

# A Case for DAG Databases

FOSDEM 2023

Nikolai Kondrashov



## Nikolai Kondrashov

    **spbnick**

Sr. Software Engineer, Red Hat, CKI Project  
KernelCI KCIDB developer & maintainer  
Electronics and embedded enthusiast

Born in Russia, living in Finland



# KernelCI Contents

- **Kernel contribution workflow**
- **Kernel testing systems**
- **What are KernelCI and KCIDB**
- **The Kernel's CI problem**
- **Our solution**
- **Our data model**
- **Queries we need**
- **Neo4j experience and ideas**
- **Fallback plan**

# Kernel contribution workflow





# KernelCI Kernel contribution workflow

- **Send patches by email to maillists and maintainers**
- **Get reviews in response**
- **Repeat until everyone is satisfied**
- **Get patches merged into a subtree**
- **Get subtree merged into the mainline**
- ***Maybe* get a report from a testing system at some point**

# Kernel testing systems

# KernelCI Kernel testing systems



...



Results from Linaro's test farm.  
No regressions on arm64, arm, x86\_64, and i386.

## Summary

```
kernel: 5.7.12-rc3
git repo: https://git.kernel.org/pub/scm/linux/kernel/git/stable/linux-stable-rc.git
git branch: linux-5.7.y
git commit: 4822ba8f16c9e980a32747f119a21a9226
git describe: v5.7.12-312-gd3223a8f6d
Test details: https://qa-reports.linaro.org/ikf/linux-stable-rc-5.7-eo/build/v5.7.12-312-gd3223a8f6d
No regressions (compared to build v5.7.12)
No fixes (compared to build v5.7.12)
```

Run 34339 total tests in the following environments and test suites.

## Environments

```
-----
- dragonboard-410c
- hi3798-hikey
- i386
- linux-r2
- juno-r2-compact
- juno-r2-kasan
- nxp-ls2088
- oemv_arm
- oemv_arm64
- oemv_i386
- oemv_x86_64
- x15
- x86
- x86-kasan
```

## Test Suites

### KMSAN: kernel-usb-infileok in hif\_usb\_send

[\[Date Prev\]](#) [\[Date Next\]](#) [\[Thread Prev\]](#) [\[Thread Next\]](#) [\[Date Index\]](#) [\[Thread Index\]](#)

- To: glider@xxxxxxxxx, gregh@xxxxxxxxxxxxxxxxx, Ingrassia@xxxxxxxxxxxxx, linux-kernel@xxxxxxxx
- Subject: KMSAN: kernel-usb-infileok in hif\_usb\_send
- From: syzbot+83a1df1ed4f7e8d8ad@xxxxxxxxxxxxxxxxxxxxxxxxxxxxx
- Date: Tue, 11 Aug 2020 10:06:18 -0700

Hello,

syzbot found the following issue on:

```
HEAD commit: c8056d1 wip: changed copy_from_user where instrumented
git tree: https://github.com/google/dmaman-git-master
console output: https://syzkaller.appspot.com/view_testcase?id=24453-08000
kernel config: https://syzkaller.appspot.com/view_config?id=3a3f089f799921f
download link: https://syzkaller.appspot.com/view_testcase?id=24453-08000
compiler: clang version 10.0.0 (https://github.com/llvm/llvm-project/commit/c24431558f245c8f172c1c72b6ea3918e68f)
userspace arch: i386
```

Unfortunately, I don't have any reproducer for this issue yet.

IMPORTANT: if you fix the issue, please add the following tag to the commit:

Reported-by: syzbot+83a1df1ed4f7e8d8ad@xxxxxxxxxxxxxxxxxxxxxxxxxxxxx

usb 4-1: at91\_ahci: Transferred FW: at91\_ahci\_9271-1.4.0.fw, size: 51008

```
-----
BUG: KMSAN: kernel-usb-infileok in kasan_handle_urb=>20/bx40 mem/kasan/kasan_books.c:207
CPU: 1 PID: 1392 Comm: kworker/1:8 Not tainted 5.8.0-rc5-syzkaller #0
Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 01/01/2011
Memory: events request, firmware_work_func
Call Trace:
```

### stable-rc/linux-4.9.y build: 197 builds: 1 failed, 196 passed, 4 warnings (v4.9.232)

kernel.org bot [hot at kernel.org](#)

Sat Aug 1 08:24:58 UTC 2020

- Previous message: [\[PATCH stable\] MIPS: Loongson: Introduce and use loongson\\_llsc\\_mb\(\)](#)
- Next message: [stable-rc/linux-4.9.y build: 197 builds: 1 failed, 196 passed, 4 warnings \(v4.9.232\)](#)
- Messages sorted by: [\[date\]](#) [\[thread\]](#) [\[subject\]](#) [\[author\]](#)

stable-rc/linux-4.9.y build: 197 builds: 1 failed, 196 passed, 4 warnings (v4.9.232)

Full Build summary: <https://kernel.org/build/stable-rc/branch/linux-4.9.y/kernel/v4.9.232/>

```
Tree: stable-rc
Branch: linux-4.9.y
Git describe: v4.9.232
Git commit: 80654129e59293b2a7b80e1c58d91abe13b
Git URL: https://git.kernel.org/pub/scm/linux/kernel/git/stable/linux-stable-rc.git
Mail: 6 unique architectures
```

Build Failure Detected:

```
mips:
  32r2el_defconfig: (gcc-8) FAIL
```

Warnings Detected:

src:

### ✓ PASS: Test report for kernel 5.7.12-3ff3d4f.cki (stable-queue)

CKI Project [cki-project at redhat.com](#)

Tue Aug 4 12:35:15 UTC 2020

- Previous message: [fixing 4.14-stable's broken DM cache writethrough support \[was: Re: \(resend\) PATCH v3: \(linux-4.14.y\) dm cache: submit writethrough writes in parallel to origin and cache\]](#)
- Next message: [stable-rc/linux-4.9.y build: 190 builds: 1 failed, 189 passed, 2 errors, 13 warnings \(v4.4.232-30-g52247eb98be\)](#)
- Messages sorted by: [\[date\]](#) [\[thread\]](#) [\[subject\]](#) [\[author\]](#)

Hello,

We ran automated tests on a recent commit from this kernel tree:

```
Kernel repo: https://git.kernel.org/pub/scm/linux/kernel/git/stable/linux-stable-rc.git
Commit: 3ff3d4f3856 - x86/IBT59: Use prints_deferred() to prevent deadlock
```

The results of these automated tests are provided below.

```
Overall result: PASSED
Merge: OK
Compile: OK
Tests: OK
```

All kernel binaries, config files, and logs are available for download here:

<https://cki-artifacts.s3.us-east-2.amazonaws.com/index.html?prefix=datawarehouse/2020/08/03/61139/>

Please reply to this email if you have any questions about the tests that we ran or if you have any suggestions on how to make future tests more effective.



### Re: [PATCH V4] hwmon: add fan/pwm driver for corsair

[\[Date Prev\]](#) [\[Date Next\]](#) [\[Thread Prev\]](#) [\[Thread Next\]](#) [\[Date Index\]](#) [\[Thread Index\]](#)

- To: Jaap aarts <[jaap.aarts1@xxxxxxxxx](mailto:jaap.aarts1@xxxxxxxxx)>, Jean Delvare <[delvare@xxxxxxxxx](mailto:delvare@xxxxxxxxx)>, Guenter Roeck <[linux@xxxxxxxxx](mailto:linux@xxxxxxxxx)>, Barnabás Pócze <[pobrn@xxxxxxxxxxxxxx](mailto:pobrn@xxxxxxxxxxxxxx)>
- Subject: Re: [PATCH V4] hwmon: add fan/pwm driver for corsair h100i platinum
- From: Kernel test robot <[clk@xxxxxxxxx](mailto:clk@xxxxxxxxx)>
- Date: Sun, 16 Aug 2020 17:34:40 +0800
- Cc: [kbuild-all@xxxxxxxxxxxxx](mailto:kbuild-all@xxxxxxxxxxxxx), jaap.aarts <[jaap.aarts1@xxxxxxxxx](mailto:jaap.aarts1@xxxxxxxxx)>
- In-reply-to: <[20200815211617.86565-1-jaap.aarts1@gmail.com](mailto:20200815211617.86565-1-jaap.aarts1@gmail.com)>
- References: <[20200815211617.86565-1-jaap.aarts1@gmail.com](mailto:20200815211617.86565-1-jaap.aarts1@gmail.com)>
- User-agent: Mutt/1.10.1 (2018-07-13)

Hi Jaap,

Thank you for the patch! Perhaps something to improve:

```
[auto build test WARNING on hwmon/hwmon-next]
[auto build test WARNING on v5.8 next-20200814]
[if your patch is applied to the wrong git tree, kindly drop us a note.
And when submitting patch, we suggest to use "--base" as documented in
https://git-scm.com/docs/git-format-patch]
```

```
url: https://github.com/8day-cy/linux/commit/jaap.aarts/hwmon-add-fan-pwm-driver-for-corsair-h100i-plat/linux/20200816-89529
base: https://git.kernel.org/pub/scm/linux/kernel/git/groeck/linux-staging.git/hwmon-next
config: nios2-randconfig-c009-20200816 (attached as .config)
compiler: nios2-linux-gcc (GCC) 9.3.0
```



# KernelCI Kernel testing systems


**LKFT** About LKFT 101 Boards Branches Tests Logs

## Test Results

linux-stable-rc-4.4-oe	142 hours ago	17799 passed	879 failed	3052 skipped	<a href="#">View Builds</a>
linux-stable-rc-4.9-oe	142 hours ago	31410 passed	927 failed	3239 skipped	<a href="#">View Builds</a>
linux-stable-rc-4.14-oe	123 hours ago	32787 passed	934 failed	3239 skipped	<a href="#">View Builds</a>
linux-stable-rc-4.19-oe	123 hours ago	42289 passed	924 failed	4239 skipped	<a href="#">View Builds</a>
linux-stable-rc-5.4-oe	135 hours ago	41554 passed	933 failed	4248 skipped	<a href="#">View Builds</a>
linux-stable-rc-5.7-oe	130 hours ago	33009 passed	1352 failed	4137 skipped	<a href="#">View Builds</a>
linux-mainline-oe	13 hours ago	13193 passed	941 failed	3441 skipped	<a href="#">View Builds</a>
linux-next-oe	7 hours ago	14015 passed	928 failed	3222 skipped	<a href="#">View Builds</a>

## LKFT 2.0 Build Status

Our second generation build and continuous integration uses GitLab pipelines and TuxBot.



## LKFT Build Status

The LKFT build uses OpenEmbedded to build a userspace image, along with the kernel, for each board and branch combination under test.



**syzbot** Linux 4.14 [sign-in](#) | [mailing list](#) | [source](#) | [docs](#)

[Open \[308\]](#) [Fixed \[100\]](#) [Invalid \[215\]](#)

Name	Active	Uptime	Corpus	Coverage	Crashes	Execs	Instances:		Kernel build		syzkaller build	
							Commit	Freshness	Status	Commit	Freshness	Status
ci2-linux-4.14	now	2621h	53390	1099200	237	1354254	14858326	104	5661352	3081h	fail1us	

open (508):

Title	Repro	Bisected	Count	Last	Reported
<a href="#">WARNING in rtl_dellink</a>	C		1	4153h	4152h
<a href="#">possible deadlock in proc_pid_personality</a>			1	21155h	21154h
<a href="#">syzkaller: make target failed: failed to run ["make" "tar...]</a>			3	1491h	2489h
<a href="#">INFO: task hung in _input_unregister_device</a>			1	3481h	2482h
<a href="#">unregister_netdevice: waiting for DPU to become free (2)</a>			1	3422h	3422h
<a href="#">INFO: task hung in do_task_dead</a>			1	5415h	5415h
<a href="#">INFO: trying to register non-static key in done (2)</a>			1	5422h	5422h
<a href="#">INFO: trying to register non-static key in shhd_char_rei...</a>	C		1	6622h	6622h
<a href="#">possible deadlock in proc_pid_syscall</a>			3	1452h	2424h
<a href="#">possible deadlock in proc_pid_stack</a>	C		3	5415h	7415h
<a href="#">INFO: trying to register non-static key in l2cap_chan_cl...</a>	592		1	7420h	2420h
<a href="#">BUG: corrupted list in bt_accept_unlink</a>			1	8481h	8480h
<a href="#">KASAN: use-after-free Read in hci_get_auth_info</a>	592		1	8494h	8492h
<a href="#">inconsistent lock state in sock_timeout</a>			1	8412h	8411h
<a href="#">INFO: rcu detected stall in tv_timer_handler</a>			1	9410h	5410h
<a href="#">general protection fault in memsw work</a>	592		1	9418h	5418h

Home Jobs Builds Tests SoCs Boots KCIDB Info

## Available Test Results

The results shown here cover the last 14 days of available data starting from **Mon, 17 Aug 2020** (time is UTC based).

