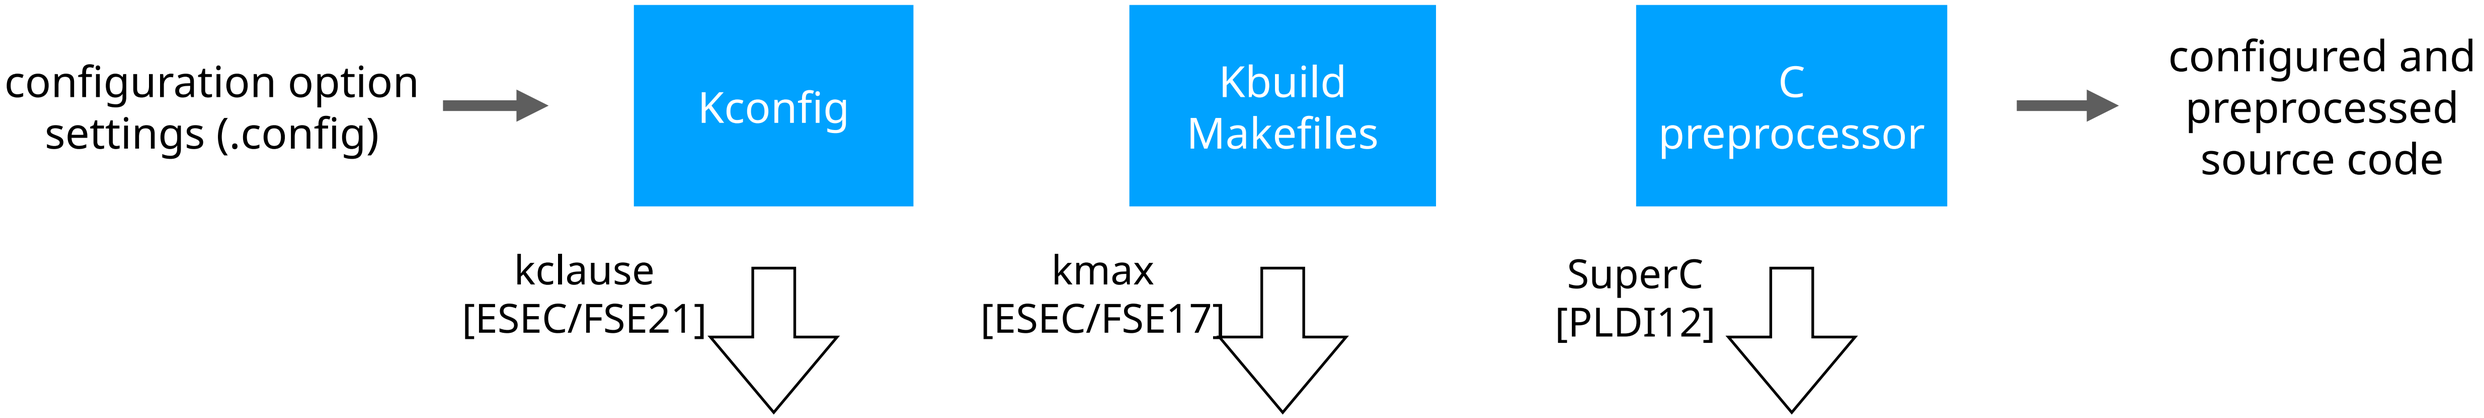

logical constraints on .config files for each file:line

.config file constraints

source code locations

Applications

- kismet [ESEC/FSE21]
  - Automatically find Linux Kconfig bugs
- krepair [FSE24]
  - Automatically change configuration files to cover patched code

Conclusion

- Current testing and analysis focuses on program code
- The software ecosystem broadens the attack surface beyond code
- Misconfigurations are one of the most critical vulnerabilities
- Our approach: formal model and test configurations
- Applications: find configuration bugs, improve testing