

XCache

(mostly at Oxford)

Itinerary

- Existing sites: Birmingham, Sheffield, Imperial
- Oxford site
- RAL config side
- Testing and results

Bham, Sheffield, Imperial

- Sheffield
 - No storage
 - Jobs run against RAL storage with no cache
 - Job mix constrained to get reasonable efficiencies.
- Imperial / QMUL
 - Reciprocal "symbiotic" relationship
 - Imperial runs ATLAS jobs v storage at QMUL
 - Vice-versa for CMS at QMUL
 - (no caches in either direction)
- Birmingham
 - Load-balanced Xrootd Disk caching proxies ("Xcache") across multiple servers, for ATLAS only.

Oxford

- Agreed to be a "test case" for exploring Caching proxies as a replacement for their existing DPM.
- Effort:
 - Vip @ Oxford
 - (honestly much of the work)
 - James Walder @ RAL
 - Testing from ATLAS side, config, liaising with rest of RAL
 - Sam Skipsey
 - Advice, planning, config
 - Rob Currie
 - (running the xrootd monitoring), config
 - Mark Slater
 - Example setup scripts, basis for original design

Oxford

- Agreed to be a "test case" for exploring Caching proxies as a replacement for their existing DPM.
- Config scripts from Mark Slater @ Birmingham
 - We simplified these a bunch, as we only have 1 Xcache server
- Monitoring from local network + job efficiencies, ATLAS monitoring, Edinburgh Xcache monitoring.
- Feb/March 2021 started building server @ Oxford
 - 720XD, removed from DPM pool
 - 4xE5-2603 v2 @ 1.80GHz
 - 12 x 3TB disks, initially in RAID6 config
- Online ~April

Security config

- Xrootd Proxies cannot* forward credentials from a client to the server they proxy.
- Security config on Oxford and RAL side needed to trust the certificate of the proxy itself for access to ATLAS data.
- (This is really more of a risk for the upstream source, as they're trusting a machine with effectively world-read-access to a VO's data.)

*in Xrootd 5.2+, in some contexts, this is now allowed. But not for our case still.

Configuration changes

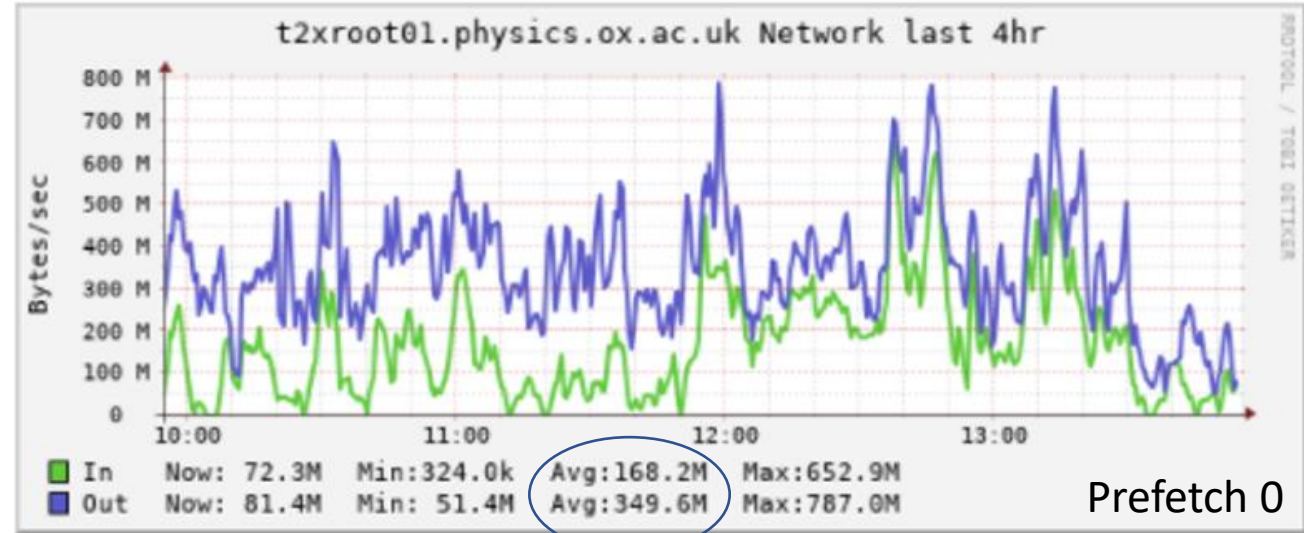

- RAID6 -> "software parallel" Xrootd "Spaces" [May 15th]
 - RAID6 configuration had insufficient IOPs to support load from WNs
 - (resilience / performance trade-off)
- Xrootd releases:
 - 4.12 (pre-production)
 - 5.0 to 5.3.1 [current]
 - (5.1 - 5.2 pretty unstable, many core dumps at load)
- Cache eviction tuning: initial config cache would fill up rapidly and not be able to purge
 - File lifetime reduced aggressively
 - High-watermark of 95% too high for cache to respond fast enough.
 - Partly due to most files being accessed only once anyway...

