

Testing a large testing software

Rémi Duraffort, Linaro Ltd.
remi.duraffort@linaro.org



Who am I?

- Rémi Duraffort
- Senior Software Engineer at Linaro
- LAVA Architect
- OSS developer since 2007
 - VLC media player
 - v8 js engine
 - PRoot/CARE
 - LAVA, lavacli, meta-lava, DummySys, lavafed, ...

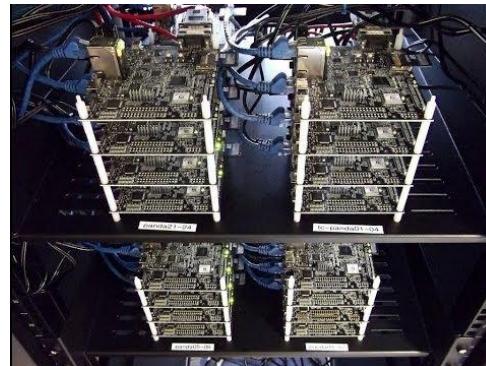


LAVA

A brief introduction

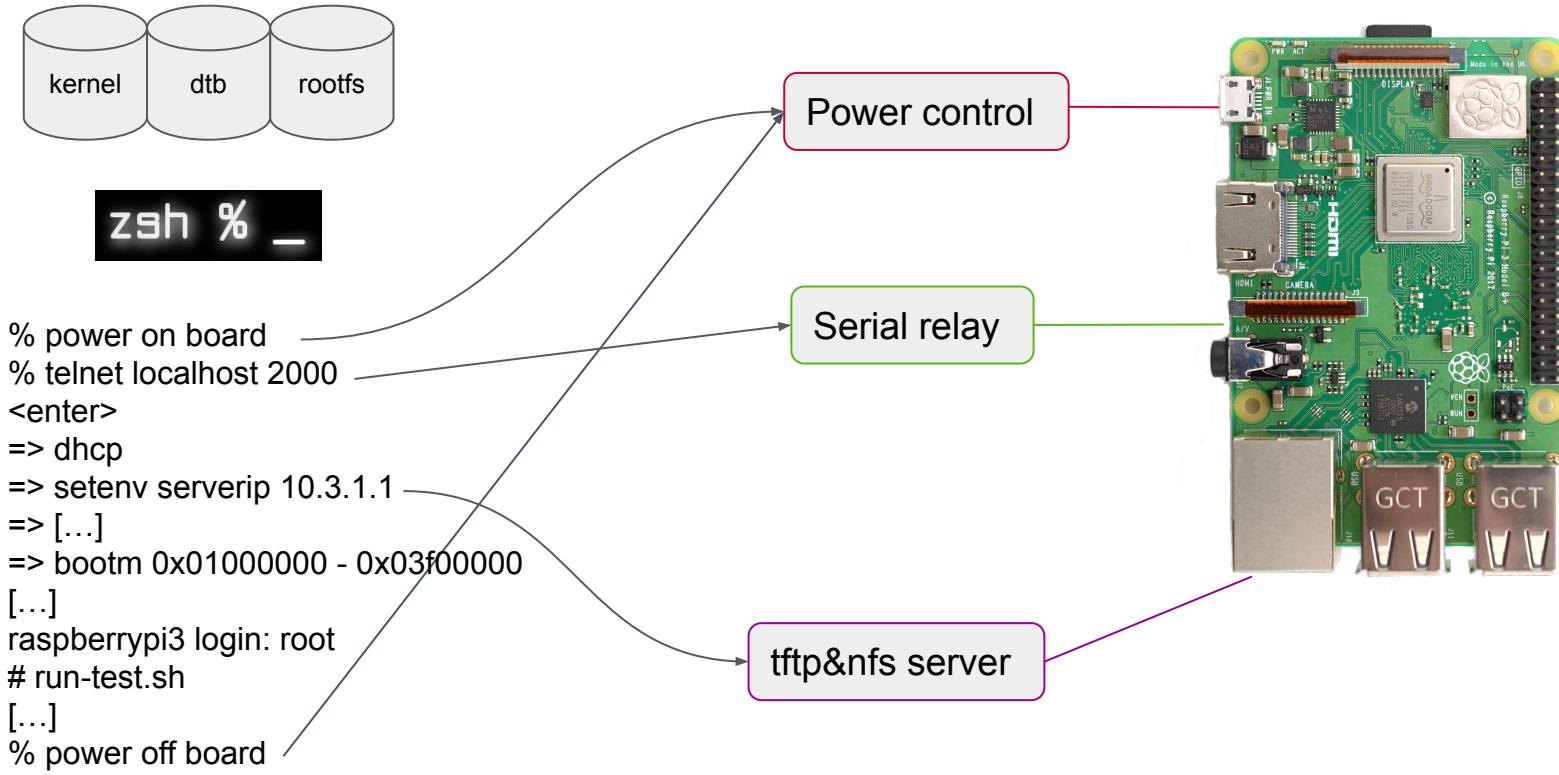
LAVA

- Linaro Automated Validation Architecture
- Test execution system: **testing software on real hardware**
 - Deploy, Boot and Test
- Usages
 - Boot testing: kernelci
 - System level testing: lkft
 - Power consumption
 - Benchmarks
 - Multinode
 - Test with many devices
 - ...

