Growing a Lab for Automated Upstream Testing: Challenges and Lessons Learned

Laura Nao, Collabora Ltd

FOSDEM'23

COLLABORA

Growing a Lab for Upstream Testing

Requirements

- Diverse ecosystem of devices
- Software for test automation
- Monitoring system
- Recovery strategies

CI Loop

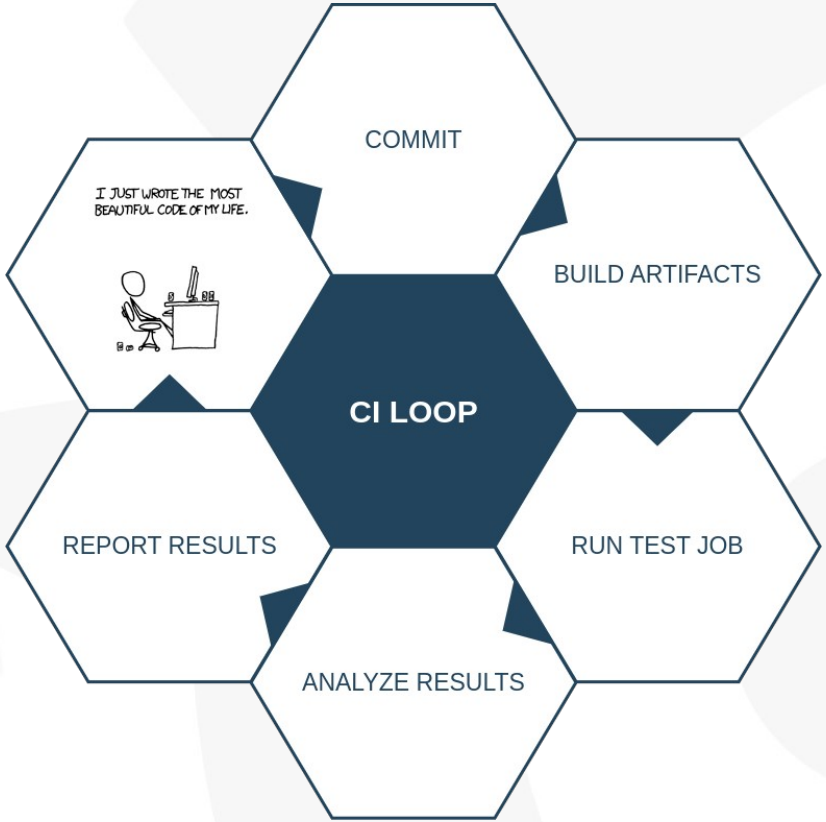
COMMIT
I JUST WROTE THE MOST BEAUTIFUL CODE OF MY LIFE.
BUILD ARTIFACTS
CI LOOP
REPORT RESULTS
RUN TEST JOB
ANALYZE RESULTS

LAVA

- **L**inaro **A**utomation and **V**alidation **A**rchitecture: https://www.lavasoftware.org/
  - Functional testing on real HW
  - Automates power control and serial access
  - Scalable scheduler
  - Results available in different formats

