



NOVEMBER 2006

[OSG HOME](#) | [SUBMIT NEWS](#) | [SUBSCRIBE](#) | [ARCHIVE](#) | [ABOUT OSG NEWS](#)

## Meetings and Events

[WLCG Collaboration Workshop](#)  
CERN, Switzerland  
January 22-26, 2007

[19th Open Grid Forum - OGF19](#)  
Chapel Hill, North Carolina  
January 29 - February 2, 2007

OSG Consortium Meeting  
San Diego, California  
March 5-8, 2007

[View Full Calendar](#)

## Spotlight on an OSG Contributor



Horst Severini is the grid computing coordinator of the Oklahoma Center for High Energy Physics at the University of Oklahoma in Norman . He is in charge of the operation of the OU part of the US ATLAS Southwest Tier 2 Center, which also includes UT Arlington and Langston University . Besides ATLAS, the OU High Energy Physics (OUHEP) group is also a member of the D0 and DOSAR VOs within OSG.

OUHEP has several OSG CEs installed, most of which run jobs for the ATLAS, D0 and DOSAR VOs. Horst and his OUHEP colleagues also maintain an OSG testbed site, always an early adopter of new OSG integration releases, that allows them to help debug and fix integration and deployment issues. Furthermore, they work with the OSG accounting and monitoring groups on software tests.


Horst is also the associate director of the OU Supercomputing Center for Education and Research, where he manages the OU Condor project. The project aims to assemble all OU campus student lab PCs into a campus Condor pool--which is also already an

## From the Executive Director

Over the past few months OSG has delivered significant resources to its VOs, and we are starting to fulfill our mission of providing the number-crunching power needed to enable innovative science. The number of concurrent jobs on OSG has fluctuated from 3,000 to almost 6,000, and CMS has run up to 55,000 test jobs a day while testing the scale and performance of the OSG infrastructure for its Computing, Software & Analysis Challenge 2006 (CSA06). GADU began a large update of its integrated public genome database in early November and is validating additional sites for execution. D0 will start its production reprocessing soon, which will serve as a good test of how this additional load can be accommodated with currently running applications such as CDF and nanoHUB. Site administrators, please contact the GOC if you need help supporting OSG applications.



We are pleased to announce that CDF is the first collaboration to agree to acknowledge OSG use in its future publications. This will help us demonstrate the scientific benefits of our infrastructure. Over the next couple of months we will work through the details of OSG acknowledgement with CDF and other collaborations.



Count of trouble tickets by Support Center over a two-week period. (Click for larger image)


Work on the next OSG release, scheduled for next February, is well underway. Testing of the first deployments of SRM/dCache storage management from VDT will begin next week. Leigh Grundhoefer continues to write a bi-weekly operations report that includes the metrics on support load from the GOC shown to the left. The GOC is there for you; please don't hesitate to contact them with problems!

Ruth Pordes

## Solving Pilot Job Security Holes Using gLExec

When a user submits a grid job to an OSG site, the job always carries the user's credentials. At the execution site, the job is assigned an appropriate userid under which to run. The security framework for this workflow has been well thought out and in place for some time.

Another option for submitting grid jobs involves the concept of a pilot job. This type of job, once it's in a site's batch slot, coordinates and calls a series of user jobs according to VO priorities at launch time. If the pilot job and the user jobs all run under the same userid, however, the pilot job framework violates the security policies of any site that requires knowledge and control of its resource users.



(Click for larger image)

gLExec, a gLite product currently used on European Computing Elements, solves this problem. gLExec is a privileged executable that, given a user credential and an execution command, obtains the appropriate Unix ID from a site's GUMS

OSG CE--and to eventually expand that project across the state of Oklahoma and beyond.

### Recent News About OSG

OSG collaborators presented at the recent Supercomputing 06 Conference in Tampa, Florida. Attendees found the presentations at the Indiana University Booth especially interesting. [More information...](#)

OSG and the Northwest Indiana Computing Grid completed an agreement to work together on the use of the VDT, giving each other opportunistic access to each other's available resources. [More information about NWICG...](#)

### Supported By



Office of Science/  
U.S. DOE

server and executes the job under that Unix ID. In order to use gLExec within OSG, VOs must configure the pilot job such that it "calls home" to get the associated user credential. The pilot then forwards the credential to gLExec, which uses it to communicate with the site security service, thus returning control to the site.

The interface to the OSG authorization and accounting infrastructure, needed to make gLExec usable on the worker nodes, has been developed jointly by gLite developers at NIKHEF and OSG developers at Fermilab. OSG is starting to deploy gLExec on its worker nodes.

Several OSG VOs have expressed interest in using pilot jobs to do resource selection on OSG, the most active of which are CDF and ATLAS. The gLExec solution in which the pilot infrastructure runs under its own userid is secure from a VO and site perspective. The separate userids protect the pilot job from malicious users, and all the users are protected from each other, just as they are when they submit jobs individually in the standard way.


Igor Sfiligoi, Fermilab

### Applications: Large Scale Event Simulation for ATLAS


As part of the preparations for real data taking next year, ATLAS is currently in the midst of a large scale Computing System Commissioning (CSC) exercise. This consists of simulating approximately 50 million "events" (collisions of protons on protons) and processing the simulated data through the entire ATLAS software chain, producing as the end result histograms that reveal new physics.

This CSC exercise is taking place on the WLCG and, in the U.S., on the OSG. The figure in the upper right shows the number of tasks that have been performed on three sub-grids of the WLCG: OSG, LCG and the NorduGrid. Each task is a collection of 10-10,000 similar compute jobs. In round numbers, the CSC so far has run 3,400 different tasks worldwide, using 10 thousand CPUs and producing 120 terabytes of data.

- Jim Shank, Boston University



The number of tasks performed as part of the ATLAS CSC on three sub-grids. PanDA is the software system used to submit ATLAS jobs to the distributed grid resources. (Click for larger image)



CPU usage (walltime) during the last three months. (Click for larger image)