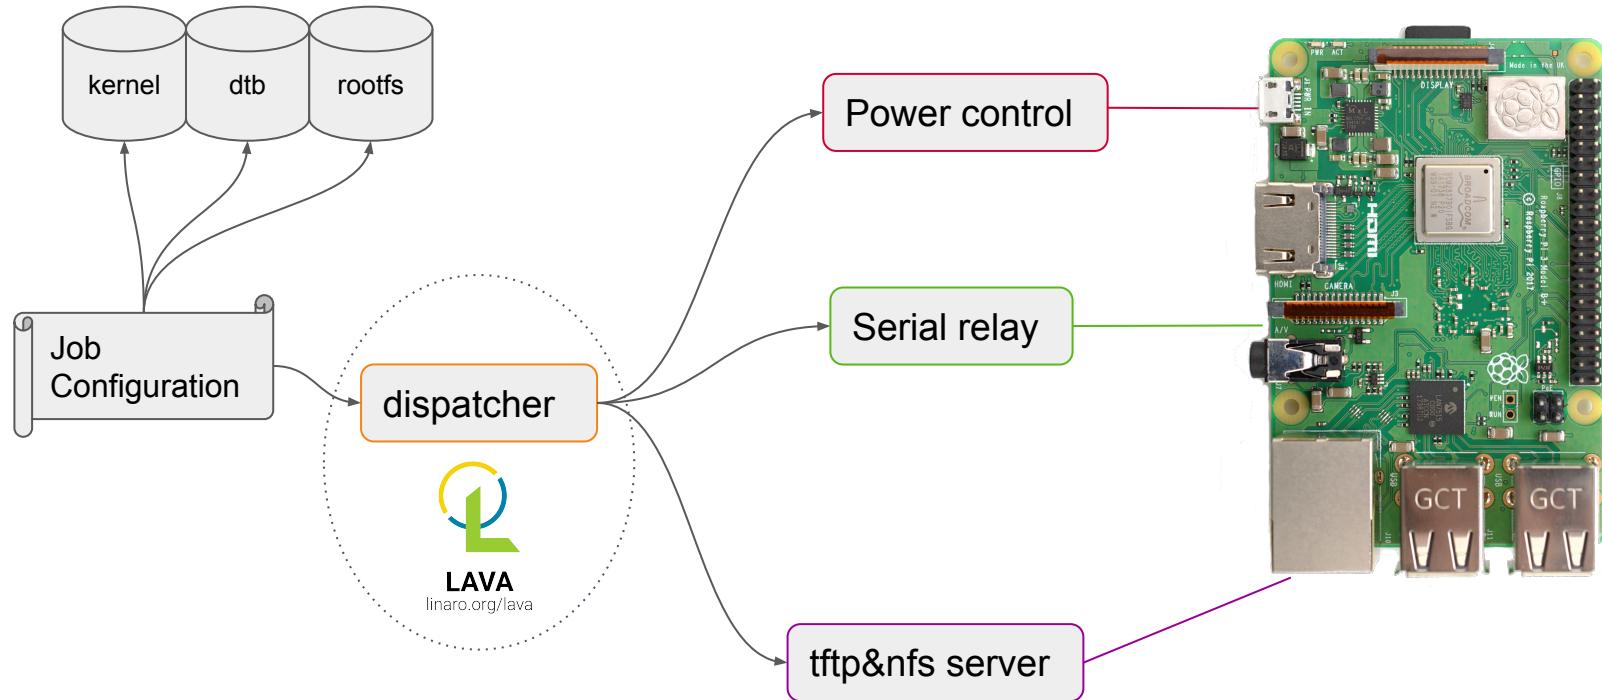


LAVA
linaro.org/lava

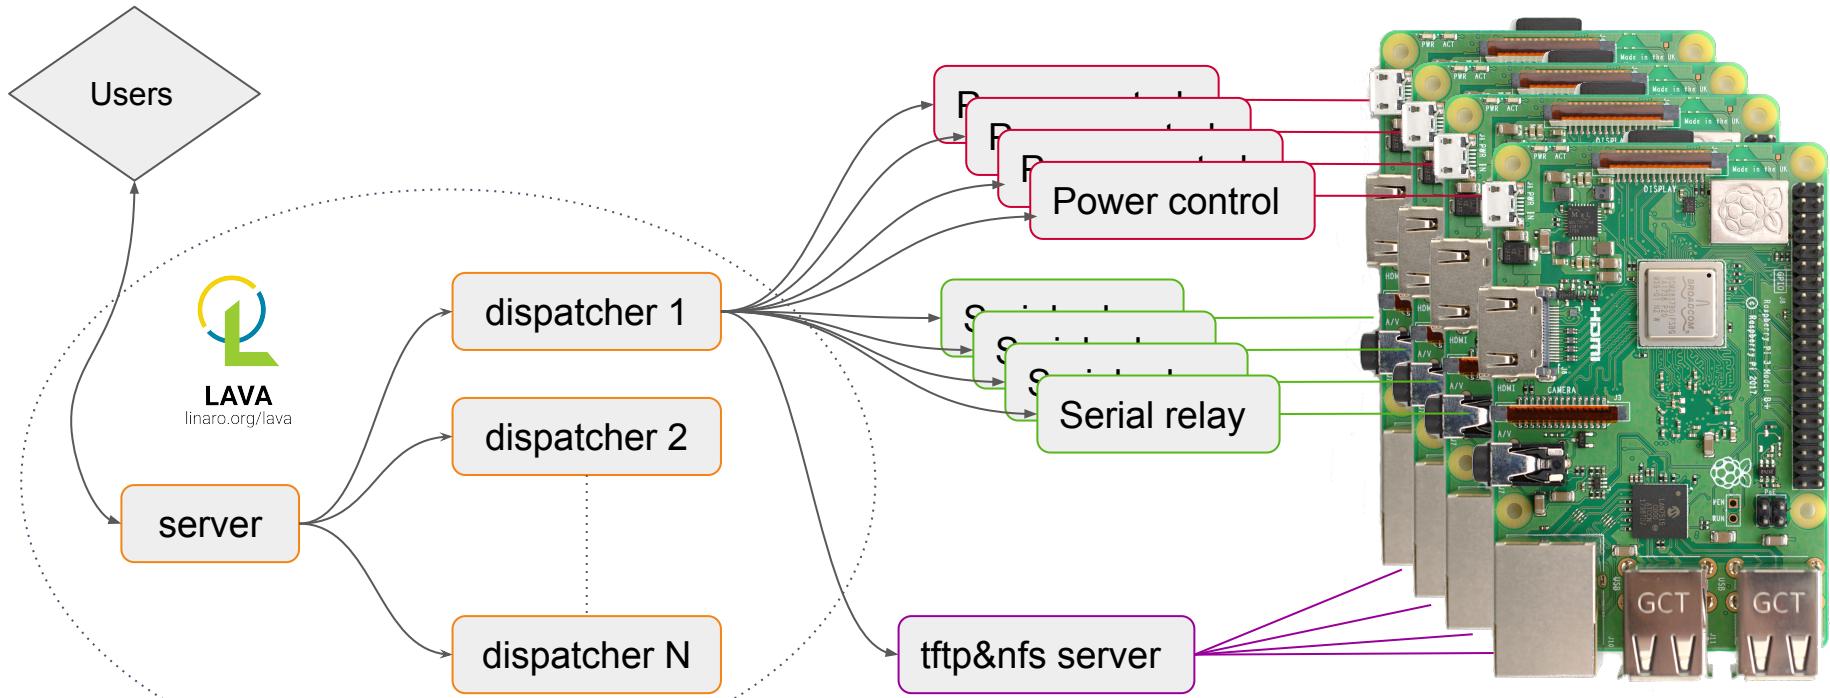
Testing without LAVA



Testing with LAVA



LAVA architecture



LAVA roles

Server

- Web UI and API
 - Submit jobs
 - Results, logs, ...
- Access control
 - Users, groups
- Scheduling jobs
 - Priority
 - Multinode jobs
- Store job logs
- Send notifications

Dispatcher

- Deploy resources
- Power on/off DUTs
- Send commands
- Parse logs
 - Kernel panic
 - Bootloader error
- Classify errors
 - Infrastructure, Bug
 - Job, Test, ...

Supported methods

deploy:

- tftp
- nbd
- flasher
- fastboot
- mps
- ssh
- docker
- vexpress ums
- ...

boot:

- cmsis
- dfu
- uboot
- pyocd
- fastboot
- docker
- qemu
- grub
- iso installer
- ...

test:

