Multi-User Pilot Jobs – GDB 7/11/2007

The MB Policy and its pre-requisites were discussed at the GDB on 7th November.

One error was spotted in the item on review of the experiment frameworks. Reviewers should consider **all** the framework, not just the distributed parts.

There was a lot of discussion but mainly for clarification. I stressed to the national representatives that they were representing all sites in their countries, not just their own. If there is great resistance to running glexec in setuid mode from the majority of sites then we should know sooner rather than later. It was stressed yet again that pilot jobs which only run under the identity of one user, the owner of the payload, are not an issue. They do not change identity. It was revealed that even single user pilot frameworks transfer and store proxies themselves so perhaps they should also be subject to a security review. This will be reported to the MWSG but ignored here for the purposes on this report.

Security Concerns Romain Wartel and Dave Kelsey reported from recent security meetings. The JSPG had discussed Multi-User Pilot Jobs . The view of the participants in the meeting room at CERN was that there are significant security risks in not switching identity. The users' workload running under the same identity as the pilot job framework would result in the ability of users to take control of the framework and to interfere with the audit logs. We should therefore <u>require</u> identity switching. OSG representatives felt that they needed to consult more widely. The logging-only mode of glexec is now considered to be unsafe for the reasons above and is not proven to be auditable especially when multiple payloads from different users run on a multi-core, multi-cpu node.

JSPG decided they should 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. They hope to get agreement on these general principles which can then be applied to the consideration of any particular service, such as pilot jobs..

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 issue if that 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.

I agree with this last point but I worry about the tactics of formulating a general policy which will then cover pilot jobs. I think this will take too long and I’d rather that pilot jobs are used as a use case to formulate a special instance of this general policy.

Review of gLExec. John White of EGEE reported that two security experts (Andrei Kruger and Alexander Yu) had reviewed gLEexec. They were JRA1 security developers in EGEE. They raised a number of issues which have been passed to the developer, SCG, and GSVG but found no showstoppers. In their opinion it is ready to start being tested by some tame sysadmins and then proceed to certification

gLEexec Certification

There is some work still required on glexec before certification.

- It needs to use syslog. This work is underway.
- YAIM needs to configure gLExec and LCAS/LCMAPS to understand and authorise gLExec, and the whitelist of accounts(s) authorized in glexec.conf

Testing is required with all batch systems. Volunteers were sought to test glexec with different batch systems and the following identified. CC-IN2P3(BQS), CERN(LSF), PBS(NIKHEF, CERN), SGE(CESGA), PBSpro(??), Condor(??)

LCAS/LCMAPS Service Version

The service version of LCAS/LCMAPS will be required for scalability before general deployment but this should not hold up testing with the shared filesystem version. JRA1 have a prototype based on alpha version of libs. Better libs by December. Shortly after this a version of the service and then one week later the client – ready for testing at that point. By end December there should be something ready for certification.

SA3 then need to figure out a deployment route. Should it run on the CE or be a new node type. Clarify with SA1. This will determine work for packaging. Testing less than 1-2 weeks. , deployment perhaps 6 weeks. This takes us to the end of February.

Review of Frameworks

The security concerns are not concerned just with glexec but with the whole framework running pilot jobs. The frameworks of the 4 LHC experiments need to be reviewed by a small panel. Points at issue include:-

- How proxies are handled and stored;
- How new jobs are launched from within the pilot job. Does this break any batch systems.
- Does the worker job tidy up after itself?

A small group of Ian Bird, Don Petravic, Dave Kelsey and John Gordon were actioned to choose a panel to review the frameworks. The first step should be for the experiments to present documentation of their architectures. The panel will then review this and then interview the relevant experts, perhaps with a questionnaire first. Having all 4 experiments on the panel might make it large but would share experiences.

Summary

Current status of the of the pre-requisites from the WLCG policy

a) glexec must be reviewed by a recognized group of security experts. **Status Done**
b) Document pilot job frameworks. **Status Not All Done**
c) Frameworks to be reviewed **STATUS Team still to be formed**
d) The frameworks should be compatible with the draft JSPG *Grid Multi-User Pilot Jobs Policy* document. **STATUS not tested**
e) *glexec* tested with the commonly used batch systems (BQS, PBS, PBS pro, Condor, LSF, SGE). **STATUS not tested**
f) *LCAS/LCMAPS*: the server version of LCAS/LCMAPS must be completed, certified and deployed. **STATUS Planned**

Progress will be reviewed at the December GDB.