



Open PBS backend bugs

- **Bug 3503** - PBS scan does not parse node information
 - submitted by DESY to GGUS
 - not enough info in the PBS logs used by ARC
- **Bug 2077** - scan-pbs-job is much too slow
 - dated back to 2010
 - pure scalability of log parsing logic
- Rewriting the parser is a good time to rewrite parsing logic
 - **Kill two birds with one stone**



Current scan-PBS-job





- LRMS job IDs and states are parsed from:
 - server_log files (Finished/Failed)
 - qstat (Queued/Running)
- Then for **each** LRMS ID `find | sed | grep` to find corresponding grid ID in control dir:
 - huge controldir I/O including stat cache killing
 - scalability issues with increasing number of jobs
 - issues with faults (parsed but not updated)



Baseline

- NSCMBR cluster with many local jobs and grid jobs (number of jobs vs. parse time)





Experimental rewrite



- Parse ***accounting*** logs instead of server logs
 - contains more info
 - more structured
 - unified information source
- Pre-fetch grid-Irms id mappings once
 - eliminate find|sed|grep scalability



How to deploy

- Experimental code requires testing
 - will be installed on the several Ukraine clusters for testing purposes
 - does not cover all cases anyway
- Implement “PBS_new” backend
 - noone will use it
- Move old backend to PBS_old and keep distrubuting for eventual need of fallback :)
 - invoke old one scan from new one with config switch