Prefetch changes


- First explored configuration changes were prefetching
- Prefetch setting is essentially the "readahead" for the Xcache, in blocks.
 - Our blocksize is 32M to match RAL ECHO blocksize on storage.
- We explored 4 prefetch settings
 - 0 (== infinity, always cache whole file on any request)
 - 1 (32MB readahead)
 - 10 (320MB readahead)
 - 20 (640MB readahead)

Prefetching and "caching efficiency"

- We can measure the effectiveness of a cache by how much it reduces the network reads "externally", versus reads "internally" from the cache.









~50% reduction in traffic (whole file caching)








Risks: Dependence on RAL config

- August: RAL gateway changes to support WebDAV broke Xroot cache @ Oxford
 - Exposed bug in XrdCeph/libradosstriper, sending more data than in file to requests, *when requested blocks are large*
 - Apparently Xcaches are vulnerable to this, despite xrdcp being immune (thanks to only requesting data up to known file size)
 - Required rollback of RAL config to "fix" Xcache (16th aug)
 - Oxford was also ticketed during this process.
- Exposes dependence of Xcache sites on their upstream sources.
- (Obviously this particular incident was unusual, but coupling is an effect)

ATLAS view of Oxford (transfers to)




ATLAS view of Oxford (transfers from jobs, via gridftp)






Slots of Running jobs by ADC activity




- MC Reconstruction
- Group Production
- MC Simulation Full
- MC Resimulation
- Data Processing

min	max	avg	current	total
0	2.81 K	738	0	2.13 Mil
0	2.72 K	550	0	1.59 Mil
0	2.76 K	156	0	450 K
0	1.82 K	125	0	361 K
0	1.73 K	92.5	0	267 K

WallClock Efficiency based on success/all accomplished jobs




Average CPU Efficiency Good jobs



avg ▾

0.949
0.918
0.865
0.847
0.775

Oxford




avg ▾

1.01
0.975
0.958
0.924
0.865

Glasgow

Average CPU Efficiency Good jobs



avg ▾

0.978
0.952
0.842
0.460

Sheffield (no cache)