- git repository
- interactive
- minimal
- multinode

Supported device-types: 197

adb-nuc alpine-db am6 apq8016-sbc-uboot ar9331-dpt-module arduino101 armada-370-db armada-370-rd armada-3720-db
armada-3720-espressobin armada-375-db armada-385-db-ap armada-388-clearfog armada-388-gp armada-398-db armada-7040-db
armada-8040-db armada-xp-db armada-xp-gp armada-xp-linksys-mamba armada-xp-openblocks-ax3-4 arndale at91rm9200ek at91sam9261ek
at91sam9m10g45ek at91sam9x25ek at91sam9x35ek at91-sama5d2_xplained at91-sama5d4_xplained b2120h410 b2260 base bcm2836-rpi-2-b
bcm2837-rpi-3-b-32 bcm2837-rpi-3-b beaglebone-black-barebox beaglebone-black beagle-xm cc13x2-launchpad cc3220SF cubietruck d02 d03
da850-lcdk disco-l475-iot1 docker dove-cubox dra7-evm dragonboard-410c dragonboard-820c dragonboard-845c frdm-k64f frdm-kw41z
fsl-imx8mm-evk fsl-imx8mq-evk hi6220-hikey-bl hi6220-hikey-bl hi6220-hikey-r2 hi960-hikey hifive-unleashed-a00 highbank hip07-d05 hsdk ifc6410
imx23-olinuxino imx27-phYTEC-phycard-s-rdk imx28-duckbill imx53-qsrB imx6dl-riotboard imx6q-nitrogen6x imx6q-sabrelite imx6ul-pico-hobbit
imx7s-warp imx8m imx8mn-ddr4-evk jetson-tk1 juno juno-uboot juno-uefi kirkwood-db-88f6282 kirkwood-openblocks_a7 kvm lava-slave-docker lxc
mediatek-8173 meson8b-ec100 meson8b-odroidc1 meson-g12a-sei510 meson-g12a-u200 meson-g12a-x96-max meson-g12b-a311d-khadas-vim3
meson-g12b-odroid-n2 meson-gxbb-nanopi-k2 meson-gxbb-p200 meson-gxl-s805x-libretech-ac meson-gxl-s805x-p241 meson-gxl-s905d-p230
meson-gxl-s905x-khadas-vim meson-gxl-s905x-libretech-cc meson-gxl-s905x-p212 meson-gxm-khadas-vim2 meson-gxm-q200 meson-sm1-sei610
mimxrt1050_evk minnowboard-max-E3825 minnowboard-turbo-E3826 moonshot-m400 mps mustang-grub-efi mustang mustang-uefi nexus10 nexus4
nexus5x nexus9 nrf52-nitrogen nucleo-l476rg nxp-ls2088 odroid-n2 odroid-x2 odroid-xu3 orion5x-rd88f5182-nas overdrive
ox820-cloudengines-pogoplug-series-3 panda peach-pi pixel poplar qcom-qdf2400 qcs404-evb-1k qcs404-evb-4k qemu-aarch64 qemu r8a7791-porter
r8a7795-salvator-x r8a7796-m3ulcb r8a7796-m3ulcb-kf rk3288-rock2-square rk3288-veyron-jaq rk3328-rock64 rk3399-gru-kevin rk3399-puma-haikou
rzn1d sama53d sama5d34ek sama5d36ek sdm845-mtp sharkl2 snow soca9 socfga-cyclone5-socrates ssh stm32-carbon stm32mp157c-dk2
sun4i-a10-olinuxino-lime sun50i-a64-bananapi-m64 sun50i-a64-pine64-plus sun50i-h5-libretech-all-h3-cc sun50i-h6-orangepi-3
sun50i-h6-orangepi-one-plus sun50i-h6-pine-h64 sun50i-h6-pine-h64-model-b sun5i-a13-olinuxino-micro sun5i-gr8-chip-pro sun5i-r8-chip
sun6i-a31-app4-evb1 sun7i-a20-cubieboard2 sun7i-a20-olinuxino-lime2 sun7i-a20-olinuxino-micro sun8i-a23-evb sun8i-a33-olinuxino
sun8i-a33-sinlinx-sina33 sun8i-a83t-allwinner-h8homlet-v2 sun8i-a83t-bananapi-m3 sun8i-h2-plus-bananapi-m2-zero sun8i-h2-plus-libretech-all-h3-cc
sun8i-h2-plus-orangepi-r1 sun8i-h2-plus-orangepi-zero sun8i-h3-libretech-all-h3-cc sun8i-h3-orangepi-pc sun8i-r40-bananapi-m2-ultra synquacer-acpi
synquacer-dtb synquacer tc2 tegra124-nyan-big thunderx upsquare vexpress x15-bl x15 x86-atom330 x86-celeron x86 x86-pentium4 x86-x5-z8350
xilinx-zcu102

Why testing a testing software?

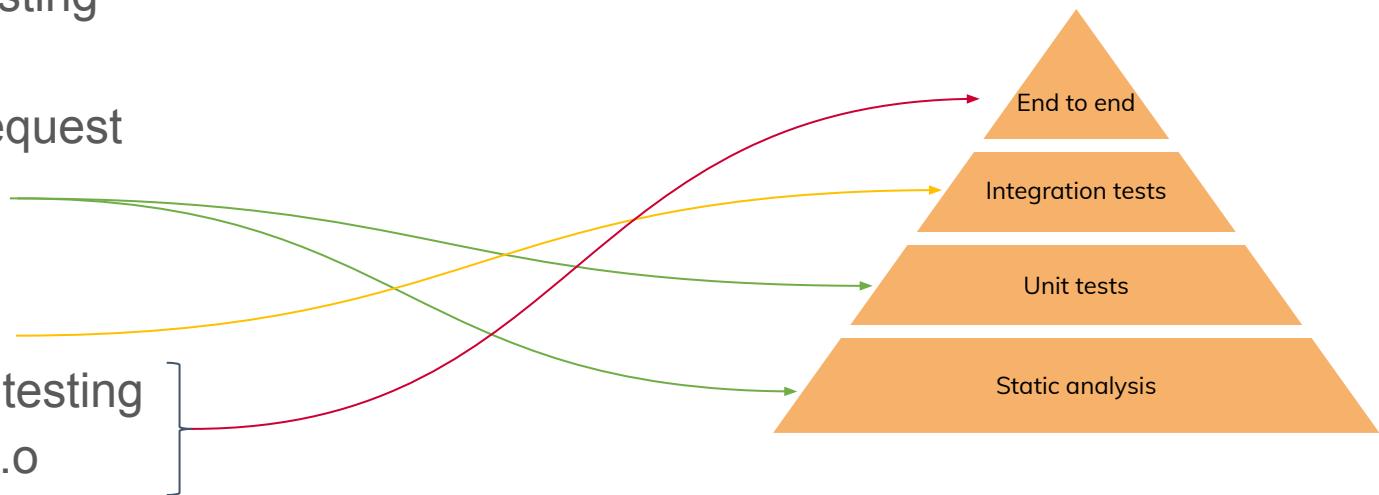
Why testing a testing softwares?

- Should be **reliable**
 - False positives
 - Lose trust in the CI
 - Developers **ignore** the CI
 - False negatives
 - Not reporting errors
 - Shipping **buggy** software
 - **Bugs and regressions**
 - Like every software
 - **Complexe** softwares

Testing LAVA

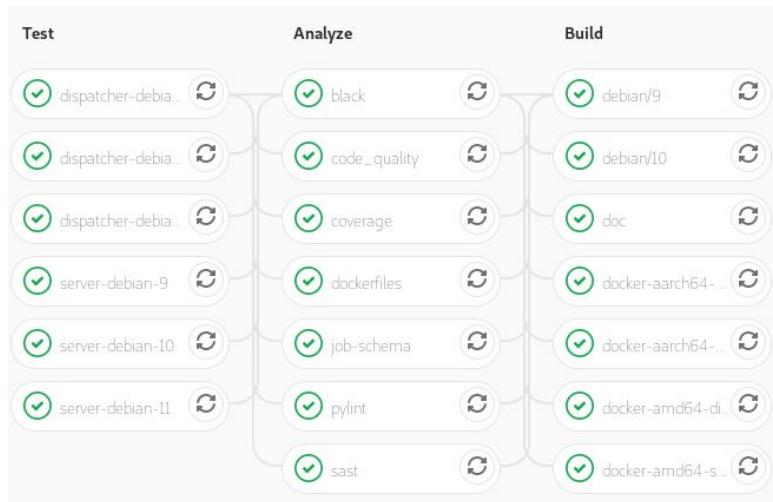
Testing LAVA

- While developing
 - Manual testing
 - ./ci-run
- Each merge request
 - Gitlab CI
- Once a day
 - meta-lava
 - federated testing
 - staging.v.l.o



Testing LAVA: GitLab CI

- Test (unittest)
 - Server: 561 tests
 - Dispatcher: 388 tests
 - On Debian 9, 10 and 11
- Analyze
 - Black
 - Static analysis: pylint, bandit
- Build
 - Debian pkg (9 and 10)
 - Docker images (amd64 and AArch64)
 - Doc

