

UK Cloud Report

Alastair Dewhurst, on behalf of the UK Cloud
Squad



Alastair Dewhurst, 27th January 2016



Overview

- T2 site evolution
- Development work
- Site feedback



UK Funding

- Roger Jones will be presenting a talk on Friday on: “UK sites evolution”
- Sites will have 90% flat cash.
 - 19% real term cut.
- Sites will be more reliant on opportunistic capital spends.
- Consolidating storage for ATLAS at T1 + 4 x T2 sites.
 - Some T2D will decommission their storage.
 - All sites will continue to contribute CPU.



Tier 2 evolution

- Challenge: How to make best use of a site with CPU resources only and 0.5FTE effort?
- Vac project[1] aims to run Grid jobs with minimal effort.
- UKI-LT2-UCL storage was completely decommissioned:
 - PanDA queue just points to QMUL storage.
 - Both sites within London.
- Aim to run HC tests at Oxford:
 - Read data via FAX from UK sites. Write locally.
- Goal: CPU resources can dynamically read and write to any UK site.


[1] <https://www.gridpp.ac.uk/vac/>

Alastair Dewhurst, 27th January 2016



Storage shrinkage

- Very difficult to shrink storage at a site:
 - ATLAS fill sites.
 - Draining data from nodes to be decommissioned very slow.
- Current method normally involves “nuking” the site and starting again.
- Effort will reduce at sites decommissioning storage:
 - Faults fixed but no new developments.
- Slow draining over months, need free space and plenty of secondary data to allow automatic cleaning of site.
- Once space drops below 300TB drop storage completely decommissioned.



Alastair Dewhurst, 27th January 2016



Tier 1 developments

- Echo
- Docker
- Load Balancing
- Preempt-able Jobs





- Echo is the name of our Ceph storage service which will replace Castor for disk only storage.
- Provides S3/Swift API:
 - Very cautiously opening up access due to security concerns.
- XrootD & GridFTP plugins also available.
- Eventually Castor will use Ceph pools for tape buffers.
- ~5PB usable storage currently.
 - Additional 13PB storage purchased.
- Aim for production status by July 2016.



S3 / Swift

- In the long term RAL would like to have an “S3 spacetoken”.
 - Short term use object store for logs.
- Testing FTS transfers to S3 endpoints:
 - Fts-devel.gridpp.rl.ac.uk – always running latest version.
 - Recent DPM and dCache storage can perform 3rd party transfers.
 - Older storage use FTS as proxy – Possible to run at RAL.
- Concern that FTS has own service for storing access credentials.
 - RAL has some development effort for S3 authentication.
- Don't forget about Swift!
 - Open standard, easier for developers.



XrootD and GridFTP

- Echo has “working” XrootD and GridFTP plugins.
 - Use gridmap file – sufficient for Castor, users mapped to atlas prod or analysis.
 - No SRM – stop asking.
- For GridFTP: 3 x Gateway machines with 4 x 10GB/s each.
- For Xrootd: Plugin on every VVN to allow efficient internal access to files.
- Still finding bugs and working on service stability.

```
# globus-url-copy file://foo.root gsiftp://gdss606.gridpp.rl.ac.uk:2811/atlasdatadisk:rucio/data15_13TeV/ab/cd/foo.root
```

```
# xrdcp root://ceph-gw1.gridpp.rl.ac.uk/atlasdatadisk:rucio/data15_13TeV/ab/cd/foo.root foo.root
```



Alias

Spacetoken

Rucio name



Docker

- Andrew Lahiff is investigating Containers to simplify WN configuration and increase flexibility.
- Testing HTCondor Docker universe on WN.
 - Successfully ran thousands of CMS jobs, ATLAS next.
 - Also testing CernVM docker container.
- Testing long running services within Containers managed by Mesos.
 - On bare metal machines rather than inside VMs.
 - Tested FTS, Bdi, MyProxy, ELK...



Load Balancing

- Current situation:
 - Hosts directly exposed to users via DNS alias if there is a problem with a host or one is rebooted for upgrades, users will see this.
- Solution is to use a load balancer between the servers and users.
 - HAProxy (load balancing)
 - Keepalived (allows for high availability)
- Benefits
 - HAProxy health checks mean that users can only be connected to healthy servers.
 - Can do transparent upgrades, hosts can fail without anyone noticing, ...
 - Can move to a more dynamic infrastructure, where the numbers of hosts behind services can vary depending on load.
- FTS3 will be the first service to move to using a load balancer.



Preempt-able Jobs

- Event Service jobs can make use of opportunistic resources where they might be killed without notice.
- Modified HTCondor to add preempt-able job classAds to certain machines:
 - Draining machines – 3% of CPU is lost to draining.
 - Cloud VMs – Significant increase in cloud resources.
- Jobs submitted to RAL-LCG2_ES will be run on preempt-able slots.
- Recently presented at ACAT 2016 [1].
 - Happy to share code with other HTCondor sites.

[1] <https://indico.cern.ch/event/397113/session/11/contribution/156>



Site Feedback

13

- Hammer Cloud
- DDM Operations



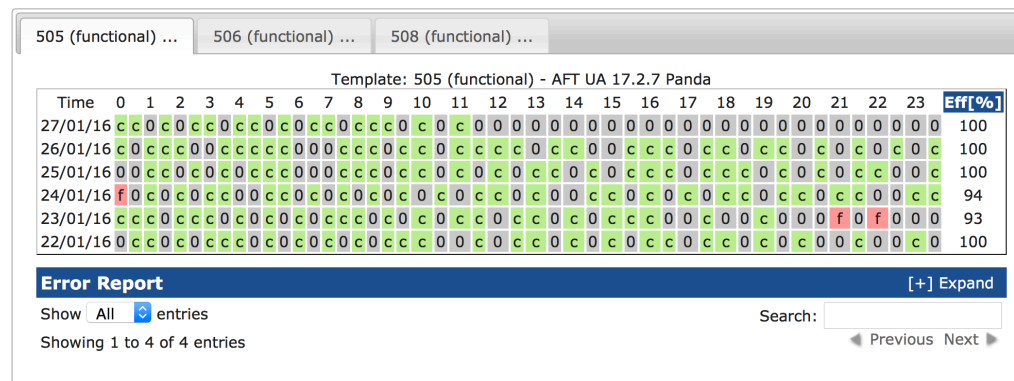
Alastair Dewhurst, 27th January 2016



Hammer Cloud

- Why is my site not back online?
 - HC is a really powerful tool for debugging site problems.
 - Want quicker response time especially after outage.
 - It would be helpful if the logic was clearer (e.g. job submitted, waiting on pilot).
- The documentation and site navigation is poor.
 - Looks like it was designed for experts, not site admins.

ANALY_RAL_SL6



DDM Operations

- Communication between sites and central support is sometimes sluggish.
 - Recent example with consistency checking GGUS tickets.
- Protocol zoo/usage – we want to simplify setup where possible.
- Groupdisk space tokens, empty directories – can we clean them up?
- Would like better understanding of T2 usage evolution – Will data be unique?



Summary

- UK Cloud is currently operating very effectively
- Significant challenges in future need to start solving them now!

