



Pilot Jobs : What Next?

Jeff Templon

NIKHEF

GDB meeting

06 december 2006



Issues

- ◆ Basic pilot jobs (single user / single task)
- ◆ Multiple tasks in pilot jobs
- ◆ Resource Commitments and Pilot Jobs
- ◆ Multiple users in pilot jobs
- ◆ Accounting and pilot jobs
- ◆ Proposals

Basic Pilot Jobs

- ◆ Rationale : insulation for user payload
- ◆ Impact : small -- only real difference is that payload is deposited on user repository rather than RB
- ◆ More frequent site floods, but not fundamentally different than "normal job"
- ◆ "Normal" jobs are also pilots
 - Send wrapper to site
 - Wrapper pulls payload from RB sandbox ...

Multiple Tasks in Pilot Jobs (*)

- ◆ Only difference is that these tend to run very long
- ◆ Ability of site to handle them depends on
 - Partitioning site into semi-dedicated resources, or
 - Dimensioning allowed run time to fit cluster size
 - 72 hrs for 350 cores is a bit long ...
- ◆ No real problems, may simply cause spectrum of operational numbers like max wall times allowed

Resource Commitments & P.J.

- ◆ Problem: computers are real time, not CPU time
- ◆ Jobs run on nodes in real time, may or may not use CPU time
- ◆ Pilot jobs seem to be worse at "apparent CPU usage"
- ◆ Running kamikaze pilots uses wall time but no CPU usage
 - "basic cost == scheduling cycle" (Lyon)
- ◆ Fair share **MUST** be based on wall time : priority or fair-sharing based on CPU time allocation, rewards the most inefficient users!!!!

Multiple Users in Pilot Jobs

From LCG User AUP

1. *You shall only use the GRID to perform work, or transmit or store data consistent with the stated goals and policies of the VO of which you are a member and in compliance with these conditions of use.*
2. *You shall not use the GRID for any unlawful purpose and not (attempt to) breach or circumvent any GRID administrative or security controls. You shall respect copyright and confidentiality agreements and protect your GRID credentials (e.g. private keys, passwords), sensitive data and files.*

Doesn't help much ...

Multiple User Pilot Jobs (ii)

The user shall take the necessary precautions to protect his personal computer or work station against unauthorized access.

The user shall also protect details of his personal account, particularly by avoiding obvious passwords and shall not divulge his passwords to any third party, unless expressly authorized by his Division Leader. Upon request from the CERN Computer Security Officer or the service manager concerned, the user shall select a new password.

CERN Circular #5

Multiple User Pilot Jobs (iii)

- ◆ Not about security, about liability
 - Trust "VO" to handle things correctly
 - We (sites) know that certain users have "given others access" and we let them ... Can we convince anyone of "due diligence"?

- ◆ Need
 - Formal agreement between site and MUPJ operators
 - Residual checks at site (otherwise we are blind!)
 - User-Signed JDLs?
 - User delegated proxies + *Glexec-on-WN* with or without *suid* bit?
 - Some other brilliant idea?

Accounting and Pilot Jobs

- ◆ See earlier comments : can't just use CPU time ... Site purchases real computers
- ◆ MUPJs essentially make fine-grained accounting a "VO" problem ... Unless one wants to declare one of the MUPJ frameworks "standard middleware"
- ◆ Other sort of pilot jobs are no problem for fine-grained accounting

Proposals

- ◆ Try to clear up confusion : “pilot jobs” have been equated with “suid on WNs” in some circles!!
- ◆ Single user pilot jobs : declare OK!!
- ◆ Single user multiple-payload jobs : declare OK, but make sites and expt's aware of consequences ... Smaller sites will need to have very short queues, or reserved resources (no sharing)
- ◆ Recognize and approve that fair sharing must be based on wall times, and this places limits on what CPU shares can be reached, based on efficiency of workload

Proposals : MUPJs

- ◆ Make decisions.
 - Decide whether we can afford (or have a choice?) to split effort in two branches:
 - ◆ MUPJs as strategy for intra-VO prioritization
 - ◆ JPWG strategy for intra-VO prioritization (almost works now but probably not enough!)
 - ◆ Other branches? (Schulz Grid Cash)
 - Then decide whether we want to allow and support them!!!
Basically a manpower question ...

Proposals : MUPJs (ii)

- ◆ Do not allow *normal users* to run MUPJ frameworks
- ◆ These frameworks are services, and a service AUP is needed
 - Define the responsibilities of both parties
 - Define the liabilities
 - Do the service providers need to sign with each site? Or with the grid infrastructure??
- ◆ Implement some reasonable (and lightweight) system of residual checks at site level (I already gave examples)
- ◆ JSPG / Lawyers / ?? -- is this enough??