Multi-User Pilot Jobs

John Gordon
Version 2.1 24/11/2009

Summary

Some experiments have started using multi-user pilot jobs without informing sites, risking breaking the trust that enables the grid to work. The relevant security policy mandates that pilot job workloads run with the identity of the originating user. This requires glexec for which there is no certified SL5 release. We thus have an impasse with experiments wanting the efficiency of multi-user pilot jobs and the sites wanting the traceability and auditability of identity change and glexec not being available for the bulk of the resources on EGEE. The WLCG Management Board should consider the options and offer guidance.

Background

At the MB of 6th November 2007 a WLCG policy on pilot jobs was agreed (an extract from the minutes showing the agreed policy is appended).
https://twiki.cern.ch/twiki/pub/LCG/MbMeetingsMinutes/LCG_Management_Board_2007_11_06.htm

This policy required 4 prerequisites before coming into force. These were:

1. A security review of glexec. (DONE) By John White of ETH
2. A review of the pilot job frameworks of the LHC experiments. (DONE) By a working group led by Maarten Litmaath.
3. glexec tested with all required batch systems. (DONE)
4. SCAS completed, certified and deployed. The first two DONE, the third started.

In 2007-08 WLCG and EGEE approved a 'Policy on Grid Multi-User Pilot Jobs' https://edms.cern.ch/file/855383/2/PilotJobsPolicy-v1.0.pdf which is optional on general EGEE sites but mandated by WLCG (above) for sites which support the LHC VOs.

The experiment have long planned to use the efficiency of 'late binding' of workload to batch worker nodes provided by multiuser pilot jobs. Such frameworks provide the experiment with the ability to make maximal use of a worker node which the pilot job has discovered to be working and correctly configured, thus reducing the number of failing jobs. It also allows implementation of grid-wide priority management bypassing the latency of batch queues and resource brokers by only deciding what to run in a job after it had started.

Current Position

At the Grid Deployment Board in November 2009 the experiments set out their positions.

LHCb

1. Production jobs: from the beginning, productions jobs were run under generic credentials (that of the production manager, used for the pilot submission) while jobs could actually be submitted by other individuals (shifters). This we don't consider as multi-user pilot jobs as it is just by chance that a job is submitted by a shifter or another.

2. Pilot jobs can be submitted in two flavours: generic (a.k.a. multi-user) pilots and private. The former are used whenever the sites accepts multi-user pilot jobs by accepting the VOMS role=pilot while the latter are used at other sites (all Tier1s are in the first category). Those pilots are running production payloads as well as user jobs, and we allow "filling mode" (i.e. getting another payload after a first one was executed, up to a certain number, and provided there is some time left). This is valid at any site that accepts role=pilot, not only at Tier1s.

If a site does not accept generic pilots (i.e. doesn't accept role=pilot) we shall not ban it but send private pilots (that can only run payloads with matching credentials). We have been waiting several years for glexec so don't expect that we are going to swap to using glexec in the first year of data taking even when it is deployed. Current experience has shown that it works with well configured sites but takes time to debug.

CMS

CMS currently uses pilots for production work which does not need to change identity. Person submits with own certificate. Analysis uses direct submission with WMS/Condor-G, but a CRAB version compatible with WMS-Glidein is being worked on so this may become an issue. (subsequent to the GDB it became known that there is a production CRAB server already submitting multi-user pilot jobs without identity change). However, WMS Glideins are already used in OSG with glexec on WN so this should be ready for EGEE sites deploying glexec. CRAB can use either: Glideins, WMS, or Condor-G so there is no need for CMS to force the use of Gliedins before sites deploy glexec.

ATLAS

ATLAS has for some time been running single-user pilot jobs where all the jobs run the work of one user who is the owner of the pilot factory. ATLAS interpreted the feedback of the GDB review as a green light to run multi-user pilot jobs without identity changing and have been doing that for several months. If sites really do not want multi-user pilots for ATLAS they should say and that site will be marked only for production (disk will not be used so ATLAS will not place data on that site).

gLexec with Panda has been shown to work; however, no site is yet using it in anger with ATLAS analysis pilots and few sites have even tried to install it

Alice

Alice track the identity of the submitting user in their own database and then submit all work under a generic identity. This breaks the Multi-User Pilot Job Security Policy

Middleware

SCAS has been certified and sites asked to deploy it in accordance with point 4 in the policy.

gLExec has been certified for SL4 but bugs were found in the SL5 port. These have been fixed and the patch was certified on 23/11. This version of glexec should appear in an EGEE release this week. Thus the majority of EGEE resources on SL5 currently do not have a certified WN release including gLExec.

