

# Storage State of the Union

GridPP 50

Sam Skipsey

(he/they)

Material and opinions from: *Simon George, Vip Davda, Raul Lopes, Matt Doidge, Gerard Hand*

# Overview

- “DPM Decommissioning” *almost* done
  - 1 site left...
- Virtual Placement *almost* works for cachey sites for ATLAS
- Cephfs/xrootd *almost* great for high storage sites
- Room for efficiency gains?

# DPM decommissioning

- DONE:
  - Edinburgh, Liverpool, RHUL, Lancaster, Bristol, ...
- Doing:
  - Brunel
    - Done for ATLAS by just making it storageless
    - Replacement Cephfs/xrootd service is now up as of this week...
  - Durham
- TBC:
  - Manchester...

# Transitions in the future

- Edinburgh
  - DPM -> dCache -> no storage
- Bristol
  - Xrootd-on-HDFS -> Xrootd-on-Cephfs?
  - Ideally needs more available effort to start this.

# Cachey sites: pairing versus VP

- Currently several Cachey/storageless sites pair with upstream “big site” for storage.
- Eg: Sheffield, Oxford, Edinburgh
- Fine, except this means we’ve had several downtimes of the big site which also brings down the tributaries.
- “Rucio” Virtual Placement should ameliorate this (at least for data inbound)

# Storage at RHUL


## Xcache experience

- **Xcache** with Virtual Placement (ATLAS Analysis data reads only)
- Direct access to **QMUL** (for ATLAS Production reads, all writes)
- Works fine on our **very old** hardware
  - 3 very old ex-DPM storage nodes, with load sharing done by ATLAS
  - 40TB RAID6 (PERC H710 mini); 64 GB RAM, 2x Xeon E5-2609 0 (2x4 cores) 10Gb/s NIC
  - No problems with i/o seen
- Services provisioned using docker-compose – very easy
- No significant increase in network traffic from JANET – to my surprise
- Reliability feels worse but it depends how you look at it

