



Test-driven kernel releases

2022-06-01

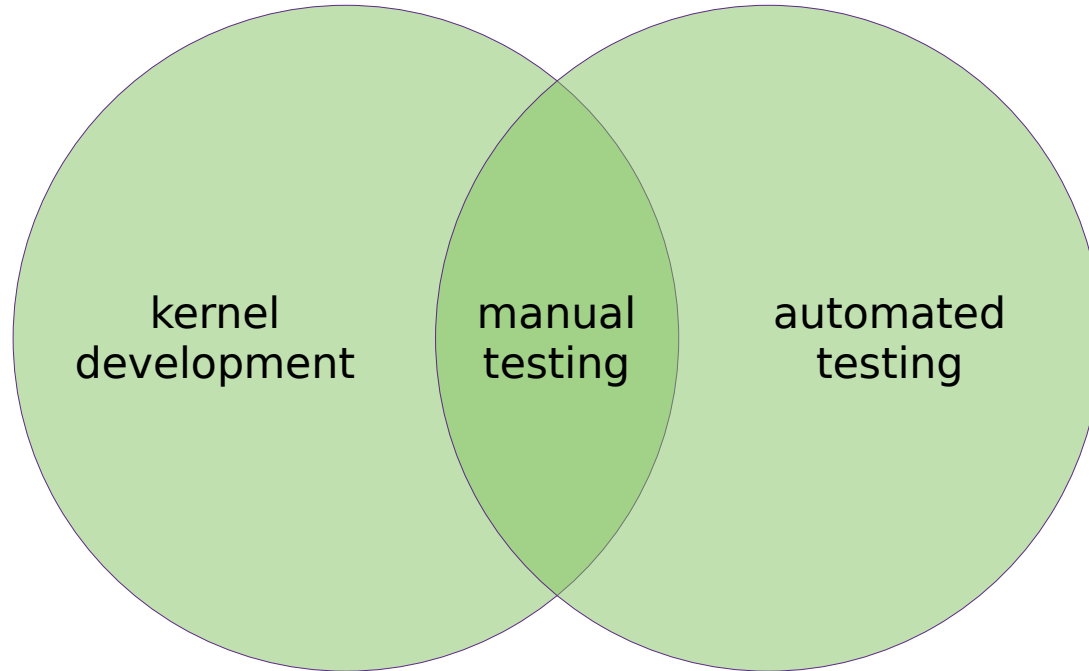
Guillaume Tucker
gtucker@collabora.com



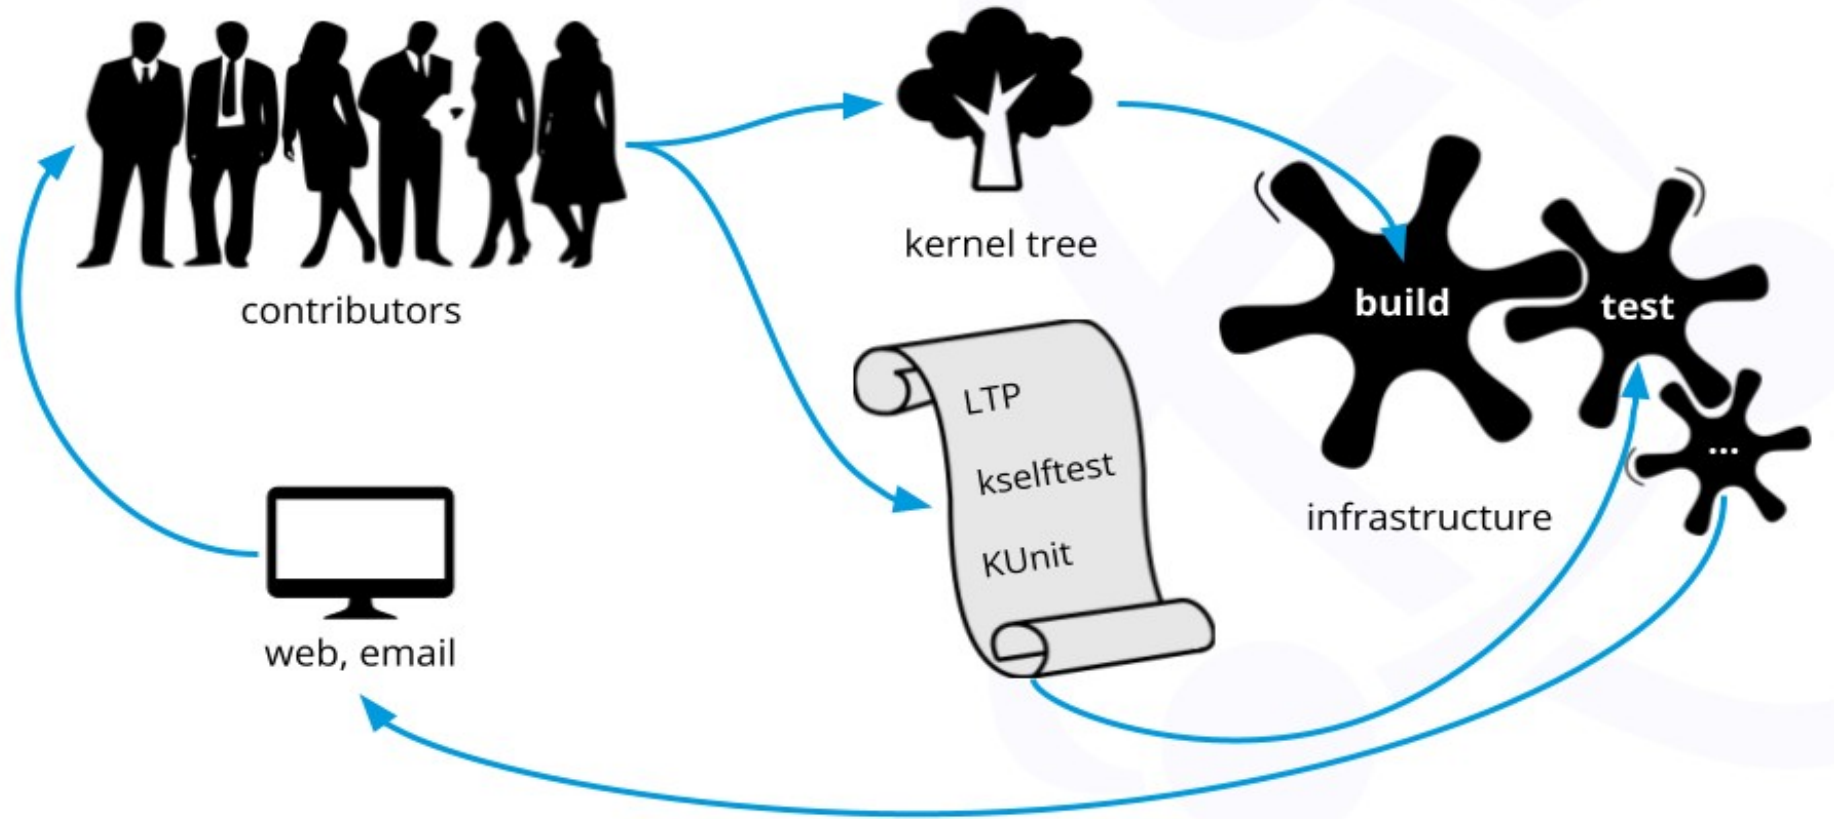
COLLABORA

Open First


Development & Testing



Automated Testing



Open Source Philosophy



Single mainline code base

Many contributors

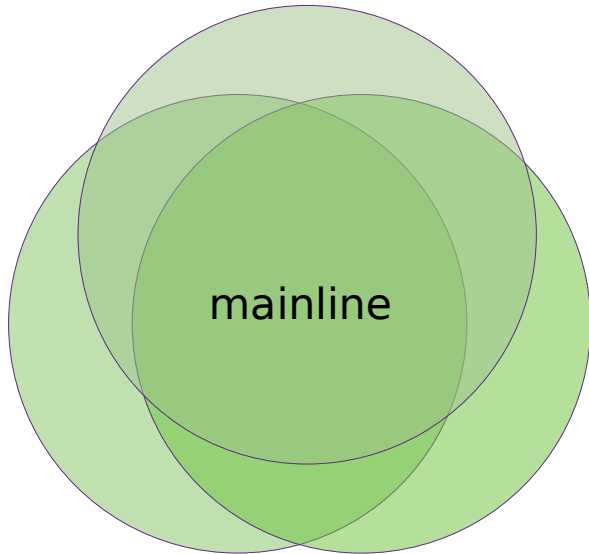
Many use-cases

Application changes are sent upstream

Reduced duplication of efforts



Open Testing Philosophy



Single mainline code base **including tests**

Many contributors **who run tests**

Test results sent upstream

Test results summary in each release

Reduced duplication of **testing efforts**



Hidden Mass of Testing

Duplicated testing efforts

No solution for tracking results upstream

Testing stays hidden as if it was downstream





COLLABORA

automated testing

syzbot

<https://syzkaller.appspot.com/>

syscall fuzzing

Automated bisection

Reproducers

Web UI

syzbot

Instances:

