

# GridPP 19 Collaboration Meeting, University of Cumbria, Ambleside

## GridPP Discussion Session 3 ~ F2F ~ 31 August 2007

**Chair: Steve Lloyd**

**Panel A:** Dave Colling, Linda Cornwall, Marian Klein, Mark Leese, Mingchao Ma

LC commenced by asking the meeting what piece of security middleware people wanted, that didn't already exist, and were there any problems to report at present? Getting rid of pool accounts was one suggestion – 8 agreed. VOMS roles was another – these should be consistent across the board. DC asked how many sites still recycled pool accounts? LC noted that the default position was to recycle. For pool accounts, did this also mean SGM and PRD? It was noted that PRD and SGM relate to static pool accounts and involve the splitting of roles, which leads to the same problems.

LC asked if there were any other issues? DC noted that glxec was an issue – this was based on software written for a different use, and it was important that a proper job is done. MM advised that this was being tested at present – not altering code, but testing whether it is secure or not and whether identity can be switched. DC reported that having glxec on the CE, it was the quality of the code that was a concern. John Walsh advised that a set UID already existed on the worker nodes – this is already there irrespective of the Grid.

MM noted that glxec can be used by any user because it is set to root. JW reported that in the early stages of EGEE II, a code audit was being done by partners outside EGEE. DC noted that they were doing a good job but it was an under-resourced activity. LC agreed that glxec was an area of concern. Dave Kant noted the impact it had on accounting. Could it be corrupted? It was noted that taking the pilot jobs approach was problematic, there were two issues: 1. was glxec well-written, and 2. it is a pilot-enabler and do we want to use it at all? There was a discussion on the use of glxec, prioritisation, design, and whether it was written securely. MM noted that the quality of the code mattered. It was noted that LHCb don't need glxec – they want pilot jobs.

LC noted that there was a gap between standards and what is written. DC advised that 3.1 WMS can give it JDL as XML – you have to ensure extra fields but you can use it like any other web service. It was noted that packages from the gLite repository (RPMs) were not signed – there was no way when downloading gLite of knowing that packages were genuine RPMs. JW reported that this topic had been raised at the ROC Manager meetings. LC asked for a show of hands of those who considered this issue important – there were 18 responses. MM reported that there was a gap in contact between middleware development and the middleware security group. LC will assist with this offline.

DC raised the issue of job requirements (18) on DB's listing. These were fixed in 3.1 but DC hadn't tested it himself – the problem was that WMS strips out requirements and doesn't pass them on. JW advised that the standard LCG CE can pass requirements to WMS. DC noted that by using Globus you can pull requirements. JW advised that it was the WMS that stripped out the code; a patch was implemented to the LCG re wallclock time and there was no modification made to the CEs.

MK raised the issue of experiments needing access to new middleware for applications testing, and noted that not many users are interested in the PPS. DC advised that experiments were encouraged to use PPS but it doesn't happen. It was noted that there were scalability problems with the PPS – this will become more important when real data is being processed and analysed. DC noted, for example, that CMS not being able to use DPM would have been found through running test jobs on

the PPS. RJ advised that ATLAS do some testing on PPS but they had encountered problems – trying to get the production system to work splits effort. For ATLAS, they prefer to test on their own PPS. DC advised that he would happily move effort away from PPS if it was not useful and was not being used, however he reported that it was more stable and more of a service now than it used to be. By the time something reaches the PPS it should have gone through integration etc, and functionality tested, the PPS should be the last stage. DC asked whether effort should continue to be allocated to the PPS.

JW asked whether PPS sites publish that they offer PPS? MK advised that the PPS uses different BDIIs. JW suggested putting a production tag on it and testing the PPS service on the production grid. SL noted that we would need to wait and see what happens when real data comes through. There ensued a discussion on dependencies and RPMS, and middleware packages. JW noted that these issues were being looked at at present but resources were an issue – it was work-in-progress.