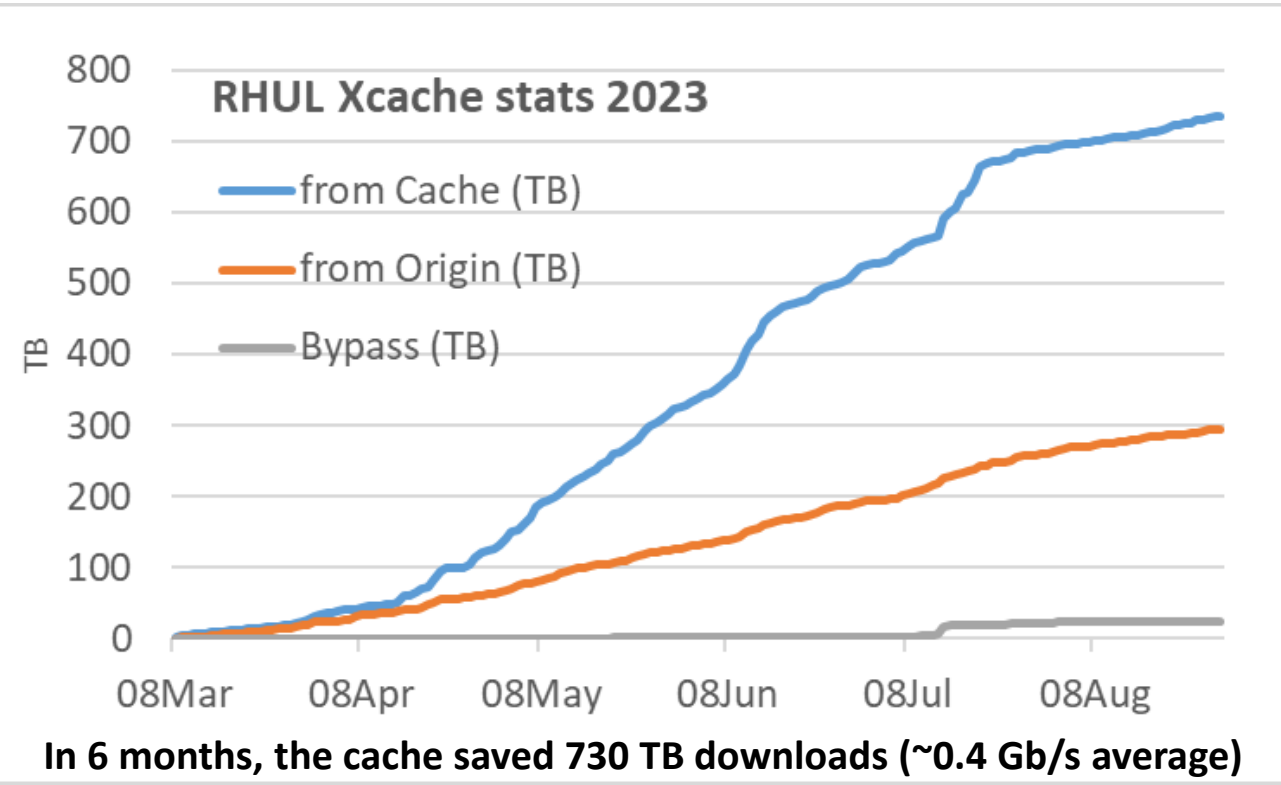
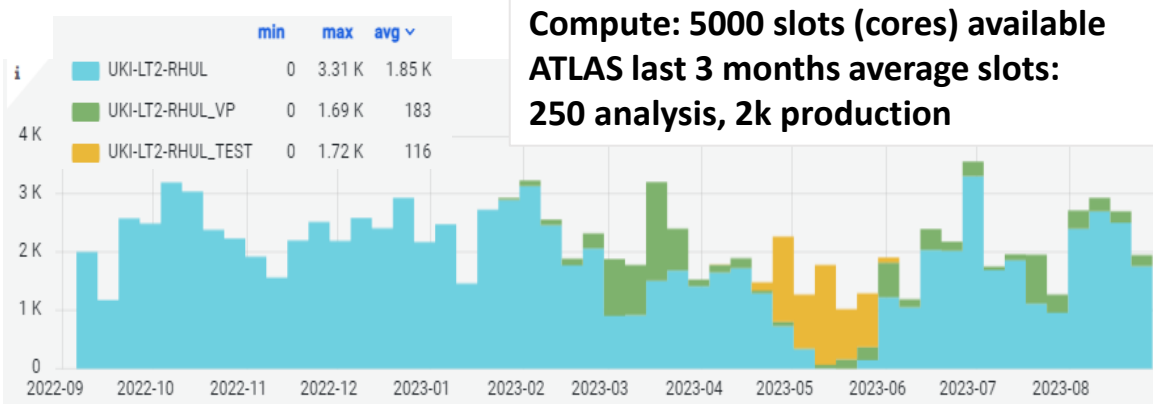
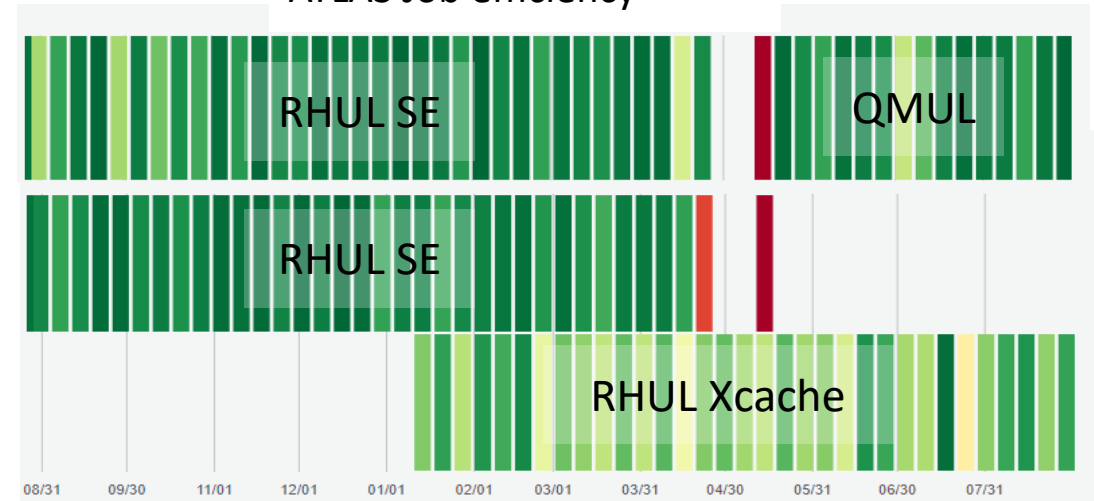
*Thanks James W, Dan T and the GridPP storage group*