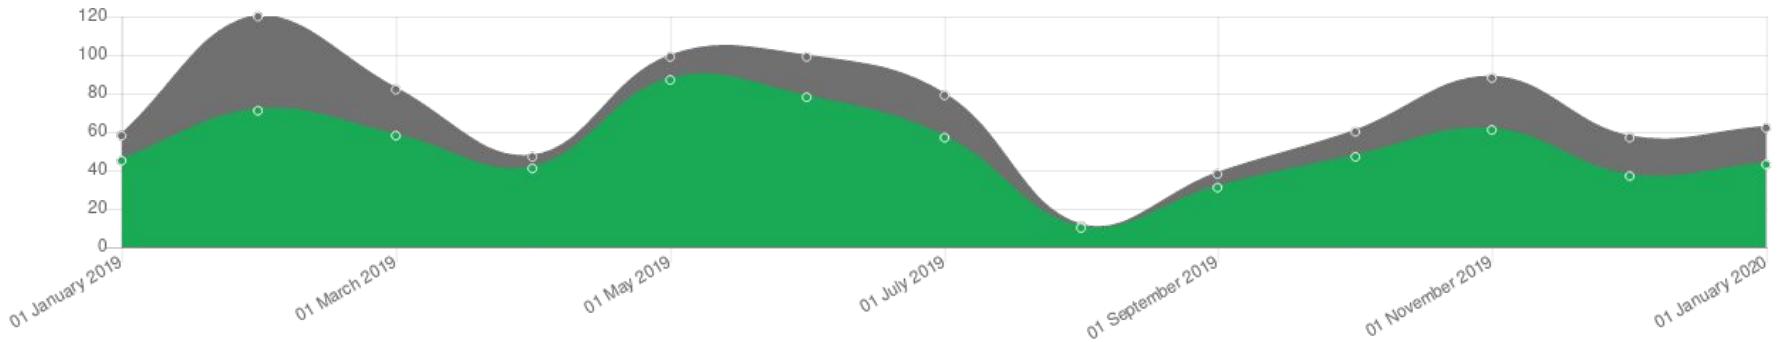


Testing LAVA: GitLab CI

- LAVA server
 - Just a django application
 - Insert data
 - Run some tests
- LAVA dispatcher
 - Board interaction
 - Master-slave protocol
 - Difficult to test
 - Use pytest mocking

Testing LAVA: GitLab CI

- Is it useful
 - Found many issues



- Is it enough?
 - NO!

Testing LAVA

meta-lava

Testing LAVA: a combinatorial issue

197 boards

- ✖ 16 deploy
- ✖ 26 boot
- ✖ 4 test

327 808 combinations*

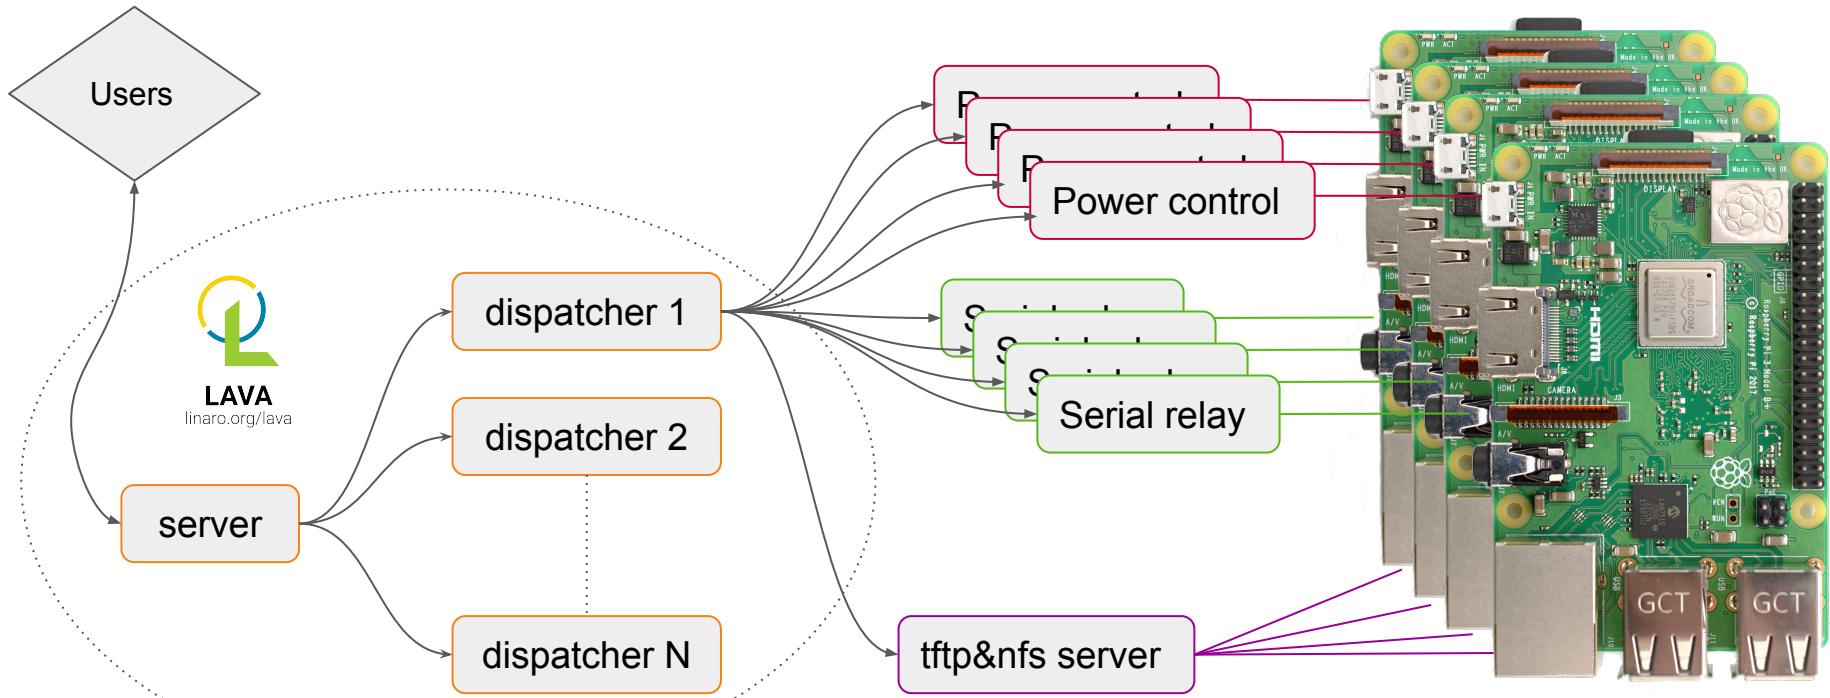
* Most case are useless but that's still huge

