

## Notes from OSG – Nuclear Physics day on Dec 6<sup>th</sup>

These are the notes and action items I took away from the meeting on Wednesday. Please correct my mistakes and add additional things to be sent to the OSG management bodies – Executive Board and Council.

Attending: Tim Hallman, Jerome Lauret, Alex Dzierba, Rob Soltz (remotely), Doug Olson, Andy Kowalski, Ruth Pordes, Miron Livny, Gene Oleynik (morning).

First I want to thank Alex, Jerome, and Ron for presenting a really useful set of slides about their projects which are a good basis for discussion of our collaborations. Many thanks to all of you who travelled in this wintery weather. I really apologise for not giving you a tour of Fermilab and visits to good local restaurants. I hope you will give us a next time with an opportunity to present experiments computing models and plans to a larger audience at Fermilab itself, talk with the CMS computing management – who were all at CERN last week and so not available to attend our meeting - and to perhaps give you a tour of the layout of the lab computing.

These last couple of months for OSG have been a time of planning and organizing, and this was the focus of our meeting.

### **STAR:**

There are short term and longer term needs and plans to be made.

A first focus is the current milestone in the OSG year 1 project plan (available from <http://osg-docdb.opensciencegrid.org/0005/000514/002/OSGProjectPlan-v17-accept.pdf> )

STAR: Migration of >80% of simulation to OSG, *Jerome Lauret*, 6/15/07

We agreed that this requires a throughput of 80 CPUdays/day on the 4 current plus one new STAR sites on OSG (NERSC/PDSF, Wayne State, Sao Paolo, Birmingham and UIC) and at least 1 site where STAR will use opportunistically available resources, proposed to be FermiGrid. Without any more stable pattern based on 80/100 jobs, the 80% milestone would be un-achievable as the burden of recovery versus accomplish would be inadequate. We need a sustained rate that will allow scaling. The STAR simulation uses a local Storage Element (SE) to directly store the data from the application and then a separate step to move the data from the local SE to the mass store at BNL. After the meeting Ruth sent mail to Jerome, Keith Chadwick and Gene to make the FermiGrid contact.

STAR jobs on OSG (from STAR and OSG monitoring about 250-300 /day when they are being run) are currently having unacceptably low efficiency of success on the NERSC and Sao Paolo, better but still inadequate efficiency on Wayne State, and are not yet running on Birmingham. We kicked off a task force to deal with this as soon as possible. It was agreed to focus on NERSC as the first site to improve. Wayne Betts, Dave Skinner and someone designated from the OSG troubleshooting team, lead by Shaowen Wang, to

## Notes from OSG – Nuclear Physics day on Dec 6<sup>th</sup>

be the people on this task force. Shaowen sent mail to Jerome, Doug and available Condor staff to get some more specifics from Jerome and plans a phone meeting to go over technical and action items. We hope that can happen Monday 11<sup>th</sup>. The people at last weeks Wednesday meeting plan to get together sometime in mid-January to review progress and needs.

STAR relies on the DRM/HRM implementation of storage management behind an SRM interface. Development of this implementation currently has insufficient sustained funding – we think there is some Earth Science Grid funding, OSG has 0.5 FTE support at LBNL (Alex Sim) for SRM testing and support within the Facility. Support for development and use for this implementation has been extremely good during the period of SRM and PPDG funding. Extensions are needed to support STAR analysis needs over the next two years. It is not clear how many of these extensions -- Role-based and later user authorization, quotas, Best data placement, data transfer queue monitoring, Policies – are available or planned for other SRM interfaced storage management systems. Options available could potentially be DPM, STORM, dCache (if installation and maintenance is made much easier). It is a goal to support dCache through the VDT over the next six months; The Birmingham site may test DPM. We did not talk about the requirement for SRM on-the-fly deployment. We agreed to think more about this and come back in a couple of months to talk.