Prod

Analy



## ATLAS Job efficiency



# Storage at RHUL

## Problems

- **Rucio geographical ordering problem**
  - Caused frequent blacklisting due to ATLAS Hammercloud failures
  - With no local SE, Hammercloud tests get input data from WAN, but rucio doesn't correctly provide a location-ordered replica list
  - HC test datasets are everywhere, so we have been downloading from Australia (alphabetical) and more recently Israel (no apparent logic) with some unreliability
  - Several attempts to get this addressed in rucio and the ATLAS pilot have not yet succeeded. Current relevant tickets: rucio issues [6141](#) [4832](#), ATLAS Jira [ATLDDMOPS-5662](#)
  - Forced online to avoid blacklisting meanwhile – not ideal
- **Xcache servers sometimes stop working without warning**
  - Only clue is jobs failing with data access timeouts and docker logs not updated for many hours
  - ATLAS monitoring is too simple: just requires a separate heartbeat container to be up; no check that xcache is responsive
- **Reliant on single maintainer: needs to be adopted by ATLAS not just Ilija**
  - ATLAS robot cert expires 22 Sep [needs interaction with Ilija to replace – happened yesterday]

# Oxford's Xcache Servers

## Setup - May 2021

### Hardware:

Dell PowerEdge R720xd  
2 x Intel Xeon E5-2603 v2 @1.80GHz  
48GB memory  
40TB Disk storage

Xrootd version 5.x

VP confirmation - 14<sup>th</sup> August 2023

OS: EL8 (Rocky 8)

## Setup (RAL Loan Server) - October 2022

### Hardware:

Dell PowerEdge R7525:  
2 x AMD 7413 2.65GHz,24C/48T,128M,180W,3200  
256GB memory  
24 x 4TB (92TB for storage) SSD Storage

Xrootd version 5.x

VP confirmation - 17<sup>th</sup> August 2023

OS: EL8 (Rocky 8)

## Network Connection


- 2x 10Gb network connection to the local switch stack.
- 20Gb to University core network router.
- 100Gb from University to JANET.

