Open Science Grid, December 2008 - Annual Report to DOE

Nuclear physics – STAR/OSG accomplishments

The STAR experiment has continued to use the OSG data movement capabilities between LBNL, BNL and new sites on the OSG at Wayne State University and/or NPI/ASCR in Prague (two fully functional Tier 2 centers) and in tests at the University of Illinois at Chicago. At Prague / Bulovka, data transfers are handled using a BeStMan SRM client interoperating with a *Disk Pool Manager* (DPM) SRM door. Data rates between BNL and Prague, reaching 300 Mb/sec at the moment, are sufficient to sustain the local needs. Local data access in Prague rests on the use of the *STAR Unified Meta-Scheduler* (SUMS), which provides users a common interface for job submission. SUMS offers a transparent submission approach to both Grid and non-Grid resources and remains at the heart of STAR's strategy to migrate an entire class of jobs to Grid resources. STAR plans to utilize the Prague resources for opportunistic Monte-Carlo event processing in 2009 as well as study and test a multi-site data transfer paradigm coordinating movement of datasets to and from multiple locations (sources) in an optimal manner that is, using a planner taking into account network and site performance. It is noteworthy to mention that analysis of data sets now entirely relies on access to Scalla/Xrootd data aggregation at BNL (since 2006) and DPM/rfio access at Prague (2007/2008). Users make extensive use of SUMS abstraction to seamlessly launch jobs on the respective farms; the same job description works on both farms.

STAR has continued to leverage and help consolidate the BeStMan SRM implementation and have engaged in active discussions, steering and integration of the messaging format from the *Center for Enabling Distributed Petascale Science* (CEDPS) Troubleshooting team, in particular targeting the use for BeStMan client/server troubleshooting for faster error and performance anomaly detection and recovery. Stability and scalability testing have been regularly conducted using BeStMan between BNL and LBNL with increasing beneficial results (downtimes are minimal and to the level of one issue per month with a few hours recovery). BNL's gatekeepers are also being upgraded and the services will be re-tuned once again upon completion (end of 2008) with the prospect of massive data transfer outside/in from the *Korea Institute of Science and Technology Information* (KISTI), a facility and team now a member of the STAR collaboration. KISTI holds potentials for serving as a Tier-1-like center for STAR. Setup and data transfer testing are in progress with a first trial in real data transfer mode planned for early 2009 (during RHIC Run-9). Access to job submission at KISTI will likely make use of standard Grid technologies.

STAR grid data processing and job handling operations have continued their progression toward a full Grid-based operation relying on the OSG software stack and the OSG Operation Center issue tracker. The STAR operation team and support has been efficiently addressing issues and stability overall of the grid infrastructure which seems to have increased. To date, STAR has however mainly achieved simulated data production on Grid resources. Since reaching a milestone in 2007, it has become routine to utilize non-STAR dedicated resources from the OSG for the Monte-Carlo event generation pass

and to run the full response simulator chain (requiring the whole STAR framework installed) on STAR's dedicated resources. STAR's primary dedicated resources in the US are at BNL's RACF and NERSC's PDSF. These and other resources are specifically allocated to and secured for STAR and are aggregated into a virtual facility via the use of the OSG interfaces and the continuous use of SUMS for job scheduling. In the case of jobs submitted to PDSF from BNL, job output is returned to BNL using BeStMan SRM for data transfers.

The relative proportions of processing contributions are at the moment at the level of 10-15% for Monte-Carlo event generation (and mainly pre-allocated use of resources from Fermi-Grid) for non-dedicated resources and 85-90% on STAR dedicated sites for full signal reconstruction. This proportion is explainable by the fact that the complete STAR software stack and environment, which is difficult to impossible to recreate on arbitrary grid resources, is necessary for the full event reconstruction and hence, access to generic and opportunistic resources are simply un-practical and not matching the realities and needs of running experiments in Physics production mode. Those like STAR have as requirements a reproducibility of all results ensured by thoroughly validating the software stack for each variant of OS and platforms via regression testing. The resources needed for running all simulated data productions (both contributions of STAR and non-STAR dedicated resources) is equivalent to 10% of a one pass real-data production and has, since 2008, been entirely handled on a Grid based operation.

STAR has also worked closely, to the extent possible, with the CEDPS Virtualization activity, seeking the benefits of truly opportunistic use of resources by creating a complete pre-packaged environment (with a validated software stack) in which jobs will run. Such approach would allow STAR to run any one of its job workflow (event generation, simulated data reconstruction, embedding, real event reconstruction and even user analysis) while respecting STAR's policies of reproducibility implemented as complete software stack validation. The multitude of combination and the fast dynamic of changes (OS upgrade and patches) make the reach of the diverse resources available on the OSG, workforce constraining and economically un-viable. While the Virtualization work has potential for dramatic positive impacts for STAR (and other VOs), our progresses were hampered due to the fact that while STAR has the knowledge and will toward the innovative, it also has limited workforce availability (none of which is OSG funded and all of which has as primary mandate to support the immediate need of RHIC operations) and hence, has not been able to sustain appropriate effort for this activity and make full use of what could become a useful resource.

Finally, since all of STAR's physics results are built on a foundation of simulated data, all STAR's physics publications acknowledge the resources provided by the OSG.