25 Tests per page  Filter the

Tree	Branch	Kernel	Test Plan	Test Results
mainline	master	v5.8-13502-gb07175d041ba	usb	54 34 4 16
rt-stable	v4.14-rt	v4.14.193-r192	baseline-rfcs	1,589 1,502 22 59
mainline	master	v5.8-13651-g84835b118ed	baseline	1,097 774 20 303
rt-stable	v4.14-rt	v4.14.193-r192	baseline	1,609 1,499 24 80
mainline	master	v5.8-13502-gb07175d041ba	baseline	1,588 1,482 26 73
mainline	master	v5.8-13502-gb07175d041ba	sleep	34 12 2 20
mainline	master	v5.8-13414-g7175e9d2792b	baseline	4,674 4,528 94 77

**CI** **2**

Go to pipeline #

**Build**

All / Revision / Build

Build status: **Passed**

Build ID: rehat956956

Valid: Yes

Origin: rehat

Start Time: Aug 17, 2020, 2:26 a.m.

Architecture: ppc64le

Command: make -j30 INSTALL\_MOD\_STRIP=1 largz-pkg

Config Name: fedora

Config URL: https://3.us-east-2.amazonaws.com/ci-artifacts/datawarehouse/2020/08/17/612139/build\_ppc64le\_rehat956956/config

Log URL: https://3.us-east-2.amazonaws.com/ci-artifacts/datawarehouse/2020/08/17/612139/build\_ppc64le\_rehat956956/build.log

Output Files

MD1.01.0rg Search this list Manage this list Sign In

**LKP** Linux Kernel Performance

Start a new thread Manage subscription

ACTIVITY SUMMARY

RECENTLY ACTIVE DISCUSSIONS

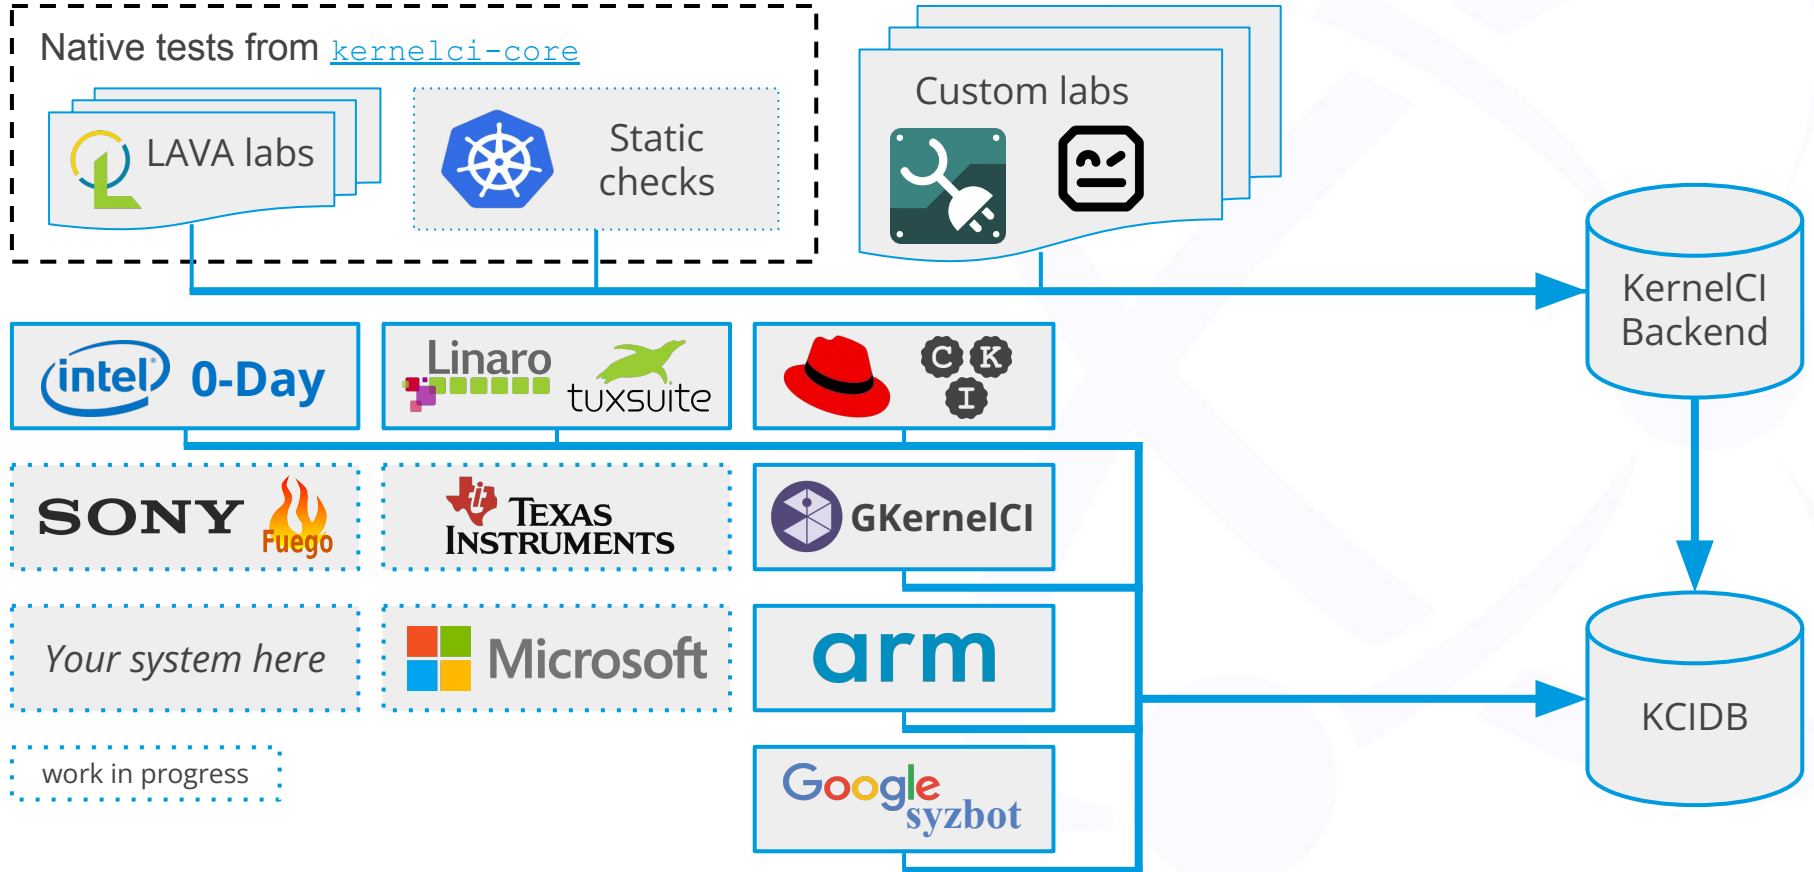
MOST ACTIVE POWERHOURS

MOST POPULAR DISCUSSIONS

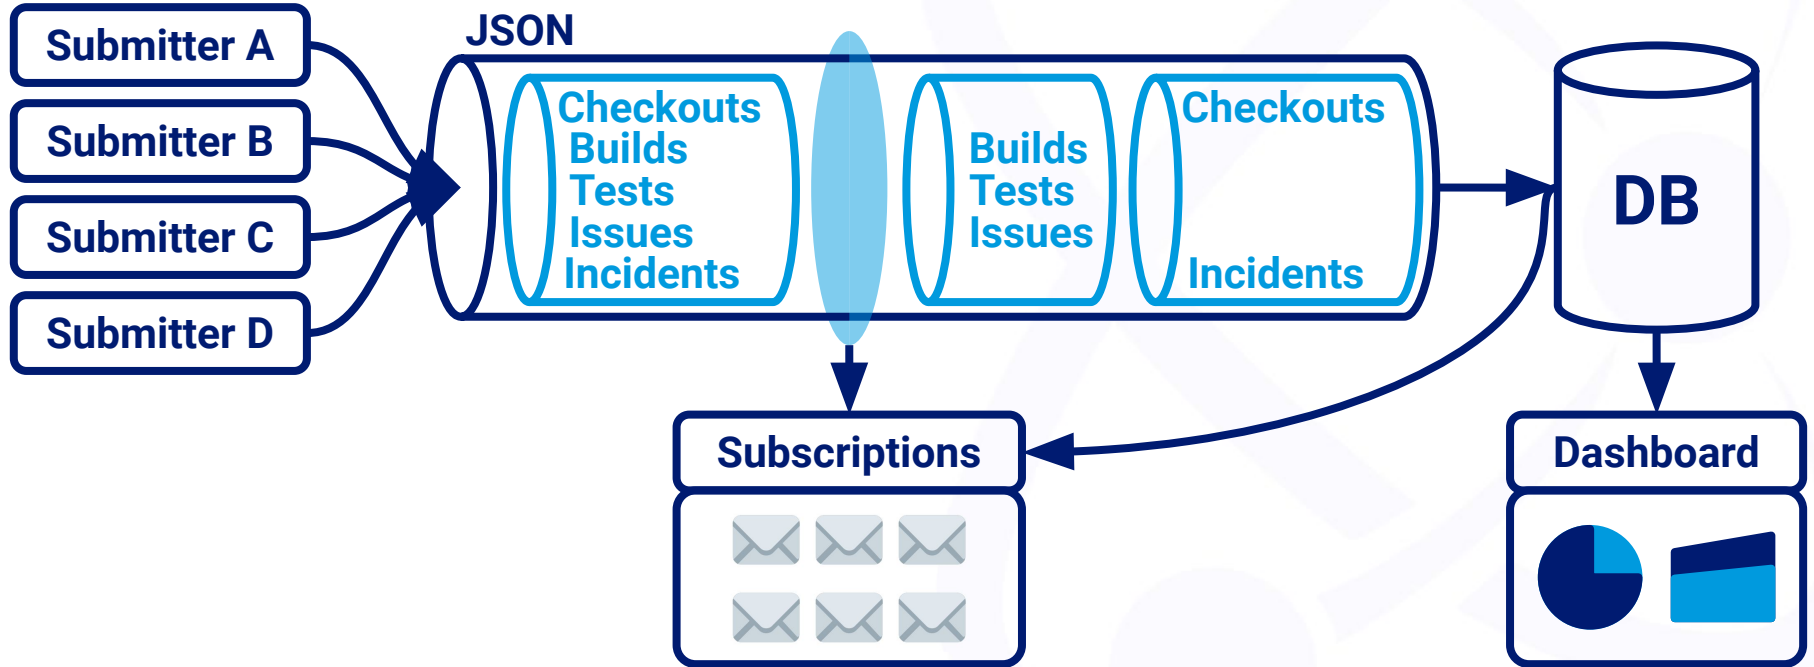
MOST ACTIVE DISCUSSIONS

# What are KernelCI and KCIDB

# KernelCI What are KernelCI and KCIDB



# KernelCI What are KernelCI and KCIDB






# KernelCI What are KernelCI and KCIDB



# KernelCI What are KernelCI and KCIDB



# KernelCI What are KernelCI and KCIDB

Home -

Welcome to Linux Kernel CI report database (KCIDB) dashboard. Here you can find testing reports aggregated from multiple CI systems ("origins"). This is a work in progress, and we welcome contributions and participation. See the source code for the [collection system](#), and for [this dashboard](#). Write to [kernelci@proton.fr](mailto:kernelci@proton.fr) if you want to start submitting results from your CI system, or if you want to receive automatic notifications of arriving results.

Repositories

URL	Branches	Revisions
https://gitlab.com/Linaro/ltfs/mimosa/stable/linux-stable-rc-queue		539
https://git.kernel.org/pub/scm/linux/kernel/git/stable/linux-stable-rc.git	queue/5.4, queue/5.14, queue/5.10 (v10)	479
https://git.kernel.org/pub/scm/linux/kernel/git/davem/net-next.git	master	347
https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git	master	448
https://android.googlesource.com/kernel/common	android/13.5.10, android/12.5.4-1b, android/12.5.4 (v16)	154
https://gitlab.com/Linaro/ltfs/mimosa/stable/linux-stable-rc		140

Revisions

Origin	Repository	Branch	Name	Hash	Discovery time	Status	Builds	Tests
kernelci	https://git.kernel.org/pub/scm/linux/kernel/git/stable/linux-stable-rc.git	linux-5.10-y	v5.10.67-1775-gb66c21646809	b66c21646809	2021-09-21 09:11:43.151+00	✓	✓	FAIL
.....	https://git.kernel.org/pub/scm/linux/kernel/git/stable/linux-stable-rc.git	queue/5.10	v5.10.67-124-g409ba09baec2	4409ba09baec2	2021-09-21 07:17:03+00	✓	✓	FAIL
fasuite	https://gitlab.com/Linaro/ltfs/mimosa/stable/linux-stable-rc-queue	-	v4.9.282-175-g41cfaebbd653	41cfaebbd653	2021-09-21 07:16:56+00	✓	✓	✓
fasuite	https://gitlab.com/Linaro/ltfs/mimosa/stable/linux-stable-rc-queue	-	v5.4.147-269-g17452f118854	17452f118854	2021-09-21 07:15:49+00	✓	✓	✓
fasuite	https://gitlab.com/Linaro/ltfs/mimosa/stable/linux-stable-rc-queue	-	v4.4.283-133-ga6ff9e9006	a6ff9e9006	2021-09-21 07:15:38+00	✓	✓	✓
fasuite	https://gitlab.com/Linaro/ltfs/mimosa/stable/linux-stable-rc-queue	-	v5.19.206-293-gc70c39e8ed9f	c70c39e8ed9f	2021-09-21 07:15:35+00	✓	✓	✓
fasuite	https://gitlab.com/Linaro/ltfs/mimosa/stable/linux-stable-rc-queue	-	v4.14.244-217-ga6ff9e9006	a6ff9e9006	2021-09-21 07:15:33+00	✓	✓	✓
fasuite	https://gitlab.com/Linaro/ltfs/mimosa/stable/linux-stable-rc-queue	-	v5.14.6-170-g6eac71c1020a	6eac71c1020a	2021-09-21 07:14:33+00	✓	✓	✓
fasuite	https://gitlab.com/Linaro/ltfs/mimosa/stable/linux-next	-	next-20210921	83fa38574812	2021-09-21 06:12:25+00	✓	✓	✓

