



# Open Science Grid

## Report to the US Department of Energy December 2010

Miron Livny	University of Wisconsin	PI, Technical Director
Ruth Pordes	Fermilab	Co-PI, Executive Director
Kent Blackburn	Caltech	Co-PI, Council co-Chair
Paul Avery	University of Florida	Co-PI, Council co-Chair

# Table of Contents

<b>1. Executive Summary .....</b>	<b>4</b>
1.1 What is Open Science Grid? .....	4
1.2 Usage of Open Science Grid .....	5
1.3 Science enabled by Open Science Grid.....	7
1.4 Technical achievements in 2009-2010.....	8
1.5 Preparing for the Future .....	11
<b>2. Contributions to Science.....</b>	<b>13</b>
2.1 ATLAS .....	13
2.2 CMS .....	20
2.3 LIGO.....	23
2.4 ALICE .....	25
2.5 D0 at Tevatron .....	26
2.6 CDF at Tevatron .....	30
2.7 Nuclear physics .....	35
2.8 MINOS .....	40
2.9 Astrophysics .....	41
2.10 Structural Biology.....	42
2.11 Multi-Disciplinary Sciences.....	44
2.12 Computer Science Research .....	44
<b>3. Development of the OSG Distributed Infrastructure .....</b>	<b>45</b>
3.1 Usage of the OSG Facility.....	45
3.2 Middleware/Software .....	47
3.3 Operations.....	49
3.4 Integration and Site Coordination .....	51
3.5 VO and User Support.....	52
3.6 Security.....	53
3.7 Content Management.....	55
3.8 Metrics and Measurements .....	56
3.9 Extending Science Applications .....	58
3.9.1 Scalability, Reliability, and Usability .....	58
3.9.2 Workload Management System .....	60
3.10 Challenges facing OSG .....	61
<b>4. Satellite Projects, Partners, and Collaborations .....</b>	<b>63</b>
4.1 CI Team Engagements .....	65
4.2 Condor.....	66
4.2.1 Support Condor.....	67
4.2.2 Release Condor .....	69
4.3 High Throughput Parallel Computing .....	70

4.4	<i>Advanced Network Initiative (ANI) Testing</i> .....	71
4.5	<i>ExTENCI</i> .....	73
4.6	<i>Corral WMS</i> .....	73
4.7	<i>OSG Summer School</i> .....	75
4.8	<i>Science enabled by HCC</i> .....	76
4.9	<i>Virtualization and Clouds</i> .....	78
4.10	<i>Magellan</i> .....	79
4.11	<i>Internet2 Joint Activities</i> .....	79
4.12	<i>ESNET Joint Activities</i> .....	80
<b>5.</b>	<b>Training, Outreach and Dissemination</b> .....	<b>82</b>
5.1	<i>Training</i> .....	82
5.2	<i>Outreach Activities</i> .....	83
5.3	<i>Internet dissemination</i> .....	83
<b>6.</b>	<b>Cooperative Agreement Performance</b> .....	<b>84</b>

Sections of this report were provided by: the scientific members of the OSG Council, OSG PI-s and Co-PIs, and OSG staff and partners. Paul Avery and Chander Sehgal acted as the editors.

# 1. Executive Summary

---

## 1.1 What is Open Science Grid?

Open Science Grid (OSG) is a large-scale collaboration that is advancing scientific knowledge through high performance computing and data analysis by operating and evolving a cross-domain, nationally distributed cyber-infrastructure (Figure 1).

Meeting the strict demands of the scientific community has not only led OSG to actively drive the frontiers of High Throughput Computing (HTC) and massively Distributed Computing, it has also led to the development of a *production quality facility*. OSG's distributed facility, composed of laboratory, campus, and community resources, is designed to meet the current and future needs of scientific operations at all scales. It provides a broad range of common services and support, a software platform, and a set of operational principles that organizes and supports scientific users and resources via the mechanism of Virtual Organizations (VOs).

The OSG program consists of a Consortium of contributing communities (users, resource administrators, and software providers) and a funded project. The OSG project is jointly funded, until late-2011, by the Department of Energy SciDAC-2 program and the National Science Foundation.



Figure 1: Sites in the OSG Facility


While OSG does not own the computing, storage, or network resources used by the scientific community, these resources are contributed by the community, organized by the OSG facility, and governed by the OSG Member Consortium. OSG resources are summarized in Table 1.

**Table 1: OSG computing resources**

Number of Grid interfaced processing resources on the production infrastructure	122
Number of Grid interfaced data storage resources on the production infrastructure	59
Number of Campus Infrastructures interfaced to the OSG	9 (GridUNESP, Clemson, FermiGrid, Purdue, Wisconsin, Buffalo, Nebraska, Oklahoma, SBGrid)
Number of National Grids interoperating with the OSG	3 (EGEE, NGDF, TeraGrid)
Number of processing resources on the Integration infrastructure	23
Number of Grid interfaced data storage resources on the integration infrastructure	14
Number of Cores accessible to the OSG infrastructure	~57,000
Size of Disk storage accessible to the OSG infrastructure	~24 Petabytes
CPU Wall Clock usage of the OSG infrastructure	Average of 47,000 CPU days/ day during Nov 2010

## 1.2 Usage of Open Science Grid


High Throughput Computing technology created and incorporated by the OSG and its contributing partners has now advanced to the point that scientific users (VOs) are utilizing more simultaneous resources than ever before. Typical VOs now utilize between 15 and 20 resources with some routinely using as many as 40 – 45 simultaneous resources. One key factor was the transition to using Pilot based job submissions. This has now become the recommended mechanism for using OSG. The overall usage of OSG has increased again by ~45% this past year and continues to grow at a steady rate (Figure 2). Utilization by each stakeholder varies depending on its needs during any particular interval. Overall use of the facility for the 12 month period ending December, 2010 was 348M hours, compared to 241M hours for the previous 12 months. (Detailed usage plots can be found in the attached document on Production on Open Science Grid.) During stable normal operations, OSG provides over 1.1M CPU wall clock hours a day (~47,000 cpu days per day) with peaks occasionally exceeding 1.3M hours a day; approximately 300K – 400K opportunistic hours (~35%) are available on a daily basis for resource sharing. Based on transfer accounting, we measure approximately 1 PetaByte of data movement (both intra- and inter-site) on a daily basis with peaks of 1.2 PetaBytes per day. Of this, we estimate 25% is GridFTP transfers between sites and the rest is via LAN protocols.



Maximum: 34,413,793 Hours, Minimum: 203,860 Hours, Average: 16,921,680 Hours, Current: 4,436,820 Hours

**Figure 2: OSG Usage (hours/month) from June 2007 to November 2010**

The number of Non-HEP CPU-hours (Figure 3) is regularly greater than 1 million CPU hours per week (average = 1.2 M) even with the LHC startup at the end of March. LIGO averaged approximately 90K hours/day (11% of the total), and non-physics use now averages 85K hours/day (10% of the total), reflecting efforts over the past year to support SBGrid and incorporate new resources such as Nebraska's Holland Computing Center.



**Figure 3: OSG non-HEP weekly usage from December 2009 to December 2010, showing more than a 2x fractional increase. LIGO (shown in red) is the largest non-HEP contributor.**

### 1.3 Science enabled by Open Science Grid

OSG's infrastructure supports a broad scope of scientific research activities, including the major physics collaborations, nanoscience, biological sciences, applied mathematics, engineering, and computer science; and, through the Engagement program, other non-physics research disciplines. The distributed facility is heavily used, as described below and by usage charts in the attachment "Production on Open Science Grid".

A strong OSG focus in the last year has been supporting the ATLAS and CMS collaborations' preparations for LHC data taking that re-started in March 2010. Each experiment ran significant preparatory workload tests (including STEP09), data distribution and analysis challenges while maintaining significant ongoing simulation processing. As a result, the OSG infrastructure has performed well during current data taking. At the same time, OSG partnered with ATLAS and CMS to develop and deploy mechanisms that have enabled productive use of over 40 U.S. Tier-3 sites that were added over the past year.

OSG made significant accomplishments in 2010 supporting the science of the Consortium members and stakeholders (Table 2). Considering first the large experiments, in late 2009 LIGO significantly ramped up Einstein@Home production on OSG to search for gravitational radiation from spinning neutron star pulsars, publishing 28 papers on this and other analyses. The D0 and CDF experiments used the OSG facility for a large fraction of their simulation and analysis processing in publishing 28 and 61 papers, respectively, over the past 12 months. The LHC experiments ATLAS and CMS also had a productive year. CMS submitted for publication 23 physics papers based on cosmic ray analyses as well as a charged particle measurement from the Decem-

ber 2009 first collision dataset. ATLAS submitted a total of 113 papers. The STAR experiment had 29 publications during this time.

Smaller research activities also made considerable science contributions with OSG support. Besides the physics communities, the 1) structural biology group at Harvard Medical School, 2) groups using the Holland Computing Center and NYSGrid, 3) mathematics research at the University of Colorado, and 4) protein structure modeling and prediction applications have sustained (though cyclic) use of the production infrastructure. The Harvard (SBGrid) paper was published in Science.

As Table 2 shows, approximately 249 papers were published over the past 12 months (listed in the attachment “Publications Enabled by Open Science Grid”). Non-HEP activities accounted for 92 (37%), a large increase from the previous 12 month period. These publications depended not only on OSG “cycles” but OSG-provided software, monitoring and testing infrastructure, security and other services.

**Table 2: Science Publications in 2010 Resulting from OSG Usage**

<b>VO</b>	<b># pubs</b>	<b>Comments</b>
ATLAS	67	
CDF	52	
CDMS	2	
CIGI	2	
CMS	33	
D0	20	+ 5 accepted
Engage	20	
GLOW	19	
HCC	1	
LIGO	6	
Mini-Boone	3	
MINOS	4	
nanoHUB	3	
NYSGRID	1	
SBGRID	3	
STAR	13	
<b>Total</b>	<b>249</b>	

#### **1.4 Technical achievements in 2009-2010**

The technical program of work in the OSG project is defined by an annual WBS created by the area coordinators and managed and tracked by the project manager. 2010 saw the following distribution of staff by area and institution (Table 3):



**Table 3: Distribution of OSG staff**

<b>Area of Work</b>	<b>OSG Staff</b>	<b>Contribution<sup>1</sup></b>	<b>Institutions (lead institution first)</b>
OSG facility Management	0.5		UW Madison
Software Tools Group	1.0	1.0	UW Madison, Indiana U, U Chicago NCSA
Production Coordination	0.7		U Chicago
Biology and Biomedical Outreach	0.85	0.21	Harvard Med School
LIGO Specific Requests	1.25	0.17	Caltech
Software	7.6	0.25	UW Madison, Fermilab, LBNL
Operations	5.5	1.15	Indiana U, Fermilab, UCSD
Integration and Sites	4.05		U Chicago, Caltech, LBNL
VO Coordination	0.45	0.1	UCSD, Fermilab
Engagement	1.3	1.35	RENCI, ISI, Fermilab, UW Madison
Campus Grids	0.4		UChicago, Clemson
Security	2.55	0.2	Fermilab, NCSA, LBNL, UCSD
Training and Content Manage- ment	1.7		Caltech, BNL, Fermilab, Indiana U
Extensions	0	0.4	UCSD, BNL
Scalability, Reliability and Usa- bility	0.95		UCSD
Workload Management Systems	2.2		BNL, UCSD
Internet2		0.27	Internet2
Consortium + Project Coordina- tion	1.08	0.2	Fermilab, Caltech, UFlorida
Metrics	0.5	0.05	UNebraska
Communications and Education	1.35		Fermilab, UFlorida, UW Madison
Project Manager	0.85		Fermilab
<b>Total</b>	<b>34.78</b>	<b>5.35</b>	

By the end of 2010 the OSG staff has decreased by 2 FTEs due to staff leaving.

More than three quarters of the OSG staff directly support the operation and software for the on-going stakeholder productions and applications (the remaining quarter mainly engages new customers and extends and proves software and capabilities; and also provides management and communications etc.).

The 2010 WBS define more than 400 tasks, including both ongoing operational tasks and activities to put in place specific tools, capabilities and extensions. The area coordinators update the WBS quarterly. As new requests are made to the project the Executive Team prioritizes these against the existing tasks in the WBS. Additions are then made to the WBS to reflect those activ-

---

<sup>1</sup> Contribution included in the WBS. Many other contributions come from external projects including the science communities, software developers and resource administrators.

ities accepted into the work-program, some tasks are then dropped and other tasks deliverable dates are adjusted according to the new priorities. In FY10 the WBS was 85% accomplished (about the same as in FY09).

In 2010, the main technical achievements that directly support science include:

- Ongoing sustained support – including response to operational issues, inclusion and distribution of new **software** and **client tools**, in support of the **LHC Operations**. Improved “ticket-exchange” and “critical problem response” technologies and procedures were put in place, and well exercised, between the WLCG, EGEE/EGI operations services, and the US ATLAS and US CMS support processes. Agreed upon SLAs were put in place for the core operational services provided by OSG.
- OSG carried out “**prove-in**” of **reliable critical services (e.g. BDII) for LHC** and operation of services at levels that meet or exceed the needs of the experiments. This effort included robustness tests of the production infrastructure against failures and outages and validation of information by the OSG as well as the WLCG.
- Success in **STAR monte-carlo production using virtual machine based software images on grid and cloud resources**.
- Significant improvements in LIGO’s ability to use the OSG infrastructure, including adapting the Einstein@Home for Condor-G submission, resulting in a greater than **5x increase in the use of OSG by Einstein@Home**. Delivery of VDT components in “native packaging” for use on the LIGO data grid.
- Two significant publications from structural biology community (SBGrid) based on production running across many OSG sites as well as a rise in multi-user access through the SBGrid portal software.
- Entry of ALICE USA to full OSG participation for their sites in the USA, following a successful evaluation activity. This includes WLCG reporting and accounting through the OSG services.
- **Sustained support and better effectiveness for Geant4 validation runs.**
- **Ongoing** support for IceCube and GlueX, and initial engagement with LSST and NEES.
- Increased opportunistic cycles provided to OSG users by our collaborators in Brazil and Clemson.
- Diverse research supported by the increasingly active and effective Campus communities at the Holland Computer Center at the University of Nebraska and the support for multi-core applications for the University of Wisconsin Madison GLOW community.
- Security program activities that continue to **improve our defenses and capabilities towards incident detection and response** via review of our procedures by peer grids and adoption of new tools and procedures.
- Better understanding of the role and interfaces of Satellite projects as part of the larger OSG consortium contributions. Increased technical and educational collaboration with TeraGrid through the NSF-funded joint OSG – TeraGrid effort, ExTENCI, which began in August 2010 (see Section 4.5) and the joint OSG-TG summer student program.

- Contributions to the WLCG in the areas of new job execution service (CREAM), use of Pilot based job management technologies, and interfaces to commercial (EC2) and scientific (Magellan, FutureGrid) cloud services.
- Extensive for US ATLAS and US CMS Tier-3 sites in security vulnerability testing, packaging and testing of XROOTD and other needed technologies, and improvements (e.g. native packaging of software components, improved documentation, support for “storage only sites”) to reduce the “barrier to entry” to participate as part of the OSG infrastructure. We have held training schools for site administrators and a storage forum as part of the support activities.
- **Contributions to the PanDA and GlideinWMS** workload management software that have helped improve capability and supported broader adoption of these within the experiments. Reuse of these technologies by other communities. Configuration of PANDA for the integration test bed automated testing and validation system.
- Increased collaboration with ESnet and Internet2 on **perfSONAR and Identity Management**.
- **Improvement** and validation of the collaborative workspace **documentation** for all OSG areas of work.
- A successful summer school for 17 students and OSG staff mentors, as well as successful educational schools in south and central America.
- Continuation of excellence in e-publishing through the collaboration with the European Grid Infrastructure represented by the International Science Grid This Week electronic newsletter ([www.isgtw.org](http://www.isgtw.org)). The number of readers continues to grow.

**In summary, OSG continues to demonstrate that national cyberinfrastructure based on federation of distributed resources can effectively meet the needs of researchers and scientists.**

## 1.5 Preparing for the Future

In October 2010 the OSG project submitted a 1 year extension request to the DOE SciDAC-2 program to enable continuation of support for HEP and NP operations and production until March 2012.

We are now writing a proposal for the future of the OSG project, again intended to be submitted to both the DOE and NSF. The vision covers: **Sustaining** the infrastructure, services and software; **Extending** the capabilities and capacities of the services and software; and **Expanding** the reach to include new resource types – shared intra-campus infrastructures, commercial and scientific clouds, multi-core compute nodes – and new user communities in early stages of testing the benefit of OSG to their science such as the NEESComm and LSST programs.

The new OSG proposal will undergo an initial review in January 2011 by external reviewers invited by Executive Director. The reviewers include the head of the WLCG project Ian Bird, the 2 XD proposal PIs John Towns and Richard Moore, and senior management of ESNet Bill Johnston. As input to this proposal the Executive Director has worked with the project leaders and

council to propose a revised organization to be put in place in September 2011. A revised management plan is in preparation and will be available as part of the new proposal.