LAVA

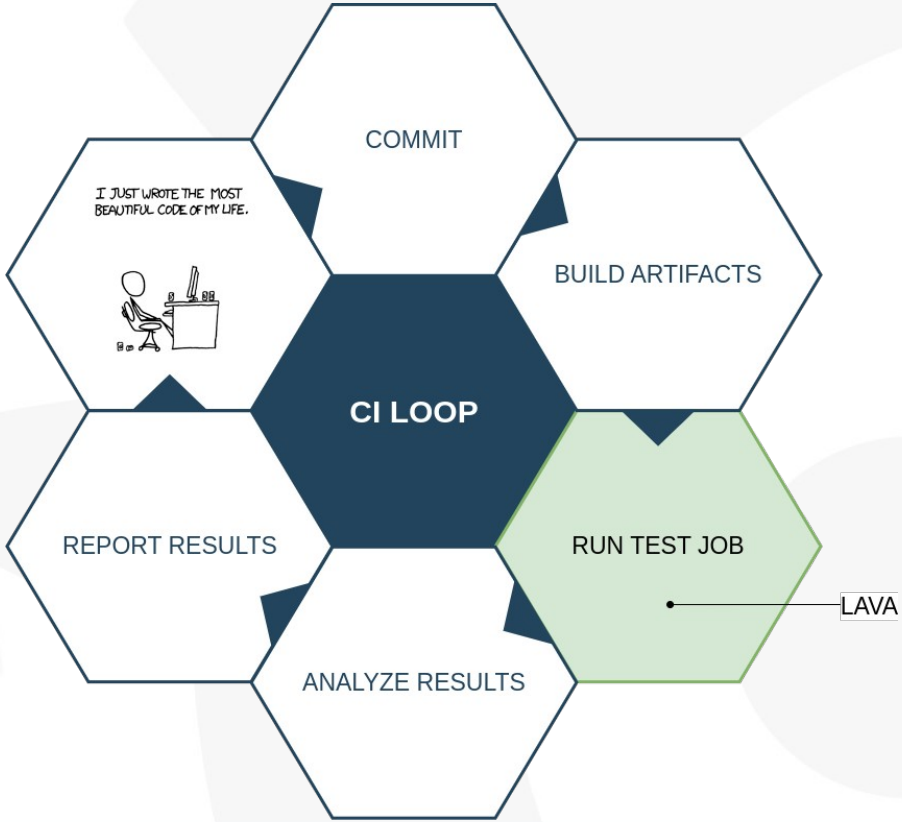
COMMIT
I JUST WROTE THE MOST BEAUTIFUL CODE OF MY LIFE.
BUILD ARTIFACTS
CI LOOP
REPORT RESULTS
RUN TEST JOB
LAVA
ANALYZE RESULTS

Devices in LAVA

- Requirements
  - Ability to be turned on/off remotely
  - Ability to access a reliable console remotely
  - Ability to boot arbitrary Kernel+(DTB)+System remotely
- Configuration (jinja2, YAML)
  - Device type template
  - Device dictionary
  - Health check

Collabora Lab

https://lava.collabora.dev/

- 16 racks
- 217 devices of 38 different types
- 16 LAVA dispatchers
- Network switches, debug interfaces, USB hubs, power supplies, tons of cables