Build status

Test status

Reports

Builds per architecture

Repository -

Home

Build status

Test status

Revisions

Origin	Branch	Name	Hash	Discovery time	Status	Builds	Tests
kernelci	queue/6.14	v4.14.244-217-ga6ff9e9006	a6ff9e9006	2021-09-21 10:28:35.278+00	✓	✓	FAIL
kernelci	queue/5.10	v5.10.67-124-g409ba09baec2	4409ba09baec2	2021-09-21 10:03:57.482+00	✓	✓	FAIL
kernelci	linux-5.13-y	v5.13.19	7feb8060c0db	2021-09-21 09:47:49.489+00	✓	✓	FAIL
kernelci	linux-5.10-y	v5.10.67-125-gb66c21646809	b66c21646809	2021-09-21 09:11:43.151+00	✓	✓	FAIL
kernelci	linux-4.19-y	v5.19.206-294-ga707776805f	9a707776805f	2021-09-21 08:53:19.501+00	✓	✓	FAIL
kernelci	linux-5.4-y	v5.4.147-261-g9f54072b6a3	9f54072b6a3	2021-09-21 08:36:16.638+00	✓	✓	FAIL
kernelci	linux-4.14-y	v4.14.244-218-g7c39e8ed9f	7c39e8ed9f	2021-09-21 08:31:30.890+00	✓	✓	FAIL
kernelci	linux-4.9-y	v4.9.282-176-g5020b6ff9f	f5020b6ff9f	2021-09-21 08:09:46.564+00	✓	✓	FAIL
kernelci	linux-4.4-y	v4.4.283-134-g3a54c909bd	3a54c909bd	2021-09-21 07:52:38.934+00	✓	✓	FAIL

Reports

Builds per architecture



Branch -

Repository URL: <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git> • Name: master • Origin: All • Build architecture: All • Build configuration: All

Home → Repository

Origins	Repository URL	Name
arm, kernelci, redhat, syzbot, tuxsuite	<a href="https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git">https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git</a>	master

Revisions

Origin	Name	Hash	Discovery time	Status	Builds	Tests
tuxsuite	v5.15-rc2-35-g05050856a77	e9b6578414c0	2021-09-21 01:00:47+00	✓	X	
arm	v5.15-rc2-2-g8f7189236e	e8f7189236e	2021-09-20 21:46:17+00	✓	✓	FAIL
tuxsuite	v5.15-rc2-2-g8f7189236e	e8f7189236e	2021-09-20 18:20:57+00	✓	X	
...	...	...	...	...	...	...
...	v5.15-rc2	e4e737bb5c17	2021-09-20 00:41:11:722+00	✓	X	FAIL
redhat	2062142f27a0	2062142f27a0	2021-09-19 20:41:12:717+00	✓	X	PASS
redhat	bc1abb9e55ce	bc1abb9e55ce	2021-09-19 19:51:28:073+00	✓	X	PASS
...	...	...	...	...	...	...
...	v5.15-rc1-239-g645016caa4b8	645016caa4b8	2021-09-18 22:06:14:758+00	✓	X	FAIL
redhat	938f913bd91	938f913bd91	2021-09-18 19:56:12:056+00	✓	X	FAIL
arm	v5.15-rc1-221-g43570536611	43570536611	2021-09-18 13:49:13+00	✓	✓	FAIL

Build status, Test status, Reports, Builds per architecture

Revision -

Origin: All • Build architecture: All • Build configuration: All

Home

Repo: <https://gitlab.com/leandro/mirrors/torvalds/linux-mainline> • Branch: master • Commit name: v5.15-rc2

Log: Status ✓ Builds X Tests FAIL

Hash: e4e737bb5c170df6135a127739a9e6148ee3da82

Description: 2021-09-20 00:41:17.381+00

Build status, Top 10 architecture build failures, Test status, Top 5 configuration build failures, Top 5 compiler build failures

Origin	Architecture	Compiler	Configuration	Start time	Duration	Description	Status	Tests
redhat	aarch64	clang version 13.0.0 (Fedora 13.0.0-rc1.1.fc35)	fedora	2021-09-20 00:54:01.94896+00	00:35:23.000		✓	
redhat	s390x	s390x-linux-gnu-gcc (GCC) 11.2.1 20210728 (Red H...	fedora	2021-09-20 00:54:03.116458+00	00:10:08.000		✓	
redhat	s390x	clang version 13.0.0 (Fedora 13.0.0-rc1.1.fc35)	fedora	2021-09-20 00:54:04.279798+00	00:12:45.000		✓	
redhat	x86_64	gcc (GCC) 11.2.1 20210728 (Red Hat 11.2.1-1)	fedora	2021-09-20 00:54:08.688925+00	00:23:37.000		✓	FAIL
redhat	ppc64le	powerpc64le-linux-gnu-gcc (GCC) 11.2.1 20210728 ...	fedora	2021-09-20 00:54:12.072940+00	00:24:02.000		✓	
redhat	aarch64	aarch64-linux-gnu-gcc (GCC) 11.2.1 20210728 (Red...	fedora	2021-09-20 00:54:13.559512+00	00:30:15.000		✓	SKIP
redhat	x86_64	clang version 13.0.0 (Fedora 13.0.0-rc1.1.fc35)	fedora	2021-09-20 00:54:16.220894+00	00:28:42.000		X	
redhat	ppc64le	clang version 13.0.0 (Fedora 13.0.0-rc1.1.fc35)	fedora	2021-09-20 00:54:19.067626+00	00:28:42.000		✓	

Patches



Below is the summary of results Kernel CI database has recorded for this revision so far. See complete and up-to-date report at:

[https://kcidb.kernelci.org/d/revision/revision?orgId=1&var-git\\_commit\\_hash=80808293888188c3923f5bd8a69370dae25ed14e5&var-patchset\\_hash=](https://kcidb.kernelci.org/d/revision/revision?orgId=1&var-git_commit_hash=80808293888188c3923f5bd8a69370dae25ed14e5&var-patchset_hash=)

## OVERVIEW

Builds: ✗ FAIL  
Tests: ✗ FAIL

## REVISION

Status ✔ PASS  
Commit  
name: v5.16-rc7-108-g8080829388818  
hash: 80808293888188c3923f5bd8a69370dae25ed14e5  
Checked out from  
<https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git> master  
<https://gitlab.com/Linaro/lkft/mirrors/torvalds/linux-mainline>  
By  
kernelci, redhat, syzbot, tuxsuite

## BUILDS

Status ✗ 56 ✔ 738  
Architectures  
arm ✗ 26 ✔ 347  
riscv ✗ 8 ✔ 21  
mips ✗ 7 ✔ 86  
powerpc ✗ 5 ✔ 47  
x86\_64 ✗ 3 ✔ 55  
arm64 ✗ 3 ✔ 48  
i386 ✗ 2 ✔ 42  
s390 ✗ 1 ✔ 17  
ppc64le ✗ 1 ✔ 1  
sh — ✔ 24  
arc — ✔ 22  
parisc — ✔ 12  
sparc — ✔ 12  
aarch64 — ✔ 2  
s390x — ✔ 2  
Failures  
✗ 4 riscv allnoconfig  
✗ 2 arm allmodconfig "v5.16-rc7-108-g8080829388818"  
✗ 2 x86\_64 allmodconfig "v5.16-rc7-108-g8080829388818"  
✗ 2 i386 allmodconfig "v5.16-rc7-108-g8080829388818"  
✗ 2 powerpc maple\_defconfig  
...  
By  
kernelci, redhat, syzbot, tuxsuite

## TESTS

Status  
✗ 23 ✔ 53 ✗ 5 🔄 6 📦 1  
Failures  
✗ baseline-qemu-docker  
✗ baseline  
✗ baseline-nfs  
✗ ltp-mm  
✗ cros-ec  
...  
By  
kernelci, redhat, syzbot

See complete and up-to-date report at:

[https://kcidb.kernelci.org/d/revision/revision?orgId=1&var-git\\_commit\\_hash=80808293888188c3923f5bd8a69370dae25ed14e5&var-patchset\\_hash=](https://kcidb.kernelci.org/d/revision/revision?orgId=1&var-git_commit_hash=80808293888188c3923f5bd8a69370dae25ed14e5&var-patchset_hash=)

## LEGEND

- ✗ ERROR - Aborted. Test, tested code, or both might be faulty.
- ✗ FAIL - Failed. Tested code is likely faulty.
- ✔ PASS - Passed. Tested code is likely correct.
- 📦 DONE - Finished. Status of tested code is unknown.
- 🔄 SKIP - Skipped. Planned, but didn't execute.
- ? UNKNOWN - In progress, or status unknown.
- 🚫 WAIVED - Waived result. Test is too new or shows known failures.
- BLANK - No data, zero.

# The Kernel's CI problem



# KernelCI The Kernel's CI problem

- **Hardware is a natural scarcity**
- **Test results depend on hardware, timing, etc.**
- **A broken change might pass and a healthy change might fail**
- **Hard to remove noise from the results**

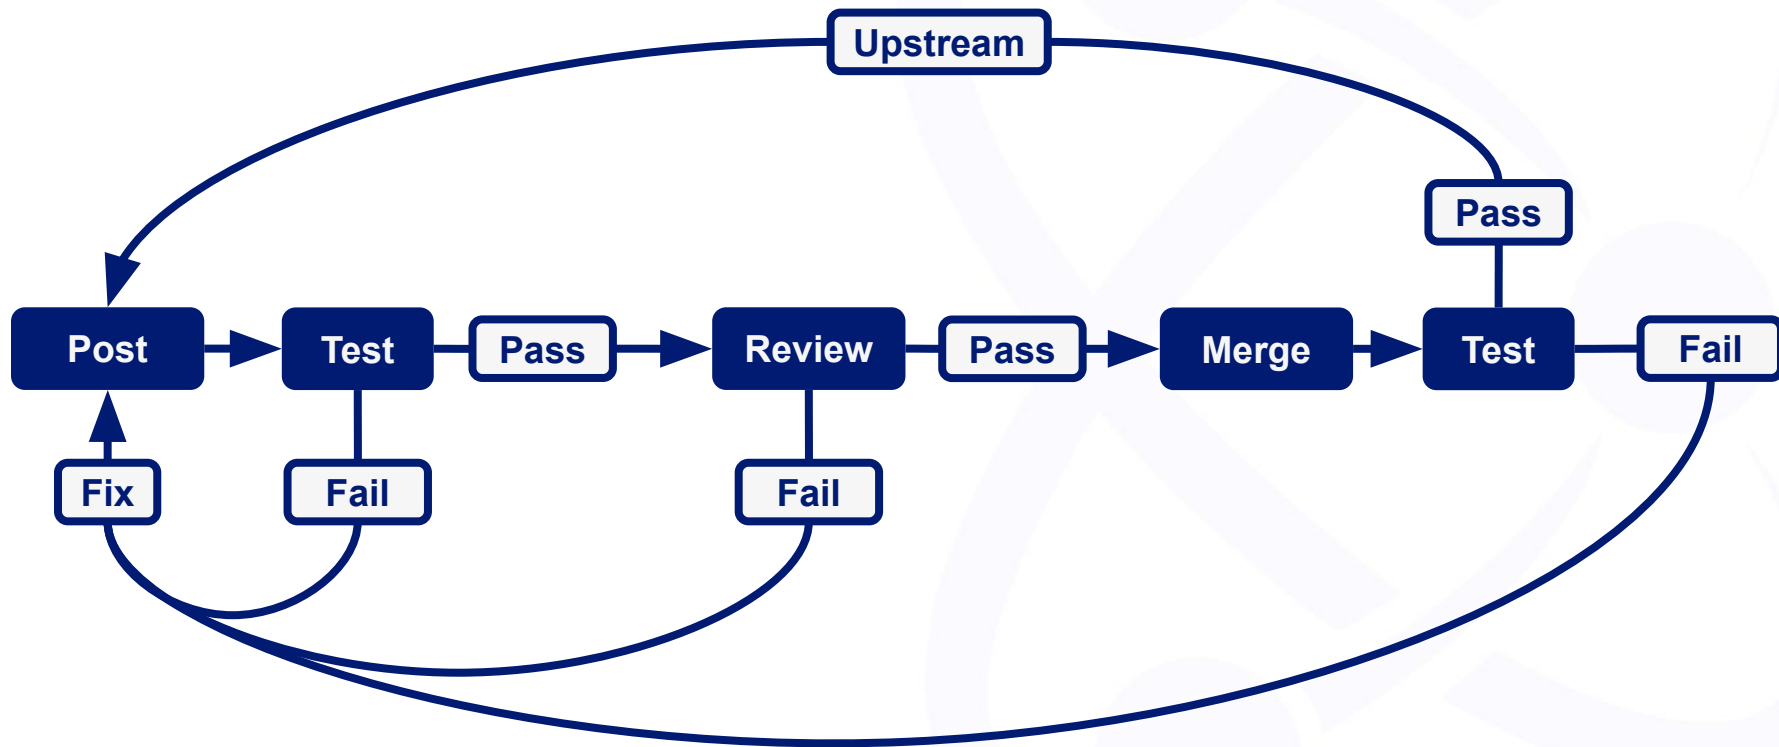


# KernelCI The Kernel's CI problem

- **Hard to investigate and reproduce**
- **Developers are skeptical of reports**
- **Reports undergo human reviews before sending**
- **Nobody stops development to fix CI**
- **Feedback loop is too long**