The following documents have been developed as input to the thinking of the OSG future programs:

- National CI and the Campuses (OSG-939)
- Requirements and Principles for the Future of OSG (OSG-938)
- OSG Interface to Satellite Proposals/Projects (OSG-913)
- OSG Architecture (OSG-966)
- Report from the Workshops on Distributed Computing, Multidisciplinary Science, and the NSF's Scientific Software Innovation Institutes Program (OSG-1002)

OSG contributed to a paper written by US ATLAS and US CMS "Assessment of Core Services provide to U.S. ATLAS and U.S. CMS by OSG" in February 2010. Since August 2010 we have included ALICE USA in the discussions and program of work. These communities have identified the following areas of need in extending the benefits provided by OSG services and capabilities:

- Configuration Management across services on different hosts – subject of a Blueprint discussion in May 2010. This is planned as future work by the software teams.
- Integration of Commercial and Scientific Clouds – Initial tests with the Magellan scientific Clouds at ANL and NERSC are underway and we have started a series of follow up technical meetings with the service groups every six weeks; explorations with EC2 are ongoing. We are using this work and the High Throughput Parallel Computing (HTPC) satellite project to understand how new resources and compute capabilities and capacities can be best integrated smoothly into the OSG infrastructure.
- Usability for collaborative analysis – evaluations of extended data management, opportunistic storage management, and the IRODS data management technologies are underway. Requirements for "Dynamic VOs/Workgroups" are being gathered from several stakeholders.
- Active management of shared capacity, utilization planning and change – an active analysis of available cycles is underway. We have started discussions on the application of more dynamic "OSG usage priorities".
- End-to-End Data Management challenges in light of advanced networks – we continue to work with Internet2 and ESNet research arms to look for opportunities for collaboration.

## 2. Contributions to Science


---

### 2.1 ATLAS


The goal of the computing effort as part of the U.S. ATLAS Operations Program is to empower U.S. physicists to address some of the most profound questions in particle physics today: what is the physical origin of mass? Do supersymmetric particles exist and will they shed light on the nature of dark matter? Does space-time have extra spatial dimensions? Answers to these questions would provide a major advance toward completing a unified view of the particles in nature, the forces with which particles interact, and their role in the past and future of our universe. This is a time when we have unusually compelling indications that the Large Hadron Collider (LHC) at CERN, with collision energies three to seven times beyond those available at previous facilities, will lead to especially important discoveries with implications over a broad field of fundamental science.

The ATLAS collaboration, consisting of 174 institutes from 38 countries, completed construction of the ATLAS detector at the LHC, and began first colliding-beam data taking in late 2009. The 44 institutions of U.S. ATLAS made major and unique contributions to the construction of the ATLAS detector, provided critical support for the ATLAS computing and software program and detector operations, and contributed significantly to physics analysis, results, and papers published.

Following the short run late 2009, LHC collider operations was resumed in March 2010. By early December 2010 the ATLAS collaboration has taken 1.2 billion events from proton-proton collisions and more than 200 million events from HI collisions. The total RAW (unprocessed) data volume amounts to almost 2 PB (1.6 PB pp and 0.3 PB HI data). While the RAW data was directly replicated to all ten ATLAS Tier-1 centers according to their MoU share (the U.S. receives and archives 23% of the total), the derived data was, after prompt reconstruction at the Tier-0 center, distributed to regional Tier-1 centers for group and user analysis and further distributed to the regional Tier-2 centers. Following significant improvements that were incorporated into the reconstruction code as well as improved calibration data becoming available, re-reconstruction of the data taken through November 2010 was conducted at the Tier-1 centers while users started to analyze the data using resources at the Tier-1 site, Tier-2 centers and their own institutional computing facilities. As the amount of initial data taking was small we observed users running data reduction steps followed by transfers of the derived, condensed data products to their interactive analysis servers, resulting in a reduced utilization of grid resources for a few months until LHC operations resumed in March 2010. However, machine luminosity ramped up rather quickly and much more data was taken in the second half of 2010, particularly in November and December.




**Figure 4: Integrated Luminosity as delivered by the LHC vs. recorded by ATLAS**




**Figure 5: Volume of RAW and derived data accumulated by ATLAS**

According to the data distribution policy as it was defined for the US region Event Summary Data (ESD) and Analysis Object Data (AOD) along with their derived versions were replicated in multiple copies to the Tier-2 centers in the U.S. Given that the replication process of several hundred terabytes of data from the Tier-1 center to the Tier-2s needed to be completed within the shortest possible period of time, the data rates the network and the storage systems had to sustain rose to an aggregate rate of 2 gigabytes per second. User analysis on the data started instantly with the arrival of the datasets at the sites. With more data becoming available the level of activity in the analysis queues at the Tier-1 and the Tier-2 centers was almost constant. A significant backlog of jobs waiting in the queues was observed at times. The workload management system (based on PanDA) distributed the load evenly across all sites able to run analysis on the required datasets. On average, the U.S. ATLAS facility contributes 30% of worldwide analysis-related data access. The number of user jobs submitted by the worldwide ATLAS community and brokered by PanDA to U.S. sites has reached an average number of 600,000 per month peaking occasionally at more than 1 million submitted jobs per month.




**Figure 6: Worldwide PanDA usage increased by a factor of three over the past year**



**Figure 7: Analysis performance of U.S. ATLAS sites in comparison to other regions – 75% of the ATLAS analysis jobs are completed by 20 sites**


Monte Carlo production is ongoing with some 50,000 concurrent jobs worldwide, and about 10,000 jobs running on resources provided by the distributed U. S. ATLAS computing facility comprising the Tier-1 center at BNL and 5 Tier-2 centers located at 9 different institutions spread across the U.S.



**Figure 8: OSG CPU hours (92M total) used by ATLAS over 12 months, color coded by facility.**

ATLAS' data distribution model has multiple replicas of the same datasets within regions (e.g. the U.S.) A significant problem was observed shortly after data taking resumed in March with sharply increasing integrated luminosity that produced an avalanche of new data. In particular, as Tier-2 sites disk storage filled up rapidly a solution had to be found to accommodate the data required for analysis. Based on job statistics that includes information about data usage patterns it was found that only a small fraction of the programmatically replicated data was actually accessed. U.S. ATLAS in agreement with ATLAS computing management consequently decided to change the Tier-2 related distribution model such that only datasets requested by analysis jobs are replicated. Programmatic replication of large amounts of ESDs was stopped, only datasets of all categories that are explicitly requested by analysis jobs are replicated from the Tier-1 center at BNL to the Tier-2 centers in the U.S. Since June 2010, when the initial version of a PanDA steered dynamic data placement system was deployed, we observe a healthy growth of the data volume on disk and are no longer facing situations where actually needed datasets cannot be accommodated.





**Figure 9: Cumulative evolution for DATADISK at Tier-2 centers in the U.S.**

Figure 9 clearly shows the exponential growth of the disk space utilization in April and May up to the point in June when the dynamic data placement system was introduced. Since then the data volume on disk is almost constant, despite the exponential growth of integrated luminosity and data volume of interest for analysis. Meanwhile the usage of the system, invented and tested by experts in the U.S. and fully transparent to users, was extended to other regions around the world. Disk capacities as provisioned at the Tier-2 centers have thus evolved from a kind of archival storage to a well-managed caching system. As a result of recent discussions it was decided to further develop the system by including the Tier-1 centers and evolve the distribution model such that it is no longer based on a strict hierarchy but allows direct access across all present hierarchy levels. A future model will include remote access to files and fractions thereof, rather than relying on formal dataset subscriptions via the ATLAS distributed data management system prior to getting access to the data.

Experience gained during the computing challenges and during the first year of ATLAS data taking gives us confidence that the tiered, grid-based computing model has sufficient flexibility to process, reprocess, distill, disseminate, and analyze ATLAS data. We have found, however, that the Tier-2 centers may not be sufficient to reliably serve as the primary analysis engine for more than 400 U.S. physicists. As a consequence a third tier with computing and storage resources located geographically close to the researchers was defined as part of the analysis chain as an important component to buffer the U.S. ATLAS analysis system from unforeseen future problems. Continued enhancement of U.S. ATLAS institutions' Tier-3 capabilities is still essential and will be based around the short and long-term analysis strategies of each U.S. group.

An essential component of this strategy is the creation of a centralized support structure to handle the increased number of campus-based computing clusters. OSG plays an important role in implementing the necessary components and helped in two key areas: packaging of batch processing (Condor) and storage management components (xrootd), both of which are easily installable and maintainable by physicists. Because this U.S. initiative (driven by Rik Yoshida from Argonne National Laboratory and Doug Benjamin from Duke University in collaboration with OSG) made rapid progress in just a few months, ATLAS Distributed Computing Management invited the initiative leaders to develop a technical and maintainable solution for the Tier-3 community. A very successful CERN workshop addressing Tier-3 issues was organized in January 2010, with good representation from around the globe. Major areas of work and interest were

identified during the meeting and short lived working groups were formed to address issues associated with in software installation, data and storage management, data replication and data access. Reports from these groups document the results relevant to their work area and provide important guidance for ongoing implementations.

Open Science Grid has organized regular Tier-3 Liaison meetings between several members of the OSG facilities, U.S. ATLAS and U.S. CMS. During these meetings, topics discussed include cluster management, site configuration, site security, storage technology, site design, and experiment-specific Tier-3 requirements. Based on information exchanged at these meetings several aspects of the U.S. ATLAS Tier-3 design were refined leading to improvements regarding the usability and maintainability.

Today U.S. ATLAS (contributing to ATLAS as a whole) relies extensively on services and software provided by OSG, as well as on processes and support systems that have been produced or evolved by OSG. OSG has become essential for the operation of the worldwide distributed ATLAS computing facility and the OSG efforts have aided the integration with WLCG partners in Europe and Asia. The derived components and procedures have become the basis for support and operation covering the interoperation between OSG, EGEE, and other grid sites relevant to ATLAS data analysis. OSG provides software components that allow interoperability with European ATLAS sites, including selected components from the gLite middleware stack such as LCG client utilities, and file catalogs.

It is vital to U.S. ATLAS that the present level of service continues uninterrupted for the foreseeable future, and that all of the services and support structures upon which U.S. ATLAS relies today have a clear transition or continuation strategy.

Based on observations U.S. ATLAS suggested to OSG to develop a coherent middleware architecture rather than continue providing a distribution as a heterogeneous software system consisting of components contributed by a wide range of projects. Difficulties we encountered included inter-component functional dependencies that require communication and coordination between component development teams. A technology working group, chaired by a member of the U.S. ATLAS facilities group (John Hover, BNL), has been asked to investigate, research, and clarify design issues and summarize technical design trade-offs such that the project teams working on component design and implementation can make informed decisions. In order to achieve the U.S. ATLAS goals, OSG needs an explicit, documented system design, or architecture, so that component developers can make compatible design decisions and virtual organizations (VO) such as U.S. ATLAS can develop their own applications based on the OSG middleware stack as a platform. The middleware architecture and the associated design roadmap are now under development. An area U.S. ATLAS is particularly interested in is Cloud Computing and Virtualization. Properties of interest include that usable resources can be made available instantaneously, once needed, at the required quantities. The foundation could be based on pre-defined virtual machines either furnished by U.S. ATLAS or the cloud provider, which establishes a homogeneous and well-defined system platform. Such a system is accessed via a clearly defined interface that supports functionality including virtual machine loading, starting, stopping, monitoring, cyber security/credential handling, node discovery and more – essentially all means ATLAS applications need to interact with cloud resources. There is a significant interest from U.S. ATLAS in running jobs supplied as virtual machines, because as experience has shown, in the classic grid context it is difficult in having the production and analysis framework to run on arbitrary systems.

Middleware deployment support provides an essential and complex function for U.S. ATLAS facilities. For example, support for testing, certifying and building a foundational middleware for production and distributed analysis activities is a continuing requirement, as is the need for coordination of the roll out, deployment, debugging and support for the middleware services. In addition, some level of preproduction deployment testing has been shown to be indispensable. This testing is currently supported through the OSG Integration Test Bed (ITB) providing the underlying grid infrastructure at several sites along with a dedicated test instance of PanDA, the ATLAS Production and Distributed Analysis system. These elements implement the essential function of validation processes that accompany incorporation of new versions of grid middleware services into the Virtual Data Toolkit (VDT), which provides a coherent OSG software component repository. U.S. ATLAS relies on the VDT and OSG packaging, installation, and configuration processes to provide a well-documented and easily deployable OSG software stack.

U.S. ATLAS greatly benefits from OSG's Gratia accounting services, as well as the information services and probes that provide statistical data about facility resource usage and site information passed to the application layer and to WLCG for review of compliance with MoU agreements.

An essential component of grid operations is operational security coordination. The coordinator provided by OSG has good contacts with security representatives at the U.S. ATLAS Tier-1 center and Tier-2 sites. Thanks to activities initiated and coordinated by OSG a strong operational security community has grown up in the U.S. in the past few years, ensuring that security problems are well coordinated across the distributed infrastructure.

No significant problems with the OSG provided infrastructure have been encountered since the start of LHC data taking. However, there is an area of concern that may impact the facilities' performance in the future. As the number of job slots at sites continues to grow the performance of pilot submission through CondorG and the underlying Globus Toolkit 2 (GT2) based gatekeeper must keep up without slowing down job throughput, particularly when running short jobs. When addressing this point with the OSG facilities team we found that they were open to evaluating and incorporating recently developed components such as the CREAM Computing Element (CE) provided by EGI developers in Italy. Intensive tests were conducted by the Condor team in Madison and numerous integration issues were identified and resolved by the VDT team.

In the area of middleware extensions, U.S. ATLAS continued to benefit from the OSG's support for and involvement in the U.S. ATLAS-developed distributed processing and analysis system (PanDA) layered over the OSG's job management, storage management, security and information system middleware and services. PanDA provides a uniform interface and utilization model for the experiment's exploitation of the grid, extending across OSG, EGEE and Nordugrid. It is the basis for distributed analysis and production ATLAS-wide, and is also used by OSG as a WMS available to OSG VOs, as well as a PanDA based service for OSG Integrated Testbed (ITB) test job submission, monitoring and automation. This year the OSG's WMS extensions program continued to provide the effort and expertise on PanDA security that has been essential to establish and maintain PanDA's validation as a secure system deployable in production on the grids. In particular PanDA's glxexec-based pilot security system developed in this program went through production readiness tests in the U.S. and Europe throughout the year.

Another important extension activity during the past year was in WMS monitoring software and information systems. During the year ATLAS and U.S. ATLAS continued the process of merging the PanDA/US monitoring effort with CERN-based monitoring efforts, together with

the ATLAS Grid Information System (AGIS) that integrates ATLAS-specific information with the grid information systems. The agreed common approach utilizes a python apache service serving json-formatted monitoring data to rich jQuery-based clients. This served as the basis for a new prototype PanDA monitoring infrastructure developed by the OSG effort, now being integrated with the CERN-based effort that also has made substantial progress during the last year. We expect that during 2011 this merge will be completed and PanDA monitoring will have a well defined evolution path for the future. The PanDA monitor began to feel the effects of Oracle scalability limitations (at least with the present configuration of the PanDA Oracle DBs) towards the end of the quarter, and planning began for a program investigating alternative back end DB technologies, particularly for the deep archive of job and file data which show the most severe scaling limitations and which have access patterns amenable to other storage approaches, in particular to the highly scalable key-value pair based systems such as Cassandra and Hive that have emerged as open source software from web behemoths such as Google, Amazon and Facebook. We expect this program to grow into a significant part of our activity in 2011.

## 2.2 CMS

During 2010 CMS has transitioned from commissioning the detector to producing its first physics results across the entire range of physics topics. The first eight scientific papers have been published in peer reviewed journals, including Phys. Rev. Letters, Physics Letters B, Eur. Phys. Journal C, and the Journal of High Energy Physics, and many tens of scientific publications are either submitted to journals, or presently in peer review within the collaboration.

Among the published results are already some first surprises, like the first “Observation of Long-Range, Near-Side Angular Correlations in Proton-Proton Collisions”, others are searches for new physics that already exceed the sensitivity reached by previous generations of experiments, still others, including the first observation of top pair production at the LHC, are major milestones that measure the cross section for the dominant Standard Model background processes to much of the ongoing, as well as future new physics searches.

Computing has proven to be the enabling technology it was designed to be, providing an agile environment for scientific discovery. U.S. CMS resources available via the Open Science Grid have been particularly important to the scientific output of the CMS experiment. The seven Tier-2 sites are among the eleven most heavily used Tier-2 sites globally, accounting for about 40% of the total data analysis volume worldwide, and among the top eight contributors to the global Monte Carlo production effort, providing roughly 50% of the total simulation volume worldwide. In aggregate, the U.S. CMS sites have received 3.2 PB of data, 1.3 PB of which was transferred from the FNAL Tier-1, 1.0 PB from other Tier-1 sites, and the rest from other Tier-2 sites. The U.S. leadership position within CMS as indicated by these metrics is attributable to superior reliability, and agility of U.S. sites. We host a complete copy of all core data samples distributed across the seven US Tier-2 sites, and due to the excellent performance of the storage infrastructures, we are able to refresh data quickly. It is thus not uncommon that data becomes available first at U.S. sites, attracting time critical data analysis to those sites.