Testing LAVA: meta-lava

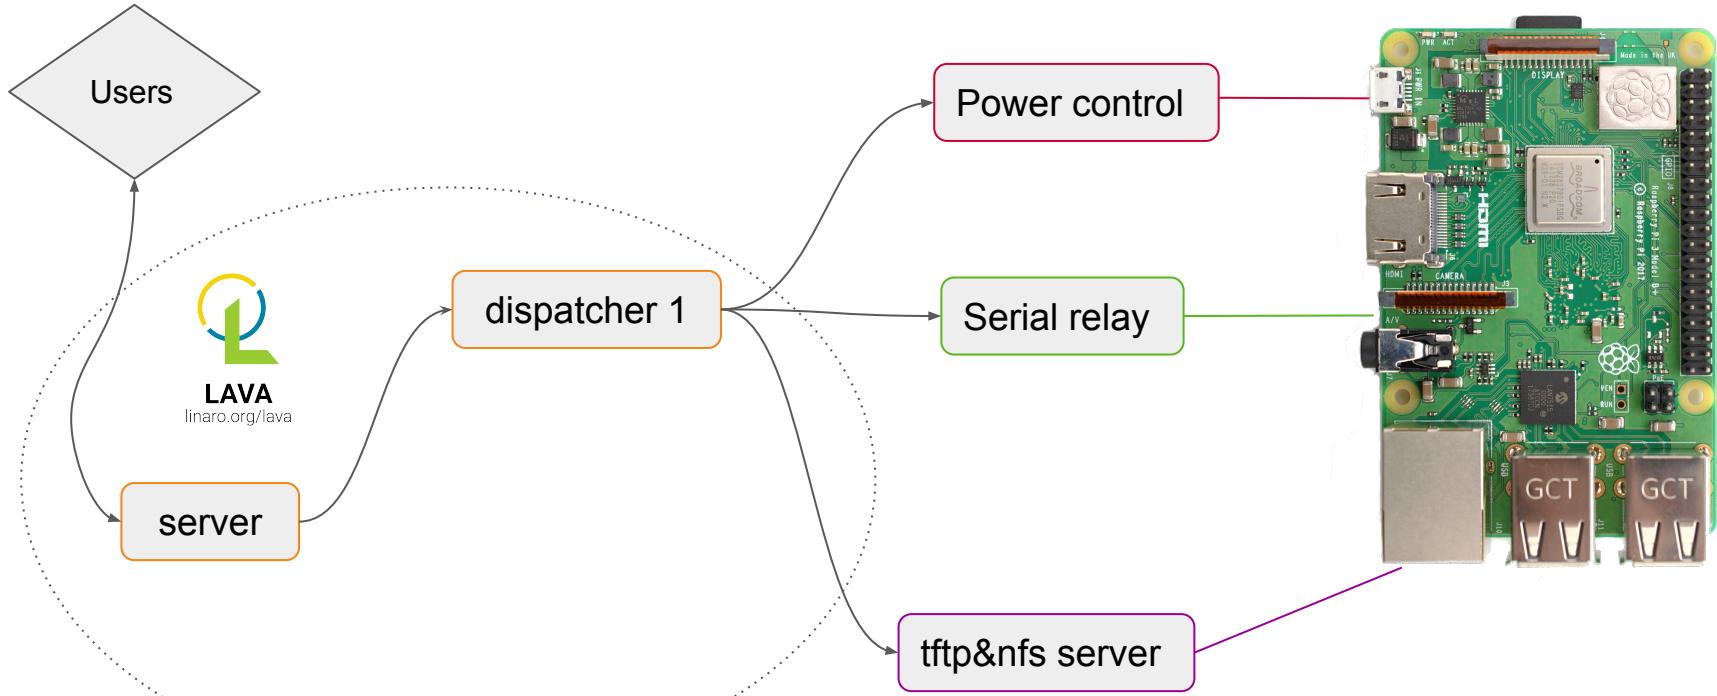
- Goal
 - Testing the full system
 - Including board interaction
 - Without any boards
 - Fast & cheap
- Solutions
 - Board emulation
 - CPU intensive
 - Expensive & slow
 - System mocking



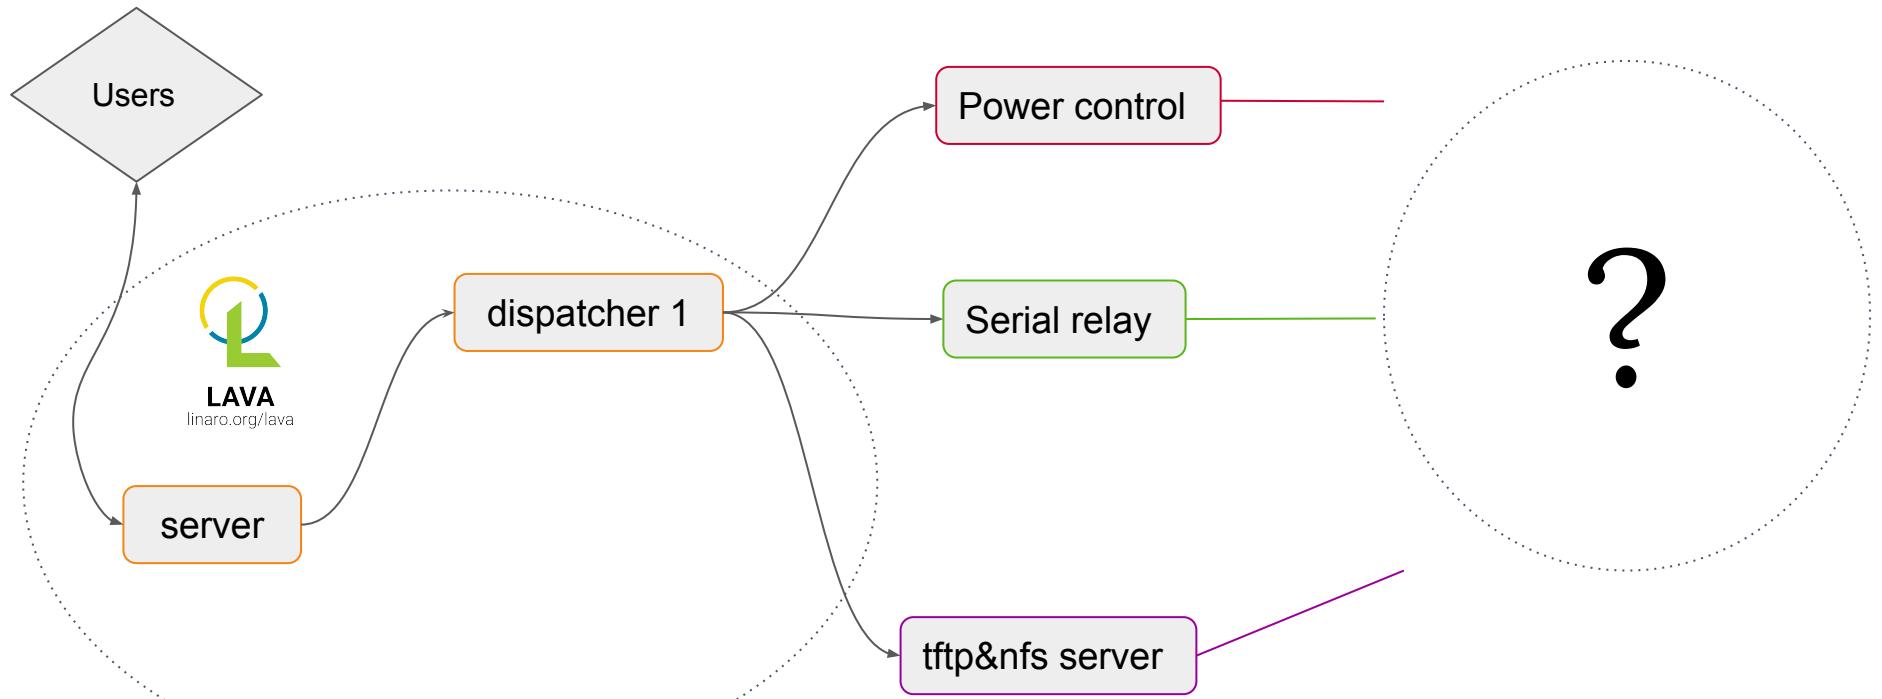
Testing LAVA: system mocking



Testing LAVA: system mocking

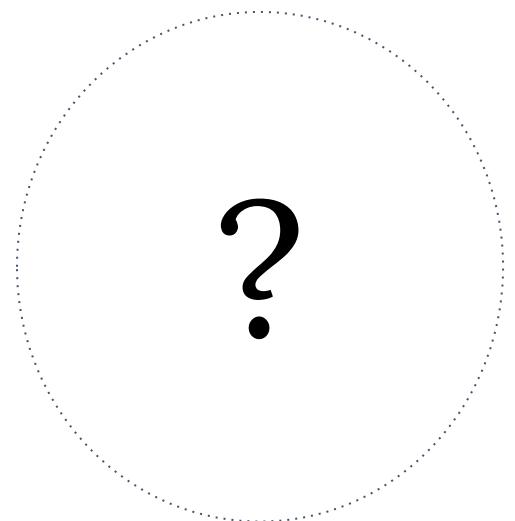


Testing LAVA: system mocking



Testing LAVA: system mocking

- Power control:
 - Just a command line
 - `/bin/true`
- Serial relay:
 - Board output
 - LAVA commands
 - DummySYS
- tftp & nfs:
 - Should check the files
 - tftp and nfs commands