Name	Last active	Uptime	Corpus	Coverage ①	Crashes	Execs	Kernel build				syzkaller build		
							Commit	Config	Freshness	Status	Commit	Freshness	Status
ci-qemu-upstream	now	59m	42456	613488	229	1016941	8ab2afa23bd1	.config	17h59m		3666edfe	now	
ci-qemu-upstream-386	now	59m	41000	569063	184	996495	8ab2afa23bd1	.config	17h59m		3666edfe	now	
ci-qemu2-arm32	now	1h05m	107869	123963	30	793004	8ab2afa23bd1	.config	17h59m		3666edfe	now	
ci-qemu2-arm64	now	1h04m	77347	89931	12	459534	8ab2afa23bd1	.config	17h59m		3666edfe	now	
ci-qemu2-arm64-compat	now	1h04m	78230	88610	15	397037	8ab2afa23bd1	.config	17h59m		3666edfe	now	
ci-qemu2-arm64-mte	now	1h05m	92952	108538	1	654586	8ab2afa23bd1	.config	17h59m		3666edfe	now	
ci-qemu2-riscv64	now	48m	10580	242971	241	39012	0966d385830d	.config	81d	failing	3666edfe	now	
ci-upstream-bpf-kasan-gce	now	2h24m	14943	330503	38	2274433	90343f573252	.config	2d02h		af70c3a9	20h38m	
ci-upstream-bpf-next-kasan-gce	now	2h25m	15457	329985	34	2930736	4b4b4f94a4f6	.config	17h02m		af70c3a9	20h38m	
ci-upstream-gce-leak	now	2h24m	45174	747405	72	1570633	8ab2afa23bd1	.config	17h59m		af70c3a9	20h38m	
ci-upstream-kasan-gce	now	2h24m	41472	601455	61	2698677	8ab2afa23bd1	.config	17h59m		af70c3a9	20h38m	
ci-upstream-kasan-gce-386	now	2h24m	36104	549381	38	957582	8ab2afa23bd1	.config	17h59m		af70c3a9	20h38m	
ci-upstream-kasan-gce-root	now	2h25m	48399	681376	169	3199595	8ab2afa23bd1	.config	17h59m		af70c3a9	20h38m	
ci-upstream-kasan-gce-selinux-root	now	2h24m	49446	734581	147	3141371	8ab2afa23bd1	.config	17h59m		af70c3a9	20h38m	
ci-upstream-kasan-gce-smack-root	now	2h24m	69473	594441	164	2619679	8ab2afa23bd1	.config	17h59m		af70c3a9	20h38m	
ci-upstream-kmsan-gce	now	2h25m	63842	373073	43	6109913	917c7d3f1a0a	.config	5d23h		af70c3a9	20h38m	
ci-upstream-kmsan-gce-386	now	2h25m	59979	405932	21	1933296	917c7d3f1a0a	.config	5d23h		af70c3a9	20h38m	
ci-upstream-linux-next-kasan-gce-root	now	2h24m	46305	703815	135	2019568	3b46e4e44180	.config	8h53m		af70c3a9	20h38m	
ci-upstream-net-kasan-gce	now	2h24m	31845	396464	122	5750513	7e062cda7d90	.config	5d18h		af70c3a9	20h38m	
ci-upstream-net-this-kasan-gce	now	2h16m	27501	380712	72	2348548	3e0b8f529c10	.config	4h09m		af70c3a9	20h38m	
ci2-upstream-kcsan-gce	now	4h34m	57241	374211	106	5311338	8ab2afa23bd1	.config	17h59m		af70c3a9	20h38m	
ci2-upstream-usb	now	4h35m	2626	67138	47	2722928	97fa5887cf28	.config	10d		af70c3a9	20h38m	

open (881):