Figure 10 and Figure 11 show example metrics for CMS use of computing. Figure 10 shows the number of CPU hours per week used by CMS on OSG. Figure 11 shows the number of pending jobs versus time for CMS worldwide. In both cases, the color coding indicates different sites. Figure 1 shows that use of OSG by CMS has reached a plateau in September 2010. Figure 2 shows that around the same time, the total number of pending jobs started to increase dramati-

ly. When ordering sites according to the number of pending jobs, we find that of the top six most heavily congested sites, four are U.S. CMS Tier-2 sites, the third site is the FNAL Tier-1, and the fifth is the CERN Tier-0. We thus conclude that U.S. CMS is presently resource constrained.


The Open Science Grid has been a significant contributor to this success by providing critical computing infrastructure, operations, and security services. These contributions have allowed U.S. CMS to focus experiment resources on being prepared for analysis and data processing, by saving effort in areas provided by OSG. OSG provides a common set of computing infrastructure services on top of which CMS, with development effort from the U.S., has been able to build a reliable processing and analysis framework that runs on the Tier-1 facility at Fermilab, the project supported Tier-2 university computing centers, and opportunistic Tier-3 centers at universities. There are currently 27 Tier-3 centers registered with the CMS data grid in the U.S., 20 of them provide additional simulation and analysis resources via the OSG. The remainder are Universities that receive CMS data via the CMS data grid, using an OSG storage element API, but do not (yet) make any CPU cycles available to the general community. OSG and U.S. CMS work closely together to ensure that these Tier-3 centers are fully integrated into the globally distributed computing system that CMS science depends on.

In addition to common interfaces, OSG provides the packaging, configuration, and support of the storage services. Since the beginning of OSG the operations of storage at the Tier-2 centers have improved steadily in reliability and performance. OSG is playing a crucial role here for CMS in that it operates a clearinghouse and point of contact between the sites that deploy and operate this technology and the developers. In addition, OSG fills in gaps left open by the developers in areas of integration, testing, and tools to ease operations.


OSG has been crucial to ensure U.S. interests are addressed in the WLCG. The U.S. is a large fraction of the collaboration both in terms of participants and capacity, but a small fraction of the sites that make-up WLCG. OSG is able to provide a common infrastructure for operations including support tickets, accounting, availability monitoring, interoperability and documentation. Now that CMS is taking data, the need for sustainable security models and regular accounting of available and used resources is crucial. The common accounting and security infrastructure and the personnel provided by OSG represent significant benefits to the experiment, with the teams at Fermilab and the University of Nebraska providing the development and operations support, including the reporting and validation of the accounting information between the OSG and WLCG.

In addition to these general statements, we'd like to point to two specific developments that have become increasingly important to CMS within the last year. Within the last two to three years, OSG developed the concept of "Satellite projects", and a notion of providing an "ecosystem" of independent technology projects that enhance the overall national computing infrastructure in close collaboration with OSG. CMS is starting to benefit from this concept as it has stimulated close collaboration with computer scientists on a range of issues including 100Gbps networking, workload management, cloud computing and virtualization, and High Throughput Parallel Computing that we expect will lead to multi-core scheduling as the dominant paradigm for CMS in a few years time. The existence of OSG as a "collaboratory" allows us to explore these important technology directions in ways that are much more cost effective, and more likely to be successful than if were pursuing these new technologies within a narrow CMS specific context.

Finally, within the last year, we have seen increasing adoption of technologies and services originally developed for CMS. Most intriguing is the deployment of glideinWMS as an OSG service, adopted by a diverse set of customers including structural biology, nuclear physics, applied mathematics, chemistry, astrophysics, and CMS data analysis. A single instance of this service is jointly operated by OSG and CMS at UCSD for the benefit of all of these communities. OSG developed a Service Level Agreement that is now being reviewed for possible adoption also in Europe. Additional instances are operated at FNAL for the Run II experiments, MINOS, and data reprocessing for CMS at Tier-1 centers.




**Figure 10: OSG CPU hours (80M total) used by CMS over 12 months, color-coded by facility.**



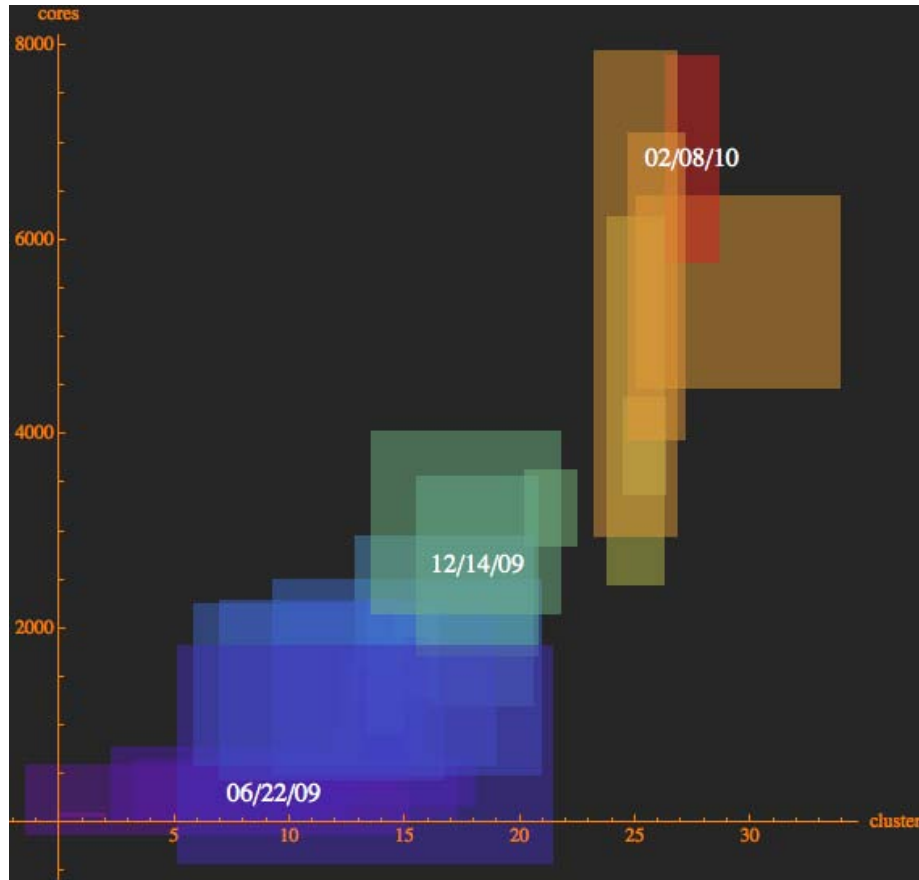
**Figure 11: Average number of pending jobs per week by CMS worldwide**

## 2.3 LIGO

LIGO continues to leverage the Open Science Grid for opportunistic computing cycles associated with its grid based Einstein@Home application, known as Einstein@OSG. This application is one of several in use for an “all-sky” search for gravitational waves of a periodic nature attributed to elliptically deformed pulsars. Such a search requires enormous computational resources to fully exploit the science content available within LIGO’s vast datasets. Volunteer and opportunistic computing based on the BOINC (Berkeley Open Infrastructure for Network Computing) has been leveraged to utilize as many computing resource worldwide as possible. Since porting to the grid based Einstein@OSG code onto the Open Science Grid more roughly two year ago, steady advances in both the code performance, reliability and overall deployment onto the Open Science Grid have been demonstrated. OSG has routinely ranked in the top two or three computational providers for this LIGO analysis worldwide. This year, more than 32 million CPU-Hours have been provided towards this search for pulsar signals by the Open Science Grid.



**Figure 12: Opportunistic usage of the OSG by LIGO’s grid based Einstein@Home application for the current year.**



**Figure 13: Scaling improvements in the utilization of the Open Science Grid by the Einstein@OSG application for the past one and a half years. Each rectangle represents the weekly view of number of sites (x-axis) versus the number of CPU cores (y-axis). We are currently saturated at roughly 30 sites, utilizing about 6000 cores on a weekly average.**

This year has also seen development effort on a variation of the search for gravitational waves from pulsars with the porting of the “PowerFlux” application onto the Open Science Grid. This is also a broadband search but utilizes a power averaging scheme to more quickly cover a large region of the sky over a broad frequency band. The computational needs are not as large as with the Einstein@Home application at the expense of lower signal resolution. The code is currently being wrapped up to provide better monitoring in a grid environment where remote login is not supported.

One of the most promising sources of gravitational waves for LIGO is from the inspiral of a system of compact black holes and/or neutron stars as the system emits gravitational radiation leading to the ultimate coalescence of the binary pair. The binary inspiral data analyses typically involve working with tens of terabytes of data in a single workflow. Collaborating with the Pegasus Workflow Planner developers at USC-ISI, LIGO continues to identify changes to both Pegasus and to the binary inspiral workflow codes to more efficiently utilize the OSG and its emerging storage technology, where data must be moved from LIGO archives to storage resources near the worker nodes on OSG sites.

One area of intense focus this year has been on the understanding and integration into workflows of Storage Resource Management (SRM) technologies used in OSG Storage Element (SE) sites



to house the vast amounts of data used by the binary inspiral workflows so that worker nodes running the binary inspiral codes can effectively access the LIGO data. The SRM based Storage Element established on the LIGO Caltech OSG integration testbed site is being used as a development and test platform to get this effort underway without impacting OSG production facilities. Using Pegasus for the workflow planning, DAGs for the binary inspiral data analysis application using of order ten terabytes of LIGO data have successfully run on three production sites. Performance studies this year have suggested that the use of glide-in technologies can greatly improve the total run time requirements for these large workflows made up of tens of thousands of jobs. This is another area where Pegasus in conjunction with its Corral glide-in features have resulted in further gains in the ability to port and effectively use a complex LIGO data analysis workflow, designed originally for running on the LIGO Data Grid, over to the Open Science Grid where there are sufficient similarities to make this possible, but sufficient differences to require detailed investigations and development activities to reach the desired science driven goals.

LIGO continues working closely with the OSG Security team, DOE Grids, and ESnet to evaluate the implications of its requirements on authentication and authorization within its own LIGO Data Grid user community and how these requirements map onto the security model of the OSG.

## **2.4 ALICE**

The ALICE experiment at the LHC relies on a mature grid framework, AliEn, to provide computing resources in a production environment for the simulation, reconstruction and analysis of physics data. Developed by the ALICE Collaboration, the framework is fully operational with sites deployed at ALICE and WLCG Grid facilities worldwide. During 2010, ALICE US collaboration deployed significant compute and storage resources in the US, anchored by new Tier 2 centers at LBNL/NERSC and LLNL. These resources, accessible via the AliEn grid framework, are being integrated with OSG to provide accounting and monitoring information to ALICE and WLCG while allowing unused cycles to be used by other NP groups.


In early 2010, the ALICE USA Collaboration's Computing plan was formally adopted and funded by DOE. The plan specifies resource deployments at both the existing NERSC-PDSF cluster at LBNL and the LLNL/LC facility, and operational milestones for meeting ALICE USA's required computing contributions to the ALICE experiment. A centerpiece of the plan is the integration of these resources with the OSG in order to leverage OSG capabilities for accessing and monitoring distributed compute resources. Milestones for this work included: completion of more extensive scale-tests of the AliEn-OSG interface to ensure stable operations at full ALICE production rates, establishment of operational OSG resources at both facilities, and activation of OSG accounting reports of utilization of these resources by ALICE to the WLCG. During this year, with the support of OSG personnel, we have met most of the goals set forth in the computing plan.

NERSC/PDSF has operated as an OSG facility for several years and was the target site for the initial development and testing of an AliEn-OSG interface. With new hardware deployed for ALICE on PDSF in June of 2010, a new set of scaling tests were carried out by ALICE which demonstrated that the AliEn-OSG interface was able to sustain job submission rates and steady-state job-occupancy required by the ALICE team. Since about mid-July of 2010, ALICE has run production at PDSF with a steady job concurrency rate of about 300 jobs consistent with the computing plan.

During the fall of 2010 a small OSG-ALICE task force was renewed to facilitate further integration with OSG. Work in the group has focused on the ALICE requirement that resource utilization is reported by the OSG to WLCG. This work has included cross-checks on accounting records reported by the PDSF OSG site as well as developing additional tools needed for deploying OSG accounting at the LLNL/LC facility. As a result of these efforts, both facilities currently report accounting records to OSG and these reports will be forwarded to WLCG via normal OSG operations as soon as the facilities are fully registered with the WLCG.

## 2.5 D0 at Tevatron

The D0 experiment continues to rely heavily on OSG infrastructure and resources in order to achieve the computing demands of the experiment. The D0 experiment has successfully used OSG resources for many years and plans on continuing with this very successful relationship into the foreseeable future. This usage has resulted in a tremendous science publication record, including contributing to improved limits on the Higgs mass exclusion as shown in Figure 14.



**Figure 14:** Plot showing the latest combined D0 and CDF results on the observed and expected 95% confidence limit upper limits on the ratios to the Standard Model cross section as a function of Higgs mass.

All D0 Monte Carlo simulation is generated at remote sites, with OSG continuing to be a major contributor. During the past year, OSG sites simulated over 500 million events for D0, approximately 1/3 of all production. The rate of production has leveled off over the past year as almost all major sources of inefficiency have been resolved and D0 continues to use OSG resources very efficiently. The changes in policy at numerous sites for job preemption, the continued use


of automated job submissions, and the use of resource selection has allowed D0 to opportunistically use OSG resources to efficiently produce large samples of Monte Carlo events. D0 continues to use approximately 31 OSG sites regularly in its Monte Carlo production. The total number of D0 OSG MC events produced over the past several years has exceeded 1 billion events (Figure 15).

Over the past year, the average number of Monte Carlo events produced per week by OSG continues to remain approximately constant. Since we use the computing resources opportunistically, it is interesting to find that, on average, we can maintain an approximately constant rate of MC events (Figure 16). Dips in OSG production are now only typically due to D0 switching to new software releases which temporarily stops our requests to OSG. Over the past year D0 has been able to obtain the necessary opportunistic resources to meet our Monte Carlo needs even though the LHC also has high demand. We have been able to achieve this by continuing to improve our efficiency and to add additional resources each year. It is hoped that the Tevatron program will continue to run for several more years, therefore D0 will continue to need OSG resources for many more years.


Last year D0 was able to use LCG resources at a significant level to produce Monte Carlo events. The primary reason that this was possible was that LCG began to use some of the infrastructure developed by OSG. Because LCG was able to easily adopt some of the OSG infrastructure, D0 was able to produce approximately 200 million Monte Carlo events last year on LCG. The ability for OSG infrastructure to be used by other grids has proved to be very beneficial.

The primary processing of D0 data continues to be run using OSG infrastructure. One of the very important goals of the experiment is to have the primary processing of data keep up with the rate of data collection. It is critical that the processing of data keep up in order for the experiment to quickly find any problems in the data and to keep the experiment from having a backlog of data. Typically D0 is able to keep up with the primary processing of data by reconstructing 6-8 million events/day (Figure 17). However, when the accelerator collides at very high luminosities, it is difficult to keep up with the data using our standard resources. However, since the computing farm and the analysis farm have the same infrastructure, D0 is able to move analysis computing nodes to primary processing to improve its daily processing of data, as it has done on more than one occasion. This flexibility is a tremendous asset and allows D0 to efficiently use its computing resources. Over the past year D0 has reconstructed nearly 2 billion events on OSG facilities. In order to achieve such a high throughput, much work has been done to improve the efficiency of primary processing. In almost all cases, only 1-2 job submissions are needed to complete a job, even though the jobs can take several days to finish, see Figure 18.


OSG resources continue to allow D0 to meet its computing requirements in both Monte Carlo production and in data processing. This has directly contributed to D0 publishing 29 papers in 2010 (with 10 additional papers submitted for publication) see [http://www-d0.fnal.gov/d0\\_publications/](http://www-d0.fnal.gov/d0_publications/).



