



GDB

Pilot Jobs + glexec

John Gordon, STFC-RAL

GDB meeting @CERN, 10 October 2007



GDB

Pilot Jobs

- **We spend a lot of time discussing glxexec**
- **But do not forget that the real user requirement is for Pilot Jobs**
- **glxexec is a way to meet the sites requirements for identifying the user who runs the payload.**
- **If sites have further reservations about pilot jobs apart from glxexec then they should raise them now, as glxexec will be a waste of time if it does not permit pilot jobs**



GDB

Pilot Jobs Policy

- See Dave Kelsey's talk today



GDB

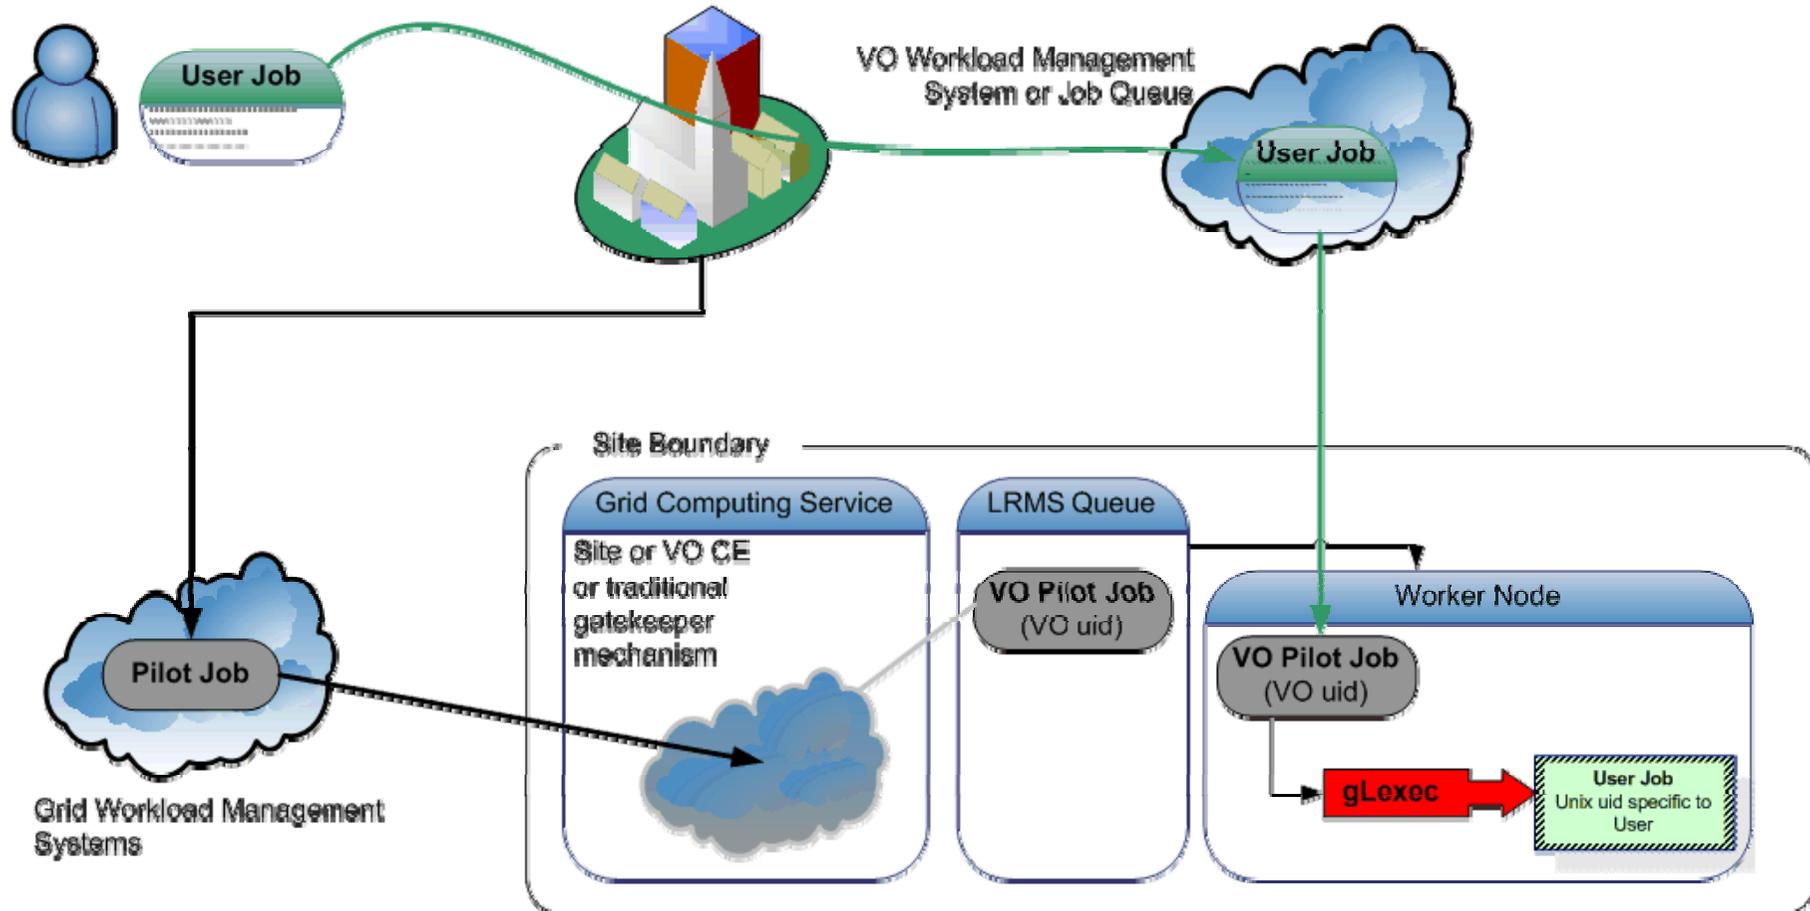
- David Groep gave two presentations at CHEP describing the technical background.
- <http://indico.cern.ch/materialDisplay.py?contribId=43&materialId=slides&sessionId=26&confId=3580>
- <http://indico.cern.ch/materialDisplay.py?contribId=92&materialId=slides&sessionId=24&confId=3580>



