

Some thoughts on accounting

LCG GDB
May 18th 2005

Note: Personal view only; there was no meeting at CERN yesterday

Tony.Cass@[CERN](mailto:Tony.Cass@CERN.ch).ch

Why accounting?

	Resource Use	Policy Adherence	Efficiency	Planning
Funding Agencies	Is our money wasted?	Are the resources being used as we expected?	?	Surely they've got enough?
Users	Where can I get a free lunch?	Why are they getting more than me?	Why do my jobs always take so long?	No we haven't! We could do so much more if only we had...
Grid Managers	Are any sites under used?	?	?	?
Site Managers	Demonstrate that we're wringing every last drop out of our underfunded kit.	Because they're supposed to. We are fulfilling the SLA	Writing small files to tape is a BAD idea!	Don't be daft. I'm cash limited

CPU related accounting

- ◆ Total capacity
- ◆ Average normalised power and spread
 - ?? Something needed here for comparison across sites
- ◆ Average usage
 - Locally submitted
 - Grid submitted
- ◆ Per VO
 - Usage, absolute & percentage
 - » Locally submitted, Grid submitted
 - Target usage percentage
- ◆ **Requires an agreed normalisation**
 - But so does anything CPU related...

Job related accounting

◆ Globally and per VO

- Number of jobs
 - » Locally & Grid Submitted
- Average duration
 - » CPU, wallclock & average ratio
- Average queue time?
 - » This will depend on queue structure at sites, but some number here is interesting

◆ Job success rates, though interesting, cannot be recorded meaningfully.

- Do all batch schedulers log an exit code?
- Successful job scheduler exit does not mean successful job completion from user viewpoint.
- Unsuccessful job scheduler exit can include user induced error status. Not correct to attribute this to a site.

Disk storage related accounting

- ◆ Total capacity
- ◆ Average occupancy?
 - Probably 100%...
- ◆ Per VO
 - Usage, absolute & percentage
 - Target usage, absolute & percentage
 - Average file size for files on disk
 - Volume read/written to/from disk by jobs?
 - » May not be easy to collect
 - ◆ Instrumentation of rfio/(x)rootd, searching of logs
 - » Even harder, but interesting, is fraction of a file that is read
 - ◆ Sparse access to large files...

Tape related accounting

- ◆ Capacity of current media
 - Used media at actual capacity, unused media at theoretical maximum
- ◆ Actual volume on tape
- ◆ Per VO
 - Volume on tape
 - Compression ratio
 - Average file size on tape
 - Volume read/written to/from tape
 - Average file size read/written
 - Tape mounts, write and, especially, read

Network related accounting

- ◆ Traffic volumes
 - Globally
 - To named sites
 - ???

- ◆ Per-VO breakdown
 - Presumably not possible to generate at network level.
 - SE generated statistics?

Granularity

◆ Weekly

◆ Daily

- Possible in principle, but necessary?
- If so, also need to collect weekly figures from sites
 - » At least for CPU figures; edge effects have greatest impact here.

Final Provocative Thought

◆ Accounting is nothing but Monitoring

- Set up a central Lemon Repository
- Sites create "sensors" to inject agreed "metrics"
 - » e.g. CPU time used per time period.
 - » can have multiple time periods (day/week) as different metrics
- Provide central extraction & display tools a la Lemon