Collabora Lab Stats

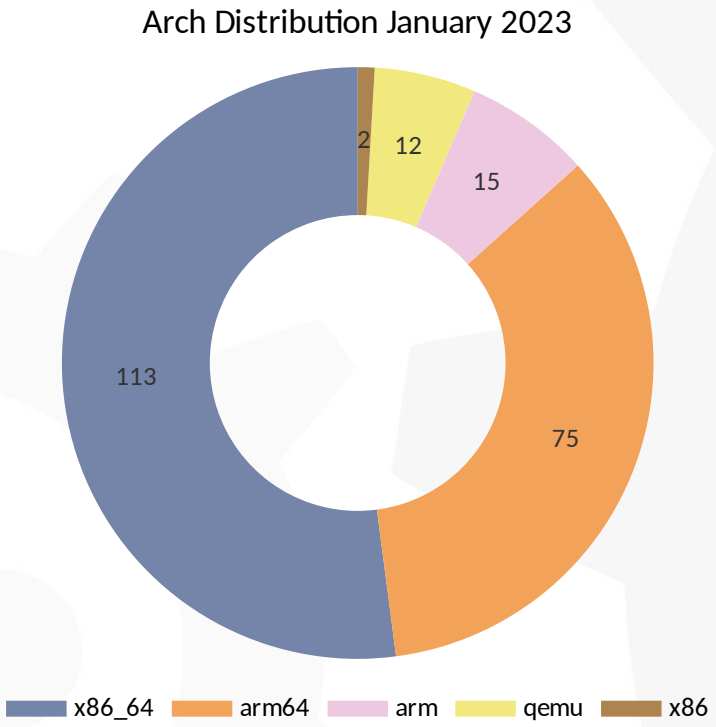

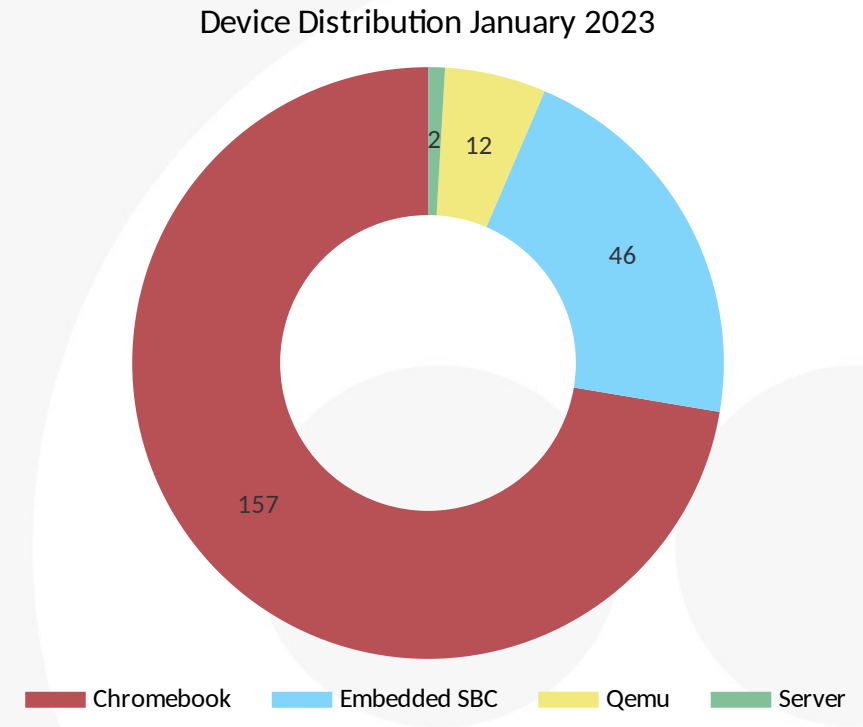

HW for Automation

- Embedded SBCs
  - Ethernet relay (e.g. Devantech ETH008, ETH8020, ETH484)
  - Ethernet power switch (e.g. Aviosys IP9850, Energenie EGPMS)
- Chromebooks
  - SuzyQable + USB Eth adapter
  - ServoV4
- Servers
  - IPMI
- All devices
  - USB regular/switchable HUB (e.g. Ykush)
  - USB cables

SW for Automation

- Power+serial control
  - PDUDaemon https://github.com/pdudaemon/pdudaemon
  - Conserver https://www.conserver.com
  - Hdctools https://chromium.googlesource.com/chromiumos/third_party/hdctools/
- Interaction w/ LAVA
  - Lavacli https://docs.lavasoftware.org/lava/lavacli.html
  - LAVA Gitlab Runner https://github.com/collabora/lava-gitlab-runner

KernelCI

- Linux Kernel continuous testing https://kernelci.org/
  - Baseline tests (e.g. bootrr, dmesg)
  - Boot tests (e.g. boot-nfs, boot-fastboot)
  - Subsystem tests (e.g. igt, v4l2-compliance)
  - Userspace tests (e.g. chromeos tast)
- Post-merge testing
  - Build reports
  - Regression reports

MesaCI

- Mesa conformance testing and performance tracking
  https://gitlab.freedesktop.org/mesa/mesa/-/pipelines
  - API
    - OpenGL, OpenGL ES, VA-API, Vulkan
  - Drivers
    - Iris, ANV, RadeonSI, RADV, Panfrost, Panvk, Freedreno, Turnip, LLVMPipe, Lavapipe, Softpipe, Etnaviv, Lima, v3d, v4c, Dozen, Virgl, Venus, Nouveau, Crocus
  - Test suites
    - dEQP, Khronos GL and VK CTS, Piglit, trace replaying for OpenGL, Vulkan and Direct3D, Skqp, va-utils
- Pre-merge conformance tests
- Post-merge performance tests

