

GridPP 19 Collaboration Meeting, University of Cumbria, Ambleside

GridPP Discussion Session 1 ~ F2F ~ 29 August 2007

Chair: Tony Doyle

Panel A: Mona Aggarwal, Chris Brew, Stephen Burke, James Catmore, Greig Cowan, Brian Davies, Peter Hobson

TD invited the panel to raise issues of concern. PH cited user issues, that of no local access. There was tension between end use of the Grid and getting to that point, especially in small Institutions (eg: Brunel) which don't have the resources to run a local farm and grid cluster. TD asked about physics analysis, and whether the grid paradigm of a Resource Broker had been adopted, eg: some of CMS users use an RB. Doing physics analysis on the grid was preferable to local use. Dave Colling advised that this varied from community to community.

JC noted that for ATLAS B physics, there was Monte Carlo and co-ordinated analysis for CSE tasks - problems encountered had not been to do with Grid infrastructure, rather, they were from the ATLAS side in relation to jobs running, bureaucracy etc, but as far as they were concerned they were happy with the Grid as users. JC also noted that Ganga gave more control, which was a significant improvement. They use distributed technology to run analysis and this works well if the catalogue is correct. Regarding their experience with the Grid, Ganga was good, ATLAS scripts less so.

BD noted that users shouldn't need to have access to low-level commands - Ganga and DQ user tools should give all the information required. RJ advised that the philosophy in ATLAS was to pre-place data where possible. BD asked how this would stop users accessing files in their own way? Stalled jobs weren't the fault of any one thing. JC noted that fewer jobs fail but it was difficult to decipher why jobs fail - he has stats on fewer jobs failing but it was not easy to tell from the log files - you had to learn by experience what the error messages meant.

GC noted that in addition to storage, being a user of LHCb and analysing events, he had found that a lot of problems didn't relate to middleware, rather, experiment configuration of software at sites. TD asked if this was a site installation problem? RN confirmed that there had been problems recently but they had been addressed now. TD asked if end users received the message as to what problems prevailed? RN noted that a site could be banned, fixed, then unbanned again - this was less of a problem at Tier-1 level. SL advised that sometimes only the Information Systems were available therefore the job wouldn't go to that site as the software wasn't recognised, but this was erratic.

SB noted that SRM2, VOMS, GLite CE, GLite RB, were all in progress but were not there yet - we were trying to build whilst using and funding for middleware has gone - he asked how this process would be completed? Quick hacks to solve immediate problems would cause greater problems later on. DC noted that we had gone from a level of 60% of jobs failing, with a lot of work improved. CB noted that we were certainly not there yet - as a site administrator and a member of an experiment, it was useful to be able to de-debug experiments, but he noted that Ganga constrains you to do something in a certain way - this was helpful, and what the user wants. Regarding local access to resources, opinion was divided depending on user requirements. It was asked how to stop users crashing the SE? It was understood that you can break the Grid simply by using it. It was agreed that if local users demand access to the SE, you can't say no. TD advised that if we work with users, every use case should be able to be handled re data placement and use of the Grid.

There was a discussion of placing of datasets and the difficulty of users pulling back their files; sites were also unstable, one day they worked, the next they didn't. This led to frustration among users, therefore we could not constrain users to a single method or a piece of software. It was noted that jobs could be lost at sites or the RB. DN noted a difference between production, analysis, and availability of events - reliability of the system was required. MA noted that it was difficult to assist users at the older sites who submit jobs, local users were less of an issue.

Panel B: Roger Jones, Akram Khan, Steve Lloyd, Raja Nandakumar, Neasan O'Neill, Glenn Patrick, Sarah Pearce, Dave Newbold

RJ noted another concern, what happens if users pull out data to Tier-3? This impedes storage elements. Should we have a rate-throttling policy? The moving of larger amounts of data should be penalised. RJ noted that what is available to users should be different to what is available to sysadmin. In a directory without permissions, or other areas, there should be experiment-acceptable use, how could abuse of the system be controlled? DN noted that control of batch-queuing systems was possible, but there was no control over storage. RJ noted the policy he would like to have: move more than 20GB get half the rate, move 50GB this goes down to 10% - a 'throttling' policy with a justifiable move of large amounts of data being planned in advance. DN noted that amongst computing, storage, and transfers, we only have control of one of these areas.

AK reported that he had moved from BaBar to CMS, had submitted analysis jobs and they had run nicely - the system does 1000 events and completes 1000, and he can access root files. The pessimism comes when jobs fail - filtering through the log files, it is difficult to see errors and difficult to know who to contact to ask about solutions - the errors are unclear. TD asked if there was an issue over error-reporting? AK asked how you could tell if a job was inefficient? SL advised that it depended how long it took for it to come back in real time.