GDB

Pilot Jobs and gLExec

Virtual Organisation



On success: the site will set the uid/gid to the new user's job

On failure, gLExec will return with an error, and pilot job can terminate or obtain other user's job



History

- Pilot jobs were designed to deliver workload to sites which are known to work for a VO and to bypass delays in job submission.
- Pilot jobs which pull in jobs from a variety of users break the current Acceptable Use Policy in EGEE/WLCG/OSG(?)
- Some sites are content that identity change is auditable and that information is available to trace the owner of a job. Some sites insist that the identity of running processes must be transparent to the site.
- Glxec was developed at NIKHEF as an acceptable, auditable and secure way to change the identity under which a job runs.
 - It can run in three modes Do nothing, Log identity change, change identity of job
 - If glxec is installed in one of these modes at every site then pilot jobs can be set to use it.



Status

- Experiments were asked at a recent GDB what they will do if they cannot use pilot jobs at particular sites.
 - ATLAS, CMS, Alice said they could live with it.
 - LHCb said it was a showstopper
- Sites have a suspicion that VOs will run (are running?) Pilot Jobs anyway.
- Some sites have listed their concerns/objections to glxec.
- TCG permission is required to add glxec to glite release
- If this is done without consensus then it will not be deployed everywhere and VOs will not be able to rely on it.



Example Site Requirements

- The UK(I) does not accept glxec in its present state and will not recommend its deployment.
- Before we will reconsider its deployment we require to see:
 - a satisfactory code inspection wrt security.
 - an AUP with VOs agreed in principle by sites (in general through GDB) and all 4 LHC VOs.
 - a satisfactory analysis of the effect on the supported batch systems and accounting of identity changes.
 - glxec should be simple and safe to configure. Simple so that all sysadmins can understand what they are doing and safe so that misconfigurations do not inadvertently lead to security exposures.
 - The three flavours of glxec should exist as separate binaries and not as one with a configuration switch. Each of these should satisfy all package dependencies. e.g., dependencies should express some requirement for an abstract "glxec" property, which is satisfied by "glxec-null" or "glxec-suid" etc.



Recent Developments

- Lots of discussions at CHEP07 and EGEE07
- Sites are gaining confidence
 - I have heard several say they may be prepared to run it now
 - Perhaps not in identity switch mode.
- David has been generous with his time to answer queries and discuss issues.



Remaining Issues

- LCAS/LCMAPS on each worker node does not scale
- Library version now ready.



GDB

What Now?

- Are we ready to recommend to TCG that this be given a deployment priority
- Pass to SA3 for certification
- PPS
- Production
- Get Pilot Jobs Policy approved

- How long will this take?