# Oxford's Xcache Servers

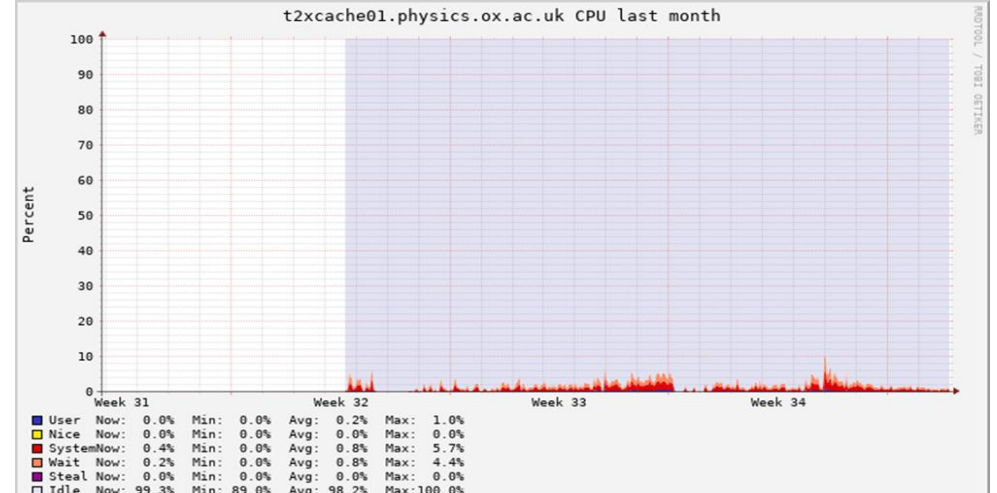
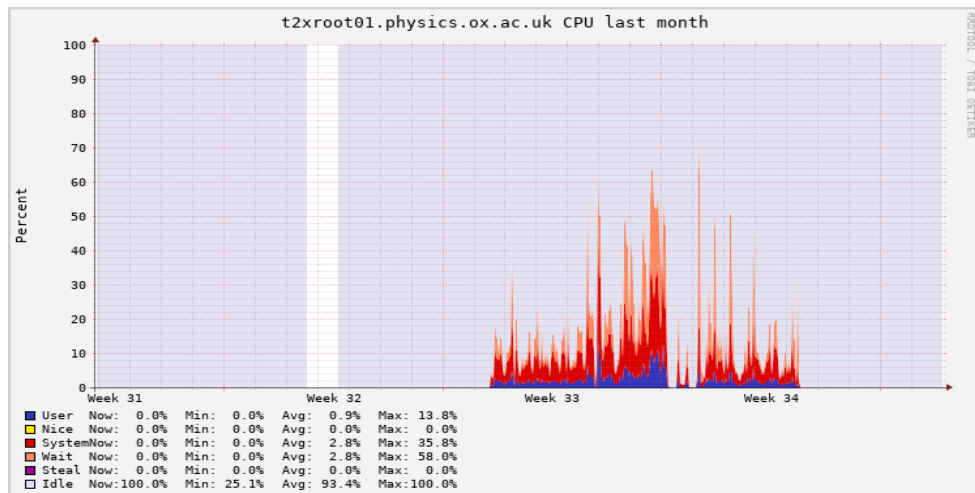
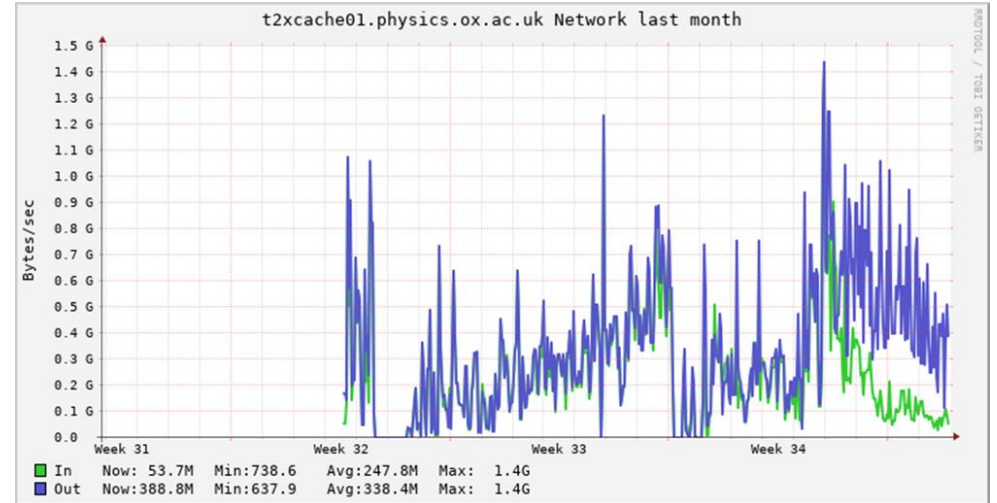
## CPU and Network Plots

(8 – 29 August 2023)