Testing LAVA: system mocking

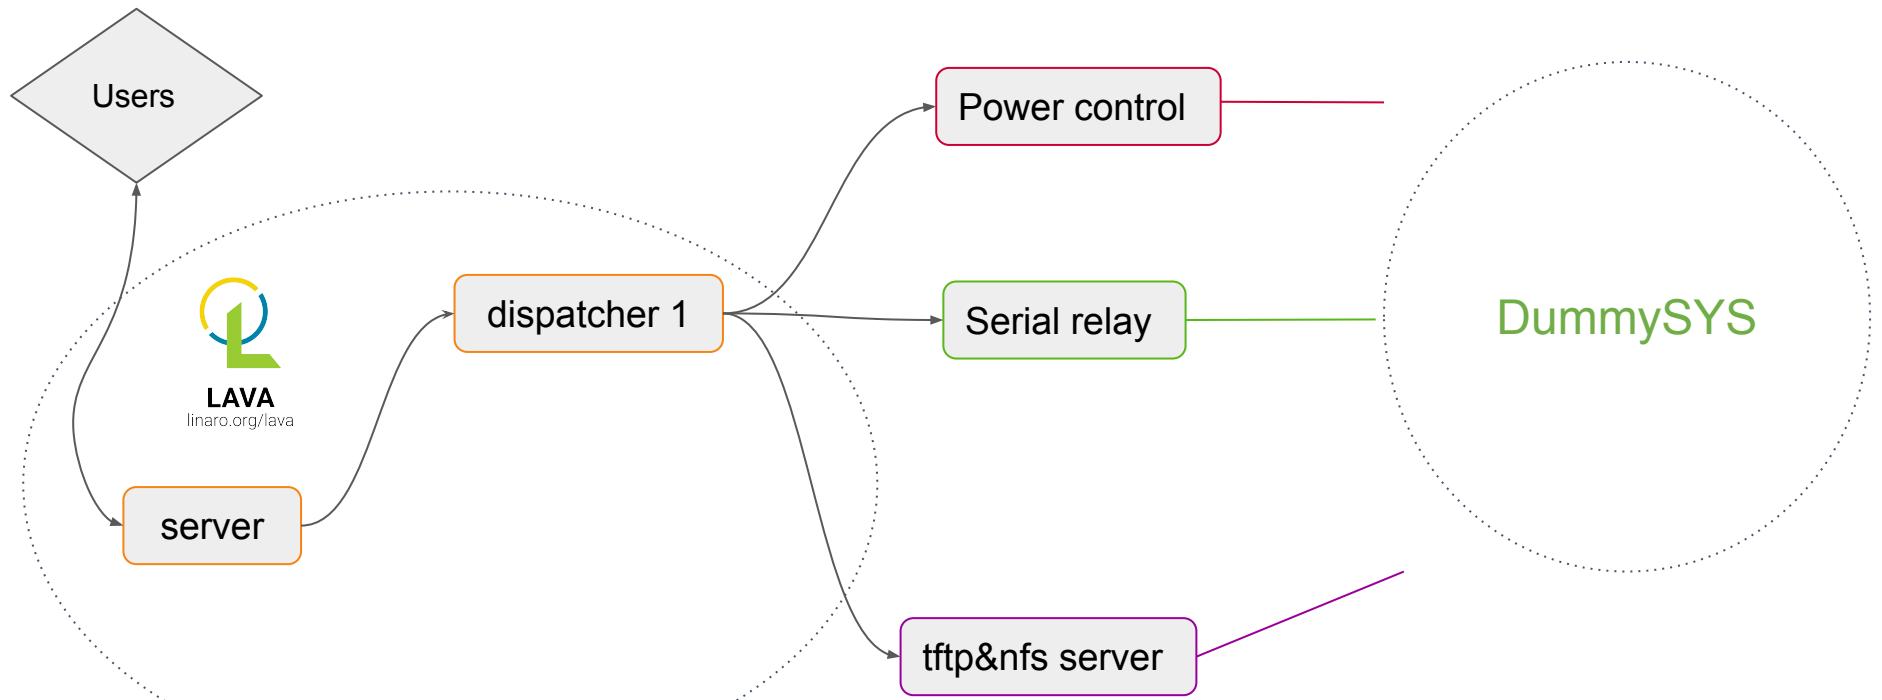
- DummySYS

- Output
 - Like a real board
- Input
 - Expect the right sequence
 - Fail if the sequence is changed
- Use tftp & nfs resources
 - Download kernel/dtb/ramdisk
 - Mount nfs rootfs
 - Checksum some files



DummySYS

Testing LAVA: system mocking



Demo

Testing LAVA: system mocking

- Meta-lava
 - Server docker container
 - Dispatcher docker container
 - With DummySYS
- Testing master every morning
 - 27 board types
 - Including board that I've never seen
 - Testing board failures
 - bootloader errors
 - dhcp failing
 - ...