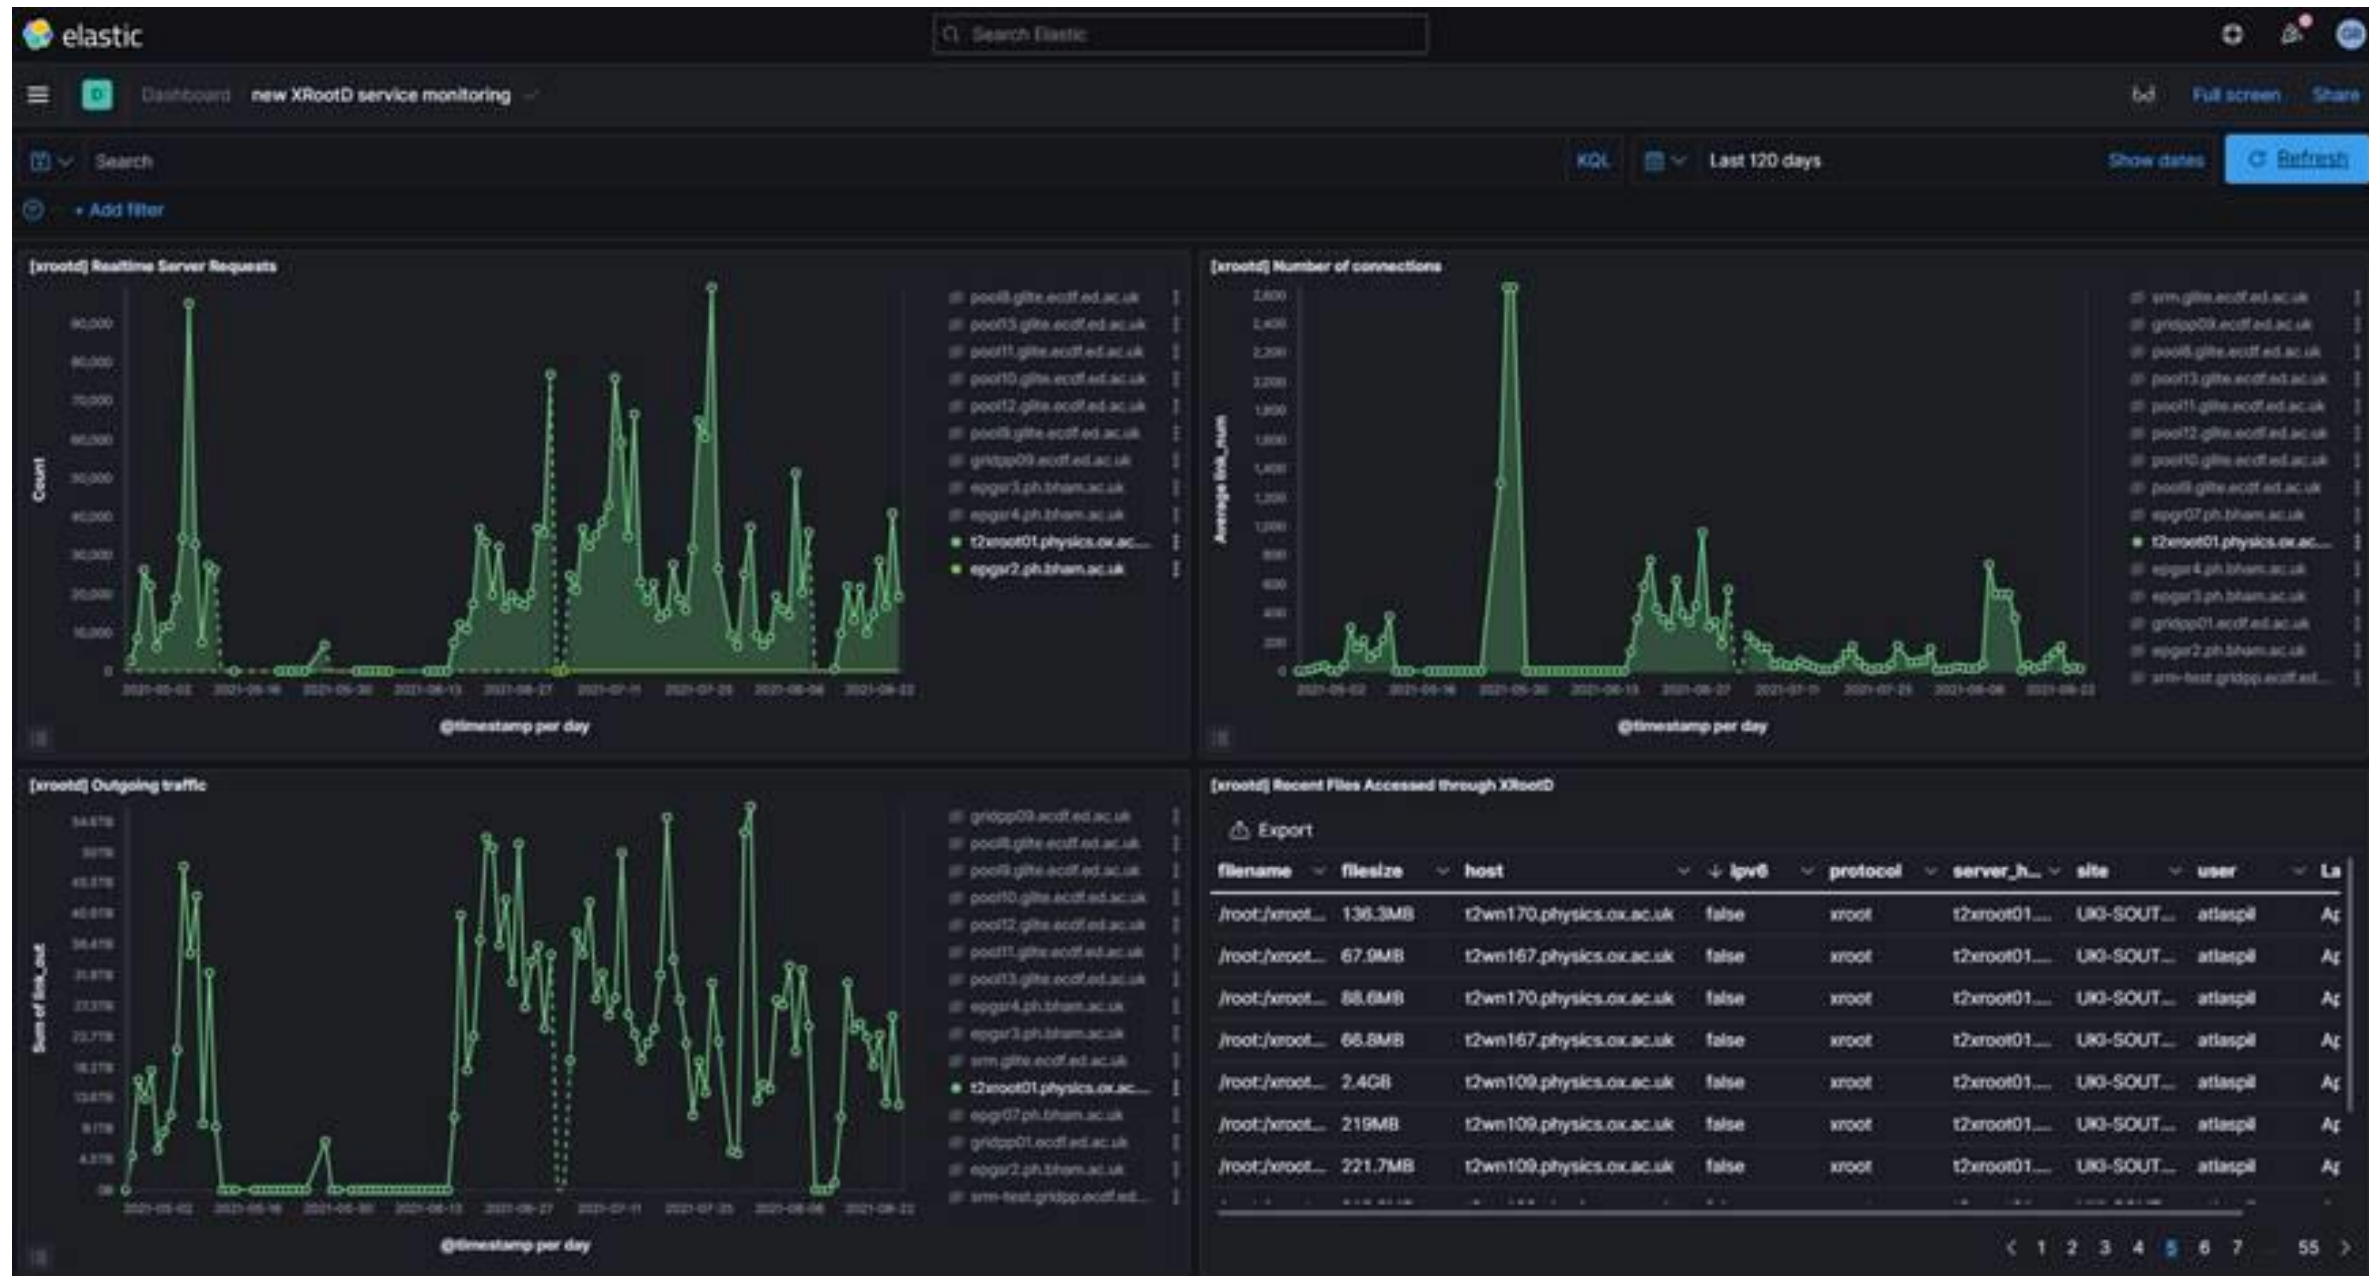
XRootD Server Monitoring

- XRootD monitoring can be very powerful.
 - Used by ALICE as it's "good enough" for their use-case.
 - WLCG would prefer it to be better for wider use.
- "*XRootD-Server*" monitoring driven by different philosophy to "*Service-level*" monitoring.
- No interest/man-power from XRootD devs in changing XRootD internals.
- WLCG working on a proposal to run an additional service at-each-site to translate XRootD to "*Service-level*" messages.
- Ideally this new service will be a lightweight and "*almost-stateless*" solution which sends 1 message per file access.

XRootD/XCache Monitoring for GridPP

- Edinburgh using SLATE/OSG-based monitoring tooling presented at vCHEP. This is “good-enough” to do some first monitoring and make plots.
- The Tier1 is already using a site-specific monitoring solution.
- Interest from Glasgow in potentially adopting similar tooling.
- Both have experienced bumpy transition XRootD 4->5.
- Edinburgh will consider migrating to the “WLCG collector” for XCache monitoring.
- Not clear yet, if or how this will meet all possible use cases.

Monitoring @ Edinburgh



Current Status

- Prefetch 0 [whole file copying]
- Xrootd 5.3.1

- Job efficiencies ~80% over last week.
- (dependent on job mix)

- Stable

Future Work

- Tune cache acceptance criteria
 - At present, caches all files requested from RAL
 - However, Teng Li's [and others] results show that 90% of files read once.
 - (And Oxford jobs *stage in* data)
 - So most of the files are "buffered" not cached
 - Can and should tune "acceptance" criteria for caching [some thought at Edinburgh on how to do this for ATLAS data]
- Direct IO versus Copy to Scratch
- Open question: how does this scale with Oxford network capacity increase? (Say, to 20Gbit/s from 10)
 - Bigger, faster cache?
 - No Cache [even at cost of higher load on RAL gateways?]