Title	Repro	Cause bisect	Fix bisect	Count	Last	Reported	Last activity
WARNING in dma_map_sgtable (2)	C	inconclusive		3	4d00h	1d00h	4h41m
INFO: task can't die in vlan_ioctl_handler				5	1d06h	1d06h	1d06h
KASAN: use-after-free Read in filp_close				2	6d13h	1d06h	1d06h
KASAN: use-after-free Read in ext4_xattr_set_entry (5)				1	7d14h	3d04h	3d04h
memory leak in xas_nomem	C			1	7d10h	3d10h	1d12h
riscv/fixes test error: lost connection to test machine				12	59m	4d00h	3d05h
WARNING in driver_unregister	C			1	8d04h	4d04h	6h04m
KASAN: use-after-free Read in macsec_get_iflink				1	8d05h	4d05h	4d05h
WARNING in floppy_queue_rq				3	2d17h	4d06h	4d06h
memory leak in hidraw_report_event	C			1	6d10h	5d23h	5d23h
KASAN: use-after-free Read in do_sync_mmap_readahead				1	9d23h	5d23h	4d01h
memory leak in autofs_wait	C			1	9d23h	5d23h	5d23h
memory leak in dvb_usb_device_init	syz			1	10d	6d15h	6d15h
memory leak in tty_init_dev (2)	C			1	10d	6d16h	6d16h



COLLABORA

Open First

KernelCI

<https://linux.kernelci.org/job/>

Tailored CI system

Web API

Distributed test labs

Kubernetes

Automated bisection

KCIDB database

The results shown here cover the last **14 days** of available data starting from **Fri, 25 Mar 2022** (time is UTC based).

25 jobs per page

Tree	Branch	Latest Build Status	Latest Test Results	Date	Status
stable-rc	linux-4.14.y	44 3 2	0 0 0	2022-03-25	
android	android12-5.4	93 45 4	24 2 0	2022-03-25	
soc	for-next	132 4 3	939 16 2	2022-03-25	
android	android-mainline	109 29 3	1,352 11 2	2022-03-25	
soc	arm/fixes	123 3 2	2,647 137 16	2022-03-25	
android	android11-5.4	87 75 12	495 21 12	2022-03-25	
stable-rc	queue/4.14	119 6 1	33 3 0	2022-03-25	
clk	clk-next	130 4 1	1,551 32 6	2022-03-25	
stable-rc	queue/5.10	129 2 3	1,391 129 12	2022-03-25	
next	master	146 9 4	3,429 211 54	2022-03-25	
stable-rc	queue/5.16	107 2 2	1,069 53 9	2022-03-25	
net-next	master	127 5 3	2,861 136 28	2022-03-25	
mainline	master	141 5 7	3,608 140 20	2022-03-25	
stable-rc	queue/5.4	115 9 1	1,044 81 8	2022-03-25	
next	pending-fixes	131 2 2	4,944 190 67	2022-03-25	
stable-rc	queue/4.19	139 10 4	1,136 118 9	2022-03-25	
stable-rc	queue/5.15	125 1 3	1,059 33 19	2022-03-25	



COLLABORA

Open First

Red Hat CKI

https://datawarehouse.cki-project.org/

Fedora kernels

Mainline kernels

Stable

LTP, kselftest etc.

KCIDDB integration

The screenshot shows the Red Hat CKI web interface. On the left is a blue sidebar with navigation icons for Checkouts, Builds, Issues, Tests, and Hosts. The main content area has a search bar at the top and a table titled 'Latest Checkouts'. The table has columns for ID, Tree, Branch, Start Time, Valid, Builds, and Tests. Each row represents a checkout job with its status and results.