Testing LAVA

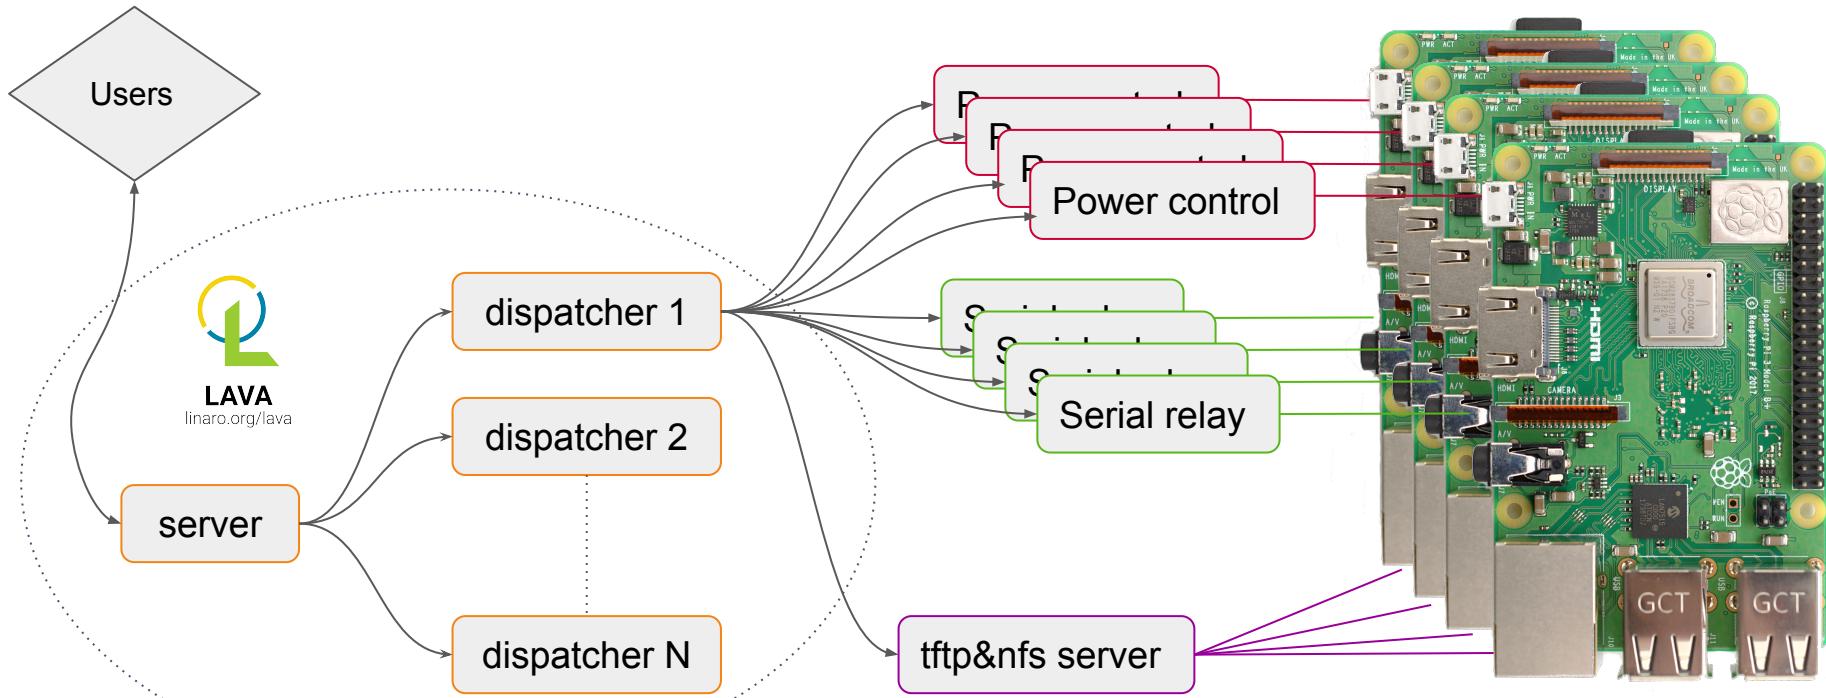
Benchmarks

Running 500 jobs in parallel?

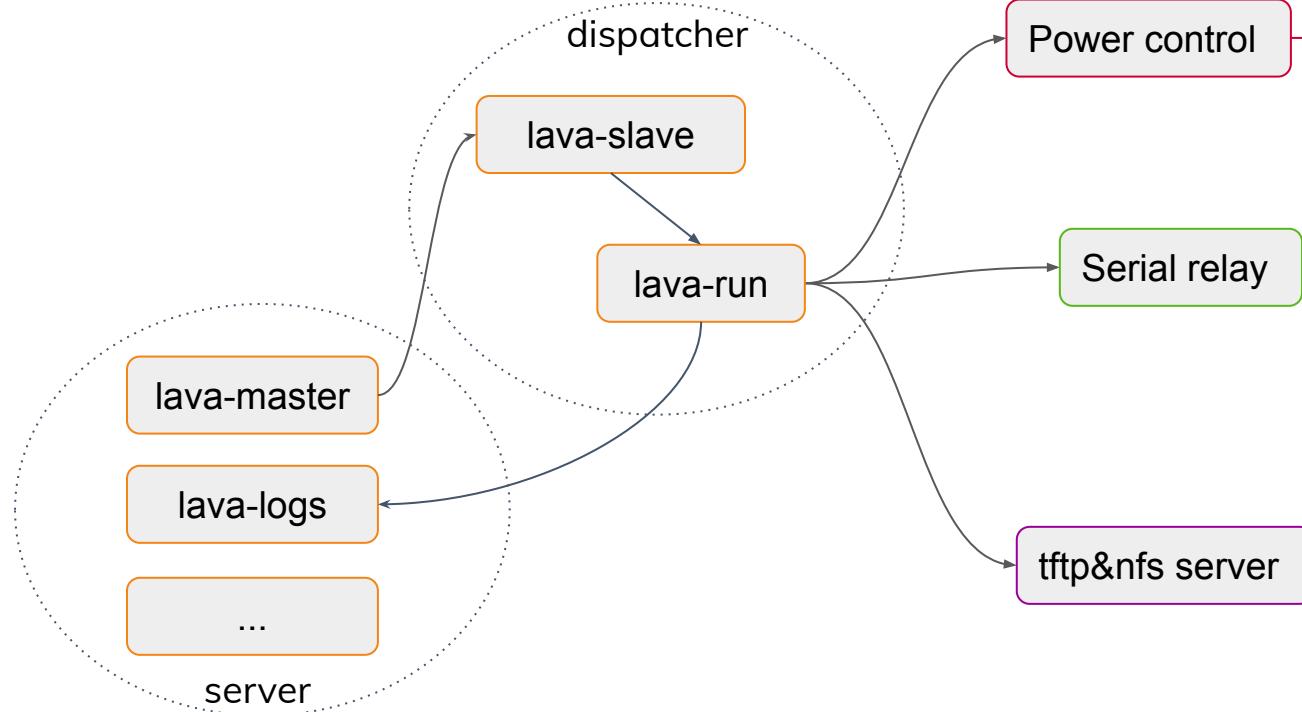
- Hardware:
 - 500 boards
 - ~100 servers
 - Many people to plug everything
 - Some board and server will fail
- Not reliable
- Mock some part of the system
 - Use only one server
 - Mock the right part
 - Keep the test effective



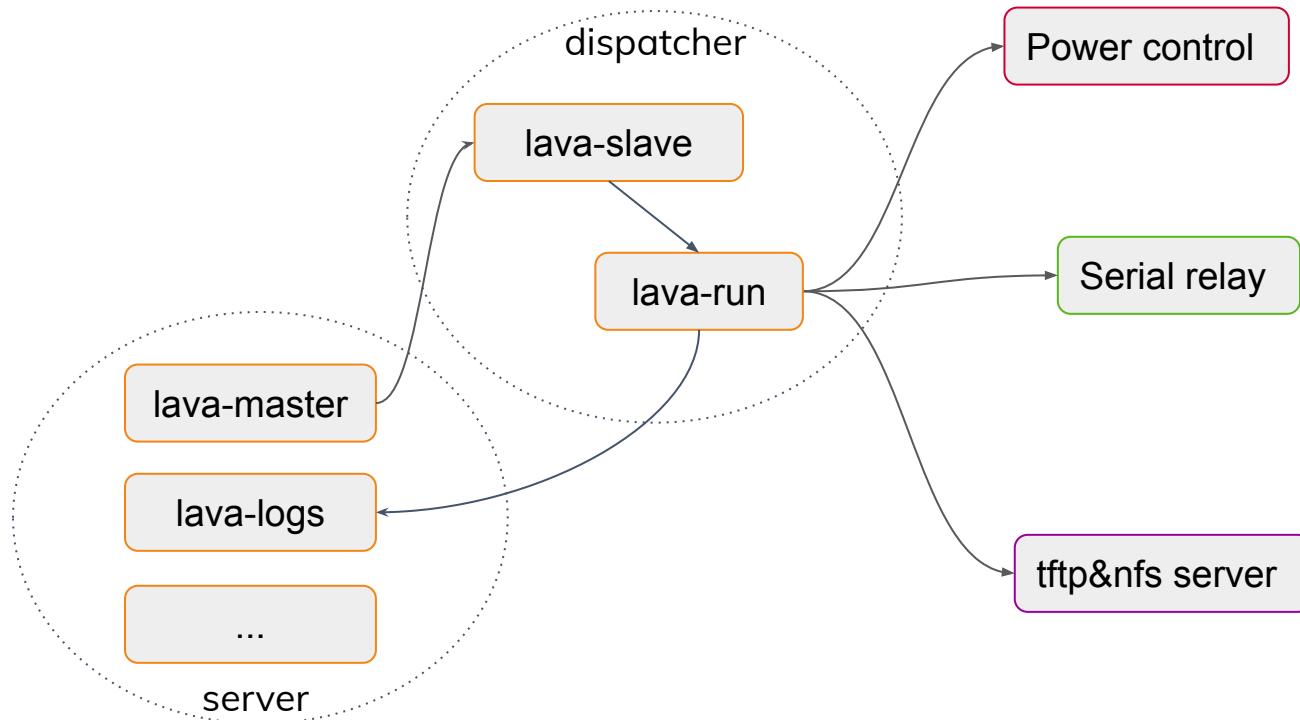
Running 500 jobs in parallel?