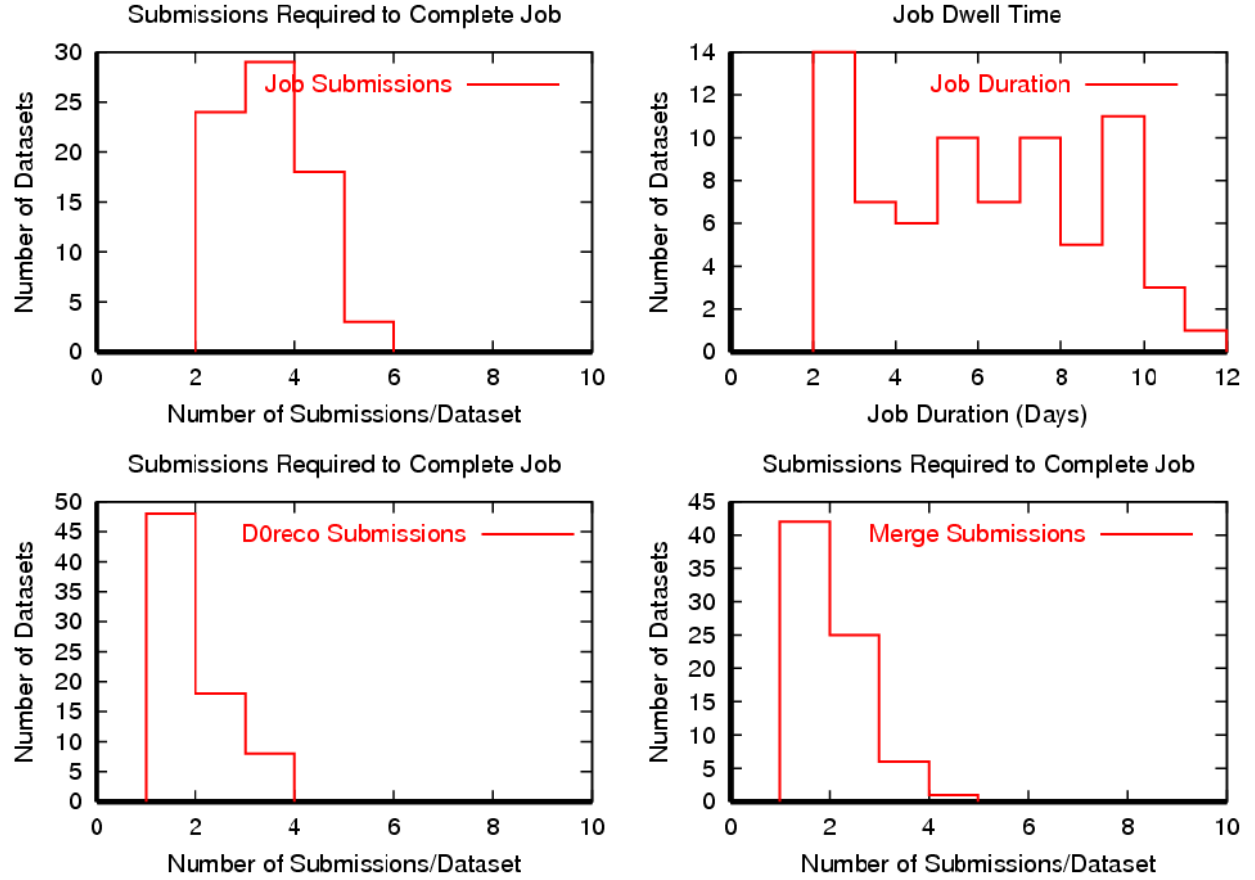
**Figure 15: Cumulative number of D0 MC events generated by OSG during the past year.**



**Figure 16: Number of D0 MC events generated per week by OSG during the past year. The dip in production in December and January was due to D0 switching to a new software release which temporarily reduced our job submission rate to OSG**




**Figure 17: Daily production of D0 data events processed by OSG infrastructure. The dips correspond to times when the accelerator was down for maintenance so no events needed to be processed.**



**Figure 18: Submission statistics for D0 primary processing. In almost all cases, only 1-2 job submissions are required to complete a job even though jobs can run for several days.**

## 2.6 CDF at Tevatron

In 2009-2010, the CDF experiment produced 48 new results for winter 2010 and then an additional 42 new results for summer using OSG infrastructure and resources. Included in these results was achieving Standard Model sensitivity at the 95% CL to a Higgs at a mass of 165  $\text{GeV}/c^2$  and, when combined with D0, excluding a Higgs with a mass between 158 and 175  $\text{GeV}/c^2$  (Figure 20).



**Figure 19: Upper limit plot of recent CDF search for the Standard Model Higgs**

The OSG resources support the work of graduate students, who are producing one thesis per week, and the collaboration as a whole, which is submitting a publication of new physics results every ten days. About 50 publications have been submitted in this period. A total of 900 million Monte Carlo events were produced by CDF in the last year. Most of this processing took place on OSG resources. CDF also used OSG infrastructure and resources to support the processing of 2.4 billion raw data events. A major reprocessing has been under way to increase the b tagging efficiency for improved sensitivity to low mass Higgs. The production output from this and normal processing was 5.4 billion reconstructed events, some of which were then processed into 4.7 billion ntuple events with the remainder to be processed with improved ntuple information in the beginning of next year. An additional 471M events were created from Monte Carlo. Detailed numbers of events and volume of data are given in Table 4 (total data since 2000) and Table 5 (data taken from January 2010 to December 2010).

**Table 4: CDF data collection since 2000**

Data Type	Volume (TB)	# Events (M)	# Files
Raw Data	1852	12519	2126833
Production	2712	18454	2520643
MC	918	6241	1062320
Stripped-Prd	96	823	90925
Stripped-MC	0	3	533
MC Ntuple	441	6494	349331
<b>Total</b>	<b>6019</b>	<b>44534</b>	<b>6150585</b>

**Table 5: CDF data collection from June 2009- June 2010**

<b>Data Type</b>	<b>Data Volume (TB)</b>	<b># Events (M)</b>	<b># Files</b>
Raw Data	400.5	2448.7	446210
Production	942.5	5410.5	776545
MC	94.38	471.07	114833
Stripped-Prd	16.52	92.545	13233
Stripped-MC	0	0	0
Ntuple	166.2	4692.3	135599
MC Ntuple	140.1	471.067	117109
<b>Total</b>	<b>1760.2</b>	<b>13586.2</b>	<b>1603529</b>

The OSG provides the collaboration computing resources through two portals. The first, the North American Grid portal (NAmGrid), covers the functionality of MC generation in an environment which requires the full software to be ported to the site and only Kerberos or grid authenticated access to remote storage for output. The second portal, CDFGrid, provides an environment that allows full access to all CDF software libraries and methods for data handling.

CDF operates the pilot-based Workload Management System (glideinWMS) as the submission method to remote OSG sites. Figure 21 shows the number of running jobs on NAmGrid and demonstrates that there has been steady usage of the facilities, while Figure 22, a plot of the queued requests, shows that there is large demand. CDF MC production is submitted to NAmGrid and use of OSG CMS, CDF, and General purpose Fermilab resources plus MIT is observed.

A large resource provided by Korea at KISTI is in operation and provides a large Monte Carlo production resource with high-speed connection to Fermilab for storage of the output. It also provided a cache that allows the data handling functionality to be exploited. The system was commissioned and 10TB of raw data were processed using SAM data handling with KISTI in the NAmGrid portal. Problems in commissioning were handled with great speed by the OSG team through the “campfire” room and through weekly VO meetings. Lessons learned to make commissioning and debugging easier were analyzed by the OSG group. KISTI is being run as part of NAmGrid for MC processing when not being used for reprocessing.






Figure 20: Running jobs on NAMGrid




Figure 21: Waiting CDF jobs on NAMGrid, showing large demand, especially in preparation for the 42 results sent to Lepton-Photon in August 2009 and the rise in demand for the winter 2010 conferences.

Plots of the running jobs and queued requests are shown in Figure 23 and Figure 24. The very high demand for the CDFGrid resources observed during the winter conference season (leading to 48 new results) and again during the summer conference season (leading to an additional 42 new results), is noteworthy. Queues exceeding 30,000 jobs can be seen. The decrease in load over the summer was due to an allocation of 15% of the CDFGrid resources for testing with SLF5. This testing period ended with the end of the summer conference season on August, 2010. At that point all of the CDFGrid and NAMGrid resources were upgraded to SLF5 and this became the default. CDF raw data processing, ntupling and user analysis has now been converted to SLF5.




Figure 22: Running CDF jobs on CDFGrid




Figure 23: Waiting CDF jobs on CDFGrid

A clear pattern of CDF computing has emerged. There is high demand for Monte Carlo production in the months after the conference season, and for both Monte Carlo and data starting about two months before the major conferences. Since the implementation of opportunistic computing on CDFGrid in August, the NAmGrid portal has been able to take advantage of the computing resources on FermiGrid that were formerly only available through the CDFGrid portal. This has led to very rapid production of Monte Carlo in the period of time between conferences when the generation of Monte Carlo datasets are the main computing demand.

A number of issues affecting operational stability and operational efficiency have been pointed out in the last report. Those that remain and solutions or requests for further OSG development are cited here.

- **Service level and Security:** Since April, 2009 Fermilab has had a new protocol for upgrading Linux kernels with security updates. While main core services can be handled with a rolling reboot, the data handling still requires approximately quarterly draining of queues for up to 3 days prior to reboots.
- **Opportunistic Computing/Efficient resource usage:**  
Preemption policy has not been revisited and CDF has not tried to include any new sites due to issues that arose when commissioning KISTI. It was found that monitoring showed that the KISTI site was healthy while we found that glideins from glideinWMS were “swallowed” leaving a cleanup operational issue. This is being addressed by OSG.
- **Management of database resources:** Monte Carlo production led to a large load on the CDF database server for queries that could be cached. An effort to reduce this load was launched and most queries were modified to use a Frontier server with Apache. This avoided a problem in resource management provided Frontier servers are provided with each site installation.
- **Management of input Data resources:** During the conference crunch in July 2009 and again in March 2010 there was huge demand on the data-handling infrastructure and the 350TB disk cache was being turned over every ten weeks. Effectively the files were being moved from tape to disk, being used by jobs and deleted. This in turn led to many FermiGrid worker nodes sitting idle waiting for data. A program to understand the causes of idle computing nodes from this and other sources has been initiated and CDF users are asked to more accurately describe what work they are doing when they submit jobs by filling in qualifiers in the submit command. Pre-staging of files was implemented but further use of file management using SAM is being made default for the ntuple analysis framework. There is a general resource management problem pointed to by this and the database overload issue.

Resource requirements of jobs running on OSG should be examined in a more considered way and would benefit from more thought by the community at large.

The usage of OSG for CDF has been fruitful and the ability to add large new resources such as KISTI as well as more moderate resources within a single job submission framework has been extremely useful for CDF. The collaboration has produced significant new results in the last year with the processing of huge data volumes. Significant consolidation of the tools has occurred. In the next year, the collaboration looks forward to a bold computing effort in the push to see evidence for the Higgs boson, a task that will require further innovation in data handling and significant computing resources in order to reprocess the large quantities of Monte Carlo and data needed to achieve the desired improvements in tagging efficiencies. We look forward to another year with high publication rates and interesting discoveries.

## 2.7 Nuclear physics

STAR’s tenth year of data taking has brought new levels of data challenges, with the most recent year’s data matching the integrated data of the previous decade. Now operating at the Petabyte scale, the data mining and production has reached its maximum potential. Over a period of 10

years of running, the RHIC/STAR program has seen data rates grow by two orders of magnitudes, yet data production has kept pace and data analysis and science productivity remained strong. In 2010, the RHIC program and Brookhaven National Laboratory earned recognition as number 1 for Hadron collider research.<sup>2</sup>

To effectively face the data challenge, all raw simulations had previously been migrated to Grid-based operations. This year, the migration has been expanded, with a noticeable shift toward the use of Cloud resources wherever possible. While Cloud resources had been an interest to STAR as early as 2007, our previous years' reports noted multiple tests and a first trial usage of Cloud resources (Nimbus) in 2008/2009 at the approach of a major conference, absorbing additional workload stemming from a last minute request. This mode of operation has continued as the Cloud approach allows STAR to run what our collaboration has not been able to perform on Grid resources due to technical limitations (harvesting of resources on the fly has been debated in length by STAR as an unreachable ideal for experiment equipped with complex software stacks). Grid usage remains restricted to either opportunistic use of resources for event generator-based production (self-contained program easily assembled) or non-opportunistic / dedicated site usage with a pre-installed software stack maintained by a local crew allowing running STAR's complex workflows. Cloud resources, coupled with virtualization technology, permit relatively easy deployment of the full STAR software stack within the VM, allowing large simulation requests to be accommodated. Even more relevant for STAR's future, recent tests successfully demonstrated that larger scale real data reconstruction is easily feasible. Cloud activities and development remain (with some exceptions) outside the scope and program of work of the Open Science Grid; one massive simulation exercise was partly supported by the ExTENCI satellite project.

STAR had planned to also run and further test the Glow resources after an initially successful reported usage via a Condor/VM mechanism. However, several alternative resources and approaches offered themselves. The use of the Clemson model especially appeared to allow for faster convergence and deliverables of a needed simulation production in support of the Spin program component of RHIC/STAR. Within a sustained scale of 1,000 jobs (peaking at 1,500 jobs) for three weeks, STAR clearly demonstrated that a full fledged Monte-Carlo simulation followed by a full detector response simulation and track reconstruction was not only possible on Cloud but of a large benefit to our user community. With over 12 billion PYTHIA events generated, this production represented the largest PYTHIA event sample ever generated in our community. The usage of Cloud resources in this case expanded the resources capacity for STAR by 25% (comparing to the resources available at BNL/RCF) and, for a typical student's work, allowed a yearlong science time wait to be delivered in a few weeks. Typically, a given user at the RCF would be able to claim 50 job slots (the facility being shared by many users) while in this exploitation of Cloud resources, all 1,000 slots were uniquely dedicated to a given task and one student. The sample represented a four order of magnitude increase in statistics comparing to other studies made in STAR with a near total elimination of statistical uncertainties which would have reduced the significance of model interpretations. The results were presented at the Spin 2010 conference where unambiguous agreement between our data and the simulation was shown. It is noteworthy that the resources were gathered in an opportunistic manner as seen in Figure 25. We would like to acknowledge the help from our colleagues from Clemson, partly funded by the ExTENCI project.

---

<sup>2</sup> <http://sciencewatch.com/ana/st/hadron/institutions/>





Figure 24: Graph of the number of available machines to STAR (in red), working machines (in green) and idle nodes (in blue) within an opportunistic resource gathering at Clemson University. Within this period, the overlap of the red and green curve demonstrates the submission mechanism allows for immediate harvesting of resources as they become available.

An overview of STAR’s Cloud efforts and usage has been presented at the OSG all hand-meeting in March 2010 (see “[Status of STAR’s use of Virtualization and Clouds](#)”) and at the International Symposium on Grid Computing 2010 (“[STAR’s Cloud/VM Adventures](#)”). Further overview of activities was given at the ATLAS data challenge workshop held at BNL that same month and finally, a summary presentation was given the CHEP 2010 conference in Taiwan in October (“[When STAR Meets the Clouds – Virtualization & Grid Experience](#)”). Based on usage trend and progress with Cloud usage and scalability, we project that 2011 will see workflow of the order of 10 to 100k jobs sustained as routine operation (see Figure 26).



**Figure 25: Summary of our Cloud usage as a function of date. As seen, the rapid progression of the exploitation and usage may indicate that a 10,000 job scale in 2011 may be at reach.**

From BNL, we steered Grid-based simulation productions (essentially running on our NERSC resources), and STAR has in total produced 4.8 Million events representing a total of 254,200 CPU hours of processing time using the standard OSG/Grid infrastructure. During our usage of the NERSC resources, we re-enabled the SRM data transfer delegation mechanism allowing for a job to terminate and pass to a third party service (SRM) the task of transferring the data back to the Tier0 center, BNL. We had previously used this mechanism but not integrated it into our regular workflow as the network transfers allowed for immediate globus-based file transfer with no significant additional time added to the workflow. However, due to performance issues with our storage cache at BNL (outside of STAR's control and purview), the transfers were recently found, at times, to add a significant overhead to the total job time (41% impact on total job time). The use of a 0.5 TB cache on the NERSC side and the SRM delegation mechanism allowed mitigation of the delay problems. In addition to NERSC, large simulation event generations were performed on the CMS/MIT site for the study of prompt photon cross section and double spin asymmetry. Forty-three million raw PYTHIA events were generated, amongst which 300 thousand events were passed to GEANT as part of cross section / pre-selection speed up (event filtering at generation), a mechanism designed in STAR to cope with large and statistically challenging simulations (cross section-based calculations require however to generate with a non-restrictive phase space and count the events passing our filter and the one being rejected). Additionally, 20 billion PYTHIA events (1 million filtered and kept) were also processed on that facility. The total resource usage was equivalent to about 100,000 hours of CPU hours spanning over a period of two months total.


STAR has also begun to test the resources provided by the Magellan project at NERSC and aims at pushing a fraction of its raw datasets to the Magellan Cloud for immediate processing via an



hybrid Cloud/Grid approach (a standard Globus gatekeeper will be used as well as data transfer tools), while the virtual machine capability will be leveraged for provisioning the resources with the most recent STAR software stack. The goal of this exercise is to provide a fast lane processing of data for the Spin working group with processing of events in near real time. While near real-time processing is already practiced in STAR, the run support data production known as “FastOffline” currently uses local BNL/RCF resources and passes over a sample of the data only once. The use of Cloud resources would allow outsourcing yet another workflow in support of the experiment scientific goals. This processing is also planned to be iterative, each pass using more accurate calibration constants. We expect by then to shorten the publication cycle of results from proton+proton 500 GeV Run 11 data by a year. During the Clemson exercise, STAR had designed a scalable database access approach which we will also use for this exercise. In essence, leveraging the capability of our database API, a “snapshot” is created and uploaded to the virtual machine image and a local database service is started. The need for a remote network connection is then abolished (as well as the possibility of thousands of processes overstressing the RHIC/BNL database servers). A fully ready database factory is available for exploitation. Final preparations of the workflow are in discussion, and if successful, this modus-operandi will represent a dramatic shift in the data processing capabilities of STAR. Raw data production will no longer be constrained to dedicated resources but allowed on widely distributed Cloud based resources).