Collabora Lab Stats

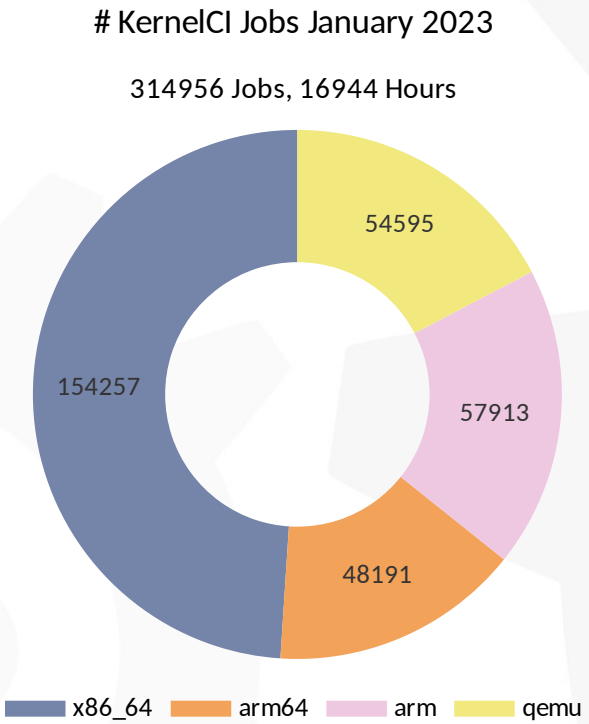

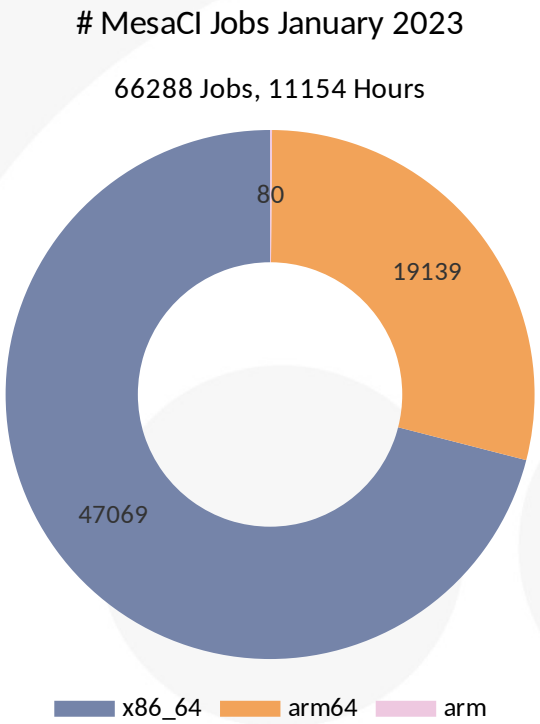

Upstream Testing

- Pre-merge tests
  - User's MR blocked until all tests succeed
  - Risk of pipeline starvation
  - Risk of wrongly rejecting MR
- Post-merge tests
  - Rely on job results to report regressions
  - Risk of detecting false regressions

Common Issues

- HW degradation
  - Faulty cables, dead battery/power supply, dead SD card
- Network
  - Connectivity issues, IP address mismatches
- Rack setup
  - Cable positioning, lid angle, overheating
- FW bugs

Pitfalls

- Serial console silence
  - Hard to determine the cause after the kernel has booted

```
Pass: 7, Duration: 2, Remaining: 11:15
Pass: 37, Duration: 4, Remaining: 4:17
Pass: 67, Duration: 6, Remaining: 3:30
Pass: 98, Duration: 8, Remaining: 3:08
Pass: 128, Duration: 10, Remaining: 2:57
Pass: 155, Duration: 12, Remaining: 2:53
Pass: 176, Duration: 14, Remaining: 2:58
Pass: 190, Skip: 1, Duration: 16, Remaining: 3:05
Pass: 202, Skip: 1, Duration: 18, Remaining: 3:16
Marking unfinished test run as failed
case: 0_mesa
case_id: 303629399
definition: lava
duration: 371.61
result: fail
uuid: 8849549_1.3.2.3.1
The job was canceled
```

