GDB

Pilot Jobs + glexec

John Gordon, STFC-RAL

GDB meeting @CERN, 7 November 2007

Recent Progress

- Multi-User Pilot Jobs Policy discussed in MB, JSPG, OSCT
- At October GDB we decided to refer the issue toMB for a decision.
- MB defined a draft policy
  - Requires sites who support LHC experiments to support their pilot jobs frameworks.
  - Mandates jobs running under the identity of the work, not some generic pilot job framework identity.
    - Glexec in setuid mode seen as the only candidate for this functionality.
  - Defines work to be done and agreed by MB before policy becomes effective.

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 jobs policy
- After discussion the following policy statement proposed and endorsed by the MB meeting on 6 November.

MB Policy

- 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:
  - 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.
  - 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.
  - *glexec* testing: *glexec* must be integrated and successfully tested with the commonly used batch systems (BQS, PBS, PBS pro, Condor, LSF, SGE).
  - *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.

JSPG(1)

- Summary of discussion on Pilot Jobs and related issues
  - (D Kelsey's personal notes - not yet approved by JSPG)

- Discussed current draft policy (v0.3) on "Multi User Pilot Jobs"
  - See https://edms.cern.ch/document/855383/1
- Discussion focussed on **whether to require Sites** to run the Grid-provided utility (today glexec) in the identity switching mode
  - or should Sites should have the right to choose
- The view of the participants in the meeting room at CERN was
  - there **are significant security risks** in not switching identity
  - the users' workload runs under the same identity as the pilot job framework resulting in the ability of users to take control of the framework and to interfere with the audit logs
- Should therefore **require identity switching**
  - But, this is contentious and may not be acceptable
- The view of the OSG participants, including Ruth Pordes,
  - had not been able to consult their Sites or discuss this suggestion internally
  - could not agree to the proposal to require identity switching at this time
  - They have concerns as to what extent they can impose such policies on their Sites.
- JSPG concluded
  - a very contentious issue
  - need to discuss with the Sites and the Grid managements
  - Can consensus be reached?

JSPG(2)

- Also discussed the requirement that pilot jobs must clean up all local data files between different user jobs within one pilot job.
  - Strong requirement coming out of discussion at the August GDB meeting
  - but Oxana noted that this may conflict with experiment requirements
  - JSPG agreed that there are security concerns, e.g. with one user job leaving infected files which the next job may pick up.
- Again we concluded that further consultation is required.
- Did not produce a new version of the policy document
  - pointless until these issues are resolved
- On the second day revisited the topic to decide best way forward
- Concentrate on the **requirements for traceability and logging**
- These are **general requirements** which apply not only to multi-user pilot jobs, but also to all other forms of job submission including, for example, Grid portals
- get agreement on these **general principles** which can then be applied to the consideration of any particular service, such as pilot jobs

JSPG(3)

- Draft words in new **"Policy on Traceability and Logging"**
  - This will replace the old policy on "Audit Requirements"
  - The words are not yet final and still need more work.
- The main points are:
  - **Risk management is crucial for Grid operations.**
  - When security incidents happen it must be possible to **identify the cause** so that it can be contained while **keeping services operational**. It must also be possible to take action to prevent the incident happening again.
  - The response to an incident needs to be commensurate with the scale of the problem. **Banning a whole large VO**, rather than just one user, **is in most cases impossible** as this affects too many users.
  - Understanding and **fixing** the cause of an incident is **essential before re-enabling access**. If a VO has been blocked in its entirety, unless there is a way of understanding what happened, it will be impossible to fix and hence **impossible to re-enable access**.
  - The minimum level of traceability for Grid usage is to be able to **identify the source** of all actions (executables, file transfers, pilot jobs, portal jobs, etc) that are part of an incident and **the individual who initiated them**. In addition, sufficiently fine-grained controls, such as blocking the originating user and monitoring to detect abnormal behaviour, are necessary for keeping services operational.
  - There are **trade-offs** between knowledge of individual user identity and controlling the code which can be executed. **Anonymous user access** may well be possible if there is **no** possibility for the **user to submit or modify executable code**, i.e. the user just provides input parameters and/or data to control a job.

Multi-User Pilot Job Policy

- Draft 0.3 https://edms.cern.ch/document/855383/1
- Should be completed
- But the ideas in it can be used to help review the experiment frameworks.

Review glexec

- 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.
- OSG - Don?
- EGEE - John?

Experiment Frameworks

- All experiments wishing to use Pilot Jobs 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.
- Do all experiments agree?
- Are frameworks sufficiently documented?
- Who will form the team?

gLExec and batch

- *glexec* must be integrated and successfully tested with the commonly used batch systems (BQS, PBS, PBS pro, Condor, LSF, SGE).
- Such testing can probably be done now using sudo.
- Who will do this?
  - *BQS – CC-IN2P3*
  - *PBS*
  - *PBS pro*
  - *Condor*
  - *LSF*
  - *SGE*
- *When?*

LCAS/LCMAPS

- *Local mode LCAS/LCMAPS does not scale to large sites*
  - *Client-server version required*
- *OSG uses GUMS – they have a common API*
  - *Only one version of glexec required*
- *LCAS/LCMAPS*: the server version of LCAS/LCMAPS must be completed, certified and deployed.
- Timescale?

- JRA1
- SA3
- SA1

gLExec deployment

- Status of certification
- PPS testing
- Safe configuration
- Does this need to wait for server mode LCAS/LCMAPS?

Timescale

- glexec review
- Framework Review
- Batch system testing
- Glexec certification and

Any Other Issues?

GDB

Science & Technology Facilities Council

Summary

- The MB Policy aims to break the current deadlock
  - One way or another
- Still much work to do before the policy becomes effective
  - Timescale
- But … want to know now if sites will not allow glexec in setuid mode under any circumstances
  - Tell Les and John
- JSPG opinions, review of glexec and of frameworks could all help persuade institute management.