

IT/CZ cluster meeting

Job Management Coordination Meeting

Data persistently stored by the Workload Manager

M. Cecchi (presenter), F. Giacomini

Workload Manager: all the file i/o we currently use:

- input (r) (submit, resubmit, _cancel_, [not the match persistently])
- output|w (ice|jc, filelist, jobdir)
- configuration (r)
- jw template (r)
- jw scripts (w)
- brokerinfo (w)
- dagman submit files (w)
- ism dump (w)
- wm_log (w)
- wms_proxy (r)
- user proxy (r)
- pid_file (rw)

Open questions

- The method must be 100% safe, of course.
- We are basically interested in:
 - Input
 - Output
 - we can also evaluate if/how to store non job-specific information, and where it makes sense
- Persistency will come in this case with atomicity, which is very good.
 - Jobdir dramatically reduces the level of contention as compared to the old filelist, how is this granted by the lbproxy/db?
 - New, old, tmp?
 - Also admin tools are required to handle the i/o
- Possible mechanisms for notification
 - Ping from each involved service (wmproxy, lm, dag planners)
 - Notifications from the lbproxy (other services simply remove code)
- Also, would it be possible – as an optimisation - to log events for multiple jobids?