ID	Tree	Branch	Start Time	Valid	Builds	Tests
redhat:551563701	mainline.kernel.or	master	May 30, 2022, 8:05 p.m.	✓	5 ✓	2 ✓ 1 ✗
redhat:551563736	mainline.kernel.or	master	May 30, 2022, 8:05 p.m.	✓	1 ✗ 4 ✓	
redhat:551558028	fedora-35		May 30, 2022, 7:49 p.m.	✓	4 ✓	
redhat:551548888	bpf-next	master	May 30, 2022, 7:30 p.m.	✓	2 ✓	4 ✓ 5 ✗
redhat:551537593	mainline.kernel.or	master	May 30, 2022, 7:10 p.m.	✓	5 ✓	2 ✓ 1 ✗
redhat:551537787	mainline.kernel.or	master	May 30, 2022, 7:09 p.m.	✓	5 ✓	3 ✓ 2 ✗
redhat:551527544	mainline.kernel.or	master	May 30, 2022, 7:04 p.m.	✓	5 ✓	3 ✓ 2 ✗
redhat:551527518	mainline.kernel.or	master	May 30, 2022, 7 p.m.	✓	5 ✓	26 ✓ 13 ✗
redhat:551516150	mainline.kernel.or	master	May 30, 2022, 6:45 p.m.	✓	5 ✓	21 ✓ 11 ✗
redhat:551516103	mainline.kernel.or	master	May 30, 2022, 6:45 p.m.	✓	5 ✓	18 ✓ 9 ✗
redhat:551471840	cki-project/kernel-	os-build	May 30, 2022, 5:23 p.m.	✓		
redhat:551387913	c9s	main	May 30, 2022, 3:21 p.m.	✓	5 ✓	115 ✗ 23 ✓ 2 ✗
redhat:551381666	c9s	main	May 30, 2022, 3:13 p.m.	✓	5 ✓	194 ✗

regzbot

<https://linux-regtracking.leemhuis.info/regzbot/mainline/>

All known regressions

Essentially manual submissions

Seamless integration with emails

Weekly report on LKML for mainline



COLLABORA

Linux kernel regression status

[\[next\]](#) [\[mainline\]](#) [\[stable/longterm\]](#) [\[dormant\]](#) [\[resolved\]](#) | [\[new\]](#) | [\[all\]](#)

current cycle (v5.18.. aka v5.18-post), culprit identified

none known by regzbot

current cycle (v5.18.. aka v5.18-post), unkown culprit

none known by regzbot

previous cycle (v5.17..v5.18), culprit identified, with activity in the past three months

- [ff042f4a9b05](#) (v5.18-rc1) **▶ mm: chiq_test runs 7 minutes instead of ~ 1 second.** by [Stefan Wahren](#) Earliest & latest [activity](#): [9](#) & [1](#) days ago. Noteworthy: [\[patch\]](#).
- [f26b3fa04611](#) (v5.18-rc1) **▶ mm: [mm/page_alloc] f26b3fa046: netperf.Throughput_Mbps -18.0% regression** by [kernel test robot](#) Earliest & latest [activity](#): [41](#) & [10](#) days ago. Noteworthy: [\[patch\]](#).

older cycles (..v5.17), culprit identified, with activity in the past three months

- [453e41085183](#) (v5.13-rc1) **▶ tboot suspend broken on Lenovo T460p** by [Derek Dolney](#) Earliest & latest [activity](#): [32](#) & [1](#) days ago. Noteworthy: [\[1\]](#).
- [44c57f205876](#) (v5.15-rc1) **▶ qla2xxx: tape drive not removed after unplug FC cable** by [Tony Battersby](#) Earliest & latest [activity](#): [5](#) & [3](#) days ago.
- [bdd8b6c98239](#) (v5.17-rc1) **▶ Xorg SEGV in Xen PV dom0 after updating from 5.16.18 to 5.17.5** by [Marek Marczykowski-Górecki](#) Earliest & latest [activity](#): [28](#) & [4](#) days ago. Noteworthy: [\[1\]](#), [\[patch\]](#).
- [a56a1138cbd8](#) (v5.17-rc7) **▶ [BUG] BLE device unpairing triggers kernel panic** by [Ahmad Fatoum](#) Earliest & latest [activity](#): [17](#) & [9](#) days ago. Noteworthy: [\[patch\]](#).
- [830aa6f29f07](#) (v5.17-rc1) **▶ pci: brcmstb: CM4 no longer boots from SD card** by [Bjorn Helgaas](#) Earliest & latest [activity](#): [28](#) & [11](#) days ago. Noteworthy: [\[1\]](#), [\[2\]](#).
- [f9aa460672c9](#) (v5.11-rc1) **▶ memory corruption on Atmel SAM5D31** by [Peter Rosin](#) Earliest & latest [activity](#): [89](#) & [13](#) days ago. Noteworthy: [\[patch\]](#).
- [a19a93e4c6a9](#) (v5.16-rc1) **▶ scsi: pm: Resume process hangs for 5-6 seconds starting sometime in 5.16** Earliest & latest [activity](#): [32](#) & [31](#) days ago. Noteworthy: [\[1\]](#).

