



Traceability and isolation within WLCG

Input from CMS



What are the logs for each job reported to your framework? (1/2)

We have **pilot logs** and **payload logs**

- ◆ pilot log goes to the pilot factory machines
- ◆ payload log -> could be called the CMS “job wrapper log” (in the following)

Pilot logs

- ◆ describe the lifetime of a pilot i.e. how many jobs it ran, what was the idle time, how did the validation scripts go, etc
- ◆ kept for 1 week, but some factories save them elsewhere before cleaning

Payload logs from **production** jobs

- ◆ job wrapper log stored on scratch disk at the execution node, collected asynchronously, tarred up with others and archived on tape
- ◆ New developments in WMAgent (WMArchive) will have easier access to logs as a goal. Still work in progress, and in any case long-term storage will only be on tape

Payloads for **analysis** jobs

- ◆ job wrapper log (from CRAB) is usually brought back to the submission node and kept on disk for 2 weeks



What are the logs for each job reported to your framework? (2/2)

A short summary of some parts of the job wrapper log is brought back by both CRAB and WMAgent

- ◆ as a JSON object, and is available on submission machine for some time (e.g. 2 weeks for CRAB).
- ◆ This part could be customised to add more info, if needed.

Finally all HTCondor ClassAds are currently collected in a test instance of ElasticSearch+Kibana at UNL (Nebraska) and kept for a month.

- ◆ This is a test project and there is no established procedure nor assigned manpower from CMS to keep it doing long term, port it to CERN ES, develop tools and dashboards etc.
- ◆ But ideally questions like "who was the owner of payload executed by a CMS pilot on node X at time T in site S" could be answered via those ClassAds

It should be noted that those ClassAds expose user DN's, directory names and all sort of information that so far WLCG has insisted must be kept private and not circulated around. Same for factory logs. So it is not known whether we can make those more generally available.



Is searching into these logs feasible? (e.g. cost of asking who was running a job on a given WN at a given time)

See previous

- ◆ current logs are kept for limited time, even when on disk, distributed on $O(20)$ submission machines located at CERN, UCSD and FNAL. Sounds like "no, not feasible".
- ◆ If someone wanted to do the search, I guess the best we can do is to upload all logs to their favorite http URL and wish good luck.
- ◆ So far this is not a priority in our work plan. Sites can find all they need to know in the **glexec** logs on their WN's and/or in their central ARGUS logs. Searching syslog of node X for time T seems much simpler and fast for a site admins



How do you isolate/protect the pilot job from the user payload?

We use glexec

- ◆ which was designed exactly for this and does its job just fine.



How can central suspension be translated to VO suspension

Revoked certificates cannot submit new jobs

glxexec to do a second check for jobs that have been on the queue before executing them.