

# A simple caching service for your CI

Rémi Duraffort  
Principal Tech Lead  
Linaro  
[remi.duraffort@linaro.org](mailto:remi.duraffort@linaro.org)



# Who am I?

- Rémi Duraffort
- Principal Tech Lead at Linaro
  
- OSS developer since 2007
  - VLC media player
  - v8 js engine
  - PRoot/CARE
  - LAVA, lavacli, meta-lava, DummySys, lavafed, ...
  - TuxRun, Tuxsuite, ...
  - Kisscache



# Kisscache

A simple caching service


# LKFT

Linux Kernel Functional Testing

# What is LKFT?

“Improve the Linux kernel quality on the Arm architecture by performing regression testing and reporting on selected Linux kernel branches and the Android Common Kernel (ACK) in real time.”

- Lead by Linaro
- Automated system to build and test a set of linux kernel trees
  - LTS trees
  - mainline
  - next
- 48 hour LTS regression reporting SLA



# LKFT 2023 numbers

## Linux Kernel


- 396 RC
- 2443 revisions
- 1.1M builds (kernels)
- 297M tests

## Android Common Kernel

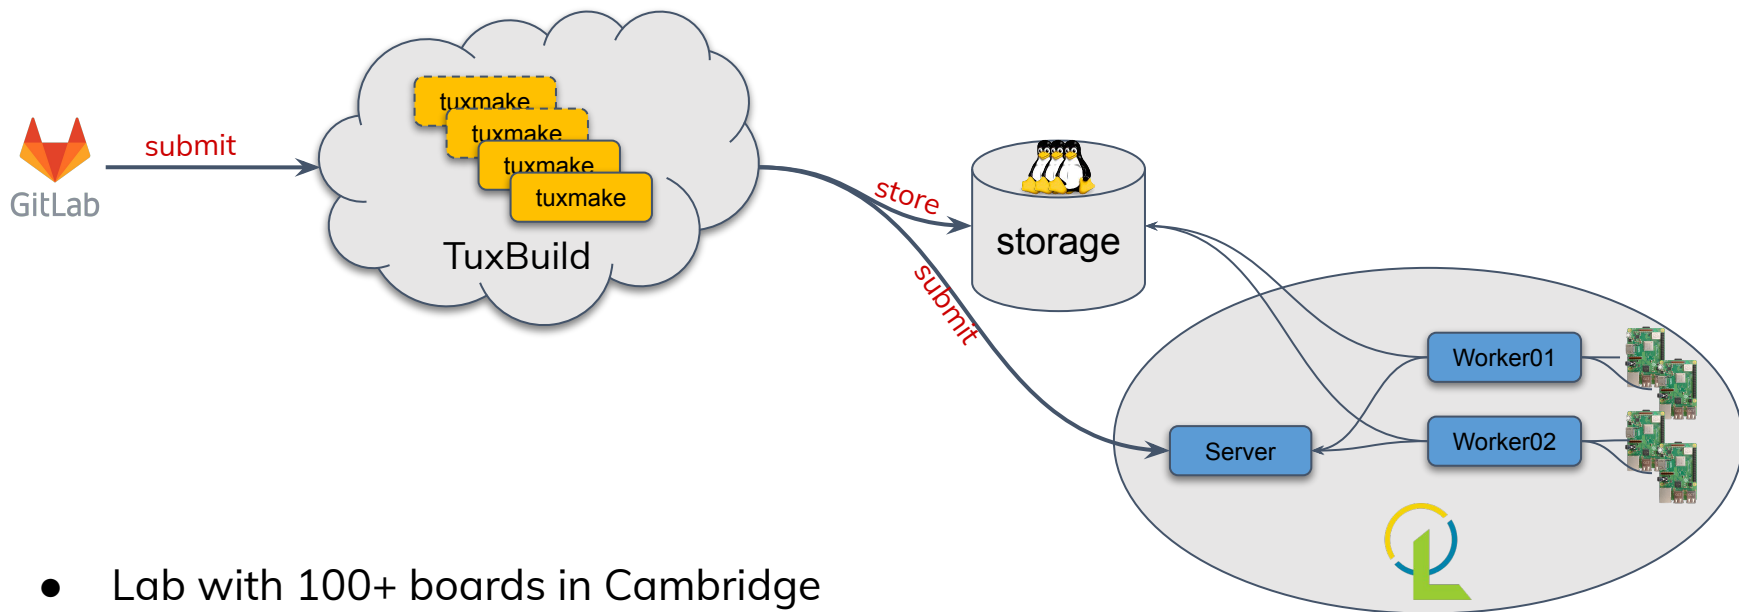
- 580M tests (VTS, CTS, ...)