The OSG infrastructure has been heavily used to transfer and redistribute our datasets from the Tier-0 (BNL) center to our other facilities. Noticeably, the NERSC/PDSF center holds full sets of analysis ready data (known as micro-DST) for the Year 9 data and, on the approach of the [Quark Matter 2011 conference](#), we plan to make available the year 10 data allowing to spread user analysis over multiple facilities (Tier2 centers in STAR typically transfer only subsets of the data, targeting local analysis needs). Up to 7 TB of data can be transferred a day and over 150 TB of data were transferred in 2010 from BNL to PDSF.

As a collaborative effort between BNL and the Prague institution, STAR is in the process of deploying a data placement planer tool in support of its data redistribution and production strategy. The planer is based on reasoning as per {from where / to where} the data has to be {taken / should be moved} to achieve the fastest possible plan, whether the plan is a data placement or a data production and processing turn-around. To get a baseline estimate of the transfer speed limit between BNL and PDSF, we have reassessed the link speed. The expected transfer profile is given by Figure 27. We expect this activity to reach completion by mid-2011.



**Figure 26: Transfer speed maximum between BNL and NERSC facility. The speed maximum is consistent with a point to point 1 Gb/sec link.**


All STAR physics publications acknowledge the resources provided by the OSG.

## 2.8 MINOS

Over the last three years, computing for MINOS data analysis has greatly expanded to use more of the OSG resources available at Fermilab. The scale of computing has increased from about 50 traditional batch slots to typical user jobs running on over 2,000 cores, with an expectation to expand to about 5,000 cores (over the past 12 months we have used 3.1M hours on OSG from 1.16M submitted jobs). This computing resource, combined with 120 TB of dedicated BlueArc (NFS mounted) file storage, has allowed MINOS to move ahead with traditional and advanced analysis techniques, such as Neural Network, Nearest Neighbor, and Event Library methods. These computing resources are critical as the experiment has moved beyond the early, somewhat simpler Charged Current physics, to more challenging Neutral Current,  $\nu+e$ , anti-neutrino and other analyses which push the limits of the detector. We use a few hundred cores of offsite computing at collaborating universities for occasional Monte Carlo generation. MINOS was also successful at using TeraGrid resources at TACC in Fall 2009 for a complete pass over our data.

MINOS recently made a disappearance anti- $\nu_\mu$  measurement (shown at Neutrino 2010) comparing the energy spectra of antineutrino interactions in a near and far target that fits well to a mass difference model.






**Figure 27: Confidence Interval contours in the fit of the MINOS Far Detector antineutrino data (red) to the hypothesis of two-flavor oscillations. The solid (dashed) curves give the 90% (68%) contours.**

## 2.9 Astrophysics

The Dark Energy Survey (DES) used approximately 40,000 hours of OSG resources during the period January 2010 – December 2010 to generate simulated images of galaxies and stars on the sky as would be observed by the survey. The bulk of the simulation activity took place during a production run which generated a total of 3.5 Terabytes of simulated imaging data for use in testing the DES data management data processing pipelines as part of DES Data Challenge 5 (DC5). The DC5 simulations consist of 2600 mock science images, covering some 200 square degrees of the sky, along with nearly another 1000 calibration images needed for data processing. Each 1-GB-sized DES image is produced by a single job on OSG and simulates the 300,000 galaxies and stars on the sky covered in a single 3-square-degree pointing of the DES camera. The processed simulated data are also being actively used by the DES science working groups for development and testing of their science analysis codes. In addition to the main DC5 simulations, we also used OSG resources to produce about an additional 1 TB of simulated images, consisting of science and calibration images for the DES weak lensing and supernova science working groups, and for 5 smaller simulation data sets generated to enable quick turnaround and debugging of the DES data processing pipelines. Figure 1 shows an example color composite image of the sky derived from these DES simulations.



**Figure 28:** Example simulated color composite image of the sky, here covering just a very small area compared to the full  $5000 \text{ deg}^2$  of sky that will be observed by the Dark Energy Survey. Most of the objects seen in the image are simulated stars and galaxies. Note in particular the rich galaxy cluster at the upper right, consisting of the many orange-red objects, which are galaxies that are members of the cluster. The red, green, and blue streaks are cosmic rays, and have those colors as they each appear in only one of the separate red, green, and blue images used to make this color composite image.

## 2.10 Structural Biology

The SBGrid Consortium, operating from Harvard Medical School in Boston is supporting software needs of ~150 structural biology research laboratories, mostly in the US. The SBGrid Virtual Organization (SBGrid VO) extends the initiative to support most demanding structural biology applications on resources of the Open Science Grid. Support by the SBGrid VO is extended to all structural biology groups, regardless of their participation in the Consortium. Within last 12 months we have significantly increased our participation in the Open Science Grid, in terms of both utilization and engagement. Specifically:

- We have launched and successfully maintained a GlideinWMS grid gateway at Harvard Medical School. The gateway communicates with the Glidein Factory at UCSD, and dispatches computing jobs to several computing centers across the US. This new infrastructure allowed us to reliably handle the increased computing workload. Within the last 12 months our VO supported ~6 million CPU hours on the Open Science Grid, and we rank as number 10 Virtual Organization in terms of overall utilization.
- SBGrid completed development of the Wide Search Molecular Replacement workflow. The paper describing its scientific impact was recently published in PNAS. Another paper presenting the underlying computing technology was presented during the 3rd IEEE Workshop on Many-Task Computing on Grids and Supercomputers, co-located with ACM/IEEE SC10 International Conference for High Performance, Networking, Storage and Analysis.
- The WS-MR portal was made publicly available in November, 2010. Since its release we have supported 35 users. The majority of users were from US academic institutions (e.g.

Yale University, Harvard, WUSTL, University of Tennessee, University of Massachusetts, Stanford, Immune Disease Institute, Cornell University, Caltech), but international research groups utilized the portal as well (including research groups from Canada, Germany, Australia and Taiwan).

- We continue planning for integration of the central biomedical cluster at Harvard Medical School with Open Science Grid resources. The cluster has been recently funded for expansion from 1000 to 7000 cores (S10 NIH award), and the first phase of the upgrade is being completed in December.
- Our VO is organizing the Open Science Grid All-Hand Meeting which is scheduled to take place in Boston in March of 2011. We have prepared a preliminary program agenda, and participated in several planning discussions.
- We successfully maintained a specialized MPI cluster in Immune Disease Institute (Harvard Medical School affiliate) to support highly scalable molecular dynamics computations. A long-term molecular dynamics simulation was recently completed on this cluster, and will complement crystal structure that was recently determined in collaboration with Walker laboratory at HMS (Nature, in press). The resource is also available to other structural biology groups in Boston area.



**Figure 29: Crystal structure of the MHC Class I molecule presenting a peptide, and bound to the T-cell receptor. MHC peptide binding domain rendered in red has a unique fold. The other six domains in the complex are all members of the immunoglobulin superfamily (IgSF). The WS-MR approach can identify a very small subset (4/~4000) of previously determined Ig structures that could be used here to bootstrap structure determination process**

## 2.11 Multi-Disciplinary Sciences

The Engagement team has worked directly with researchers in the areas of: biochemistry (Xu), molecular replacement (PRAGMA), molecular simulation (Schultz), genetics (Wilhelmsen), information retrieval (Blake), economics, mathematical finance (Buttimer), computer science (Feng), industrial engineering (Kurz), and weather modeling (Etherton).

The computational biology team led by Jinbo Xu of the Toyota Technological Institute at Chicago uses the OSG for production simulations on an ongoing basis. Their protein prediction software, RAPTOR, is likely to be one of the top three such programs worldwide.

A chemist from the NYSGrid VO using several thousand CPU hours a day sustained as part of the modeling of virial coefficients of water. During the past six months a collaborative task force between the Structural Biology Grid (computation group at Harvard) and OSG has resulted in porting of their applications to run across multiple sites on the OSG. They are planning to publish science based on production runs over the past few months.

## 2.12 Computer Science Research

OSG continues to provide a laboratory for research activities to deploy and extend advanced distributed computing technologies in the following areas:

A continuing collaboration between the Condor project, US ATLAS, and US CMS is using the OSG to test the scalability and methods for “just-in-time” scheduling across the OSG sites using “glide-in” methods. Glideins introduce new challenges, like two-tiered matching, two-tiered authorization model, network connectivity, and scalability. The two-tiered matching is being addressed within the glideinWMS project sponsored by US CMS. The two-tiered authorization is addressed by the gLExec component developed in Europe by NIKHEF, and partially supported by Fermilab for OSG. The network and scalability issues are being addressed by Condor.

Cybersecurity is a growing concern, especially in computing grids, where attack propagation is possible because of prevalent collaborations among thousands of users and hundreds of institutions. The collaboration rules that typically govern large science experiments as well as social networks of scientists span across the institutional security boundaries. A common concern is that the increased openness may allow malicious attackers to spread more readily around the grid. Mine Altunay of OSG Security team collaborated with Sven Leyffer and Zhen Xie of Argonne National Laboratory and Jeffrey Linderth of University of Wisconsin-Madison to study this problem by combining techniques from computer security and optimization areas. The team framed their research question as how to optimally respond to attacks in open grid environments. To understand how attacks spread, they used OSG infrastructure as a testbed. They developed a novel collaboration model observed in the grid and a threat model that is built upon the collaboration model. This work is novel in that the threat model takes social collaborations into account while calculating the risk associated with a participant during the lifetime of the collaboration. The researchers again used OSG testbed for developing optimal response models (e.g. shutting down a site vs. blocking some users preemptively) for simulated attacks. The results of this work has been presented at SIAM Annual Conference 2010 at Denver, Colorado and also submitted to the Journal of Computer Networks.


In addition applied research has been done to deploy the Hadoop storage system and XROOTD data caching systems into the US CMS and US ATLAS distributed systems.

### 3. Development of the OSG Distributed Infrastructure

#### 3.1 Usage of the OSG Facility

The OSG facility provides the platform that enables production by the science stakeholders; this includes operational capabilities, security, software, integration, testing, packaging and documentation as well as VO and User support. Scientists who use the OSG demand stability more than anything else and we are continuing our operational focus on providing stable and dependable production level capabilities.


The stakeholders continue to increase their use of OSG. The two largest experiments, ATLAS and CMS, after performing a series of data processing exercises last year that thoroughly vetted the end-to-end architecture, were ready to meet the challenge of data taking that began in March 2010. The OSG infrastructure has demonstrated that is up to the challenge and continues to meet the needs of the stakeholders. Currently over 1 Petabyte of data is transferred nearly every day and more than 4 million jobs complete each week.



**Figure 30: OSG facility usage vs. time broken down by VO**

During the last year, the usage of OSG resources by VOs increased by ~45% from about 5.5M hours per week to about 8M hours per week; additional detail is provided in the attachment entitled “Production on the OSG.” OSG provides an infrastructure that supports a broad scope of scientific research activities, including the major physics collaborations, nanoscience, biological sciences, applied mathematics, engineering, and computer science. Most of the current usage continues to be in the area of physics but non-physics use of OSG is a growth area with current usage approximately 570K hours per week (averaged over the last year) spread over 17 VOs.





**Figure 31: OSG facility usage vs. time broken down by Site.**  
(Other represents the summation of all other “smaller” sites)

With over 85 sites, the production provided on OSG resources continues to grow; the usage varies depending on the needs of the stakeholders. During normal operations, OSG provides more than 1.1 M CPU wall clock hours a day with peaks occasionally exceeding 1.3 M CPU wall clock hours a day; between 300K and 400K opportunistic wall clock hours are available on a daily basis for resource sharing.

This year, OSG provided significant effort and technical planning devoted to enabling both new and upgraded CMS and ATLAS Tier-3 sites that have been funded and are continuing to come online. The new Tier-3 sites are notable since many of their administrators do not have formal computer science training and thus special frameworks were developed to provide effective and productive environments. To support these sites (in collaboration with ATLAS and CMS), OSG focused on creating both documentation as well as a support structure suitable for these sites and continues to address the ongoing needs of the ATLAS and CMS Tier-3 Liaisons. To date the effort has addressed:

- Onsite help and hands-on assistance to the ATLAS and CMS Tier-3 coordinators in setting up their Tier-3 test sites including several multi-day meetings to bring together the OSG experts needed to answer and document specific issues relevant to the Tier-3s. OSG hosts regular meetings with these coordinators as well to discuss issues and plan steps forward.
- OSG packaging and support for Tier-3 components such as Xrootd that are projected to be installed at over half of the Tier-3 sites (primarily ATLAS sites). This includes testing and working closely with the Xrootd development team via bi-weekly meetings.

- OSG support for the Canadian and WLCG clients that have been selected as the mechanism for deploying ATLAS software at T3 sites. This involved adding features to the VDT to meet the ATLAS requirement of strict versioning, as well as features to the WLCG Client tool to support specific directory and log file changes to support ATLAS.
- Many OSG workshops have been updated to draw in the smaller sites by incorporating tutorials and detailed instruction. A site admins workshop was held in August 2010. One new feature was a tutorial on the Hadoop file system.
- OSG documentation for Tier-3s has been extended to support T3s beginning with installation on the bare hardware. Sections for site planning, file system setup, basic networking instructions, and cluster setup and configuration were updated and maintained together with more detailed explanations of each step (<https://twiki.grid.iu.edu/bin/view/Tier3/WebHome>). This documentation is used directly by CMS, and serves as the reference documentation that was used by ATLAS to develop more specific documentation for their T3s.
- OSG worked directly with new CMS and ATLAS site administrators as they started to deploy their sites. We made arrangements to work directly with local site administrators to work through security issues and barriers that many T3 sites were beginning to encounter as they attempted to setup T3 sites for the first time.
- The OSG Security team has set up a PAKITI server that is centrally monitoring and enabling all the CMS T3 sites to identify and fix security loopholes.
- Regular site meetings were geared toward Tier-3s in conjunction with the ongoing site coordination effort including office hours held three times every week to discuss issues that arise involving all aspects of the sites.

In summary OSG has demonstrated that it is meeting the needs of US CMS and US ATLAS stakeholders at all Tier-1's, Tier-2's, and Tier-3's. OSG is successfully managing the uptick in job submissions and data movement when LHC data taking resumed in 2010. And OSG continues to actively support and meet the needs of a growing community of non-LHC science communities that are increasing their use and reliance on OSG.

### **3.2 Middleware/Software**

In 2010, our efforts to provide a stable and reliable production platform have continued and we have focused on support and incremental, production-quality upgrades. In particular, we have focused on support for the relatively new Tier-3 sites, native packaging, and storage systems.

As in all software distributions, significant effort must be given to ongoing support. We have focused on continual, incremental support of our existing software stack release, OSG 1.2. Between January 2010 and mid-December 2010, we released 12 minor updates, or approximately one per month. These included regular software updates, security patches, bug fixes, and new software. We will not review all of the details of these releases here, but we instead wish to emphasize that we have invested significant effort in keeping all of the software up to date so that OSG's stakeholders can focus less on the software and more on their science. This general maintenance consumes roughly 50% of the effort of the OSG software effort.

There have been several software updates and events in the last year that are worthy of deeper discussion. As background, the OSG software stack is based on the VDT grid software distribu-

tion. The VDT is grid-agnostic and used by several grid projects including OSG, WLCG, and BestGrid. The OSG software stack is the VDT with the addition of OSG-specific configuration.