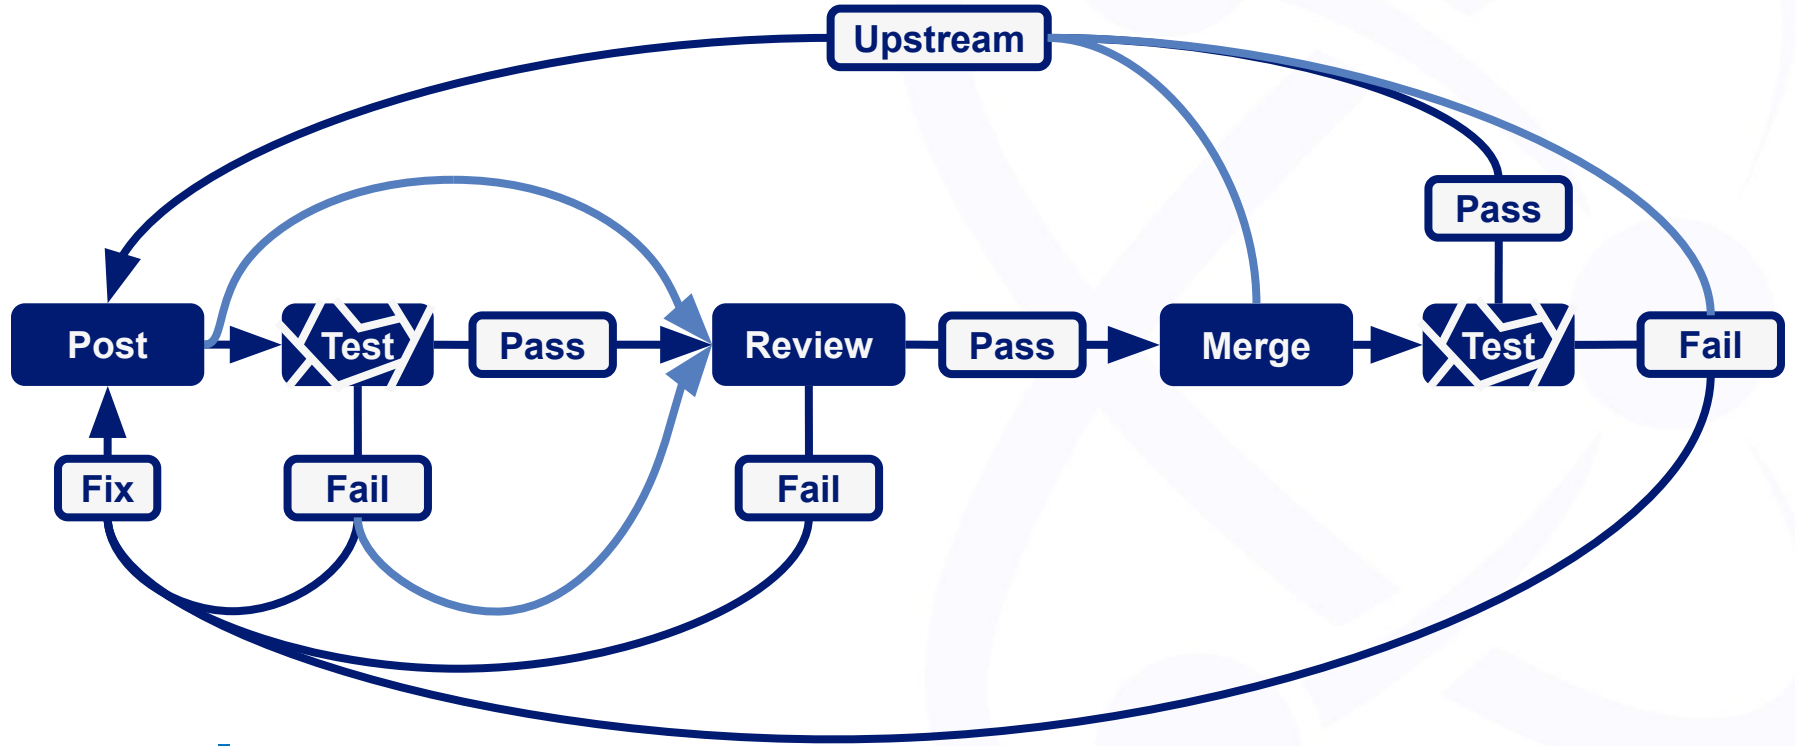


# KernelCI The Kernel's CI problem



**Ideal**



# KernelCI The Kernel's CI problem



**Actual**


# Our solution













# KernelCI Our solution

The screenshot shows a web interface for a bug report. On the left is a blue sidebar with navigation links: Builds, Issues, CONFIDENCE, Tests, and Hosts. The main content area has a header with the bug title and a 'Kernel Bug' label. Below the header, there are fields for 'Kind: Kernel bug', 'Resolved: No', 'Policy: public (read: None, write: None)', 'Ticket URL', and 'Origin Tree: stable'. A 'Regexes' table follows, with columns for Text, File Name, and Test Name. The 'Affected hosts' section is currently empty. The 'Affected Checkouts' table lists three entries with columns for ID, Tree, Branch, and Start Time. Three large, stylized letters 'C', 'K', and 'I' are overlaid on the right side of the page.

**Kernel Bug** Storage block - storage fio numa: Performance comparison: min:40535 \* 1.15 < max:81259 stable Created by:

**Kind:** Kernel bug **Resolved:** No **Policy:** public (read: None, write: None)

**Ticket URL:** [https://bugzilla.redhat.com/show\\_bug.cgi?id=2032094](https://bugzilla.redhat.com/show_bug.cgi?id=2032094)

**Origin Tree:** stable

### Regexes

Text	File Name	Test Name
<code>\[s+FAIL\s+\] :: Performance comparison: min:\d+ \* 1\.\d+ &lt; max:\d+</code>	taskout\.log	Storage block - storage fio

### Affected hosts

>

### Affected Checkouts

ID	Tree	Branch	Start Time
redhat:449002104	mainline.kernel	master	Jan. 16, 2022, 8:26 a.m.
redhat:448605532	kernel-rt-rhel9		Jan. 15, 2022, 12:54 a.m.
redhat:448600281	rhel9		Jan. 15, 2022, 12:33 a.m.

**C** **K** **I**

# KernelCI Our solution

Edit filter

Short description CFL: igt@kms\_flip@2x-flip-vs-expired-vblank\* - fail - Failed assertion: timercmp(&reply.ts, &o->flip\_state.last\_ts, ==), Last errno: 25

Select the tags

Ignored

Empty list

Filter

>> >

Leave empty to select all tags

Selected

Showing all 1

Filter

< <<

DRM-TIP

Select the machines

Ignored

Showing all 186

Filter

>> >

Tag: APL  
Tag: BDW  
Tag: BLB  
Tag: BSW  
Tag: BWR  
Tag: BXT  
Tag: BYT  
Tag: CHAMELI  
Tag: CNL  
Tag: CTG

Leave empty to select all machines

Selected

Showing all 1

Filter

< <<

Tag: CFL

Select the tests

Ignored

Showing all 2963

Filter

>> >

IGT: igt@kms\_frontbuffer\_tracking@fbc-2p-primscrn-pri-indfb-draw  
IGT: igt@kms\_chamelium@hdmi-hpd  
IGT: igt@kms\_flip@dprms-off-confusion  
IGT: igt@gem\_persistent\_relocs@forked-faulting-reloc  
IGT: igt@kms\_chv\_cursor\_fail@pipe-b-256x256-top-edge  
IGT: igt@kms\_cursor\_legacy@basic-flip-before-cursor-legacy  
IGT: igt@kms\_atomic\_transition@5x-modeset-transitions-fencing  
IGT: igt@gem\_persistent\_relocs@forked-thrashing  
IGT: igt@kms\_draw\_crc@draw-method-xrgb8888-render-xtiled  
IGT: igt@kms\_pipe\_crc\_basic@hang-read-crc-pipe-b

Leave empty to select all tests

Selected

Showing all 1

Filter

< <<

IGT: igt@kms\_flip@2x-flip-vs-expired-vblank

Select the statuses

Ignored

Showing all 14

Filter

>> >

IGT: pass  
IGT: warn  
IGT: dmesg-wa  
IGT: dmesg-fail  
IGT: incomplete  
IGT: notrun  
IGT: timeout  
Piglit: warn  
IGT: crash  
Piglit: crash

Leave empty to select all statuses

Selected

Showing all 1

Filter

< <<


IGT: fail

Stdout regex Regular expression that needs to be matched on the test's standard output. Leave empty to ignore.

Stderr regex Failed assertion: timercmp(&reply.ts, &o->flip\_state.last\_ts, ==)in.\*Last errno: 25, Inappropriate ioctl for device

Dmesg regex Regular expression that needs to be matched on the kernel logs. Leave empty to ignore.

Matches 0/564 unknown failures: 0 tag(s), 0 machine(s), 0 machine tag(s), 0 test(s), and 0 status(es) - Apply filter





# Our data model




# KernelCI Our data model

Checkout


Checkout

Checkout

# KernelCI Our data model

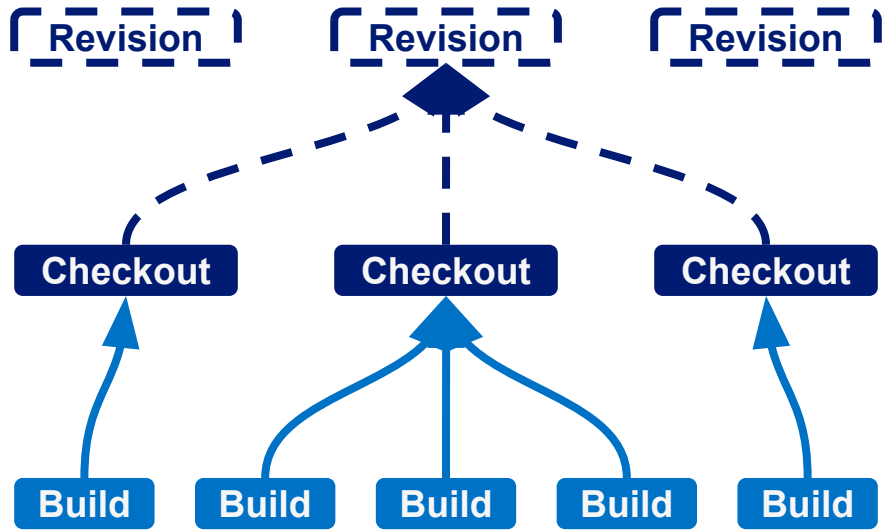


# KernelCI Our data model



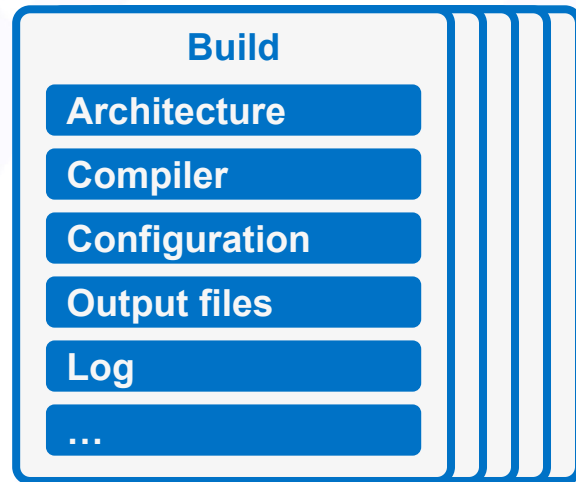
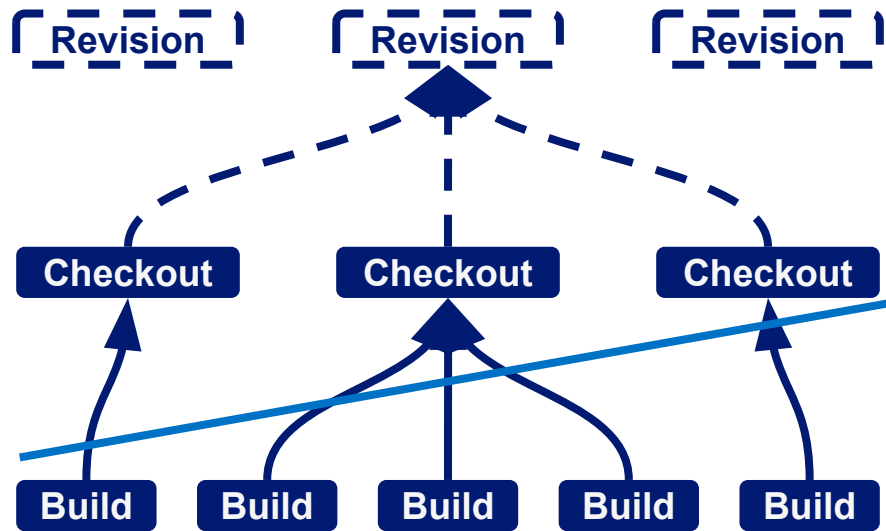


