



WLCG Demonstrator

WLCG storage, Cloud Resources and volatile storage into
HTTP/WebDAV-based regional federations

Report: GDB November 2016



Project motivation

- The project addresses two related topics
 - Exploitation of cloud storage (AWS S3, Ceph S3, MS Azure)
 - Consolidation of existing storage via a local federation
 - within a single site, OR
 - between a few ‘friend’ sites showing up as one

Teams and Experiments

- ATLAS Canada: UVIC and others
 - Evaluated dynafed in the last 2-3 years with small pilot projects
 - Got approval for funding a regional/national project involving 3rd party Cloud storage
- Belle-II
 - R&D project on HTTP and federations
 - Also UVIC is involved as one of the sites
- ATLAS Italy: Naples+Frascati
 - Previous experience (using PROOF+Dynafed) described at CHEP 2015
 - Evaluating various technical possibilities for Tier-1 and Tier-2s

- Belle-II : The R&D has full endorsement, some results have already been circulated
- Some ATLAS teams participate with their own funding, likely to increase
- ATLAS central ADC is an observer, will work with the new macro-sites

Why a demonstrator?

- Promote the visibility of the participating projects
 - Attract other contributions and create a focus for related work
- Promote information exchange on these topics between different groups of people
- Publically address the relevant questions
- Why should experiments care?
 - Help projects which are trying to help them
 - Low-cost initiative (will in fact attract resources)

Targets

- Understand if local HTTP federations can be advantageous
 - Usual fed advantages in a robust and user friendly way
 - Storage efficiency, fallback, redundancy, regional consolidation
 - **Non-intrusive**; no need to reconfigure participating endpoints, which can also be accessed as before
 - Very low cost for joining sites or remote services
- Understand how cloud storage can be exploited
 - Following the model of a bridge which makes it look like any other (HTTP-enabled) grid storage
 - No need to change applications
 - No distribution of sensitive keys
 - Could support site or in-cloud CPU
 - Or following the model of a cache in front of it ?
- What if one wants to support the current SRM-based apps ?
- Focus on ‘opportunistic’ storage, can change at the pace of shorter term contracts

Tech choices

- The teams work autonomously exchanging information often
- The teams chose to evaluate two technologies: Dynafed (Belle-II, Uvic) and DPM 1.9 (the latest one with volatile pools capabilities)
 - Both techs share a common DMLite framework
- Dynafed does federations of S3/Azure/WebDAV, aggressive in performance and management simplicity
- DPM does site storage
 - Servers can be distributed across sites since ever, never exploited so far AFAIK
 - A side effect of the last release makes it easier than before, clustering through HTTPS
 - It could use cloud storage as a backend in various way, or it can be installed in the cloud itself
 - Potential full set of Grid features over multi-site and clouds, at the price of a more “static” data management approach
- The two techs can be mixed into a fed of multisite DPMs

About this report

- Each of the three teams contributed a short presentation on the status of their explorations and plans
 - The correlation between their works is pretty evident
- For practical reasons I will read their slides and push questions to the end
- The authors are available here at the GDB, either in person or through Vidyo



★ Somehow volatile

Regional/National multi-tier HTTP storage