- Since summer 2009 we have been focusing on the needs of the ATLAS and CMS Tier-3 sites. In particular, we have focused on Tier-3 support for new storage solutions. In the last year, we have improved our packaging, testing, and releasing of BeStMan, Xrootd, and Hadoop, which are a large part of our set of storage solutions. We have released several iterations of these.
- We have emphasized improving our storage solutions in OSG. This is partly for the Tier-3 effort mentioned in the previous item, but is also for broader use in OSG. For example, we have created new testbeds for Xrootd and Hadoop and expanded our test suite to ensure that the storage software we support and release are well tested and understood internally. We have monthly joint meetings with the Xrootd developers and ATLAS to make sure that we understand how development is proceeding and what changes are needed. We have also provided new tools to help users query our information system for discovering deployed storage systems, a task which has traditionally been hard in OSG. We expect these tools to be particularly useful to LIGO and SCEC, though other VOs will likely benefit as well. We have recently tested the new Bestman2 which has improved scalability and will be moving it to production in the near future.
- We have continued our intense efforts to provide the OSG software stack as so-called “native packages” (e.g. RPM on Red Hat Enterprise Linux). With the release of OSG 1.2, we have pushed the packaging abilities of our infrastructure (based on Pacman) as far as we can. While our established users are willing to use Pacman, there has been a steady pressure to package software in a way that is more similar to how they get software from their OS vendors. With the emergence of Tier-3s, this effort has become more important because system administrators at Tier-3s are often less experienced and have less time to devote to managing their OSG sites. We have wanted to support native packages for some time, but have not had the effort to do so, due to other priorities; but it has become clear that we must do this now. We initially focused on the needs of the LIGO experiment and in April 2010 we shipped to them a complete set of native packages for both CentOS 5 and Debian 5 (which have different packaging systems), and they are now in production. The LIGO packages are a small subset of the entire OSG software stack, and we are now phasing in complete support for native packages across the OSG software stack. We are currently focusing our efforts on providing Xrootd as native packages for ATLAS Tier 3 sites, and we expect they will be ready for production use in early 2011.
- We are currently preparing a major software addition: ATLAS has requested the addition of the gLite CREAM software, which has similar functionality to the Globus GRAM software; that is, it handles jobs submitted to a site. Earlier this year it was evaluated by OSG and it is likely to scale quite well. We are far along with our work to integrate it into the OSG Software Stack, and expect it to go into testing early in 2011.
- We have added a new software component to the OSG software stack, the gLite FTS (File Transfer Service) client, which is needed by both CMS and ATLAS.



- We made improvements to our grid monitoring software – RSV – to make it significantly easier for OSG sites to deploy and configure it. This is currently in testing, with release expected in January 2011.
- We worked with the Network for Earthquake Engineering Simulation (NEES) project to provide them with the capability of archiving their data at the reliable tape facility at Fermilab.
- We have expanded our ability to do accounting across OSG by implementing mechanisms that perform accounting of file transfer statistics and storage space utilization.
- We have worked hard on outreach through several venues:
  - a. We also conducted an in-person Storage Forum in September 2010 at the University of Chicago, to help us better understand the needs of our users and to directly connect them with storage experts.
  - b. We participated in three schools/workshops (OSG Summer School, OSG Site Administrator Forum, and Brazil/OSG Grid School) to assist with training with grid technologies, particularly storage technologies.
  - c. We participated in the OSG documentation effort to significantly improve the OSG technical documentation.

The VDT continues to be used by external collaborators. EGEE/WLCG uses portions of VDT (particularly Condor, Globus, UberFTP, and MyProxy). The VDT team maintains close contact with EGEE/WLCG via the OSG Software Coordinator's engagement with the gLite Engineering Management Team. EGEE is now transitioning to EGI, and we are closely monitoring this change. TeraGrid and OSG continue to maintain a base level of interoperability by sharing a code base for Globus, which is a release of Globus, patched for OSG and TeraGrid's needs. The VDT software and storage coordinators are members of the WLCG Technical Forum, which is addressing ongoing problems, needs and evolution of the WLCG infrastructure in the face of data taking.

### **3.3 Operations**

The OSG Operations team provides the central point for operational support for the Open Science Grid and provides the coordination for various distributed OSG services. OSG Operations publishes real time monitoring information about OSG resources, supports users, developers and system administrators, maintains critical grid infrastructure services, provides incident response, and acts as a communication hub. The primary goals of the OSG Operations group are: supporting and strengthening the autonomous OSG resources, building operational relationships with peering grids, providing reliable grid infrastructure services, ensuring timely action and tracking of operational issues, and assuring quick response to security incidents. In the last year, OSG Operations continued to provide the OSG with a reliable facility infrastructure while at the same time improving services to offer more robust tools to the OSG stakeholders.

OSG Operations is actively supporting the LHC and we continue to refine and improve our capabilities for these stakeholders. We have supported the additional load of the LHC start-up by increasing the number of support staff and implementing an ITIL based (Information Technology Infrastructure Library) change management procedure. As OSG Operations supports the LHC data-taking phase, we have set high expectations for service reliability and stability of existing and new services.

During the last year, the OSG Operations continued to provide and improve tools and services for the OSG:

1. Ticket Exchange mechanisms were updated with the WLCG GGUS system, ATLAS RT system, and the VDT RT system to use a more reliable web services interface. The previous email based system was unreliable and often required manual intervention to ensure correct communication was achieved. Using the new mechanisms, tickets opened by the WLCG are in the hands of the responsible ATLAS representative within 5 minutes of being reported.
2. The OSG Operations Support Desk regularly responds to ~160 OSG user tickets per month.
3. A change management plan was developed, reviewed, and adopted to insure service stability during WLCG data taking.
4. The BDII (Berkeley Database Information Index), which is critical to CMS and ATLAS production, is now functioning with an approved Service Level Agreement (SLA) which was reviewed and approved by the affected VOs and the OSG Executive Board. The BDII performance has been at 99.91% availability and 99.94% reliability during the 8 preceding months.
5. A working group was created to deploy a WLCG Top Level BDII for use by USCMS and USATLAS. Recommendations have been submitted to the OSG Blueprint Group.
6. The MyOSG system was ported to MyEGEE and MyWLCG. MyOSG allows administrative, monitoring, information, validation and accounting services to be displayed within a single user defined interface.
7. Using Apache Active Messaging Queue (Active MQ) we have provided WLCG with availability and reliability metrics.
8. The public ticket interface to OSG issues was continually updated to add requested features aimed at meeting the needs of the OSG users.
9. We have completed the SLAs for all Operational services.

And we continued our efforts to improve service availability via the completion of several hardware and service upgrades:

- A third instance of BDII was prepared and tested and upgraded to the new software version.
- The OSG collaborative documentation environment was moved to a dedicated host to improve performance.
- Monitoring of OSG Resources at the CERN BDII was implemented to allow end-to-end information system data flow to be tracked and alarmed on when necessary.
- A migration to a virtual machine environment for many services is now complete to allow flexibility in providing high availability services.

**Table 6: OSG Service Availability and Reliability (8 Months Beginning May 1, 2010)**

Service	Availability	Reliability
OSG BDII	99.91%	99.94%
MyOSG	99.91%	99.91%
GOC Ticket	99.92%	99.93%
OIM	99.87%	99.88%
RSV Collector	99.8%	99.8%
Software Cache	99.55%	99.55%
OSG TWiki	99.89%	99.92%

Service reliability for OSG services remains excellent and we now gather metrics that can quantify the reliability of these services with respect to the requirements provided in the Service Level Agreements (SLAs). SLAs have been finalized for all OSG hosted services. Regular release schedules for all OSG services have been implemented to enhance user testing and regularity of software release cycles for OSG Operations provided services. It is the goal of OSG Operations to provide excellent support and stable distributed core services the OSG community can continue to rely upon and to decrease the possibility of unexpected events interfering with user workflow.

### **3.4 Integration and Site Coordination**

The OSG Integration and Sites Coordination activity continues to play a central role in helping improve the quality of grid software releases prior to deployment on the OSG and in helping sites deploy and operate OSG services thereby achieving greater success in production. For this purpose we continued to operate the Validation (VTB) and Integration Test Beds (ITB) in support of updates to the OSG software stack that include compute and storage element services. In addition to these activities, there were three key areas of focus involving sites and integration undertaken in the past year: 1) provisioning infrastructure and training materials in support of a second OSG sponsored workshop in Colombia, as part of the launch of the Grid Colombia National Grid Infrastructure (NGI) program, discussed in Training below; 2) continued development and an automated workflow system for validating compute sites in the ITB; and 3) directed support for OSG sites, in particular activities targeted for the ramp-up of US-LHC Tier-3 centers with particular focus on XRootd storage systems.

The “ITB Robot” is an automated testing and validation system for the ITB which has a suite of test jobs that are executed through the pilot-based Panda workflow system. The test jobs can be of any type and flavor. The current set includes simple ‘hello world’ jobs, jobs that are CPU-intensive, and jobs that exercise access to/from the associated storage element of the site. Importantly, ITB site administrators are provided a command line tool they can use to inject jobs aimed for their site into the system, and then monitor the results using the full monitoring

framework (pilot and Condor-G logs, job metadata, etc) for debugging and validation at the job-level. In addition, a web site to provide reports and graphs detailing testing activities as well as results on multiple sites for various time periods was developed. This web site will also allow ITB site administrators to schedule tests to run on their site and allow them to view the daily testing that will run on their ITB resources. In the future, we envision that additional workloads will be executed by the system, simulating components of VO workloads.

As new Tier-3 facilities came online we found new challenges in supporting systems administrators. Often Tier-3 administrators were not UNIX computing professionals but post-docs or students working part time on their facility. To better support these sites we installed a virtualized Tier-3 cluster using the same services and installation techniques that are being developed by the ATLAS and CMS communities. An example is creating user-friendly instructions for deploying an XRootd distributed storage system.

Finally in terms of site and end-user support, we continue to interact with the community of OSG sites using the persistent chat room (“Campfire”) that has now been in regular operation for nearly 18 months. We offer three hour sessions at least two days a week where OSG core Integration or Sites support staff are available to discuss issues, troubleshoot problems or simply “chat” regarding OSG specific issues; these sessions are archived and searchable.

### **3.5 VO and User Support**

The VO and User Support team coordinates and supports the “at-large” Science VOs in OSG which covers all user communities except for ATLAS, CMS, and LIGO which are directly supported by the OSG Executive Team. At various times through the year, we provided assistance to science communities in leveraging and improving their use of OSG. We conducted a range of activities to jump-start and help VOs that are new to OSG or changing/increasing their use of OSG.

- The D0 collaboration continues to leverage OSG at ever increasing scales for Monte Carlo simulation; the peak weekly OSG production was about 16 million events with over 507 million events for the year.
- The IceCube Neutrino Observatory started grid operations using resources at GLOW and 5 remote sites.
- GLUE-X was started up as a VO and is sharing usage of its resources by 5 other VOs in production. Recently it has adopted the pilot jobs framework for submitting its analysis jobs within OSG.
- GridUNESP in Brazil achieved full functionality with active support from the DOSAR community and end-user MPI applications were submitted through its full regional grid infrastructure using OSG.
- Molecular dynamics simulations of mutant proteins run by CHARMM group at NHLBI/NIH and JHU were re-established on OSG using PanDA.
- SBGrid transitioned to a pilot job based system and was able to achieve their goal of over 4000 simultaneous jobs over 48 hours; this has enabled SBGrid to broaden the access to OSG by other life science and structural biology researchers in the Boston area.

- The Large Synoptic Survey Telescope (LSST) collaboration has undergone two cycles of community engagement in OSG, targeting image simulation as a pilot application. The pilot program has used 40,000 CPU hours per day for a month, with peaks of 100,000. The activity included a result validation process and has demonstrated the viability of the OSG platform to the LSST community.
- The ALICE-USA community has been ramping up their use of OSG in the last year. Two sites have been added to OSG and use the standard OSG methods and systems for usage reporting to WLCG.
- The GEANT4 Collaboration's EGEE-based biannual validation testing production runs were efficiency-improved and expanded onto the OSG, assisting in its toolkit's quality releases for MINOS, ATLAS, CMS, and LHCb.
- The Holland Computing Center (HCC) provides resources and expert user support for selected computational challenges. The communities involved are granted the usage of a well supported Glidein job submission system that allows access to an OSG-wide overlay batch system infrastructure. The system is fully integrated with a POSIX-enabled highly-available Hadoop-based storage element. (See also Section 4.8.)

Another key contribution is the weekly VO forum teleconferences that promote regular interaction between representatives of VOs and staff members of OSG. These meetings focus on

- in-depth coverage of VO and OSG at-large issues
- community building through shared experiences and learning
- expeditious resolution of operational issues
- knowledge building through tutorials and presentations on new technology

As an outcome of the VO forum, VOs identify areas of common concern that then serve as inputs to the OSG work program planning process; some recent examples include: dynamic mechanisms to manage and benefit from public storage availability; accounting discrepancies; exit code mismatches in Pilot-based environments; need for real-time job status monitoring; and the need for more accurate site-level advertisement of heterogeneous sub-cluster parameters.

We continued our efforts to strengthen the effective use of OSG by VOs. D0 increased to 75-85% efficiency at 80K-120K hours/day and this contributed to new levels of D0 Monte Carlo production. Work continues to support and enable CDF in broader use of remote sites in the USA and South Korea. The Fermilab VO with its wide array of more than 12 science communities continued efficient operations. In addition, we continue to actively support additional new communities in joining and leveraging the OSG.

### **3.6 Security**

The Security team continued its multi-faceted approach to successfully meeting the primary goal of maintaining operational security, developing security policies, acquiring or developing necessary security tools and software, and disseminating security knowledge and awareness.

During the past year, we increased our efforts on assessing the identity management infrastructure and the future research and technology directions in this area. Towards the end of 2009, in collaboration with ESN Net Authentication and Trust Fabric Team, we had planned a series of se-

curity workshops. Two of the workshops, Living in an Evolving Identity World Workshop that brought technical experts together and OSG Identity Management Requirements Gathering Workshop that brought the VO security contacts together, were held at the end of the 2009. We continued our work in 2010 and organized a Security and Virtual Organizations Workshop, held during OSG All Hands Meeting. This was a follow-up to the issues identified at the two previous workshops and brought technical experts and VOs together to discuss the current state of the security infrastructure and necessary improvements. We conducted a detailed survey with our VO security contacts to pinpoint the problem areas. The workshop reports and the survey results are available at <https://twiki.grid.iu.edu/bin/view/Security/OsgEsnetWorkshopReport>.

Usability of the identity management system has surfaced as a key element. Obtaining, storing, and managing certificates by the end user have significant usability challenges and thus easy-to-use tools for the end user are a critical need. Most end user software products lack the native support for PKI mainly because the majority of vendors do not favor PKI as their security mechanism. Moreover, the widespread adoption of vendor-products within the science community makes it inevitable that OSG adopts and integrates diverse security technologies with its infrastructure.

Our community feedback indicated that solving these problems is a priority. We have started identifying both short-term and long-term solutions. The short-term plans included quick fixes on the most urgent problems while we are working towards the long-term solutions that can re-structure our security infrastructure on the basis of usability and diverse security technologies.

Among our short-term accomplishments, we have designed a new web site for end users to manage the certificate lifecycle and integrated the web site with DOEGrids CA. The page is customized for OSG users and simplifies the request/renew/revoke operations. From the design phase, we worked with our VOs and end users since usability of the web site was our number one concern. We had the site continuously tested by 20 to 30 users and re-designed the page according to their feedback. After the release, we received very positive feedback from the VOs and end users. Our future work is to promote adoption of this web site by our user community.

Another short-term fix was to improve the certificate issuance process. We developed and deployed a monitoring tool into the DOEGrids CA infrastructure that showed us the time spent in various steps of the process. Based on the collected data, we worked with our VOs and certificate-issuing agents, identified the bottlenecks, and modified the process to remove these bottlenecks.

Our final accomplishment in short-term goals was working with SBGrid in re-designing a user registration workflow that would meet SBGrid's needs. The workflow includes all of the steps from the moment a user decides to join SBGrid till the user has access privileges to run jobs in the Grid. Originally, this took 8 steps for the users. We reduced the process to a 3 to 4-step process for the end user (depending on VO tools). We identified the tools that would streamline the registration process and also help the end user with this experience. Our work with SBGrid is applicable to other small non-LHC VOs. We are in the process of applying our outcomes from this work to other VOs in OSG.

Of our long-term goals, we started designing a security architecture that would not only meet the current needs, but also will have flexibility and openness to meet the future needs. Providing diverse identity services, besides certificates, was one of the first steps. We started working with the CILogon project team at NCSA, which provides a Shibboleth-enabled certificate authority.

We strongly supported CILogon CA to be accredited by IGTF TAGPMA. With CILogon CA, we identified pilot universities that can utilize CILogon CA technology and worked with the university contacts to set up a test bed in OSG. We also teamed with a small business, Galois, who won a DOE-grant to deliver identity management solutions to OSG. OSG provided the necessary guidance for Galois to define the problem space.

We discovered a high number of kernel vulnerabilities during the last 12 months in operational security. Although these vulnerabilities are not specific to our infrastructure or to grid software, their severity led us to inform our member communities of the risks and the mitigation methods. We contacted each of the Tier-3 sites and helped their personnel patch and test their systems as there was a clear need for this type of support for smaller sites. We had one attack reported in our community, but the attacker did not target the grid infrastructure and only compromised the local systems.

During the vulnerability monitoring, we realized individualized care for the multitude of OSG sites was not sustainable. We searched for automated vulnerability monitoring products and found one from EGI operational security team, Pakiti. We obtained the software in March 2010, installed an instance and demonstrated it to our community in April. Based on the positive feedback and further feature requests from our community, we decided to continue our work, and relayed these requests to EGI team. The EGI team was receptive to our input, and delivered the requested features in late May. Since June 2010, we have operated this service in OSG.

We started a security blog that we use as a security bulletin board to improve our communication with our community. Our long-term goal is to build a security community, where each site administrator is an active contributor (<https://twiki.grid.iu.edu/bin/view/Security/FY10SecuritySupport>).

