

Storage Resource Reporting

Draft proposal v0.1
Alessandro Di Girolamo, Oliver Keeble

Introduction

Storage accounting for WLCG is under review as part of a more general reconsideration of accounting and the accounting portal. This document proposes a small set of requirements for storage providers, intended to satisfy the needs of the accounting system, and at the same time fulfil the related storage resource reporting requirements of the experiments.

Alessandro's presentation at the GDB

https://docs.google.com/presentation/d/1hJ_bh8TzBdxLsi2WDIXaIzoUtySg_OHn2t4OtVMtg8w/edit#slide=id.g12418cc379_0_68

The GDB agreed to that the approach should be to finalise consumer side requirements (from the experiments and WLCG management) and then take this to the storage providers. The stakeholders "just" need to say what they need, and how often.

Proposal

The interface and structure are not defined here, this is left to the storage providers, with the constraint that for any given protocol all should implement the same interface.

1 High level overview of the total and available space.

- Total used and total available space for the following
 - All distinct "Storage Areas" available to the experiment
 - This will allow calculation of total volume provided by the site to the experiment
 - How do clients know where these are? (e.g. if 10TB is "mounted" deep in the tree?)

- Any entity on which a quota can be applied (typically a directory or a space token)
 - This will allow clients to understand if they can write
- The top level of the experiment part of the namespace
 - tbc – in the HTTP Task Force this was considered to be a requirement
- It can be at the level of the high level quota nodes, grouping the spaces with thousands of them if needed (e.g. EOS user areas)
- Number of Files would be interesting but is not required.
- What is required for tape?
- Query frequency - can be up to half an hour (not Hz)
- Accuracy
 - volume - order of tens of GB (i.e. experiments are not picky on super precisions, Storage providers should comment on what's doable with a limited amount of complication).
 - time – a freshness of tens of minutes or so is acceptable
- CLI - is a command line client required, or is the API sufficient?

2 More detailed storage information

Full storage dump enumerating each file

- path
- size
- accesstime
- mtime/ctime?
- Checksums?

To be provided on weekly timescale when not possible with higher frequency (e.g. EOS find allows, in a couple of hours, of having the full dump. CMS run it daily).

Here the aim is to allow a single utility per storage system which will work for all interested experiments. Are the requirements above sufficient for all?

Usage for Accounting

Only scenario 1 above is required for accounting.

Experiments will be provided of a client (as today lcg-stmd, gfal-getattr) which will allow to get the information described in the scenario 1. The workflows of the experiments for what concern the space accounting won't change, but agreeing on the format, structure and content of the information exposed will enable the possibility to setup a WLCG collector in parallel to the experiments workflows to collect the Storage Resources accounting information.

For transport to the accounting aggregator (currently APEL), there are a couple of options

- Storage systems emit accounting records and send them to APEL
 - There is a varying amount of existing support for StAR accounting records, and this standard is under review by EGI who will provide the accounting portal
- A “pull” model where APEL or an intermediate service gets the data from the endpoints and creates the relevant records