Pitfalls

- Serial connection shared between kernel and test shell
  - Possible interference between LAVA signals, kernel and test output

```
/lava-8850597/bin/lava-test-runner /lava-8850597/0
+ export TESTRUN_ID=0_v4l2-compliance-uvc
+ cd /lava-8850597/0/tests/0_v4l2-compliance-uvc
+ cat<4>[   24.613693] mtk_vpu 10020000.vpu: Direct firmware load for mediatek/mt8173/vpu_p.bin failed with error -2
 uuid
+ UUID=88<6>[   24.624411] mtk_vpu 10020000.vpu: Failed to load mediatek/mt8173/vpu_p.bin, -2, retry
50597_1.5.2.3.1
<4>[   24.633804] mtk_vpu 10020000.vpu: Direct firmware load for vpu_p.bin failed with error -2
+ set +x
<LAVA<3>[   24.643670] mtk_vpu 10020000.vpu: Failed to load vpu_p.bin, -2
_SIGNAL_STARTRUN<3>[   24.651139] mtk_vpu 10020000.vpu: Failed to request vpu_p.bin, -2
 0_v4l2-complian<3>[   24.658858] mtk-mdp 14001000.rdma: vpu_load_firmware failed -2
ce-uvc 8850597_1<4>[   24.666678] mtk_vpu 10020000.vpu: Direct firmware load for mediatek/mt8173/vpu_p.bin failed with error -2
.5.2.3.1>
```

```
Unknown test uuid. The STARTRUN signal for this test action was not received correctly.
Invalid TESTCASE signal
```

Pitfalls

- Network drop during bootloader phase
  - Needs to be marked as infrastructure error

```
Waiting for the transfer...
00000000 ################################################################
00080000 ################################################################
00100000 ################################################################
00180000 ################################################################
00200000 ################################################################
00280000 ################################################################R8152: Bulk read error 0xfffffffc
Receive failed.
R8152: Bulk read error 0xfffffffc
Receive failed.
R8152: Bulk read error 0xfffffffc
matched a bootloader error message: '(Bulk read error(.*)){3}' (9)
end: 2.2.4 bootloader-commands (duration 00:00:06) [common]
case: bootloader-commands
case_id: 303466122
definition: lava
duration: 6.01
extra: ...
level: 2.2.4
namespace: common
result: fail
depthcharge-retry failed: 1 of 3 attempts. 'matched a bootloader error message: '(Bulk read error(.*)){3}' (9)'
```

Pitfalls

- Network drop during test phase
  - Critical on tests that rely on a NFS rootfs

```
[  244.548249] usb 2-3: USB disconnect, device number 2
[  244.554103] xhci_hcd 0000:00:14.0: WARN Set TR Deq Ptr cmd failed due to incorrect slot or ep state.
[  244.799342] usb 2-3: new SuperSpeed USB device number 3 using xhci_hcd
[  244.820061] usb 2-3: New USB device found, idVendor=0bda, idProduct=8153, bcdDevice=30.00
[  244.829229] usb 2-3: New USB device strings: Mfr=1, Product=2, SerialNumber=6
[  244.837216] usb 2-3: Product: USB 10/100/1000 LAN
[  244.842482] usb 2-3: Manufacturer: Realtek
[  244.847069] usb 2-3: SerialNumber: 000001
[  245.009642] usb 2-3: reset SuperSpeed USB device number 3 using xhci_hcd
```

```
[  427.097864] nfs: server 192.168.201.1 not responding, still trying
[  427.104850] nfs: server 192.168.201.1 not responding, still trying
[  427.111838] nfs: server 192.168.201.1 not responding, still trying
[  433.104998] nfs: server 192.168.201.1 not responding, still trying
```

Best Practices

- Write robust health checks
  - Ensure faulty devices are taken down automatically
- Monitor LAVA Infrastructure Error exceptions
  - Spot issues with specific racks or device types
- Ensure device redundancy
  - Monitor the device's health and job queue
- Isolate test shell output and kernel messages when possible
  - e.g. docker container + SSH connection to the DUT

Next Steps

- Keep adding new devices
  - Increase lab capacity + cover variety of platforms from different vendors
- Keep improving the infrastructure and our monitoring tools
- Increase the coverage of test suites

Thank you!

We are hiring - col.la/careers