previous cycle (v5.17..v5.18), unkown culprit, with activity in the past three weeks

none known by regzbot

older cycles (..v5.17), unkown culprit, with activity in the past three weeks

- v5.16..v5.17 **▶ fs-cache: file contents are corrupted** by [Max Kellermann](#) Earliest & latest [activity](#): [48](#) & [0](#) days ago. Noteworthy: [\[patch\]](#).
- v5.10..v5.15 **▶ net: netfilter: Intermittent performance regression related to ipset** by [McLean, Patrick](#) Earliest & latest [activity](#): [76](#) & [0](#) days ago.
- v5.12..v5.13 **▶ NFSD stops serving clients** by [Ho Ming Shun](#) Earliest & latest [activity](#): 10 days ago. Noteworthy: [\[1\]](#).
- v5.16..v5.17.4 **▶ Hang in 5.17.4+ that appears to be due to Xen** by [Demi Marie Obenour](#)



show me the results

Focusing on the results

Manual runs

Maintainer scripts

Automated systems

Focusing on the results

Manual runs

Maintainer scripts

Automated systems



**Results are the
least common
denominator**



Benefits of results in releases

Valuable for users in general

Canonical way to keep track of code quality

Essentially, avoiding the “works for me” syndrome

```
[10:45:37] =====  
[10:45:37] Testing complete. 115 tests run. 0 failed. 0 crashed.  
[10:45:37] Test results stored in /root/linux/kunit.json  
[10:45:37] Elapsed time: 265.516s total, 0.007s configuring, 41.982s building, 0.000s running
```

Challenges

Shift in workflow: results are needed before the release

- Similar to how -rc works for stable and mainline

- Expect positive results rather than solely look for regressions

- Additional step for maintainers

Keeping it simple and not disruptive

- Optional

- Up to each maintainer to decide which results to include



COLLABORA

in practice

Where to start?

Results reproducible on any hardware

Tests included in the kernel source tree

Plain builds with reference toolchain binaries, Docker images

Builds with sparse enabled `make C=1`

coccicheck

KUnit

Device tree validation

Where to start?

Results reproducible on any hardware

Tests included in the kernel source tree

Plain builds with reference toolchain binaries, Docker images

Builds with sparse enabled `make C=1`

coccicheck

KUnit

Device tree validation

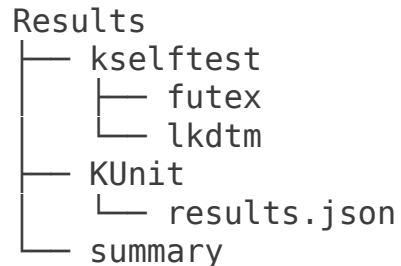
Documentation: <https://docs.kernel.org/>

RFC 1: Test results in-tree

Similar to linux-next merge logs

Updated for each release (stable, mainline, -next)

Rely on Git history for older results



RFC 2: Test-link in commit

```
commit b843adea138649361a8d68b4c9d5824c33fb0146
Author: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
Date:   Wed May 25 09:18:02 2022 +0200

    Linux 5.10.118

    Link: https://lore.kernel.org/r/20220523165812.244140613@linuxfoundation.org
    Test-link: https://testing.kernel.org/stable/5.10.118
    Tested-by: Florian Fainelli <f.fainelli@gmail.com>
    Tested-by: Shuah Khan <skhan@linuxfoundation.org>
    Tested-by: Fox Chen <foxhlchen@gmail.com>
    Tested-by: Salvatore Bonaccorso <carnil@debian.org>
    Tested-by: Sudip Mukherjee <sudip.mukherjee@codethink.co.uk>
    Tested-by: Pavel Machek (CIP) <pavel@denx.de>
    Tested-by: Hulk Robot <hulkrobot@huawei.com>
    Signed-off-by: Greg Kroah-Hartman <gregkh@linuxfoundation.org>
```



RFC 3: Git meta-data

Tied to Git history

Separate from commit merge workflow

Similar to Git notes

```
git results show REVISION
```

Some thoughts

Subsystem-specific results in separate location?

Integration results for mainline / stable / linux-next

Subsystem results could be pulled in alongside code

Follow regular email workflow for adding results

Keep in-tree result summaries in plain text

Extra data can be hosted on separate systems

RFC: <your idea here>

How does the concept sound?

Has this been tried or discussed before?

Does it seem worth the effort?

Time for an RFC on LKML to go through some details?





Thank you!



COLLABORA

Open First