# KernelCI Our data model

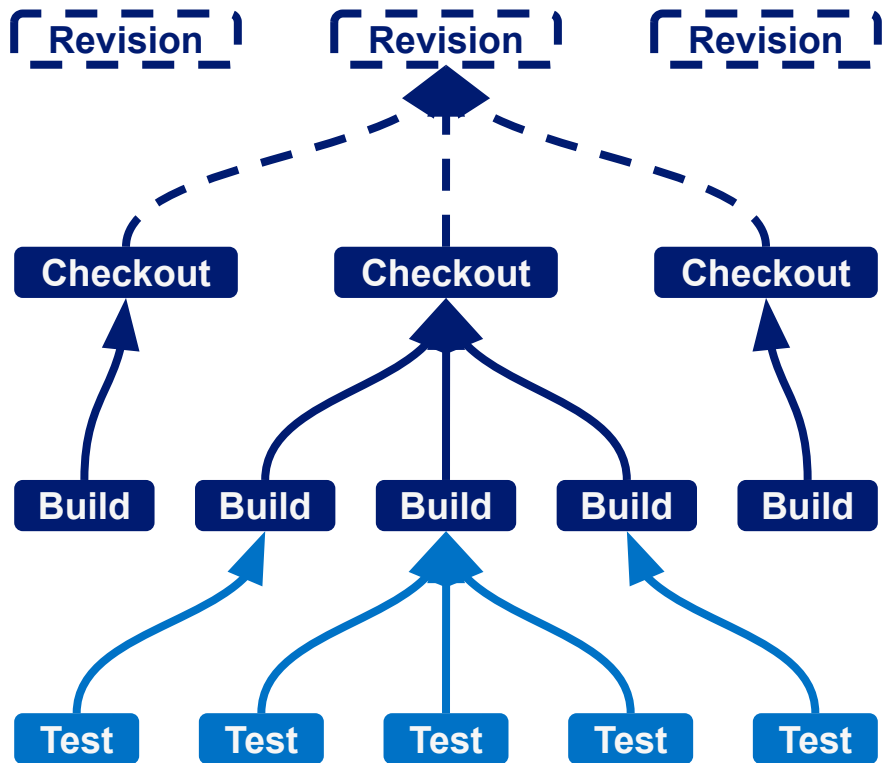




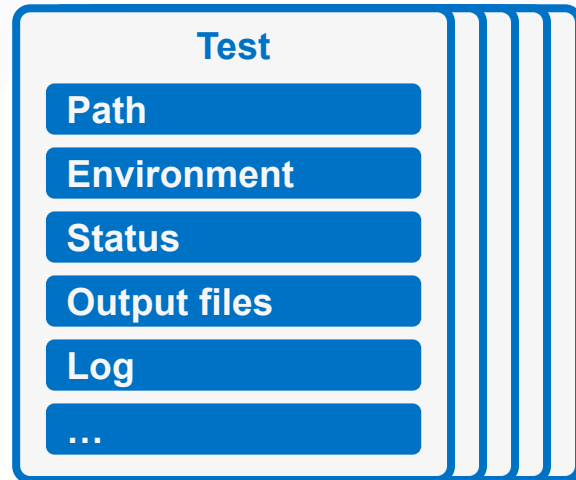
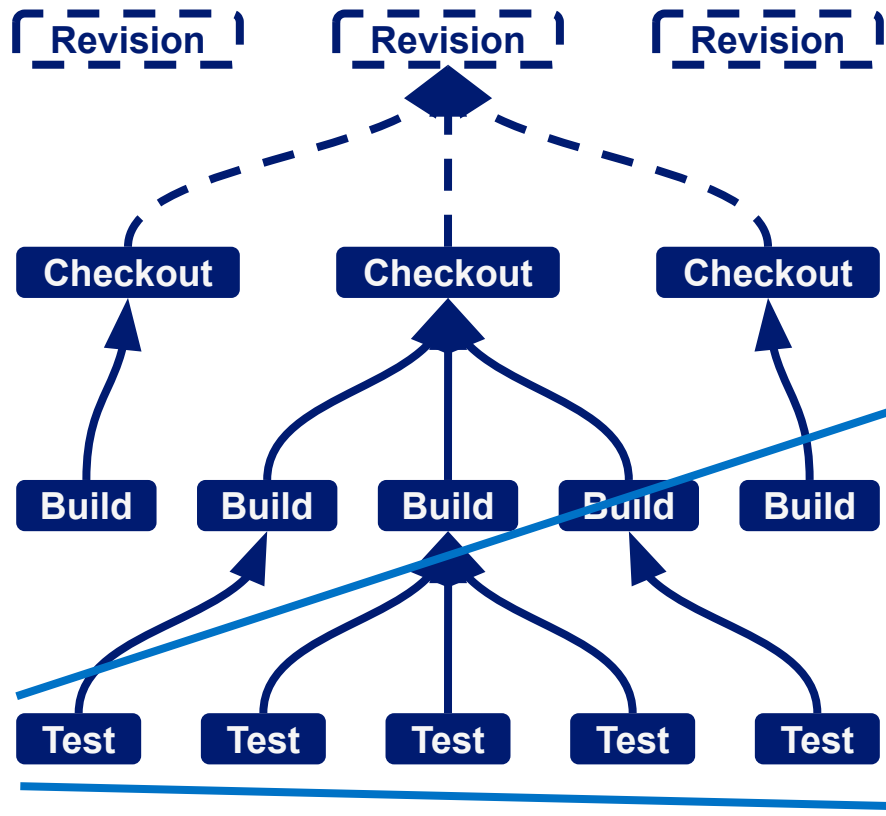
# KernelCI Our data model