### Dell PowerEdge R720xd



### Dell PowerEdge R7525



# Oxford's Xcache Servers

## VP Data sources

atlasxrootd-kit.gridka.de

bohr3226.tier2.hep.manchester.ac.uk

ccxrdatlastpc.in2p3.fr

clrlcgse01.in2p3.fr

dcgftp.usatlas.bnl.gov

eosatlas.cern.ch

eos.grif.fr

fax.mwt2.org

ftp1.ndgf.org

grid-dc.rzg.mpg.de

gridftp.atlas-swt2.org

lapp-dccentral01.in2p3.fr

lcg-dp01.uw.computecanada.ca

lcg-se01.icepp.jp

lcg-se0.ifh.de

mover.pp.rl.ac.uk

se01.grid.cyfronet.pl

tbit00.nipne.ro

xrootd01.esc.qmul.ac.uk

xrootd-at1-door.pic.es

xrootd-at2-door.pic.es

xrootd.echo.stfc.ac.uk

# Experiences with Ceph

- Diverse set of ceph installs at Tier-2s:
- Glasgow
  - Ceph Nautilus w/ xrootd-on-RADOS *and* xrootd-on-Cephfs
- Lancaster
  - Ceph Pacific w/ xrootd-on-cephfs
- Brunel
  - Ceph Quincy w/ xrootd-on-cephfs
    - In production... this week!

Also

Durham

Ceph Quincy w/ xrootd on cephfs  
(commissioning)

# Experiences with Ceph

## Glasgow

- Ceph Nautilus w/ xrootd-on-RADOS *and* xrootd-on-Cephfs

## Ceph “generally good”

Some minor issues due to age of Nautilus codebase

No data loss so far, despite various moments of excitement over last 4 years.

Currently “saving up” big administrative changes:

We need to increase PG count from 1024 to 4096 for the atlas pool

We need to transition to (at least) Pacific for future proofing.

# Experiences with Ceph

## Lancaster

- Ceph Pacific w/ xrootd-on-cephfs
- Generally happy... except:
- Unhelpfully “helpful” orchestration
  - Several interventions have been made more complex by the automation not allowing direct control [cf Glasgow, where I just turn things off directly...]
- Sudden SSD failure for RocksDB storage [half of OSDs on one node down as result]
  - Long tail of recovery issues due to geometry of cluster representation, needing a lot of manual intervention and hand-holding.

# Experiences with xrootd (GLA + LANC)

- Glasgow and Lancaster have xrootd “gateways” fronting ceph storage
- Mostly good, once you internalise configuration for xrootd...
- Shared issues:
  - Xrootd releases seem to often mix critical bugs with needed fixes
    - 5.6.x series *fix* handshakes between RHEL7 and RHEL9 OpenSSLs
    - *but* also break EGI availability tests due to incorrect handling of chunked transfers over HTTP.
  - Xrootd “stream” checksums are the missing feature for efficient transfers for Lancs (especially), and Glasgow (less so)
    - Little momentum on this despite prodding from James W and I
    - Workaround @ Lancs is throwing increasing numbers of servers at the problem (Not good for power efficiency!)

# Possible Efficiency Gains

- ARM-based Storage servers?
  - CPU load of Ceph OSD nodes is not zero, but spinning disks also a considerable component of power draw.
- More efficient operations
  - See: stream checksums (probably halves number of gateway nodes @ Lancs)
- More parsimonious storage
  - Deduplication across ceph object store?
  - There's a considerable memory tradeoff here...
  - Better Erasure Coding? [we mostly just use Reed-Solomon atm...]
- SSDs v HDDs for storage
  - Less energy when running from SSDs (?)
  - More embodied carbon from SSDs

# Summary

- Mostly DPM free
- Past the honeymoon phase of Ceph
  - We need to sync Tier-2 releases and processes better over GridPP7
- Virtual Placement remains important to us (and we need to make sure it is important to ATLAS + Rucio)
- Consolidation of storage brings some questions about Core Tier-2 scaling
- It is good that we have Xrootd dev effort @ RAL
  - Xrootd needs to work for us for GridPP7 and beyond.