

UK DPM Feedback

DPM Workshop 2019.

Thanks to Raul, Sam, Simon, Alessandra, Daniela, Kashif and Rob for their input.

UK Overview

- **11 DPM Sites in the UK (1 recently retired)**
- **Number of SEs in the UK contracting due to various pressures.**
 - Decreased Staff Time
 - Less storage hardware investment
 - Desire for fewer endpoints

Site by Site

- **Birmingham - decommissioned DPM for an EOS SE (at ALICE request).**
- **Sheffield, Cambridge - plan to decommission rather than upgrade.**
- **Bristol - dependent on HDFS support.**
- **Glasgow - migrating away from DPM for ATLAS.**

Site by site continued.

- **Liverpool, ECDF, RHUL - waiting on the dust to settle before migrating to DOME.**
 - RHUL uncertain of long term future of their SE, could make use of HDFS support.
- **Oxford - DOME'd with some effort**
 - As with RHUL uncertain of long term future, no storage hardware is in warranty.
- **Brunel, Lancaster, Manchester - DOME'd and not going anywhere.**

DOME upgrade status

- **If you're optimistic the UK is (nearly) halfway to being all upgraded.**
 - 4 Sites upgraded
 - 2 Sites decommissioning
 - 1 Site dependent on HDFS so can't upgrade.
 - Leaving 4 Sites to upgrade
- **As no UK site had a 100% smooth upgrade there is some reluctance from the rest to move.**
 - Putting out the upgrade deadline so soon has been described as “too much stick and not enough carrot”.

DOME Upgrade Experiences

- **Brunel, Lancaster**

- Well documented – we went a little early and things were a bit rough, but it was a necessary experience.
- But we can't go through that very often.

- **Manchester**

- Smoother, had some problems as dual-stacked at the same time.
- Found the documentation a little lacking in structure and distinct follow up recipes.

DOME Experience Continued.

- **Oxford**

- Found some conceptually holes in the documentation delayed the upgrade.
- A few issues after the fact (enabling gridftp redirection).
- Benefited from being able to fall back on other's experiences.

UK Storage Users

- **Main consumers are ATLAS, CMS, LHCb**
 - No surprises there.
- **But influx of new user groups (some not even HEP)**
 - Neutrino groups – T2K, Sno+, DUNE, LZ
 - Astronomers – LSST, SKA, CTA
 - Many of these user groups use DIRAC.

Moving Users away from SRM

- **CMS - don't need it**
- **ATLAS - switching it off for DOME sites**
 - No problems yet.
- **LHCB - consulting UK DPM sites**
 - See GGUS tickets.
- **Other (new) VOs - more complicated**
 - Tend to use SRM out of “tradition” - as this is what everyone else used to use.
 - Efforts by UK DIRAC team to move to root/http.
 - Hampered by lack of xroot publishing in C7 headnodes.
 - Also see recent problems with gsiftp on single-stack nodes talking to dual-stacked servers.

General Feedback

- **Documentation could be improved.**
 - Where to start (<https://wlcg-ops.web.cern.ch/dpm> is a deadend).
 - Wiki-centric in the UK, so like that approach.
 - Of course you can ask us to help with this!
- **Lancaster/Brunel/Manchester are willing to keep “early-adopting” DPM.**
 - But this is not without risks
 - Periods of bad running tarnish both the site’s and DPM’s reputation.

Feedback In Depth: Glasgow

Glasgow is moving the majority of its storage provision (the ATLAS part) to Ceph/Xrootd; and off-DPM. (In fact, the first storage in Ceph is new, so it's not "off" DPM so far, technically, but we will be migrating some of our existing storage sideways in future). This is to provide both additional non-grid interfaces (object storage, for example; and allow efficient use of a single distributed storage pool for multiple services), and also to gain technological improvements (resilience across server boundaries, automatic rebalancing, data striping...). The opportunity was provided by the inevitable disruption due to moving to a new facility.

We will probably be retaining the DPM service long enough to transition it to DOME (and have been watching the UK experience in general of this); we have not decided, yet, if we will retire the remaining DPM storage or fold it into Ceph/Xrootd.

Feedback In Depth: Brunel

- **DOME 1.12**

- happy that I moved to DOME
- had a lot of trouble moving from 1.9 to 1.10.
- - not moving out of DOME, unless my PI starts pressing me because David C is pressing him.

- **Issues in DOME**

- checksums deadlock;
- dmlite-shell doesn't deserve the name of shell
- dmlite-shell help has shoddy syntax and doesn't agree with dmlite-shell wiki.
- minimal changes in dmlite-shell: new name, new help.
- ideal changes in dmlite-shell: compile to an AST (like oil shell) that can be easily extended by the community.

- **Question to devs: Can DOME be front end to CEPH or BeeGFS?**

- **Improvements in the documentation would help us:**

- **Database**

- I'd like to see the whole lifecycle of a file in the DB?
- Interaction with DOME through the RESP may not be possible in an emergency situation or with HN under heavy load. Access to the DB helps.

- **Many of us want to monitor the many protocols for errors. We need at least a wiki with pointers to:**

- How to parse/understand your dpm-gsiftp?
- How to parse/understand your xroot logs?

- **Final observations:**

- DOME 1.12 is good

- **The dev's support has been very good!**

Other feedbacks

- **It would be nice if a DPM could self-test (run `dpm-tester.py` on the host).**
- **Given many factors, the end-of-support for legacy was considered too short notice.**
 - Responses range from a lack of sense of urgency to a reduction in trust in the service.
- **Lack of xroot publishing, noted problems with gridftp on dual-stacked servers worrying for new user groups.**