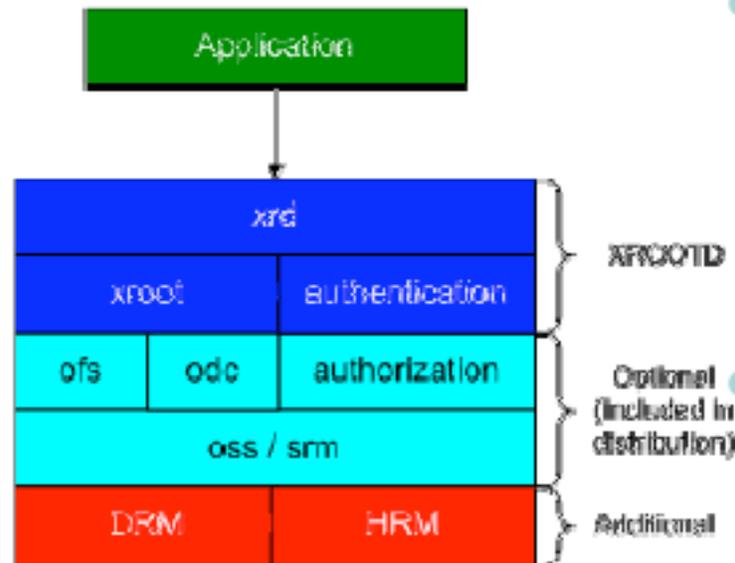
STAR object level data access and UJDL. Neither of these are currently in the OSG Year 1 scope. The Object Level access is un-funded at this stage and STAR believes it to be direly important to the experiments success. STAR will have to seek funding for this and would surely welcome agency support via support statements from partners. Job/workflow tracking and monitoring. STAR has a phase-II SBIR with Tech-X on this. The Scidac-2 Center for Distributed Science (CEDS) is doing work in this area. We left it that Jerome would contact Brian Tierney, who is leading the CEDS work here, and talk to Tech-X about their plans to see what level of collaboration/communication would be useful.

A longer term milestone for STAR is to be able to support all STAR analysis on OSG, 10,000 jobs/day in time for the next STAR run in 2009 Q2, where there is a factor of 2 increase in luminosity from the current run. STAR analysis relies on ROOT being installed on each local storage element and XROOTD for the management and interface of storage local to each compute (worker) node. An 2<sup>nd</sup> prototype of the XROOTD-SRM interface, written by STAR, is in test at BNL.

Alice depends on Root/Xrootd and there are evaluations for its use for the other LHC experiments either happening or in the works. All LHC experiments have their data in Root I/O format (as far as I know). OSG is interested in discussing collaboration for this, but technical work could not start until the summer of 2007. People felt that this would be in sufficient time to meet the STAR extensions and production needs if we knew what activities to start by then. This would need extended support in both STAR and OSG to ensure meeting the STAR needs and milestones and could be a basis for some extended work in Nuclear Physics as part of the OSG program.

## Notes from OSG – Nuclear Physics day on Dec 6<sup>th</sup>

STAR will write up their user analysis milestones (requirements, how to, ...) STAR for discussion in March 2007, and OSG and STAR will re-assess the plan in September



**Figure 1: Software layers on each local Storage Element - expected on each Worker Node, in the STAR model**

### Alice:

Alice goals in the US are to provide the hardware and software resources to allow US Alice to fully partner in the science of the experiment. Concrete plans for collaborating with OSG are waiting for the results of the Physics at the Information Frontier (PIF) proposal to NSF sometime February/March of next year. The goal is to interface Alien to the OSG infrastructure – we don't yet understand exactly the technical details of the work, what changes might be useful for the OSG services and/or for Alien itself to have the best integration.

An Alien installation is on a VO-Box at the edge of a grid computing resource. This installation can manage one or potentially many sites. Alien uses Root. However at the moment we don't understand that there are requirements that for software to be installed permanently on the worker nodes or compute site other than the VO-Box. (.e.g is the file system on the VO-Box shared across all worker-nodes?).

There are 4 sites where Alice has or will have resources in the US: Ohio Supercomputer Center, LBNL, LLNL, Texas Learning and Computation Center at the University of Houston. All these sites have mass storage systems (HPSS). Ron has successfully worked with LLNL to support GridFTP services for MIPP, and MiniBoone also uses the facility.

## **Notes from OSG – Nuclear Physics day on Dec 6<sup>th</sup>**

There has been no discussion or expectation that the other sites than NERSC/PDSF join the OSG or would offer opportunistic or shared resources to other experiments or VOs, and no expectation that Alice would want to use other OSG resources. OSG will see if the Facilities are interested in their own right.

### **Glue-X**

Concrete plans for collaborating with OSG are waiting for the results of the PIF proposal and hiring of post-docs. JLAB will provide remote access to mass storage for Glue-X and at the moment it is being discussed what other grid services JLAB will provide. Alex will contact Miron and Ruth for some technical discussion when Glue-X is ready to start specific planning for simulations and computing that would use local clusters and/or remote sites. The visualization requirements are new for OSG and we would like to work with Glue-X at the technical planning stage to best deliver what is needed.