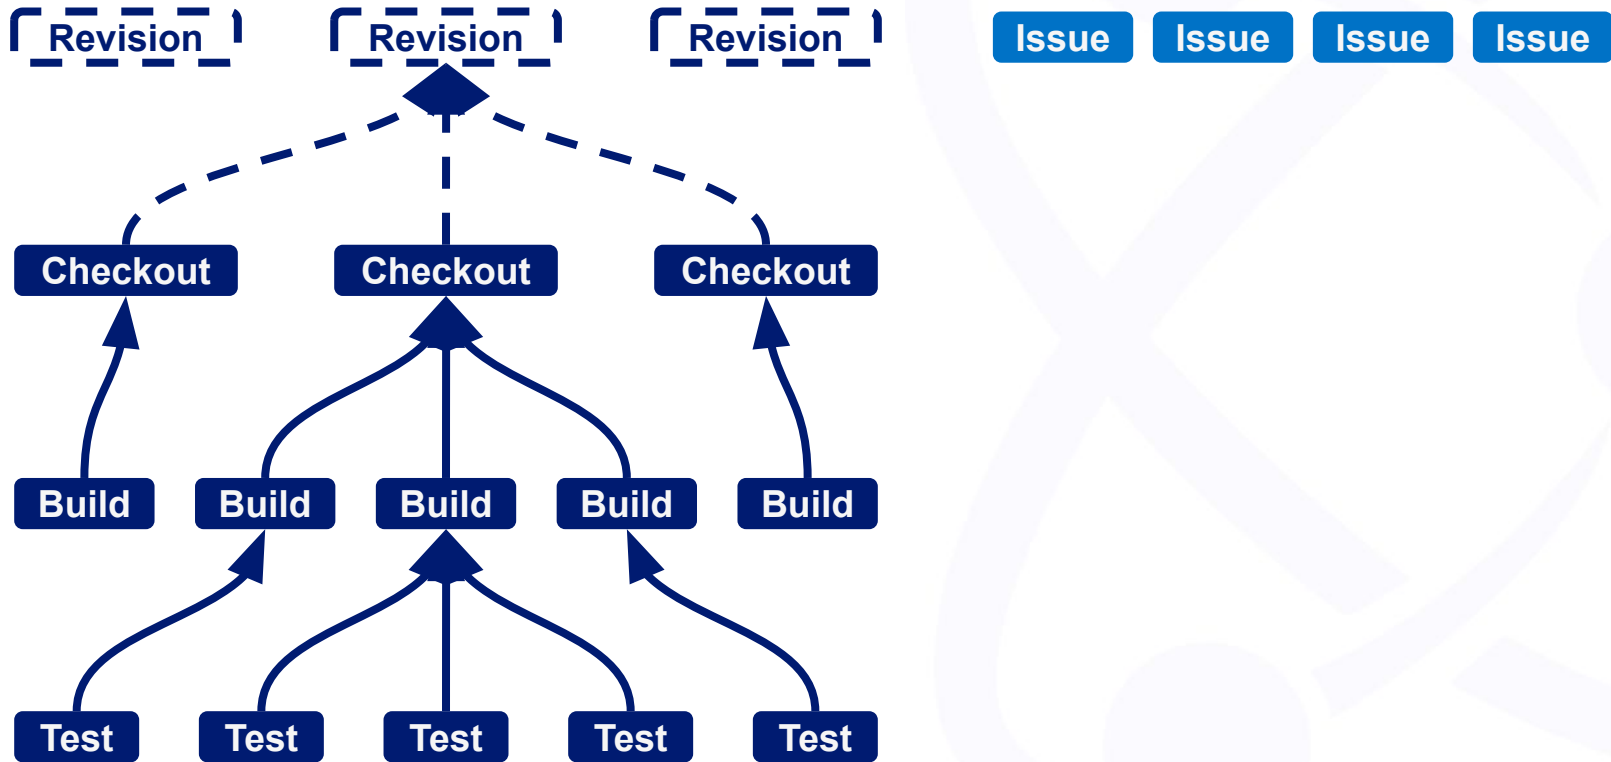
# KernelCI Our data model



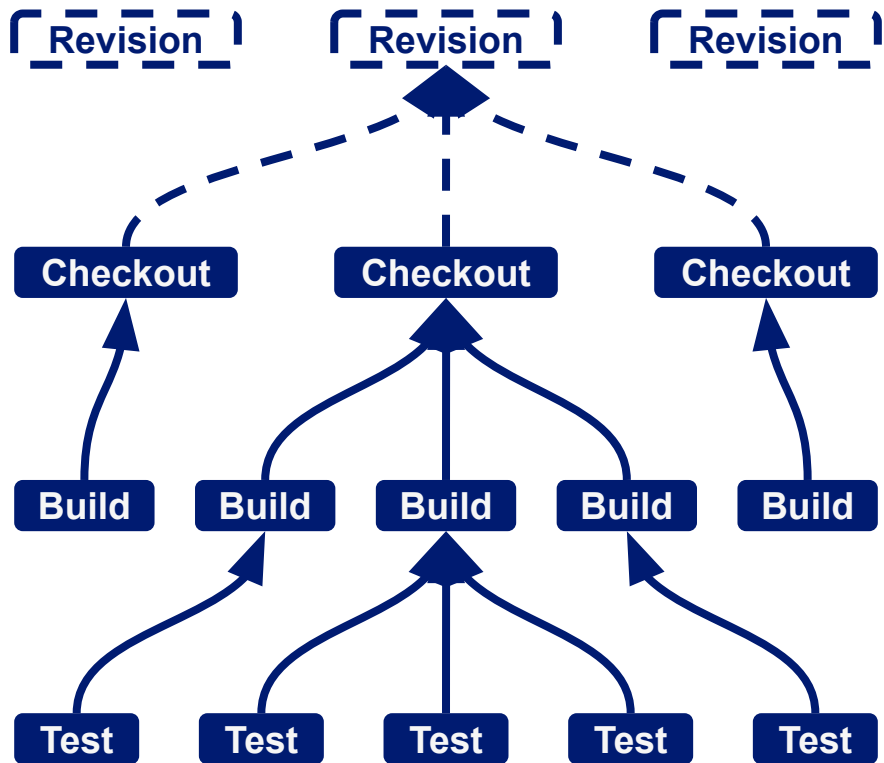
# KernelCI Our data model



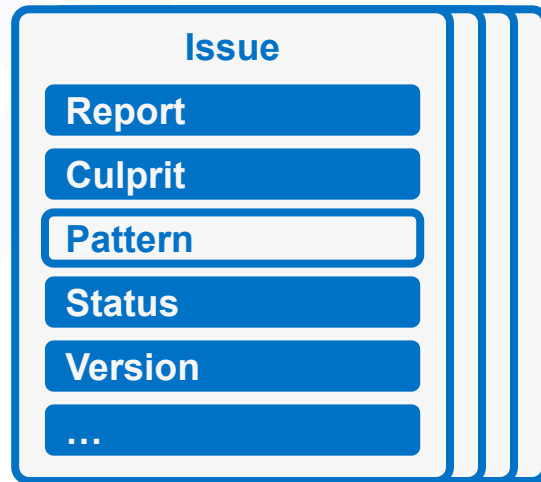
# KernelCI Our data model



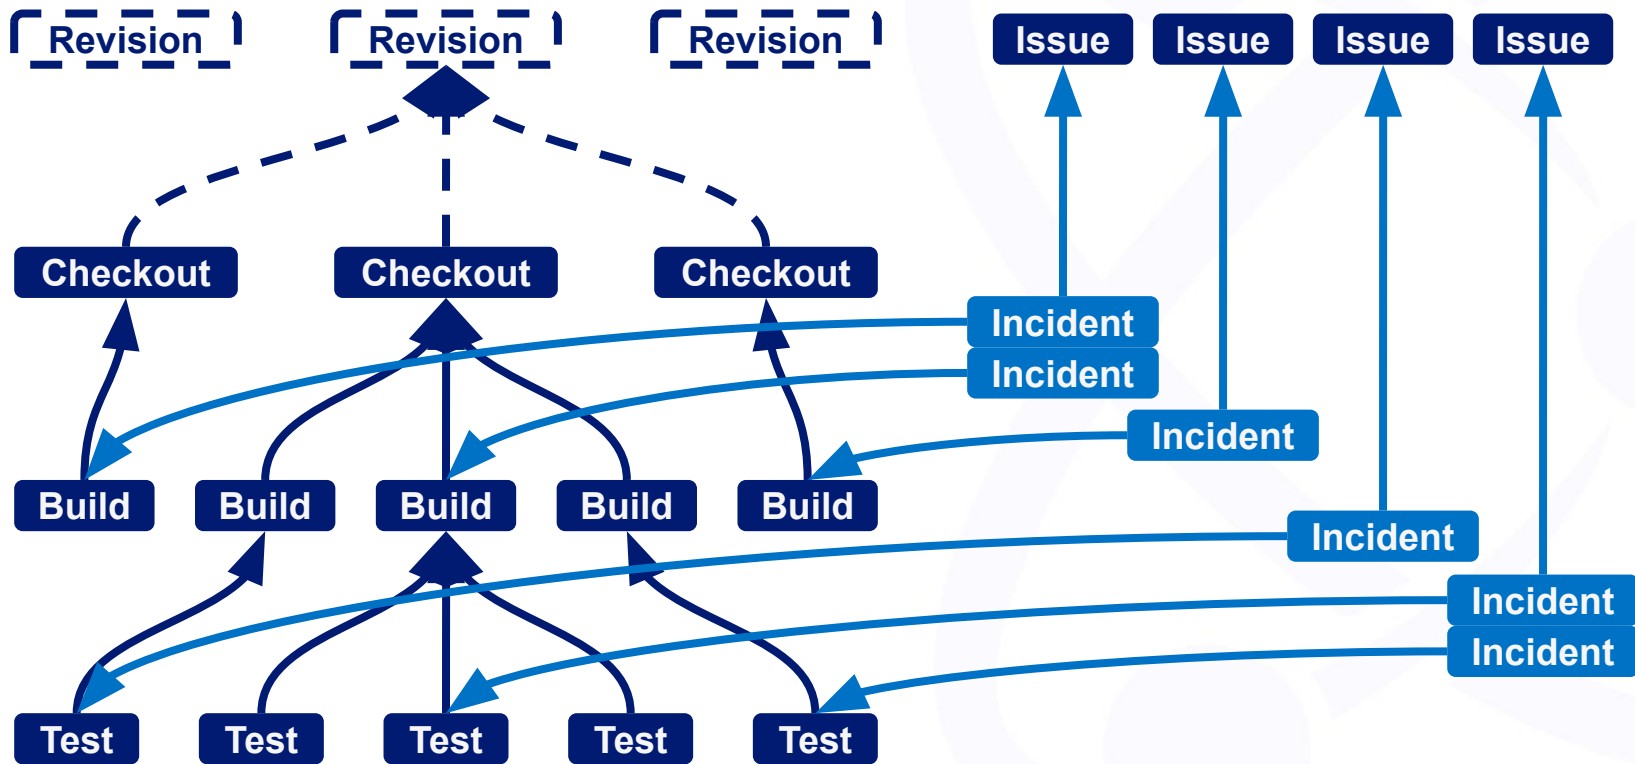
# KernelCI Our data model



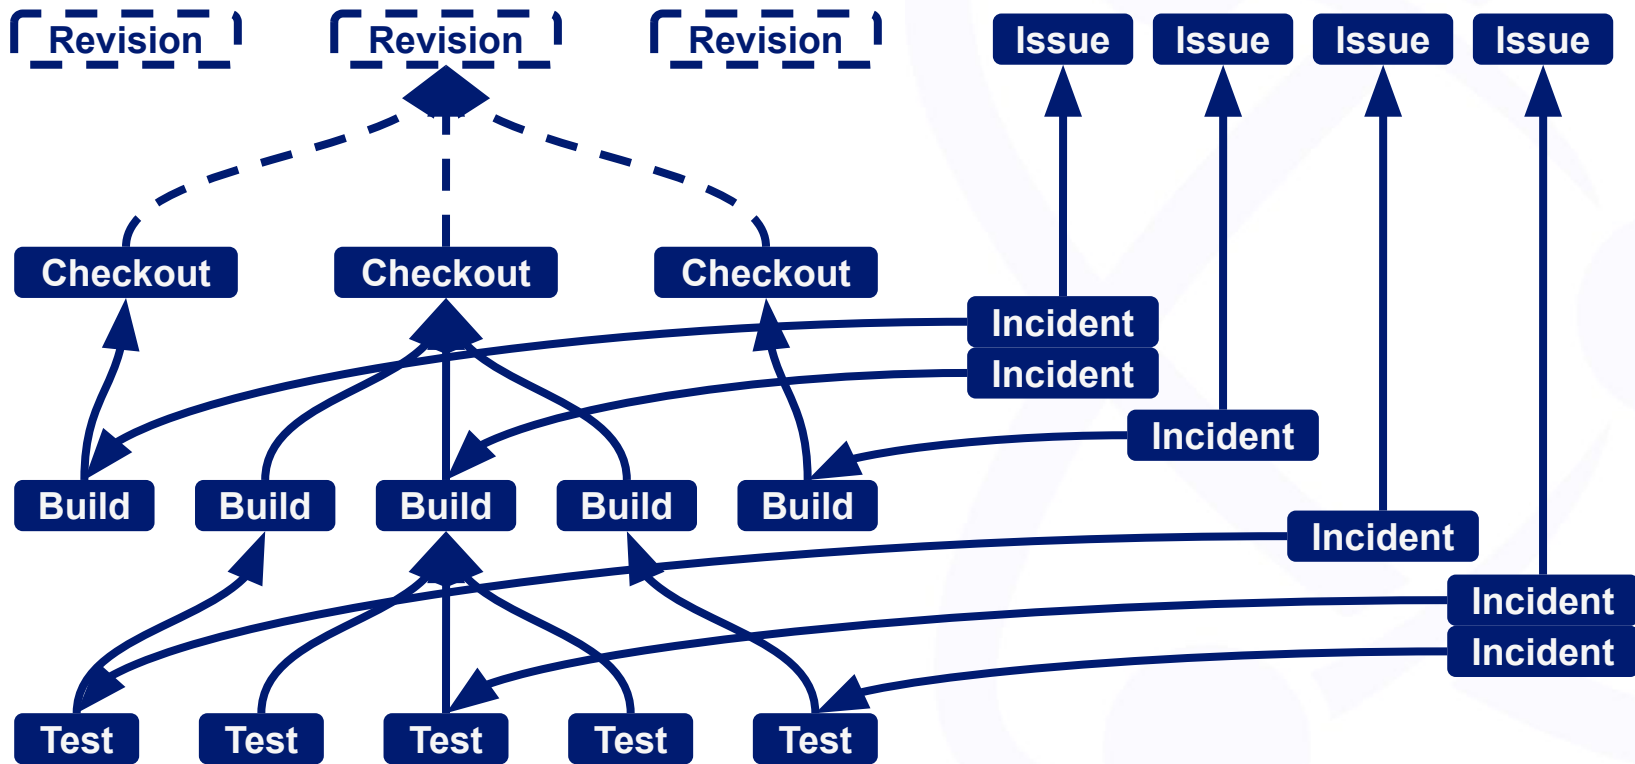
Issue Issue Issue Issue



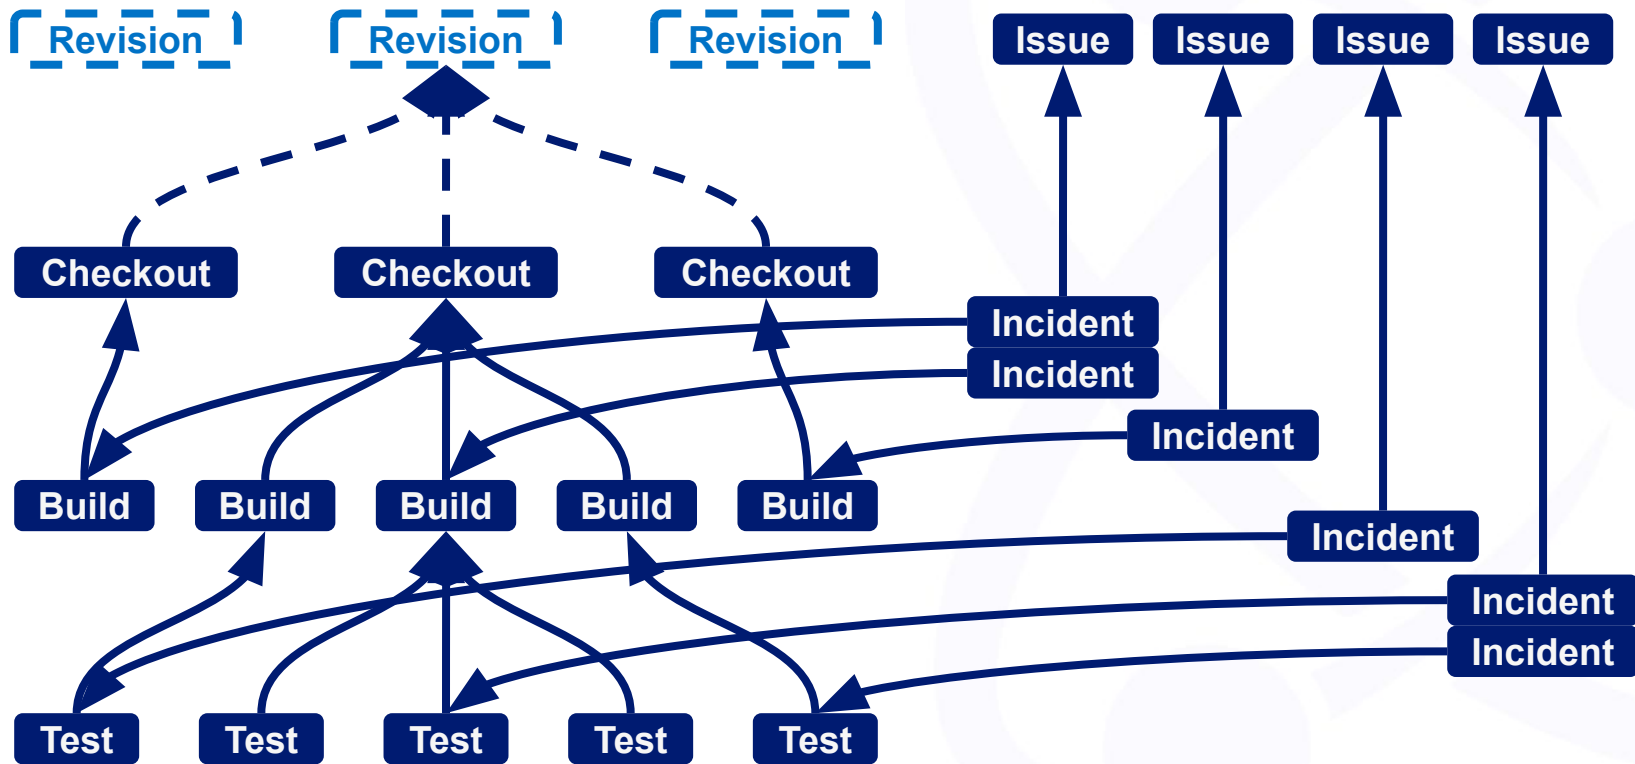
# KernelCI Our data model



# KernelCI Our data model

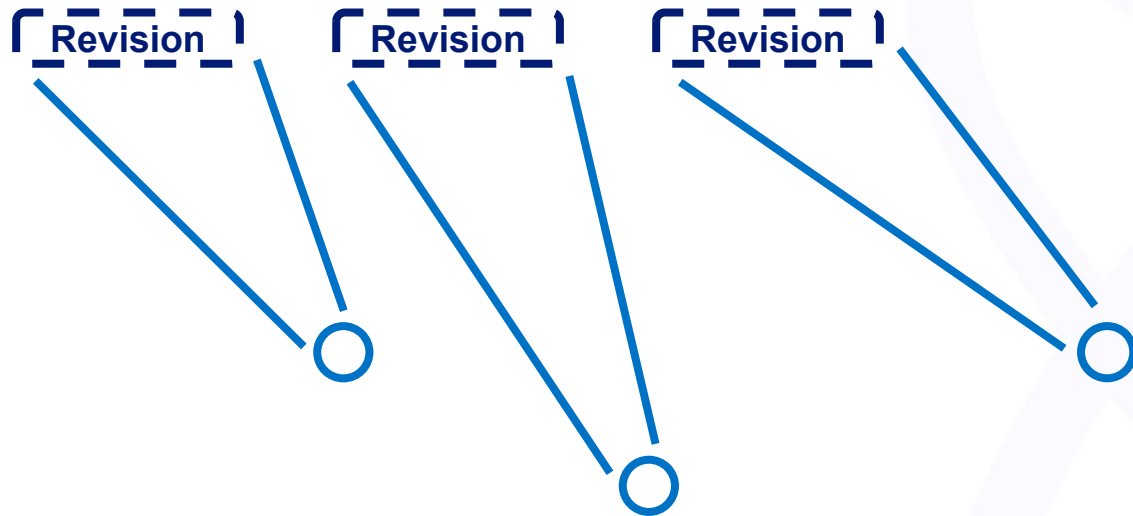


# KernelCI Our data model






# KernelCI Our (planned) data model



# Our (planned) data model




# KernelCI Our (planned) data model



 Commit parent


# KernelCI Our (planned) data model



← Commit parent

◁ Implicit parent


# KernelCI Our (planned) data model



← Commit parent


◁ Implicit parent

# KernelCI Our (planned) data model



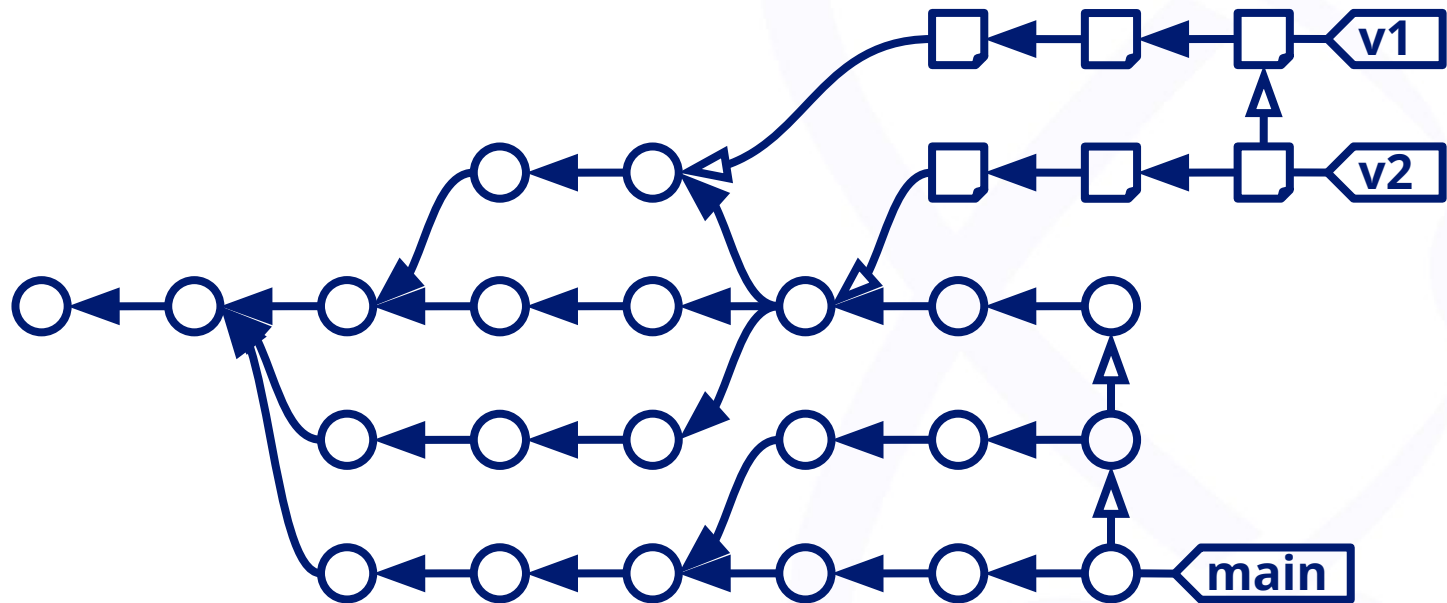
- ← Commit parent
- ◁ Implicit parent

# KernelCI Our (planned) data model



- ← Commit parent
- ↳ Implicit parent

# KernelCI Our (planned) data model



- ← Commit parent
- ◁ Implicit parent



# Directed Acyclic Graph (DAG)



- ← Commit parent
- ⤴ Implicit parent



# KernelCI Our (planned) data model

- **There's likely too many build and test results for a Graph Database.**
- **We might be able to keep issues there, but likely not incidents.**

# Queries we need



# KernelCI Queries we need

- **Which issues should we triage for this revision?**
  - We have to prioritize issues to triage.
- **What were last X test results for last Y commits?**
  - We gotta alert on changing result patterns.
  - Performance trends along revision history.
- **Can we release this branch?**
  - Which issues were found since last release?
- **Which branch revisions have a particular issue?**

# Neo4j experience and ideas



```
#!/usr/bin/env python3

import sys
import io
import csv

LOG = io.TextIOWrapper(sys.stdin.buffer, errors="replace")

def readfield(first=False):
    field = LOG.readline()
    if field == '':
        sys.exit(int(not first))
    return field[:-1]

with open(sys.argv[1], "w") as commits_file, \
    open(sys.argv[2], "w") as relations_file:
    commits_writer = csv.writer(
        commits_file, doublequote=False, escapechar="\\"
    )
    commits_writer.writerow([
        "hash",
        "author_name", "author_email", "author_timestamp",
        "subject"
    ])
    relations_writer = csv.writer(
        relations_file, doublequote=False, escapechar="\\"
    )
    relations_writer.writerow(["hash", "parent_hash"])
    while True:
        commit_hash = readfield(True)
        parent_hashes = readfield()
        author_name = readfield()
        author_email = readfield()
        author_timestamp = readfield()
        subject = readfield()
        commits_writer.writerow([
            commit_hash,
            author_name, author_email, author_timestamp,
            subject
        ])
        for parent_hash in parent_hashes.split():
            relations_writer.writerow([commit_hash, parent_hash])
```




```
#!/usr/bin/env python3

import sys
import io
import csv

LOG = io.TextIOWrapper(sys.stdin.buffer, errors="replace")

def readfield(first=False):
    field = LOG.readline()
    if field == '':
        sys.exit(int(not first))
    return field[:-1]

with open(sys.argv[1], "w") as commits_file, \
    open(sys.argv[2], "w") as relations_file:
    commits_writer = csv.writer(
        commits_file, doublequote=False, escapechar="\\"
    )
    commits_writer.writerow([
        "hash",
        "author_name", "author_email", "author_timestamp",
        "subject"
    ])
    relations_writer = csv.writer(
        relations_file, doublequote=False, escapechar="\\"
    )
    relations_writer.writerow(["hash", "parent_hash"])
    while True:
        commit_hash = readfield(True)
        parent_hashes = readfield()
        author_name = readfield()
        author_email = readfield()
        author_timestamp = readfield()
        subject = readfield()
        commits_writer.writerow([
            commit_hash,
            author_name, author_email, author_timestamp,
            subject
        ])
        for parent_hash in parent_hashes.split():
            relations_writer.writerow([commit_hash, parent_hash])
```



```
nkondras@ruby:linux$ git log --format='%H%nP%n%an%n%ae%n%aI%n%s' | ./log2csv commits.csv relations.csv
```




```
#!/usr/bin/env python3

import sys
import io
import csv

LOG = io.TextIOWrapper(sys.stdin.buffer, errors="replace")

def readfield(first=False):
    field = LOG.readline()
    if field == '':
        sys.exit(int(not first))
    return field[:-1]

with open(sys.argv[1], "w") as commits_file, \
     open(sys.argv[2], "w") as relations_file:
    commits_writer = csv.writer(
        commits_file, doublequote=False, escapechar="\\"
    )
    commits_writer.writerow([
        "hash",
        "author_name", "author_email", "author_timestamp",
        "subject"
    ])
    relations_writer = csv.writer(
        relations_file, doublequote=False, escapechar="\\"
    )
    relations_writer.writerow(["hash", "parent_hash"])
    while True:
        commit_hash = readfield(True)
        parent_hashes = readfield()
        author_name = readfield()
        author_email = readfield()
        author_timestamp = readfield()
        subject = readfield()
        commits_writer.writerow([
            commit_hash,
            author_name, author_email, author_timestamp,
            subject
        ])
        for parent_hash in parent_hashes.split():
            relations_writer.writerow([commit_hash, parent_hash])
```



```
nkondras@ruby:linux$ git log --format='%H%n%P%n%an%n%ae%n%aI%n%s' | ./log2csv commits.csv relations.csv
nkondras@ruby:linux$ wc -l commits.csv
1154921 commits.csv
```





```
#!/usr/bin/env python3

import sys
import io
import csv

LOG = io.TextIOWrapper(sys.stdin.buffer, errors="replace")

def readfield(first=False):
    field = LOG.readline()
    if field == "":
        sys.exit(int(not first))
    return field[:-1]

with open(sys.argv[1], "w") as commits_file, \
     open(sys.argv[2], "w") as relations_file:
    commits_writer = csv.writer(
        commits_file, doublequote=False, escapechar="\\"
    )
    commits_writer.writerow([
        "hash",
        "author_name", "author_email", "author_timestamp",
        "subject"
    ])
    relations_writer = csv.writer(
        relations_file, doublequote=False, escapechar="\\"
    )
    relations_writer.writerow(["hash", "parent_hash"])
    while True:
        commit_hash = readfield(True)
        parent_hashes = readfield()
        author_name = readfield()
        author_email = readfield()
        author_timestamp = readfield()
        subject = readfield()
        commits_writer.writerow([
            commit_hash,
            author_name, author_email, author_timestamp,
            subject
        ])
        for parent_hash in parent_hashes.split():
            relations_writer.writerow([commit_hash, parent_hash])
```

```
nkondras@ruby:linux$ git log --format='%H%nP%n%an%n%ae%n%aI%n%s' | ./log2csv commits.csv relations.csv
```

```
nkondras@ruby:linux$ wc -l commits.csv
```

```
1154921 commits.csv
```

```
nkondras@ruby:linux$ head commits.csv
```

```
hash,author_name,author_email,author_timestamp,subject
```

```