RN noted this was a good point: 80% of jobs could be done, and 20% disappear, however when you submit jobs they are tracked. With Ganga you can easily run a test job on a local machine, but the first issue is what happens when a job has problems? There is no report and it is not simply explainable. Also, how does a new user get up to speed? TD asked whether RN's 'knowledge' existed on a wiki page? RN noted that he could put together some information but pointed out that users don't know what JDL is or what SAM boxes are. The second issue was which data was at what site, and how to access it. TD noted that if file catalogues were consistent there would be fewer problems? RN advised that file catalogues were reasonably consistent. It was noted that CMS and ATLAS don't regularly check their file catalogues and there was no single process for reconciliation. RN noted that if a job is submitted to a site that is working fine, then it all comes back fine.

Another issue was raised: if you are running something, the CPU writes-out data, then it reads data and writes-out more data. SB commented that we needed a mindset that expected things to go wrong - there were components, network links, the organisation was vast as a whole, it was expected that the bottom layer would fail. DN disagreed, stating that robustness should be built-in to a system - for example, the network was reliable but the hardware was not. DN noted that even if you get logging info back when something does go wrong, you can't read it. More importantly than the thing going wrong, is getting comprehensible return information. There was general consensus on this from those present.

Regarding the storage systems, DN noted that we hadn't exercised them with bulk analysis. There had been a use case at Tier-2 but we hadn't had multiple users using the Tier-2 storage, and a lack of robustness was expected. Tests were required. GC advised that this issue was coming up at CHEP, and they were trying to measure tests. DN wanted to know that the Tier-1 will survive with

multiple users - and at present the feeling was that it wouldn't. DN noted that the classic test at any site was to open thousands of files and see what happens to the storage system - CASTOR dies. GC noted that they had not looked at disaster scenarios yet, they had only looked at users acting irresponsibly - it was the place to start. It was suggested that FTS could be used for this. DN countered that it was easy to break FTS.

SL noted that the main thing of interest from a user perspective was an easily-installable UI - we need to make it easier to install and maintain these. JC added we need to do the same with tools. TD advised that there was documentation on the GridPP website which points to configuration of UIs. SB noted that there was no mechanism to solve this as no-one was specifically responsible for it. RJ noted that, on the contrary, we have the WLCG framework and the dashboard. SL noted that when the RB goes down, the user is stuck, and to switch RBs is difficult. SB added that further, users don't know which RB to use and asked whether the site admins were updating the UI anyway? It was often seen as a low priority. There was a discussion on RBs and UIs.

SP reported that she often spoke to potential Grid users (eg: AstroGrid), and asked if someone wished to use the Grid, where would she point them? GS noted that he had also spoken to AstroGrid and it might be advantageous for them to have local site help - establishing contacts between people running the sites and local user groups is very useful. TD advised that potential communities can be pointed to Grid tools and assisted, however it is better for them to have an afternoon with an experienced user. DB noted that all levels of enquiry needed some kind of support, right up to the formal level of the DB, but local site help was invaluable. It was understood that at the beginning, a new user wouldn't know what questions to ask - there was no replacement for someone on the ground. DN commented that good documentation was required. Training was also important.

CB advised that there was a Grid Users mailing list and this was also a forum for questions - initially this might help? SB noted that it would depend what users wanted to do. SL suggested a document for generic use cases. There ensued a discussion regarding new clients and their requirements, VOs, and how to get any interested parties started. TD noted that DK was involved in authentication issues, and backup from higher-level Experiments was also required. GP commented that the biggest issues were data access, documentation (the latter was given low priority but is a requirement, and a critical mass of users merits that kind of information), then there was Tier-3 and a sub-culture of people based on LXplus. He noted that there had to be a cultural shift at sites.

NO reiterated SP's question and pointed-out that it hadn't been answered. If someone contacts us and wants to be involved in the Grid, what do they do? TD advised that the best course of action was to ask them where they're from, and point them towards someone in that Institution region who could be an immediate initial point of contact. It was understood how difficult it was to initially get going and hand-holding was required.

SP asked whether we needed a list of individuals at sites who are contactable by potential new clients? SL advised that there is the GridPP Users mailing list, but it is not used. It is still best to contact someone at the site initially, following which a visit would be required, and the willingness of someone to spend an afternoon helping an enquirer get up to speed. The consensus of the meeting was that local expert help would be the best route. RJ noted that there had been a user 'show and tell' meeting. TD advised that this was difficult to advertise. TD noted that we do have documentation on 'how to get going', as NGS does. Dave Colling commented that there was no mailing list at the moment for someone with a user question. The existing gridpp-users@jiscmail.ac.uk was available but had not been used. [This forum included experts and

potential users]. SP noted that EGEE had also set-up a users' forum. TD brought the meeting to a close and thanked the panels for their contributions.