

Storage Group and Tier-2 Developments

New Old Things

dCache 4.1x + dCache in general

- GridFTP TCP connection reuse (like GridFTP pipelining).
- Improved messaging and logging (for Elasticsearch compatibility)
- dCache's external communications have been emphasising "distributed storage" more than anything else.
 - see, later in talk, "Data Lakes" and other buzzwords

DPM 1.10, 1.x

- Bugfixes...
- DOME is now 100% complete dmlite replacement for rfio.
 - "distributed DPM" now more possible [see Italy]
 - SRMless DPM [DPM devs request a site to test this against ATLAS]
- xrootd is now a required component (and must be $\geq 4.8.0$)
- Caching "Volatile" storage pools [transparent proxy to a remote store... like a "Data Lake"]

The Future?

"Data Lakes"

- Big central (distributed? federated?) repositories of data, often stored as objects.
 - In danger of becoming a buzzword, caveat emptor.
 - (old idea, repackaged - see ZEPHYR etc)
- Aims: Reduce number of endpoints
- Must: Rely on smart caching to work well with remote sites.
- Must: Rely on intelligent authentication
- "Static" data on Lake can be intelligently stored - for ex. erasure coding, not wasteful replicas.

Capabilities not Identities

- Problems with Identity Delegation-based security systems
 - see, e.g. RFCs on X509; our talk from several years back; Brian Bockleman's SciTokens talks etc
- Capability-based models more widespread in real world
 - AWS etc
- SciTokens are part of the post-GridFTP transfer model (which also feeds into those Data Lakes)

T2 Models

Types of Tier 2

- No Storage
 - seems to be ATLAS model for several small T2s
 - (has always been LHCb, CMSs)
- Caching
- "Full Fat" storage
 - Internal cache for performance?

ARC Caching

- Durham (and poss. Sussex?) are the only ARC Caching systems in the UK
- Works well for them, for ATLAS work. (Needs support from Atlas-hosted ARC Control Tower to shuck workload from pilots.)
- Unusual -> harder to support.
 - Recent issue with ACT config caused confusion in tickets
 - Shifters not familiar with ARC Cache sites, don't know how to ticket appropriately (which isn't their fault)

Other Caching approaches

- Internal Xrootd Cache
 - ECDF (see Teng's talk) for ATLAS + other VOs
- External Xrootd Cache
 - RALPP (see Chris' talk) for CMS
- In both cases, hit rate is important to judge effectiveness (unlike preemptive caching)