



Pilot jobs & policy

WLCG MB

24th Nov 2009

Ian Bird
LCG Project Leader





Policy

- Two years ago we (WLCG) agreed:
 - “WLCG sites must allow job submission by the LHC VOs using *pilot jobs* that submit work on behalf of other users. It is mandatory to change the job identity to that of the real user to avoid the security exposure of a job submitted by one user running under the credentials of the pilot job user.”
- But there were pre-requisites for the policy to come into force:
 - Security review of glxexec: **done**
 - Review of pilot frameworks: **done**
 - Glxexec tested with all needed batch systems: **done**
 - SCAS completed, certified, deployed: **not deployed**
 - NB. forgot glxexec deployed: **not done!**
- There is a MUPJ JSPG policy:
 - <https://edms.cern.ch/file/855383/2/PilotJobsPolicy-v1.0.pdf>
 - Approved in Aug and Sep 2008 by WLCG and EGEE respectively
 - This states that the VO must obtain Site approval before submitting MUPJ (i.e. optional from the site point of view)
 - **It also states that the pilot framework must meet the requirements of fine-grained traceability, e.g. via identity switching.**



Implementation & Deployment

- Takes forever ...
- glexec:
 - Has been available in SL4 for some time, but we need SL5
 - Several glexec-SL5 patches failed in certification and had to be sent back to developers
 - glexec/SL5 is now certified
- SCAS on SL5:
 - Has been a long time coming
 - Passed certification and started roll-out (but very recent!)
- While sites have been encouraged to deploy glexec/SCAS they have also been urged to move to SL5 and so lack of deployment of glexec/SCAS is understandable



Situation now

- Sites have not been able to deploy glxexec/SCAS until now to support the implementation of the policy
 - Not their fault
 - Anyway, is now the time to add such critical services?
- Experiments need to run MUPJ to support their analysis
 - But adaptations to glxexec/SCAS have not been fully implemented/tested (see above!)
 - Anyway, is now the time to make such changes in essential software?
- But:
 - The experiments need to process and analyse data according to the ways that they have been preparing
 - The sites need to be able to manage their security



Proposal

- In the absence of any deployable fine-grained authentication mechanism that the experiments can make use of now I propose:
- Establish a policy that:
 - Permits the experiments to run pilot jobs that run workloads of other members of the VO
 - Agrees that the owner of the pilot job – or the VO itself – is responsible for all work run by that pilot job
 - i.e. in case of problems the entire VO may be banned at a site
 - The VO framework should provide the fine-grained traceability
 - This policy must be reviewed in **xx** months, or earlier if there are operational issues
 - During this time the existing JSPG policy would need to be suspended (?)
- We must continue to push the deployment of glexec/SCAS as rapidly as possible to be in a situation to implement the agreed policy at the earliest opportunity.