## Test run on:

- Virtual devices (qemu, FVP) in tuxsuite
- Physical devices in Linaro Lab by LAVA



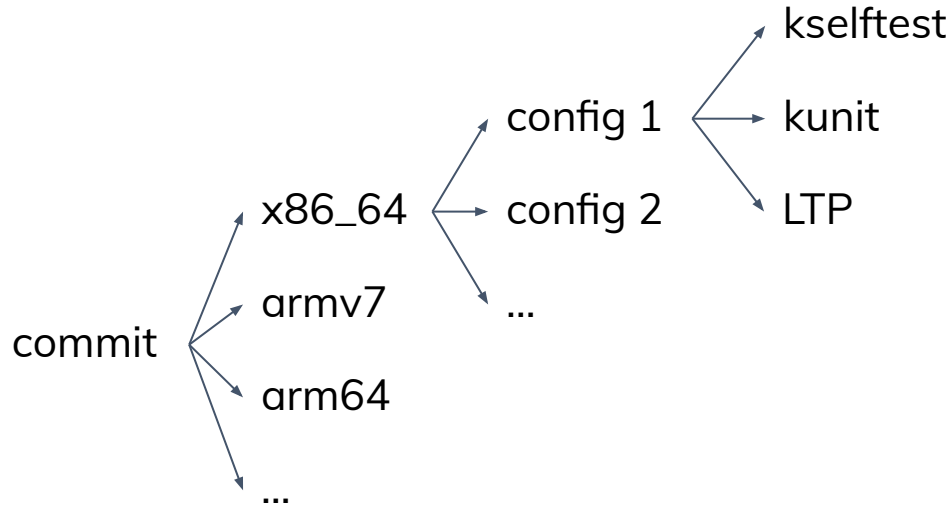
# LKFT (really) simplified architecture



- Lab with 100+ boards in Cambridge
  - RPi, dragonboard, hikey, juno, x15, rockpi4b, ...
- Can all run in parallel
- Downloading from the same storage

# Artifacts duplication

- CI system build and test kernel/dtb/rootfs
- Each kernel/dtb/rootfs is reused multiple times by the CI



- High network usage and a lot of duplication



# Just add caching

- LAVA (CI workers) downloads a lot of artifacts
  - Multiple times
  - In parallel (almost exact same time)
- SQUID should fix this?




# Just add caching

- LAVA downloads a lot of artifacts
  - Multiple times
  - In parallel (almost exact same time)
- SQUID should fix this?
  - Short answer, **NO!**



# Just add caching

- LAVA downloads a lot of artifacts
  - Multiple times
  - In parallel (almost exact same time)
- SQUID should fix this?
  - Short answer, **NO!**
  - Artefacts are served over **https**
    - Requires to fake SSL certificates
      - Create a wildcard certificate (for every domains)
      - Install on every clients



# Just add caching

- LAVA downloads a lot of artifacts
  - Multiple times
  - In parallel (almost exact same time)
- SQUID should fix this?
  - Short answer, **NO!**
  - Artefacts are served over **https**
    - Requires to fake SSL certificates
      - Create a wildcard certificate (for every domains)
      - Install on every clients
  - Multiple concurrent downloads of the same artefacts
    - SQUID will download multiple times the same artefacts
    - Cache only when a first download is completed