We were also faced with a significant change in the third-party security libraries (openssl v1.0) that we use within OSG software stack. These changes caused us to modify the way we distribute certificates within the VDT stack. We built a security testbed to test the modified certificate distribution service and the effects on the OSG software. The testbed results guided our modifications to the OSG code to make it compatible with the libraries and was fundamental for making our infrastructure work.

### **3.7 Content Management**

A project was begun in late 2009 to improve the collaborative documentation process and the organization and content of OSG documentation. The project was undertaken after an initial study using interviews of users and providers of OSG documentation and a subsequent analysis of the OSG documentation system. Implementation was begun in 2010, an overall leader was assigned and the documentation was divided into nine areas, each with a responsible owner.

A workshop was held in February 2010 to report on the results of a triage of existing documents and to train the area owners on an improved process for producing and reviewing documents, including new document standards (with tools to simplify application of the standards) and templates defined to make the documents more consistent for readers. The new content management process includes (1) ownership of each document, (2) a formal review of new and modified documents, and (3) testing of procedural documents. A different person fills each role, and all of them collaborate to produce the final document for release. Several new tools were created to facilitate the management and writing of documents in this collaborative and geographically dis-

persed environment. For example, individual contributors can see the real time status of all documents for which they have a role and the project manager can view the status of each document area or look at documents by the role of the targeted reader.

The team had several areas of focus: user and storage documentation and documentation specific to bringing up new Tier-3 sites. By the end of 2010, the team, with collaborators in many of the OSG VOs, had:

- Reduced the number of documents by 1/3 by combining documents with similar content and making use of an “include mechanism” for common sections in multiple documents.
- Produced a new navigation system that is now in the early testing stage.
- Provided reviewed documentation specific to establishing Tier-3 sites.
- Revamped documentation of the storage technologies in OSG.
- Improved and released 60% of the documents overall.

Two document areas were delayed because of staffing issues that were later addressed, and we are still working on the largest area, where work has been slow because of the volume of documents it contains. Overall, however, progress has been significant.

During the first quarter of 2011, the review and release process will continue to improve the remaining 40% of the documents. At about the same time, the prototype navigation will have been implemented and integrated into the OSG TWiki, which will allow users to navigate and search the content more easily. Once these milestones have been achieved the document process will change its current focus from improving individual documents to improving the navigation for readers in different roles such as users, system administrators, and VO managers. This process will test the connectivity and improve the usability of a set of documents when carrying out typical OSG user tasks. In the second half of 2011 the document process will be integrated into the OSG release process to ensure that software and services provided by the OSG are in sync with the corresponding documentation.

### **3.8 Metrics and Measurements**

OSG Metrics and Measurements strive to give OSG management, VOs, and external entities quantitative details of the OSG Consortium’s growth throughout its lifetime. The focus in the last year was maintenance and stability of existing metrics projects, as well as the new “OSG Display” built for the DOE.


The OSG Display (<http://display.grid.iu.edu>) is a high-level, focused view of several important metrics demonstrating the highlights of the consortium. It is meant to be a communication tool that can provide scientifically savvy members of the public a feel for the services that OSG provides.

The OSG Metrics area converted all of its internal databases and displays to the Gratia accounting system. This removes a considerable amount of legacy databases and code from the OSG’s “ownership,” and consolidates metrics on one platform. Continued report tasks include nightly RSV reports, metric thumbnails, monthly eJOT reports, and the CPU normalization performance table. This year, we have produced an updated “Science Field” report classifying OSG usage by science field. OSG Metrics has all of its applications running as pre-production services at the



GOC, and will likely turn them into production January 2011. We believe that our goal of making our work “maintainable” was largely accomplished in 2010.

The new “science field” report was created in response a request from the OSG consortium members, specifically the owners of large sites. It categorizes the large majority of OSG CPU usage by the science domain (physics, biology, computer science, etc) of the application. While this is simple for VOs within a single domain (LHC VOs), it is difficult for community VOs containing many diverse users, such as HCC or Engage. The current solution is a semi-manual process, which we are working on automating. This analysis (Figure 33) shows a dramatic increase in the non-HEP usage over the past 23 months.



**Figure 32: Monthly non-HEP Wall hours for different science fields**

The OSG Metrics team continues to perform as the liaison to the Gratia accounting project, maintaining a close working relationship. We have worked to implement service state accounting, allowing the OSG to record the historical status of batch systems. This probe has provided a monitoring-like capability for the OSG to measure the recent history of sites’ batch systems.

The collaboration between OSG and Gratia has been important as the LHC turn-on greatly increased the number of transfer records collected. We had several unplanned outages in our accounting system, and OSG has worked with Gratia on reliability issues for the next release. Collaborating with the OSG Storage area, we have re-written and deployed the dCache transfer probe to reduce the number of records, alleviating some of the load. In the next year, we foresee continuing to apply the Gratia core technology to new use cases (HTPC, new batch systems, and revisiting network metrics).

We have started discussions with Gratia about better recording “internal” information in the VO pilot systems. Gratia records batch system data, but pilot frameworks hide many details from the batch system, reducing the level of detail we have access to.

For non-accounting data, OSG Metrics has delivered a new monitoring<sup>6</sup> probe to verify the consistency of the OSG Information Services. The new monitoring probe is a piece of the OSG site monitoring framework (RSV), but deployment has been delayed until the next major release of RSV. We hope to see this at many sites in early 2011, as coherent information services need improvement.

The Metrics area continues to be heavily involved with the coordination of WLCG-related reporting efforts. Items continued from last year include installed capacity reporting, upgrading of the reports to a new, WLCG-specific, benchmark, and transferring of accounting data from the OSG to the WLCG. Installed capacity reporting has now become “official” in the WLCG, and we have installed corresponding verification reports to check the OSG is fairly reflected in the global numbers.

In FY11, the OSG will again focus on maintaining the services and reports the OSG Executive Team depends on. Most upcoming changes in metrics and accounting, while important, appear to be minor. The new metrics task for 2011 will be to try and derive a few high-quality, high-level observations of the data we collect.

### **3.9 Extending Science Applications**

In addition to operating a facility, the OSG includes a program of work that extends the support of Science Applications in terms of both the complexity as well as the scale of the applications that can be effectively run on the infrastructure. We solicit input from the scientific user community both as it concerns operational experience with the deployed infrastructure, as well as extensions to the functionality of that infrastructure. We identify limitations, and address those with our stakeholders in the science community. In the last year, the high level focus has been threefold: (1) improve the scalability, reliability, and usability as well as our understanding thereof; (2) evaluate new technologies, such as GRAM5 and CREAM, for adoption by OSG; and (3) improve the usability of our Work Load Management systems to enable broader adoption by non-HEP user communities.

We have continued with our previously established processes that help us to understand and address the needs of our primary stakeholders: ATLAS, CMS, and LIGO. The OSG has designated certain members (“senior account managers”) of the executive team to handle the interface to each of these major stakeholders and meet, at least quarterly, with their senior management to go over their issues and needs. Additionally, we document the stakeholder desired work lists from OSG and cross map these requirements to the OSG WBS; these lists are updated quarterly and serve as a communication method for tracking and reporting on progress.

#### **3.9.1 Scalability, Reliability, and Usability**

As the scale of the hardware that is accessible via the OSG increases, we need to continuously assure that the performance of the middleware is adequate to meet the demands. There were three major goals in this area for the last year and they were met via a close collaboration between developers, user communities, and OSG.

- At the job submission client level, the CMS-stated goal of 40,000 jobs running simultaneously and 1M jobs run per day from a single client installation has been achieved, with a success rate in excess of 95%. The job submission client goals were met in collaboration with CMS, Condor, and DISUN, using glideinWMS. This was done in a controlled environment, using the “overlay grid” for large scale testing on top of the production infrastructure developed the year before. To achieve this goal, Condor was modified to drastically reduce the number of ports used for its operation as by making many operations asynchronous. The glideinWMS is also used for production activities in several scientific communities, the biggest being CMS, D0, CDF and HCC, where the job success rate has constantly been above the 95% mark.
- At the storage level, the present goal is to have 50Hz file handling rates with hundreds of clients accessing the same storage area at the same time, and delivering at least 1Gbit/s aggregate data throughput. The two versions of BeStMan SRM, based on two different java container technologies, Globus and Jetty, both running on top of the HadoopFS has shown to easily scale to about 100Hz. It can also handle in the order of 1M files at once, with directories containing up to 50K files. There was no major progress on performance of the dCache based SRM which never exceeded 10Hz in our tests. On the throughput front, we achieved a sustained throughput of 15 Gbit/s over wide area network using BeStMan and FTS.
- At the functionality level, this year’s goal was to evaluate and facilitate the adoption of new Gatekeeper technologies in order to replace the Globus preWS Gatekeeper (GRAM2) currently in use on OSG. This is particularly important due to the fact that Globus has deprecated the WS Gatekeeper (GRAM4) that was to be the successor of the preWS method which had been tested in OSG over the past years. The chosen candidates to replace GRAM2 have been Globus GRAM5, INFN CREAM and Nordugrid ARC. GRAM5 and CREAM have been tested at the functionality level and recommended for adoption in OSG. ARC testing did not produce satisfactory results, however, and was thus not pursued further. Significant effort has also been invested in testing and facilitating the improvement of the client side, namely Condor-G, which is needed to use these new gatekeepers, both in OSG and in the partner Grids.

In addition, we have continued to work on a number of lower priority objectives:

- A package containing a framework for using Grid resources for performing consistent scalability tests against centralized services, like CEs and SRMs. The intent of this package is to quickly “certify” the performance characteristics of new middleware, a new site, or deployment on new hardware, by using thousands of clients instead of one. Using Grid resources allows us to achieve this, but requires additional synchronization mechanisms to perform in a reliable and repeatable manner.
- A package to monitor a certain class of processes on the tested nodes. Existing tools typically only measure system wide parameters, while we often need the load due to a specific class of applications. This package offers exactly this functionality in an easy to install fashion.
- We have been involved in tuning the performance of BeStMan SRM by performing a configuration sweep and measuring the performance at each point.

- We have been evaluating the capabilities of a commercial tool, CycleServer, for submitting jobs and monitoring Condor pools. Given that Condor is at the base of much of OSG infrastructure, having a commercially supported product could greatly improve the usability of OSG. We established that the product provides interesting features, but is missing features in key areas. We provided feedback to the developers, and will continue the evaluation.
- We have been working with CMS to understand the I/O characteristics of CMS analysis jobs. We helped by providing advice and expertise. Changes have been made to all layers of the software stack to improve the management of data I/O and computation. This work has resulted in improved CPU efficiencies of CMS software on OSG sites.
- We continuously evaluate new versions of OSG software packages, in the aim of both discovering eventual bugs specific to the OSG use cases, as well as comparing the scalability and reliability characteristics against the previous release.
- We have been involved with other OSG area coordinators in reviewing and improving the user documentation. The resulting improvements are expected to increase the usability of OSG for both the users and the resource providers.

### 3.9.2 *Workload Management System*

The primary goal of the OSG Workload Management System (WMS) effort is to provide a flexible set of software tools and services for efficient and secure distribution of workload among OSG sites. In addition to two Condor-based suites of software previously utilized in OSG, Panda and glideinWMS, the OSG Match Maker (based directly on Condor) has reached significant usage level.

The Panda system continued to be supported by OSG as a crucial infrastructure element of the ATLAS experiment at LHC as we entered the critically important period of data taking and processing in both proton and (recently) heavy ion programs. With more experience in continuously operating Panda itself as well as the monitoring services attached to it, we gained better insight into the direction of the Panda monitoring upgrade, choice of technologies and integration options. We have created a prototype of an upgraded Panda monitor based on a modern technology platform (framework-based web service as a data source, and a rich AJAX-capable client). Migration to this application will allow us to simplify application code, make full use of data caching techniques which is critical for optimal performance, leverage open source for tools such as authentication and authorization mechanisms, and provide a richer and more dynamic user experience. In addition, we started work on providing a data feed from Panda to a major LHC monitoring system, the Global Job Dashboard.

This reporting period saw continued utilization of Panda by the CHARMM collaboration working in the field of structural biology. We have also collaborated with the BNL research group active in Daya Bay and LBNE (DUSEL) neutrino experiments and have started running their Monte Carlo simulation software on Panda, with plans for expansion to multiple sites.

The ITB (Integration Testbed) activity in OSG has benefited from using Panda, which allows site administrators to automate test job submission and monitoring and have test results documented and accessible through a Web portal (see Section 3.4).

With the glideinWMS system, we continued stable operation across global large-scale resources of the CMS experiment (with an instance hosted at FNAL), and deployed a newer version capa-

ble of serving multiple virtual organizations from a single instance (these include CMS, HCC, GlueX, SBBGrid and GLOW/IceCube) at UCSD. Other applications of this system include D0, CDF and MINOS. There have been important improvements of the glideinWMS security model, as well as added support for NorduGrid and CREAM. In addition, the glideinWMS has been expanded with added ability to run in virtualized environments, notably on Amazon and Magellan resources. Work continued on improvements in documentation, installation, scalability, diagnostics and monitoring areas.

We continued the maintenance of the gLExec (user ID management software), as a project responsibility.

One of the issues we had in the previous reporting period, the lack of awareness of potential entrants to OSG of capabilities and advantages of OSG Workload Management Systems, was addressed by creation of a document that contains comparative analysis of features and characteristics of the systems used, such as depth of monitoring provided and ease of installation and maintenance.

This program will continue to be important for the science community and OSG. First, Workload Management Systems supported by OSG continue to be a key enabling factor for large science projects such as ATLAS and CMS and have been proven as such in the challenging and important period of data taking and processing at the LHC. Second, OSG continues to draw new entrants who benefit greatly by leveraging stable and proven Workload Management Systems for access to opportunistic resources and also job submission and monitoring tools created by OSG.

### 3.10 Challenges facing OSG

The current challenge facing OSG is how to sustain, improve and extend our support for our existing communities, while engaging with and expanding our support for additional communities that can benefit from our services, software and expertise. We continue to face the challenge of providing the most effective value to the wider community, as well as contribute to the XD, SciDAC-3 and Exascale programs, in the areas of:

- Utilizing **shared storage** with other than the owner group - not only more difficult than (the quantized) CPU cycle sharing, but also less well supported by the available middleware.
- **Federation of the local and community identity/authorization attributes** within the OSG authorization infrastructure.
- Interfacing and transparency for **sharing of resources across and between campuses**.
- **Validation, analysis** and active response to of all types of measurements and metrics - availability and reliability testing and accounting and monitoring information, error conditions, cases of retry etc.
- The **scalability and robustness of the infrastructure** has not yet reached the scales needed by the LHC for analysis operations in the out years. The US LHC software and computing leaders have indicated that OSG needs to provide  $\sim 2\times$  in interface performance over the next year or two, and the robustness to upgrades and configuration changes throughout the infrastructure needs to be improved.

- **Full usage of available resources.** The GlideinWMS factory is now a production OSG service supported for about 4 communities. We need to work on more automated selection of resources to match the needs of the applications.
- **User and customer frameworks** are important for engaging non-Physics communities in active use of grid computing technologies; for example, the structural biology community has ramped up use of OSG enabled via portals and community outreach and support.
- A common operations infrastructure across heterogeneous communities can be brittle. Efforts to improve the **early detection of faults and problems** before they impact the users help everyone.
- **Making a useful mentorship program that best leverages the expertise of the OSG staff mentors and helps transition students from the classroom to be users of and/or contributors to distributed computing.**
- **Use of new virtualization, multi-core and job parallelism techniques, scientific and commercial cloud computing.** We are making some progress through two satellite projects funded in these areas. The first on High Throughput Parallel Computing (HTPC) on OSG resources for an emerging class of applications where large ensembles (hundreds to thousands) of modestly parallel (4- to ~64- way) jobs. The second a research project to do application testing over the ESNet 100-Gigabit network prototype, using the storage and compute endpoints supplied by the Magellan cloud computing at ANL and NERSC.

**These challenges are not unique to OSG.** Other communities are facing similar challenges in educating new entrants to advance their science through large-scale distributed computing resources.

## 4. Satellite Projects, Partners, and Collaborations

---

The OSG coordinates with and leverages the work of many other projects, institutions, and scientific teams that collaborate with OSG in different ways. This coordination varies from reliance on external project collaborations to develop software that will be included in the VDT and deployed on OSG to maintaining relationships with other projects where there is a mutual benefit because of common software, common user projects, or expertise in areas of high throughput or high performance computing.

In 2010, we looked at our many external relationships with the objective of classifying them in order to match the importance to our stakeholders with the level of our effort in maintaining and benefitting from the relationship. The external projects with which we are actively involved are Satellite Projects and others that we work with are Partners.

Projects are Satellite Projects if they meet the following pre-requisites:

1. OSG was involved in the planning process and there was communication and coordination between the proposal's PI and OSG Executive Team before submission to agencies.
2. OSG commits support for the proposal and/or future collaborative action within the OSG project.
3. The project agrees to be considered an OSG Satellite project.

Satellite Projects are independent projects with their own project management, reporting, and accountability to their program sponsors; the OSG core project does not provide oversight for the execution of the satellite project's work program. OSG does have a close working relationship with Satellite Projects and a member of our leadership is involved.