c0b67534c95c537f7a506a06b98e5e85d72e2b7d,Linus Torvalds,torvalds@linux-foundation.org,2023-01-31T16:02:05-08:00,Merge tag 'c
e5ae8803847b80fe9d744a3174abe2b7bfed222a,Waiman Long,longman@redhat.com,2023-01-31T10:48:03-05:00,cgroup/cpuset: Fix wrong c
58706f7fb045b7019bada81fa17f372189315fe5,Linus Torvalds,torvalds@linux-foundation.org,2023-01-31T11:39:08-08:00,Merge tag 's
88b356e19199dca700b1180adc0b843a3228d4a0,Linus Torvalds,torvalds@linux-foundation.org,2023-01-31T11:17:38-08:00,Merge tag 'm
22b8077d0fcec86c6ed0efce9f7e7e5a4c2d56a,Linus Torvalds,torvalds@linux-foundation.org,2023-01-30T11:58:53-08:00,Merge tag 'f
3288666c72568fe1cc7f5c5ae33dfd3ab18004c8,Hou Tao,houtao1@huawei.com,2023-01-13T19:52:11+08:00,fscache: Use clear_and_wake_up
8226e37d82f43657da34dd770e2b38f20242ada7,Hou Tao,houtao1@huawei.com,2023-01-13T19:52:10+08:00,fscache: Use wait_on_bit() to
6d796c50f84ca79f1722bb131799e5a5710c4700,Linus Torvalds,torvalds@linux-foundation.org,2023-01-29T13:59:43-08:00,Linux 6.2-rc
ab072681eabe1ce0a9a32d4baa1a27a2d046bc4a,Linus Torvalds,torvalds@linux-foundation.org,2023-01-29T11:26:49-08:00,Merge tag 'i
```




```
#!/usr/bin/env python3

import sys
import io
import csv

LOG = io.TextIOWrapper(sys.stdin.buffer, errors="replace")

def readfield(first=False):
    field = LOG.readline()
    if field == "":
        sys.exit(int(not first))
    return field[:-1]

with open(sys.argv[1], "w") as commits_file, \
     open(sys.argv[2], "w") as relations_file:
    commits_writer = csv.writer(
        commits_file, doublequote=False, escapechar="\\"
    )
    commits_writer.writerow([
        "hash",
        "author_name", "author_email", "author_timestamp",
        "subject"
    ])
    relations_writer = csv.writer(
        relations_file, doublequote=False, escapechar="\\"
    )
    relations_writer.writerow(["hash", "parent_hash"])
    while True:
        commit_hash = readfield(True)
        parent_hashes = readfield()
        author_name = readfield()
        author_email = readfield()
        author_timestamp = readfield()
        subject = readfield()
        commits_writer.writerow([
            commit_hash,
            author_name, author_email, author_timestamp,
            subject
        ])
        for parent_hash in parent_hashes.split():
            relations_writer.writerow([commit_hash, parent_hash])
```



```
nkondras@ruby:linux$ git log --format='%H%nP%n%an%n%ae%n%aI%n%s' | ./log2csv commits.csv relations.csv
nkondras@ruby:linux$ wc -l commits.csv
1154921 commits.csv
nkondras@ruby:linux$ head commits.csv
hash,author_name,author_email,author_timestamp,subject
c0b67534c95c537f7a506a06b98e5e85d72e2b7d,Linus Torvalds,torvalds@linux-foundation.org,2023-01-31T16:02:05-08:00,Merge tag 'c
e5ae8803847b80fe9d744a3174abe2b7bfed222a,Waiman Long,longman@redhat.com,2023-01-31T10:48:03-05:00,cgroup/cpuset: Fix wrong c
58706f7fb045b7019bada81fa17f372189315fe5,Linus Torvalds,torvalds@linux-foundation.org,2023-01-31T11:39:08-08:00,Merge tag 's
88b356e19199dca700b1180adc0b843a3228d4a0,Linus Torvalds,torvalds@linux-foundation.org,2023-01-31T11:17:38-08:00,Merge tag 'm
22b8077d0fcec86c6ed0efce9f7e7e5a4c2d56a,Linus Torvalds,torvalds@linux-foundation.org,2023-01-30T11:58:53-08:00,Merge tag 'f
3288666c72568fe1cc7f5c5ae33dfd3ab18004c8,Hou Tao,houtao1@huawei.com,2023-01-13T19:52:11+08:00,fscache: Use clear_and_wake_up
8226e37d82f43657da34dd770e2b38f20242ada7,Hou Tao,houtao1@huawei.com,2023-01-13T19:52:10+08:00,fscache: Use wait_on_bit() to
6d796c50f84ca79f1722bb131799e5a5710c4700,Linus Torvalds,torvalds@linux-foundation.org,2023-01-29T13:59:43-08:00,Linux 6.2-rc
ab072681eabe1ce0a9a32d4baa1a27a2d046bc4a,Linus Torvalds,torvalds@linux-foundation.org,2023-01-29T11:26:49-08:00,Merge tag 'i
nkondras@ruby:linux$ wc -l relations.csv
1247839 relations.csv
```



```
#!/usr/bin/env python3

import sys
import io
import csv

LOG = io.TextIOWrapper(sys.stdin.buffer, errors="replace")

def readfield(first=False):
    field = LOG.readline()
    if field == "":
        sys.exit(int(not first))
    return field[:-1]

with open(sys.argv[1], "w") as commits_file, \
     open(sys.argv[2], "w") as relations_file:
    commits_writer = csv.writer(
        commits_file, doublequote=False, escapechar="\\"
    )
    commits_writer.writerow([
        "hash",
        "author_name", "author_email", "author_timestamp",
        "subject"
    ])
    relations_writer = csv.writer(
        relations_file, doublequote=False, escapechar="\\"
    )
    relations_writer.writerow(["hash", "parent_hash"])
    while True:
        commit_hash = readfield(True)
        parent_hashes = readfield()
        author_name = readfield()
        author_email = readfield()
        author_timestamp = readfield()
        subject = readfield()
        commits_writer.writerow([
            commit_hash,
            author_name, author_email, author_timestamp,
            subject
        ])
        for parent_hash in parent_hashes.split():
            relations_writer.writerow([commit_hash, parent_hash])
```



```
nkondras@ruby:linux$ git log --format='%H%nP%n%an%n%ae%n%aI%n%s' | ./log2csv commits.csv relations.csv
nkondras@ruby:linux$ wc -l commits.csv
1154921 commits.csv
nkondras@ruby:linux$ head commits.csv
hash,author_name,author_email,author_timestamp,subject
c0b67534c95c537f7a506a06b98e5e85d72e2b7d,Linus Torvalds,torvalds@linux-foundation.org,2023-01-31T16:02:05-08:00,Merge tag 'c
e5ae8803847b80fe9d744a3174abe2b7bfed222a,Waiman Long,longman@redhat.com,2023-01-31T10:48:03-05:00,cgroup/cpuset: Fix wrong c
58706f7fb045b7019bada81fa17f372189315fe5,Linus Torvalds,torvalds@linux-foundation.org,2023-01-31T11:39:08-08:00,Merge tag 's
88b356e19199dca700b1180adc0b843a3228d4a0,Linus Torvalds,torvalds@linux-foundation.org,2023-01-31T11:17:38-08:00,Merge tag 'm
22b8077d0fcec86c6ed0e0fce9f7e7e5a4c2d56a,Linus Torvalds,torvalds@linux-foundation.org,2023-01-30T11:58:53-08:00,Merge tag 'f
3288666c72568fe1cc7f5c5ae33dfd3ab18004c8,Hou Tao,houtao1@huawei.com,2023-01-13T19:52:11+08:00,fscache: Use clear_and_wake_up
8226e37d82f43657da34dd770e2b38f20242ada7,Hou Tao,houtao1@huawei.com,2023-01-13T19:52:10+08:00,fscache: Use wait_on_bit() to
6d796c50f84ca79f1722bb131799e5a5710c4700,Linus Torvalds,torvalds@linux-foundation.org,2023-01-29T13:59:43-08:00,Linux 6.2-rc
ab072681eabe1ce0a9a32d4baa1a27a2d046bc4a,Linus Torvalds,torvalds@linux-foundation.org,2023-01-29T11:26:49-08:00,Merge tag 'i
nkondras@ruby:linux$ wc -l relations.csv
1247839 relations.csv
nkondras@ruby:linux$ head relations.csv
hash,parent_hash
c0b67534c95c537f7a506a06b98e5e85d72e2b7d,58706f7fb045b7019bada81fa17f372189315fe5
c0b67534c95c537f7a506a06b98e5e85d72e2b7d,e5ae8803847b80fe9d744a3174abe2b7bfed222a
e5ae8803847b80fe9d744a3174abe2b7bfed222a,58706f7fb045b7019bada81fa17f372189315fe5
58706f7fb045b7019bada81fa17f372189315fe5,88b356e19199dca700b1180adc0b843a3228d4a0
58706f7fb045b7019bada81fa17f372189315fe5,15600159bcc6abbae6b33a849bef90dca28b78f
88b356e19199dca700b1180adc0b843a3228d4a0,22b8077d0fcec86c6ed0e0fce9f7e7e5a4c2d56a
88b356e19199dca700b1180adc0b843a3228d4a0,89c08aef8f8dcb5fc98ad8db7ca92c9dab1f26b0
22b8077d0fcec86c6ed0e0fce9f7e7e5a4c2d56a,6d796c50f84ca79f1722bb131799e5a5710c4700
22b8077d0fcec86c6ed0e0fce9f7e7e5a4c2d56a,3288666c72568fe1cc7f5c5ae33dfd3ab18004c8
nkondras@ruby:linux$
```