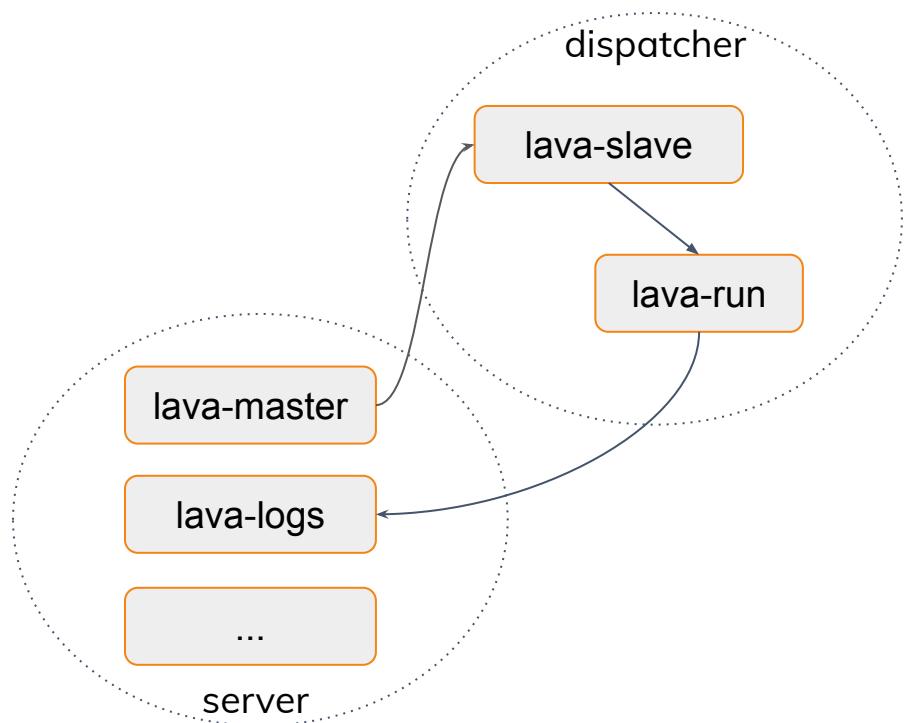
Running 500 jobs in parallel?



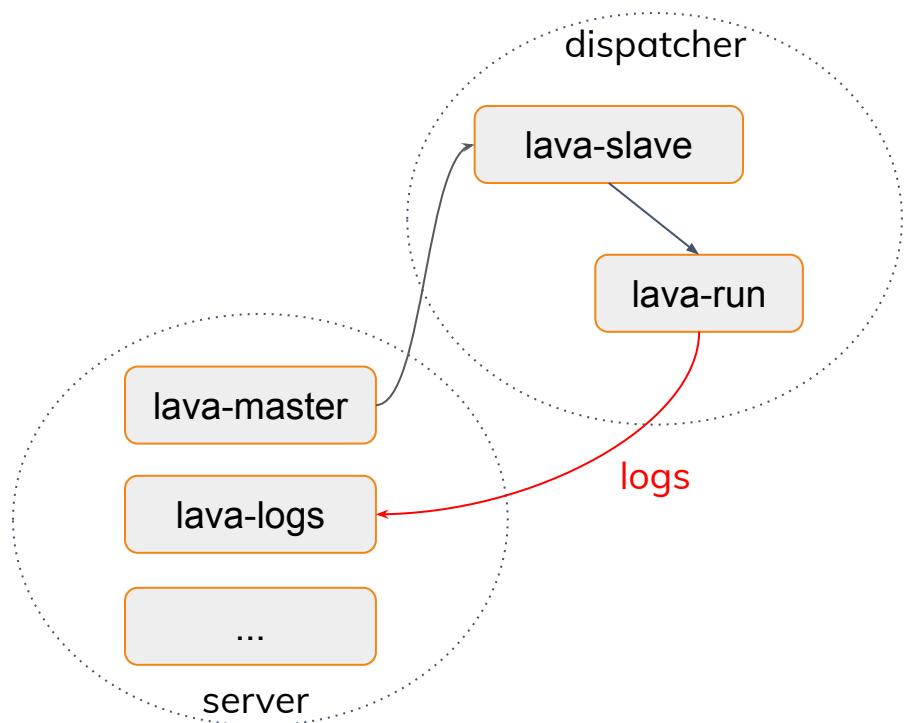
Running 500 jobs in parallel?



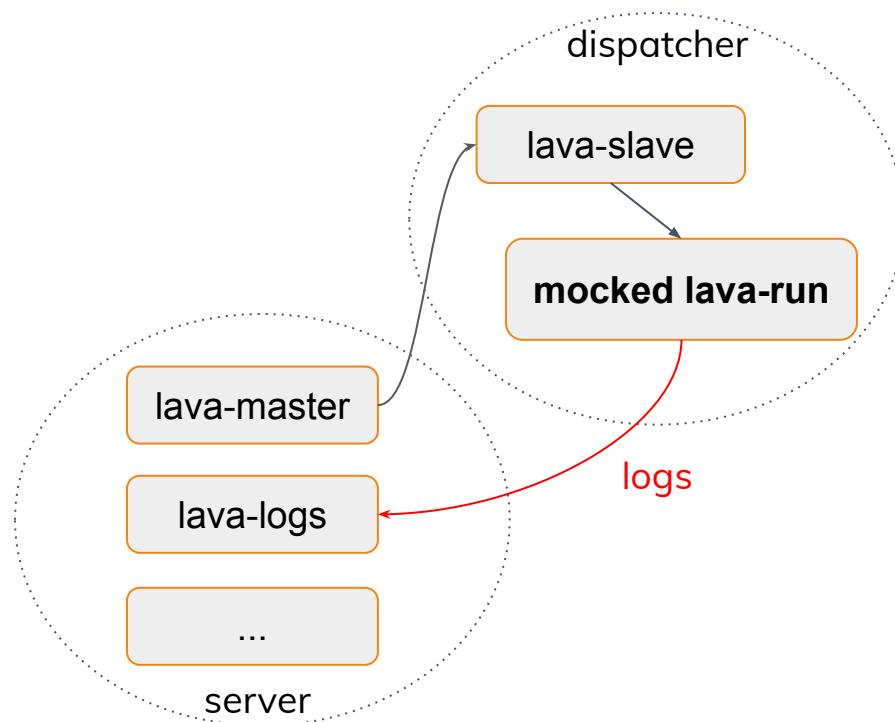
Running 500 jobs in parallel?



Running 500 jobs in parallel?



Running 500 jobs in parallel?



- Looks like lava-run
 - Same command line
 - Handle signals
 - Return value
 - Sending logs
 - Right format
 - Similar speed
- Not using CPU/RAM/IO/...
- Every services running normally

Conclusion

- System mocking is fun
 - Test on fake hardware
 - Even hardware that you don't own
 - Run benchmarks
 - ...
- Can find many bugs that unittest won't
- Not that difficult
 - Be creative!
 - Look at the boundaries

Question?