# XCache and Fidium-2

- Background/Disclaimer
- Overview
- Some ideas
- Further info

# Background/Disclaimer

- LMU Group was active during ErUM-IDT-UM project on Xcache setup, optimisation and integration
  - N. Hartmann as main actor
  - setup of several XCache server nodes (old dcache pool, dedicated SSD server)
  - did exercise several tests on functionality, speedup, scaling, …
  - integrated our XCache into ATLAS virtual placement service
  - in operation since many years
- But not experts or developers in Xrootd/Xcache
- Several further Xcache setups in Fidium/DE
  - Not covered here, please comment...

# XCache overview - 1

- Xrootd based “smart” caching system
  - basic functionality similar to web-proxy
  - 2 modes to get remote data
    - forwarding mode: specify source with request, e.g.
      root://lcg-lrz-xcache0.grid.lrz.de:1094//root://dcache-atlas-xrootd.desy.de:1094//pnfs/desy.de/atlas/dq2/atlasdatadisk/rucio/data18_13TeV/e2/6f/data18_13TeV.00349263.physics_Main.merge.AOD.f937_m1972._lb0154._0001.1
    - Direct/global mode: specify global file name and let DMS (Rucio) find out concrete source → world-wide distributed storage with global name-space:
      root://lcg-lrz-xcache0.grid.lrz.de:1094//atlas/rucio/data15_13TeV:AOD.07709524._000001.pool.root.1
    - Direct mode conceptually very nice but in practice forwarding mode mostly used
- Principal setup rather straightforward
  - typically use some old disk server
  - though devil in details concerning config, remote authentication, monitoring, ...

# Xcache overview - 2

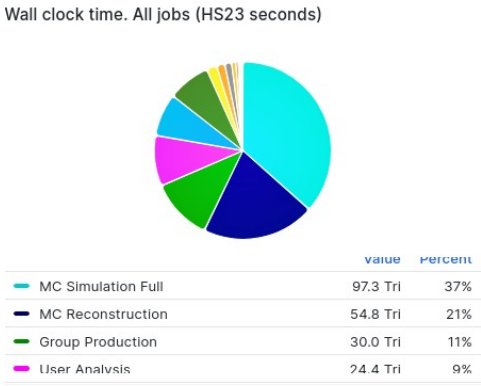


- Caching where/how to use (in HEP environment)
  - relatively straightforward as add-on to local analysis cluster
    - users don't need to download datasets but rather use global/remote name
    - caching helps with iterative analysis cycles
  - much less obvious in production environment, e.g. ATLAS: Panda/Rucio manage job&data placement
    - (almost) central principle: job goes to data
    - standard production workflows (in ATLAS) not really suited for caching
      - event-generation and simulation rather little IO (remote reading usually ok)
      - reconstruction/derivation usually single-pass
      - pileup mixing randomly selects files out of ~2 PB pool → too large
    - Analysis on Grid: special setup in ATLAS with analysis VP (virtual placement):
      - data gets virtually assigned in Rucio to cache sites
      - job scheduling follows that
      - Works in principle fine and in regular use since several years
      - But real caching benefit difficult to judge

# Some ideas for xcache@fidium2

- Test&exploit new/recent xcache features
  - pss.dca – redirect access to local read, could be very beneficial for cluster-fs
  - async read for optimised streaming mode
  - blockwise checksums
  - ...
- ATLAS VP related:
  - principal functionality fine, but issues with robustness & trouble communication
  - test/implement download from webdav endpoints
  - ...

# Further info

- Main site: https://xrootd.slac.stanford.edu/
- Regular FTS/Xrootd/XCache workshops
  - Ljublijana Sep23: https://indico.cern.ch/event/875381
  - STFC/UK Sep 24: https://indico.cern.ch/event/1386888
- IDT-UM Meeting LMU Feb 2022, XCache talk by Nikolai:
  - https://indico.physik.uni-muenchen.de/event/166/contributions/705/attachments/378/761/nikolai_xcache_idt_um_14.02.2022.pdf
- ATLAS S&C, June 24, I. Vukovic on VP:
  - https://indico.cern.ch/event/1407709/contributions/5974978/attachments/2871069/5027198/Mothballing%20Virtual%20Placement.pdf