# Kisscache

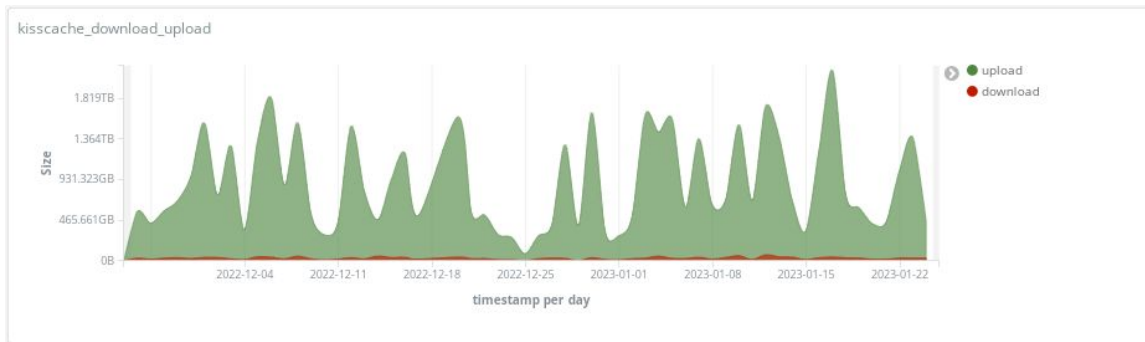
A simple caching service

# KissCache

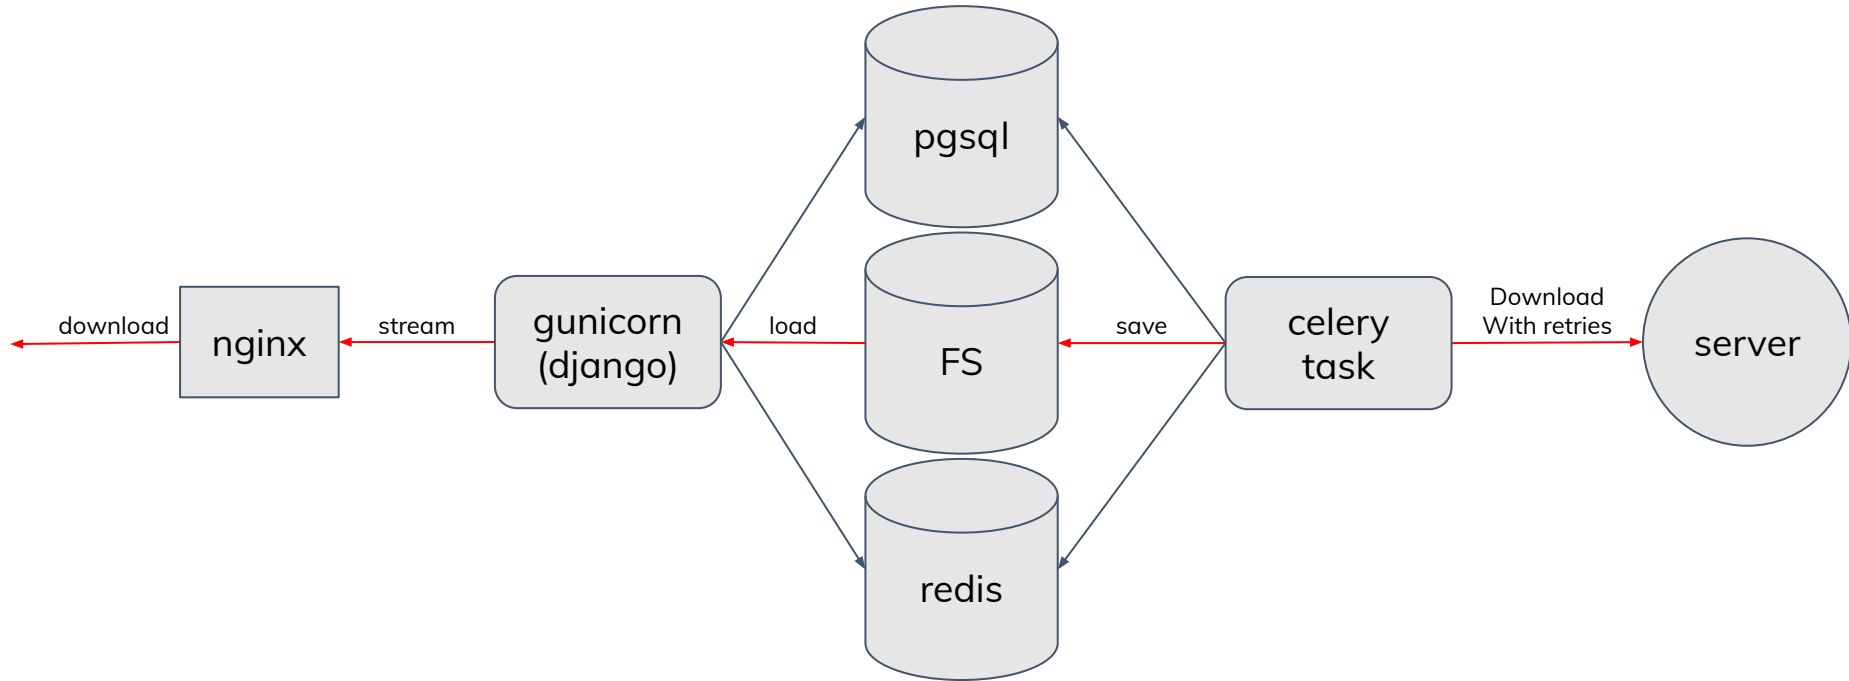
- A simple and stupid caching service
  - Cache HTTPS resources
  - Download **once** while streaming to **multiple** clients
  - <https://gitlab.com/Linaro/kisscache>
- Not transparent (prefix based), not a proxy
  - `https://kisscache/api/v1/fetch/?url=https://example.com/rootfs.ext4.zs`
    - no need for fake SSL certificates
  - Need support in the clients
- Automatic retries on multiple errors
  - 408, 413, 420, 425, 429, 430, 500, 502, 503, 504, 507, 509, 529 and 598
  - Partial download
    - Will use range request to download remaining content

# KissCache

- Over 3.5 years
  - 32 TB downloaded (from internet)
  - 1.6 PB served (local network) by KissCache
  - Network usage divided by 51x
  - Improved stability
  - Lower the s3 egress fees
    - ~3k\$ vs ~150k\$



# KissCache Architecture





Questions?

# Thank you

<https://linaro.org>