# KernelCI Neo4j experience and ideas

```
1 // Load commits
2 :auto LOAD CSV WITH HEADERS FROM 'file:///commits.csv' AS commit
3 CALL {
4     WITH commit
5     CREATE (c:Commit)
6     SET c = commit
7 } IN TRANSACTIONS
```

Added 1154920 labels, created 1154920 nodes, set 5774600 properties, completed after 46692 ms.

Table

Warn

Code

Added 1154920 labels, created 1154920 nodes, set 5774600 properties, completed after 46692 ms.

# KernelCI Neo4j experience and ideas

```
mainline$ CREATE TEXT INDEX FOR (c:Commit) ON (c.hash)
```

Added 1 index, completed after 24 ms.

Table

Code

Added 1 index, completed after 24 ms.

# KernelCI Neo4j experience and ideas

```
1 // Load relations
2 :auto LOAD CSV WITH HEADERS FROM 'file:///relations.csv' AS relation
3 CALL {
4     WITH relation
5     MATCH (child:Commit {hash:relation.hash})
6     MATCH (parent:Commit {hash:relation.parent_hash})
7     CREATE (child)-[:PARENT]→(parent)
8 } IN TRANSACTIONS
```

Table

Code



# Neo4j experience and ideas



# KernelCI Neo4j experience and ideas

```
1 MATCH (s:Commit {hash:'3123109284176b1532874591f7c81f3837bbdc17'})
2 MATCH (e:Commit {hash:'d0d96121d03d6d9cf608d948247a9f24f5a02da9'})
3 WHERE exists((s)-[:PARENT*0..]->(e))
4 RETURN true
```







```
1 MATCH (s:Commit {hash:'3123109284176b1532874591f7c81f3837bbdc17'})
2 MATCH (e:Commit {hash:'68e77ffbfd06ae3ef8f2abf1c3b971383c866983'})
3 CALL apoc.path.expandConfig(s, {
4   relationshipFilter: 'PARENT>',
5   uniqueness: 'NODE_GLOBAL',
6   terminatorNodes: [e],
7   limit: 1
8 })
9 YIELD path
10 RETURN length(path)
```



# KernelCI Neo4j experience and ideas

```
1 MATCH (s:Commit {hash:'68e77ffbfd06ae3ef8f2abf1c3b971383c866983'})
2 MATCH (e:Commit {hash:'1e20904e417738066b26490de2daf7ef3ed34483'})
3 CALL apoc.path.expandConfig(s, {
4     relationshipFilter: 'PARENT>',
5     uniqueness:'NONE',
6     terminatorNodes: [e]
7 })
8 YIELD path
9 UNWIND nodes(path) AS node
10 RETURN count(DISTINCT node)
```





KernelCI

# Neo4j experience and ideas

```
nkondras@ruby:linux$ time git log --oneline 1e20904e417738066b26490de2daf7ef3ed34483..68e77ffbfd06ae3ef8f2abf1c3b971383c866983 | wc -l  
283
```

```
real    0m0.302s  
user    0m0.268s  
sys     0m0.038s
```

```
nkondras@ruby:linux$
```





# KernelCI Neo4j experience and ideas

- **Add support for DAG-restricted databases**

- It is also possible to check whether a given directed graph is a DAG in linear time, either by attempting to find a topological ordering and then testing for each edge whether the resulting ordering is valid<sup>[18]</sup> or alternatively, for some topological sorting algorithms, by verifying that the algorithm successfully orders all the vertices without meeting an error condition.<sup>[17]</sup>

– [Wikipedia](#)

- **Gradually add optimizations**

- [Topological sorting](#) is the algorithmic problem of finding a topological ordering of a given DAG. It can be solved in linear time.<sup>[16]</sup>

– [Wikipedia](#)

# Fallback plan



# KernelCI Fallback plan

- **Store all revisions as commits in a giant git repo**
  - **Branches**
    - **And their force-pushes**
  - **Patchsets**
    - **And their versions**
- **Query with [libgit2](#)**
- **Shuttle commit hashes to/from the relational DB**


**Thank you!**



**Join Us!**



# KernelCI Main repos on GitHub



## KernelCI

Upstream Linux Kernel Validation Project

<https://kernelci.org> [kernelci@groups.io](mailto:kernelci@groups.io)

Repositories 41 Packages People 22 Teams 2 Projects 1

Search:  Type: All Language: All New

3 results for repositories matching **kcidb** Clear filter

### kcidb

kernelci.org common database tools

Python GPL-2.0 7 5 38 1 Updated 6 days ago

### kcidb-io

kernelci.org common reporting - I/O data library

Python 2 1 7 2 Updated 8 days ago

### kcidb-grafana

KCIDB Grafana docs and files


1 2 5 0 Updated on Jan 13

#### Top languages

- Python
- Shell
- HTML
- Makefile
- SaltStack

#### People

22 >





is:issue is:open label:"good first issu" 7 Pull requests Issues Marketplace Explore

Repositories	2
Code	0
Commits	0
Issues	8
Discussions <span>Beta</span>	0
Packages	107K
Marketplace	0
Topics	815K
Wikis	0
Users	0

<b>States</b>	
Open	8

<b>Languages</b>	
Python	8

[Advanced search](#) [Cheat sheet](#)

## 8 issues or view all results on GitHub

- kernelci/kcidb #157

[bigquery: Add and use a single function for queries](#)

To reduce code duplication implement a single querying function to be called by dump()/query()/oo\_query()/etc. in the BigQuery database code. Add DEBUG-level logging of queries and their parameters.

good first issue  
spbnick
- kernelci/kcidb-io #9

[Mark all symbols not supposed to be exposed as private](#)

Figure out the correct way for, and mark all symbols not explicitly supposed to be exposed as private to avoid unintended use and unnecessary coupling.

good first issue  
spbnick
- kernelci/kcidb #111

[Upgrade google-cloud-secret-manager dependency](#)

The latest release of the google-cloud-secret-manager package breaks the existing KCIDB code, and was pinned to versions <2.0.0dev. Fix the incompatibilities and unpin the package. Or perhaps re-pin ...

good first issue  
spbnick
- kernelci/kcidb #110

[Upgrade google-cloud-pubsub dependency](#)

The latest release of the google-cloud-pubsub package breaks the existing KCIDB code, and was pinned to versions <2.0.0dev. Fix the incompatibilities and unpin the package. Or perhaps re-pin to <3.0.0dev. ...

good first issue

## #kernelci at libera.chat

#kernelci The upstream kernel validation project: <http://kernelci.org> | Mailing list: <https://lists.linaro.org/mailman...> 71

Tuesday, February 16th, 2021 about 21 hours ago

**A** **alicef** none  
spbnick: sending it 21:31  
→ FFY00 nipped out 22:41

Wednesday, February 17th, 2021

→ sestebanz and kverz\_broke joined ← alexandra, shuah and kveremitz quit → selvak1631189 popped in → csd nipped out 01:10

**S** **spbnick** Nikolai Kondrashov  
alicef: Yep, I see it now, awesome :D 10:52  
→ mwasilew joined ← veremitz quit → selvak1631189, gaulishcoin9 and gaulishcoin3 popped in • kverz\_broke → kveremitz 10:57

**S** **spbnick** Nikolai Kondrashov  
Everyone who's around, say hi for a channel screenshot for my talk on Friday 🙌! 12:19

**G** **gtucker** Guillaume Tucker  
o/ KernelCI rules :) 12:33  
→ vkabatova joined (~vkabatova@176-74-128-122.netdatacomm.cz) 12:40

**M** **montjoie**  
hello 12:49  
→ monstr joined 15:50

**S** **spbnick** | 18:26