

“Data Management”

Some updates from GridPP and SWIFT-HEP related work

Sam Skipsey (he/they)

On behalf of: Tim Noble, Raul Lopes, Matt Doidge and others

Data Management and Storage Topics

- Storage:
 - DPM Migration
 - “Xrootd matters” [see Jyothish’s talk immediately preceding]
 - **Ceph care-and-feeding an increasingly important aspect of UK storage sites**
 - (outside London...)
- Data Management:
 - SWIFT-HEP
 - **work on Data Lakes with Rucio led by Tim Noble @ RAL**
 - **work with Xcache led by Raul Lopes, Ivan Reid @ Brunel**

Swift-HEP data management

Work Package 1.1 objective is to create and develop a UK based data lake using the data management software Rucio. The data lake will incorporate different types of storage endpoints.

Some of the storage endpoints need to include flash based storage for high performance that the data management can be developed to utilise performance for job efficiency.

In order to be able to maximise flash based storage endpoints; monitoring needs to be in place to be able to compare and contrast the performance between the endpoint technologies.

UK data lake monitoring

The CERN Monit FTS monitoring was used as a gold-standard monitoring page for data movement to be build for Rucio. The brief for the monitoring solution was to design a solution that used Rucio components, a dashboard creation service like Grafana, and minimal steps in between. Thereby making it easy to maintain and troubleshoot.

The features designed into the monitoring are:

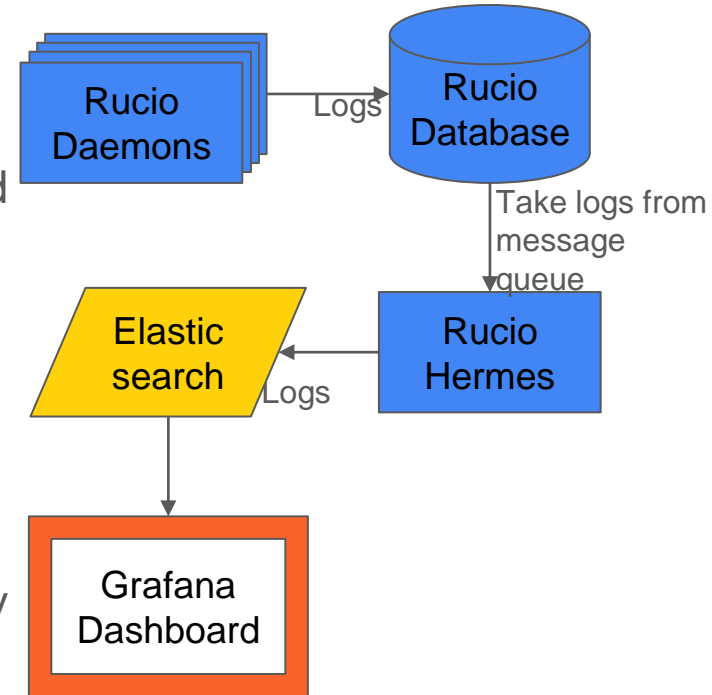
- Data rate and total
- Transfer efficiency
- Deletion Rate and total
- Breakdown of site to site metrics
- Customisable filters
- Transfer Error metrics
- Direct links to FTS error logs
- Flexibility to add more types of monitoring

Keeping the required software number low allow the monitoring to be deployed for communities outside of CERN to have monitoring of their dataflow for campaign and functional tests.

The designed solution

The monitoring solution Tim designed, implemented, and bug fixed:

1. Rucio logs are created by the daemons and stored in the database.
2. The messaging daemon then retrieves the messages from the database.
3. The messages are then sent to and ingested into Elasticsearch.
4. The data is then parsed in Elasticsearch to make the data searchable and filterable.
5. The indexed data is then used as a data source by Grafana to build dashboards.



Data lake Monitoring

These figures are from the deployed monitoring solution at RAL. A similar one is deployed for LSST and is waiting for data to populate it.

Top figure is designed to replicate the FTS monitoring most HEP communities know and use.

Lower figure expands on this information to deletions and errors from Rucio to identify issues and have direct links to the error logs.

Flexibility in filters with plans to add more to facilitate flash based/ HDD comparisons for data management work.

Helm charts created for Multi-VO and LSST and in the process of being created for the Rucio repo to allow all Rucio users access.



Next Steps for data lake

A robust, and flexible monitoring solution has been deployed for the UK data lake.

Since the monitoring has been deployed in stages a baseline of functional test monitoring data based on HDD storage endpoints has accumulated to be able to now work on incorporating flash based storage endpoints for direct comparison.

Data from the comparison will be used to inform the development of Rucio to best utilise the storage endpoints capabilities.

The general purpose helm chart will also be completed to allow all communities that use Rucio a monitoring dashboard for the Rucio internals as well as now the externals.

Brunel XCache Work

Measuring efficiency of (CMS) analysis through a cache versus direct access.

- **CMS xcache: dc2-grid-srv-009.brunel.ac.uk**
 - 270TB cache on ZFS
 - packages from OSG

xcache-3.6.0-1.osg36.el9.x86_64

cms-xcache-3.6.0-1.osg36.el9.x86_64

- Alma 9

- **XRootD/CEPH: xrootdgw.brunel.ac.uk**
 - 2.5PB CephFS

- 2 XRootD endpoints

xrootd-5.6.8-1.1osg36

- Alma 9

Thanks to GridPP Storage Group
(and especially Rob Currie @ Edinburgh)
for extensive discussion and configs for Xcache
Install/configuration/management

Next steps:

**Ivan to run extensive analysis to compare
access through cache
access directly to xrootdgw.**

**create a multi-server cache and declare in the
CMS AAA.**

**Consider adding to the WLCG Xcache network
that is being created by CMS.**

Lancaster Storage Status (1)

- In generally working okay, most of the time.
 - XRoot definitely has its moments, and CEPH generally works, but when it doesn't...
- XRootD checksumming a significant performance impact
 - Forces scaling horizontally
 - ~doubles stress caused by the gateways to the CEPH cluster
 - ~doubles the number of gateway servers required
 - requires significant gateway hardware investment per gateway (memory as well as network)
- The key to running a successful system is extensive monitoring, for both CEPH and XRoot
 - XRoot monitoring is hard to tune, Shoveller potentially powerful but finicky.
 - Steven has worked integrating this into our monitoring
 - CEPH has a metric for everything, but needs correlation to identify issues
 - Gerard's working on this. It involves piping stuff into SQLite databases...

Lancaster Storage Status (2)

- A quick case study in CEPH “moodyness”, as touched on by Jyoti in his talk we had ~2 weeks of sub-optimal performance.
 - Smoking gun was ridiculously long time taken for some CEPH ops during “bad” periods.
 - The dreaded “slow ops”
 - We thought one contributor was being overfull, but clearing space actually made things worse
 - Was the stress from deletions pushing us over? Hard to identify these operations (which is its own problem) – ceph deletions are effectively “shreds” in posix terms
 - Concurrent FTS connections were set high - lowering these to the defaults made things get “less bad”.
 - Was it a bad mix of user jobs hitting us in just the wrong way? The job mix changed at about the time, but it doesn’t perfectly line up.
 - Was there some bad disks/sectors? There was a bad disk at the time.
 - Was it just that Gerard was trying to have a week off and the CEPH missed him?
- The issue fixed itself over the weekend, without any direct intervention.
 - Been running fine since, with occasionally blips but nothing that the users/xroot have noticed.

Graph showing the bad time.



Note the log scale!