ML asked whether anyone had any issues or questions about the network, in terms of performance or monitoring performance. It was reported that 3 boxes were not running at present. A document was being compiled comprising ‘real life’ examples using data to define problems. There was an issue with the Sheffield box which was not yet solved. It was suggested that bandwidth reservation would be good to have as this would assist with data transfers. It was noted that site network traffic was affected by testing. DC reported that there was a different attitude at London colleges, the networking people were surprised at how much bandwidth people wanted to use. It was noted that networks are there to be used – if you fill up the network bandwidth, this tends to lead to expansion. There ensued a discussion about packages and dependencies.

**Panel B:** Robin Middleton, Alvin Tan, James Thorne, John Walsh, Paul Trepka

AT advised that he was coming from the applications side and the users of middleware – the issue he wished to raise was the problem of UI not being able to be installed easily, was any work being done on size and portability? JW noted that one of the problems of UI on a desktop is configuration. AT noted that you want a minimal system to run what you want it to run. JW advised starting with Globus. It was noted that no-one wants to take out and repackage Globus, and it was acknowledged that it provides more than is needed for users. There ensued a discussion on UI access and security issues. It was noted that RBs are not a site resource, they are a grid resource – supposed to be available from anywhere, however you needed flexibility of switchover and flexibility of resource. It was noted that there are more than 3 RBs in the UK, and 100 worldwide. Regarding UI, 1 per site was the minimum consideration. SL noted that this is a distributed system, communication is necessary therefore such that the ‘edges’ be kept secure.

The next issue raised was 64 bit-ness (3) on DB’s listing. It was reported that the disk servers at RAL are being tested – what were other experiences? RJ advised that ATLAS code would need twice as much memory. It was noted that the recommendation from the experiments was to go to 64 bit, but on a longer timescale. From the site admin point of view, were experiments ready to have 64 bit SL4? It was reported that, mainly, yes they were, however there existed a mix of native 32 bit and 32 bit on 64. It was noted that you can run a 32 application on a 64 bit-ness, however handling 32 or 64 bit kernels depends on hardware and therefore memory capacity.

The next issue was lack of source RPMs – there was no route of feedback to middleware from site admin. JW noted this was a question of resources – and this needs to be implemented in the build system. Middleware develop and produce a tag, this is passed onto the integration team, and anyone can take that tag and build RPMs – that is a route for source RPMs. SL noted that this had been an issue a few years ago. Access to source RPM is required - why can’t we have it? JW

advised that consortium members have a right to demand access and source code – a complaint could justifiably be made.

JW raised the issue of software licences (11). In Ireland they don't own the batch systems – these are owned by the Universities. One complaint about the Grid is that it doesn't work because of software licensing issues, which are restricted to Universities only – there are resources there, they have storage, but these can't be used because licence agreements need to be signed. It was noted that VOs can negotiate for licences. DC reported that this was a company by company process, as communities come on board we can negotiate with companies and their software. DB advised that the GridPP view was that NGS is more active in this area, and GridPP shouldn't become directly involved.

The next issue raised was the LCG monitoring working group – there was a session coming up at the EGEE conference relating to SA1 operations monitoring. Each Grid component would produce monitoring data and working groups have plugins so that you should be able to plug into your site. More information would be available at the EGEE conference regarding dependencies and what works on the UI.

PT raised the issue of site infrastructure, and were there any suggestions regarding enlarging storage elements and switching services/decommissioning? JW advised that a document had been produced relating to this, where you announce you are decommissioning the SE and give VOs a timeframe. It was noted that there is a documented procedure (but it can often be ignored). It was advised that information has to percolate from site admin to all users on a VO. Re the ATLAS issue last year, this involved an unclear message being received. It was noted that an individual user doesn't know where his files are. This should be up to the VO. One person can do transfers within a site for the whole VO. RJ advised that the Tier-2s have a custodial role – if the experiment pins it, then the site is expected to be responsible for that data. It was reported that 99% of SE space in the UK is reporting as non-volatile. RJ noted that a substantial portion is not expected to be volatile, but notice is required before taking it down. RJ noted another issue, that of changing names of storage elements. It was agreed that the question in hand had been generally experienced and was well-documented by the storage group.

SL brought the meeting to a close, thanked the panels, all delegates, and the University of Cumbria.