Security

The policy document referenced above is very clear

Even a single user pay load pulled down by a pilot job running under another identity counts as a MUPJ. The policy definition says "when this Grid job begins to execute at a Site, it pulls down and executes workload, hereafter called a user job, owned and submitted by a different member of the VO or multiple user jobs owned and submitted by multiple different members of the VO."

It says: "Before submitting pilot jobs to a Site the VO must have approval from the Grid and from that Site."

"Each pilot job must be the responsibility of one of a limited number of authorised and registered members of the VO. The VO is responsible for implementing a process for authorising pilot job owners and ensuring that they accept the conditions laid down here."

The first point has obviously been broken and I have my doubts about the second. We should be honest about it.

1. Experiments will often tell you that they trust all their users. Fine. But user accounts are compromised. This happens regularly. Operational security has to maintain the ability to monitor and control individual user activities, otherwise when an attack happens they may have to ban the whole VO and if they cannot work out what happened they can never fix the problem and perhaps never be able to re-enable the VO.
2. If the user job runs under the same identity as the pilot job framework, the user then owns the framework. They can modify its behaviour in any way they wish. They can tamper with audit logs. They may enable back-doors which are only activated weeks later, etc etc.
3. Point 6 of the Pilot Job security policy says...
   "The pilot job must respect the result of any local authorisation and/or policy decisions, e.g. blocking the running of the user job."
   If a site has banned a particular user, they don't want the experiment to run work for that user.

Sites

The sites wish to know which individuals are using their resources. This is not only mandated by the policies above but is a legal requirement in many countries. Multiuser pilot jobs with identity changing of each payload by use of a setuid program called gLExec would seem to meet this requirement. The use of a setuid program is itself contentious and some sites may not have sufficient control of the resources they make available to WLCG to implement this (for example a physics department with resources owned and managed by an IT department on behalf of the whole university). It was to force this issue that the WLCG policy mentioned above was formed.
Until today the sites were in the position of being asked to deploy gLExec but not having an SL5 release.

What Next?

By moving from single-user pilot mode to multi-user mode, the experiments are in breach of trust with the sites. The sites should be told formally that this is happening now. If we wish to suspend the policy for a fixed period we must announce this.

The sites have the right to ask the experiments to cease and desist running multi-user pilots on their site. The experiments should either make this happen or publish how a site disables itself. Given the closeness to data taking, disabling multi-user pilots should not disable all use of that site by an experiment.

EGEE SA3 should be asked to treat the certification of the SL5 version of gLExec as top priority. (certification was completed on 23/11, this will be in a new release next week).

The sites should be encouraged to continue deployment of SCAS and gLExec and tell the sites when they can start using multi-user pilots with identity changing at their site.

The experiments should switch on identity changing in their frameworks as soon as possible.

WLCG policy on pilot jobs submitting work on behalf of third parties

The topic of pilot jobs has been discussed several times in the GDB, and in particular at the last two meetings. At the October meeting it was agreed to make a proposal to the MB to adopt a policy requiring that sites support pilot jobs submitting work on behalf of third parties.

A summary note was prepared by J.Gordon (17/10/07) and presented to the MB on 23 October. This identified a number of issues and made recommendations for a pilot job policy. After discussion the following policy statement proposed and endorsed by the MB meeting on 6 November.

> 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.
>
> Implementation of this policy is subject to the following pre-requisites:
>
> 1. The identity change and sub-job management must be executed by a commonly agreed mechanism that has been reviewed by a recognized group of security experts. At present the only candidate is *glexec*, and a positive review by the security teams of each of the grid infrastructures (OSG, EGEE) would be sufficient.
> 2. All experiments wishing to use this service must publish a description of the distributed parts of their pilot job frameworks. A positive recommendation to the MB on the security aspects of the framework by a team of experts with representatives of OSG and EGEE is required. The frameworks should be compatible with the draft JSPG *Grid Multi-User Pilot Jobs Policy* document.
> 3. *glexec* testing: *glexec* must be integrated and successfully tested with the commonly used batch systems (BQS, PBS, PBS pro, Condor, LSF, SGE).
> 4. *LCAS/LCMAPS*: the server version of LCAS/LCMAPS must be completed, certified and deployed.

The policy will come into effect when the MB agrees that all of the above pre-requisites have been met.