



DECEMBER 2006

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

[2007 American Association for the Advancement of Science \(AAAS\) Annual Meeting](#)
Long Beach, California
February 15-19, 2006

[OSG Consortium All Hands Meeting](#)
San Diego, California
March 5-8, 2007

[View Full Calendar](#)

CE Installation



Run II Computer Farms at Fermilab for DZero Experiment. Image courtesy Reidar Hahn

The first step to making your institution's resources available to the Open Science Grid is to install a Compute Element (CE). The CE is the gateway machine which accepts jobs and file transfers from the rest of the grid and forwards them to your batch system. It also runs monitoring software that tells the rest of the grid how many CPUs you have, how many resources are free, and how to access them. [Installation instructions](#) for any production release of software for the OSG are always available. The current release version is OSG 0.4.1.

Two keys to successfully installing a CE machine are a large and robust file system to store the intermediate files, and a machine with sufficient CPU speed and memory. Buy as much of both as you can afford. At the August 2006 Consortium meeting, there was a session on [Best Practices for](#)

Open Science Grid News

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

From the Executive Director

Most of the OSG technical work is performed through small group activities with a few coordinating meetings covering the major areas of work. The Operations meeting, held on Mondays at 1 pm Eastern time, is a key venue for bringing any and all kinds of problems and issues to the table. If you are trying to get work done or making resources accessible to the OSG, it is expected that you will send your Support Center representative to this meeting on a regular basis. If the Operations meeting is not addressing your needs, please help us make it better. Other meetings of particular interest are the Integration Meeting on Thursdays at 1:30 PM Eastern and the biweekly Storage and Monitoring Technical Group meetings on Mondays.



In the last couple of weeks, the OSG Troubleshooting activity has gotten underway. Shaowen Wang from the University of Iowa is the interim coordinator for this activity. There is a lot of initial information on the [OSG at Work twiki site](#). Problems are handed to the troubleshooting team from the Operations, Virtual Data Toolkit (VDT) and/or User Liaisons, when there are issues that need some in depth digging and triaging to resolve.

Last week OSG applications clocked over 6,000 CPU Weeks, or about 120 CPU Years within the week. OSG resources are not full yet! We encourage you to fill the pipes!

Ruth Pordes

Spotlight on an OSG Contributor

Tim Cartwright is a member of the VDT development team at the University of Wisconsin—Madison. He works on building, configuring, and testing the software components that go into the VDT distribution and on the VDT infrastructure itself. Also, he supports the VDT user community, especially OSG, with troubleshooting problems that are discovered in the field.



As a member of the larger Condor team in Madison, Tim is also the liaison between the ATLAS project and the rest of the Condor team. In addition, he coordinates efforts on the high-availability subsystem of Condor, making sure it is used locally and monitoring its behavior for possible future enhancements.

CSA06: Full Service Challenge for CMS

interfacing your CE to your batch system and getting things to operate smoothly. Many of the suggestions from this session have now been incorporated into the Integration web. If you examine the Best Practices, you will see there are some that have yet to be filled in. Knowledgeable system administrators from across the OSG are encouraged to share their experiences and ideas!

If you have questions, there are always a number of OSG mailing lists which can help you out, such as [goc](#), [osg-general](#) or [osg-int](#).

-Steven Timm & Leigh Grundhoefer

Supported By




From October to November 2006, CMS tested its software and computing infrastructure during CMS Simulation and Analysis 2006, (CSA06) a complete service challenge. In the CMS computing model, data access is location-driven while a bookkeeping and transport system takes care of distribution and discovery of experimental and Monte Carlo simulation data samples.

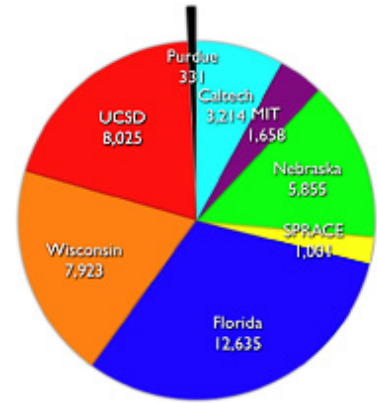
CSA06 consisted of a MC simulation part and a reconstruction/analysis part that represented 25% of the full capacity needed for CMS operation in 2008. The goal was to exercise the full event flow, starting with simulated data taking and reconstruction at CERN to analysis on Tier-2 resources.

Before the challenge began, simulated data was produced using the EGEE and OSG infrastructures. Approximately 65 million events were used for the challenge, of which 27% were produced on OSG resources. During the challenge, the simulated data samples were reconstructed at CERN and transported to Tier-1 sites, where parts of the samples were extracted (also called skimming) and re-reconstructed.

The network utilization of the US Tier-1 for CMS, located at Fermilab, reached a peak of over 275 MBps to and from more than two dozen T2 sites. Rates of over 3 Gbps to and from CERN and the eight OSG Tier-2 sites were observed, which shows a good performance of the network infrastructure.



Reconstruction of the Z mass in the decay channel $Z \rightarrow \mu + \mu$.
Image Courtesy Chang Liu, CERN
(Click for larger image)



Breakdown of CSA06 jobs completed on November 9 using OSG resources.
(Click for larger image)

Subsequent access to the samples on Tier-1 and Tier-2 level was user-driven, supplemented by automatic submission and retrieval of short test jobs (30-45 minute job duration) using dedicated job robots. The goal of 50,000 jobs per day on EGEE and OSG resources was reached and sustained over several days. For example, on November 9, about 59,000 jobs were submitted of which 70% were executed on OSG resources.

One example for the many analysis exercises carried out in the context of CSA06 is the reconstruction of the Z mass in the decay channel into two muons. The di-muon mass was reconstructed using skimmed datasets at the T1 at Fermilab and the OSG T2 at Purdue. Different muon reconstruction possibilities were studied, from using exclusively the muon system to combining it with the tracking system.

Additional information:

[Dashboard plot](#)

- Oliver Gutsche