The OSG Satellite Projects in 2010 were:

- CI-TEAM: Cyberinfrastructure Campus Champions (CI-CC)
- ExTENCI: Extending Science Through Enhanced National Cyberinfrastructure is a new project that began in August, 2010. It jointly serves OSG and TeraGrid by providing mechanisms for running applications on both architectures.
- High Throughput Parallel Computing (HTPC) on OSG resources for an emerging class of applications where large ensembles (hundreds to thousands) of modestly parallel (4- to ~64-way) jobs.
- Application testing over the ESNet 100-Gigabit network prototype, using the storage and compute end-points supplied by the Magellan cloud computing at ANL and NERSC.
- CorralWMS to enable user access to provisioned resources and "just-in-time" available resources for a single workload integrate. It builds on previous work on OSG's GlideinWMS and Corral, a provisioning tool used to complement the Pegasus WMS used on TeraGrid.
- VOSS: "Delegating Organizational Work to Virtual Organization Technologies: Beyond the Communications Paradigm" (OCI funded, NSF 0838383) studies how OSG functions as a collaboration.

We maintain a partner relationship with many other organizations that are related grid infrastructures, other high performance computing infrastructures, international consortia, and certain

projects that operate in the broad space of high throughput or high performance computing. These collaborations include:

- Community Driven Improvement of Globus Software (CDIGS),
- Condor
- Cyber Infrastructure Logon service (CILogon),
- European Grid Initiative (EGI),
- European Middleware Initiative (EMI),
- Energy Sciences Network (ESNet),
- FutureGrid study of grids and clouds,
- Galios R&D company in the area of security for computer networks,
- Globus
- Colombian National Grid (GridColombia),
- São Paulo State University's statewide, multi-campus computational grid (GridUNESP),
- Internet2,
- Magellan,
- Network for Earthquake Engineering Simulation (NEES),
- National (UK) Grid Service (NGS), NYSGrid,
- Open Grid Forum (OGF),
- Pegasus workflow management,
- TeraGrid, and the
- WLCG.

The OSG is supported by many institutions and experiments including:

- Boston University
- Brookhaven National Laboratory
- California Institute of Technology
- Clemson University
- Columbia University
- Distributed Organization for Scientific and Academic Research (DOSAR)
- Fermi National Accelerator Laboratory
- Harvard University (Medical School)
- Indiana University




- Information Sciences Institute (USC)
- Lawrence Berkeley National Laboratory
- Purdue University
- Renaissance Computing Institute
- Stanford Linear Accelerator Center (SLAC)
- University of California San Diego
- University of Chicago
- University of Florida
- University of Illinois Urbana Champaign/NCSA
- University of Nebraska – Lincoln
- University of Wisconsin, Madison
- University of Buffalo (council)
- US ATLAS
- US CMS
- STAR
- LIGO
- CDF
- D0

Selected Satellite Projects and Partnerships and their work with OSG are described below.

#### **4.1 CI Team Engagements**

During calendar year 2010, we have seen consistent usage of the OSG by users from various science domains as a direct result of the Engagement program. The Engage team also provides job submission infrastructure and VO services for other projects including the ExTENCI effort, which is utilizing 2000-3000 opportunistically available CPUs over 20+ computer resources in the Engage VO for a post-processing workflow that computes synthetic seismograms and peak ground motions from all related rupture variations to produce hazard curves for the Southern California Earthquake Center (SCEC).



**Figure 33: Calendar year 2010 CPU hours per engaged user**

The Engage VO use of OSG depicted in Figure 34 represents a number of science domains and projects including: Biochemistry (Zhao, Z.Wang, Choi, Der), theoretical physics (Bass, Peterson, Bhattacharya, Coleman-Smith), Mechanical Engineering (Ratnaswamy), Earthquake Engineering (Espinosa), RCSB Protein Data Bank (Prlic), Wildlife Research (Kjaer) and Oceanography (Thayre). We note that all usage by Engage staff depicted here is directly related assisting users, and not related to any computational work of the Engage staff themselves. This typically involves running jobs on behalf of users for the first time or after significant changes to test wrapper scripts and probe how the distributed infrastructure will react to the particular user codes. In an effort to increase the available cycles for Engage VO users, RENCi has made available two clusters to the Engage community including opportunistic access to the 11TFlop BlueRidge system.

In February 2010, James Howison and Jim Herbsleb of Carnegie Mellon University conducted a survey of the OSG Engagement Program as part of their VOSS SciSoft research project funded by NSF grant number 0943168. The full 17 page report is available upon request, and indicates that the OSG Engagement Program is effective, helpful, and appreciated by the researchers relying on both the human relationship based assistance and the hosted infrastructure which enables their computational science.

In December of 2010, Mats Rynge of the Engagement team provided content development and instruction at the Sao Paulo Brazil: <http://indico.fnal.gov/conferenceDisplay.py?confId=3586>

The Engage VO team, infrastructure, and efforts are funded under NSF award number 0753335.

## 4.2 Condor


The OSG software platform includes Condor, a high throughput computing (HTC) system developed by the Condor Project. Condor can manage local clusters of computers, dedicated or

cycle scavenged from desktops or other resources, and can manage jobs running on both local clusters and delegated to remote sites via Condor itself, Globus, CREAM, and other systems.


The Condor Team provided support to the OSG and HTC community around Condor, and performed release engineering to deliver updated versions of the Condor software. Support activities frequently involved identifying problems or opportunities for improvements to the Condor software, which would be tested and released.

#### 4.2.1 Support Condor

Users received support directly from project developers by sending email questions or bug reports directly to the *condor-admin@cs.wisc.edu* email address. All incoming support email was tracked by an email-based ticket tracking system running on servers at UW-Madison. From January 2010 through mid-December 2010, over 1,300 email messages were exchanged between the project and users towards resolving 353 support incidents (see Figure 35). Support is also provided on various online forums, bug tracking systems, and mailing lists. For example, approximately 10% of the of the hundreds of messages monthly on Condor Users email list originated from Condor staff; see Figure 36.



**Figure 34: Number of tracked incidents and support emails exchanged per month**



**Figure 35: The condor-users email list receives hundreds of messages per month; project staff regularly contribute to the discussion.**

The Condor team provides on-going support through regular phone conferences and face-to-face meetings with OSG collaborations that use Condor in complex or mission-critical settings. This includes monthly meetings with USCMS, weekly teleconferences with ATLAS, and biweekly teleconferences with LIGO. Over the last year Condor's email support system has been used to manage 20 issues with LIGO, resolving 12 of them. The Condor team also uses a web page system to track ongoing issues. This web system is tracking 21 issues associated with ATLAS, of which 15 are resolved, 67 issues associated with CMS, of which 43 are resolved, 60 tickets for LIGO, of which 25 are resolved, and 24 for other OSG users, of which 15 are resolved.

Support work for LIGO has included improving scalability problems in Condor DAGMan, helping to debug problems in their cluster, testing DMTCP's suitability for use, and fixing numerous bugs uncovered by the work of LIGO and other OSG groups. Work with ATLAS revealed bugs and scalability bottlenecks in Condor-G, which have been fixed. For groups including ATLAS and LIGO the Condor team provides advice, configuration support, and debugging support to meet policy, security, and scalability needs.

The Condor team has crafted specific tests as part of our support for OSG associated groups. For example, LIGO's work pushes the limits of Condor DAGMan, so specific extremely large scale tests have been added and are regularly run to ensure that new releases will continue to meet LIGO's needs.


The Condor team has developed and performed multiple stress tests to ensure that Condor-G's ability to manage large GRAM and CREAM job submissions will scale to meet ATLAS needs. We did multiple tests where 20,000 tests were submitted at once, monitoring for correctness under load and scaling bottlenecks. Several bugs were discovered (not only in Condor, but also in CREAM) and scalability improved.

#### 4.2.2 Release Condor

This activity consisted of the ongoing work required to have regular new releases of Condor. Creating quality Condor releases at regular intervals required significant effort. New releases fixed known bugs; supported new operating system releases (porting); supported new versions of dependent system software and hardware; underwent a rigorous quality assurance and development lifecycle process (consisting of a strict source code management, release process, and regression testing); and received updates to the documentation.

From January 2010 through mid-December 2010, the Condor team made 8 releases of Condor, with at least one more planned before the end of December 2010. During this time the Condor team created and code-reviewed 72 publicly documented bug fixes. Condor ports are maintained and released for about a dozen operating systems and different Linux distribution combinations.

We continued to invest significant effort to improve our automated test suite in order to find bugs before our users do, and continued our efforts to maximize our leverage of the NMI Build and Test facility and the Metronome framework.<sup>3</sup> The number of automated builds we perform via NMI averages over 80 per day. This allows us to better meet our release schedules by alerting us to problems in the code or a port as early as possible. We currently perform approximately 30,000 tests per day on the current Condor source code snapshot (see Figure 37).



**Figure 36: Number of daily automated regression tests performed on the Condor source**


In the course of performing our Condor *release* activities, in a typical month we:

- Released a new version of Condor to the public
- Performed over 150 commits to the codebase
- Modified over 275 source code files
- Changed over 9,500 lines of code (Condor source code written at UW-Madison now sits about 812,000 lines of code)
- Compiled about 2,300 builds of the code for testing purposes
- Ran about 900,000 regression tests (both functional and unit)

The 7.4 stable series remains popular, with over 10,000 downloads since January 2010, bringing lifetime downloads to over 13,000. The new 7.5 development series was first released in January 2010 and has over 4,500 downloads. The stable series contains only bug fixes and ports, while

3 See <http://nmi.cs.wisc.edu/> for more about the NMI facility and Metronome.

the development series contains new functionality. Throughout 2003, the project averaged over 400 downloads of the software each month; in the past year, that number has grown to over 1,000 downloads each month as shown in Figure 38. Note this graph only depicts downloads from the Condor Project homepage; we do not have statistics on downloads made from distributions that include Condor, such as the VDT (Open Science Grid software stack), Red Hat MRG, or Fedora Linux. Although Linux and Windows are the dominant platforms downloaded by users, we invested effort in Unix platforms such as AIX and Solaris in order to meet the needs of certain collaborations or user groups.



**Figure 37: Stacked graph depicting number of downloads per month per Condor version**

Besides the Condor software itself (covered in the section below), the project maintained the Condor web site at URL <http://www.condorproject.org>, which contains links to the Condor software, documentation, research papers, technical documents, and presentations. The project maintained numerous public project related email lists (such as condor-users, condor-devel) and their associated archives on-line. The project also maintained a publicly accessible version control system allowing access to in-development releases, a public issue tracking system, and a public wiki of documentation on using, configuring, maintaining, and releasing Condor.

A primary publication for the Condor team is the Condor Manual which currently stands at over 1,000 pages. In the past year, the manual has been kept up to date with changes and additions in the stable 7.4 series as well as the development 7.5 series in preparation for the 7.6 release. The most recent stable release is the *Condor Version 7.4.4 Manual* which can be found at <http://www.cs.wisc.edu/condor/manual/>.

### 4.3 High Throughput Parallel Computing

With the advent of 8-, 16- and soon 32-cores packaged in commodity CPU systems, OSG stakeholders have shown an increased interest in computing that combines small scale parallel appli-

cations with large scale high throughput capabilities, i.e. ensembles of independent jobs, each using 8 to 64 tightly coupled processes. The OSG “HTPC” program is funded through a separate NSF grant to evolve the technologies, engage new users, and support the deployment and use of these applications.

HTPC is an emerging paradigm that combines the benefits of High Throughput Computing with small way parallel computing. One immediate benefit is that parallel HTPC jobs are far more portable than most parallel jobs since they do not depend on the nuances of parallel library software versions and machine specific hardware interconnects. For HTPC, parallel libraries are packaged and shipped simultaneously with job. This pattern allows for two additional benefits: First, there are no restrictions as to the method of parallelization, these can be MPI, OpenMP, Linda, or utilize other parallelization methods. Second, the libraries can be optimized for on-processor communication so that these jobs can run optimally on Multi-core hardware.

The work advanced significantly this year as the groundwork was laid for using the OSG Glide-in mechanism to submit jobs. The implication is that users will soon be able to submit and manage HTPC jobs as easily as they do ordinary HTC jobs via GlideinWMS.

The focus of the program has been to:

- Bring the MPI and other specific libraries from the client to the remote executive site as part of the job – thus removing the dependence on the different libraries invariably found on different sites.
- Adapt applications to only use the number of cores available on a single CPU.
- Extend the OSG information services to advertise support for HTPC jobs.
- Extend Glide-in technology so that users can use this powerful mechanism to submit and manage HTPC jobs.

To date applications from the fields of chemistry, weather, and computational engineering have been run across 6 sites – Oklahoma, Clemson, Purdue, Wisconsin, Nebraska and UCSD. We have logged nearly 10M hours since the first HTPC jobs ran in November of 2009. The work is being watched closely by the HTC communities who are interested in taking advantage of multi-core while not adding a dependency on MPI.

#### **4.4 Advanced Network Initiative (ANI) Testing**

The ANI project's objective is to prepare typical OSG data transfer applications for the emergence of 100Gbps WAN networking. This is accomplished in close collaboration with OSG, contributing to OSG in terms of testing OSG software, and benefitting from OSG. Thus we look to shrink the time between 100Gbps network links becoming available, and the OSG Science Stakeholders being able to benefit from their capacity. This project is focusing on the following areas:

- Creation of an easy to operate "load generator" that can generate traffic between Storage Systems on the Open Science Grid.
- Instrumenting the existing and emerging production software stack in OSG to allow benchmarking.
- Benchmark the production software stack, identify existing bottlenecks, and work with the external software developers on improving the performance of this stack.

The following is a summary of achievements in various aspects related to ANI project that we accomplished in 2010.

1. Architecture design of OSG/HEP Application with 100Gb/s connection

Through intensive discussion with other groups the ANI project and OSG colleagues, a high level design of OSG/HEP application architecture was documented that is consistent with both the envisioned ANI and the LHC data grid architectures. This was documented in *Specification of OSG Use Cases and Requirements for the 100 Gb/s Network Connection* ([OSG-doc-1008](#)).

2. Building hardware test platform

We built a test cluster at UCSD to conduct various tests which involve the core technologies for the OSG/HEP applications for the ANI project. The cluster has all the necessary components to function as a SE, including BeStMan, HDFS, GUMS, GridFTP, FUSE, and also used for other type of data transfer tool, e.g. Fast Data Transfer (FDT). This test platform has been used for transaction rate tests described in [OSG-doc-1004](#), as well as the 2009 and 2010 Supercomputing bandwidth challenges. We are presently upgrading this testbed to be capable of 40Gb/s bi-directional data transfers. Detailed bandwidth tests are scheduled for the December 2010 vacation period as this is a low activity period for CMS, and we thus can use the UCSD production cluster as source or sink against the test platform with the goal of reaching 40Gb/s transfers between OSG storage elements for the first time.

3. Validating Hadoop Distributed File System

We previously validated the Hadoop Distributed File System (HDFS) as a key technology of an OSG Storage Element (SE). We gave a presentation on this at the International Symposium on Grid Computing (ISGC) 2010 with the title: *Roadmap for Applying Hadoop distributed file system in Scientific Grid Computing*.

4. Test of Scalability of Storage Resource Manager (SRM) System

We conducted transaction rate scalability tests of two different versions of BeStMan SRM. These two versions are based on two different java container technologies, Globus and Jetty. Our tests contributed to the release of the new Jetty-based BeStMan-2 by improving the configuration and documenting the corresponding performance as compared with the previously available Globus based implementation on the same hardware. We worked closely with various parties: Hadoop distribution and packaging from OSG storage, BeStMan development team from LBNL, and used the new glidein based scalability tool from OSG Scalability. The results of BeStMan test were documented in *Measurement of BeStMan Scalability* ([OSG-doc-1004](#)). The use of glideinWMS and glideinTester is documented in *Use of Late-binding technology for Workload Management System in CMS* ([OSG-doc-937](#)).

5. Test of WAN data transfer tools and participation at Supercomputing 2010

Various WAN data transfers have been tested between UCSD and other CMS Tier-2 sites. We present our results in networking configuration, storage architecture and data transfer in 18<sup>th</sup> International Conference on Computing in High Energy Physics (CHEP) 2010 with title: *Study of WAN Data Transfer with Hadoop-based SE*.

6. Automated deployment on Magellan



The ANI architecture assumes that the endpoint for the 100Gbps testbed will be provided by the Magellan sites at Argonne and NERSC. We have thus worked with the OSG-Magellan coordination activity to understand the modus operandi for deploying an OSG-SE dynamically on Magellan. The status of dynamic deployment using glideinWMS on cloud resources, including Magellan was presented at CHEP 2010.

#### **4.5 ExTENCI**

ExTENCI (Extending Science Through Enhanced National CyberInfrastructure) is a joint project of OSG and TeraGrid, funded by the NSF OCI. The two-year project began in August 2010 with the goal to develop and provide production quality enhancements to the National CyberInfrastructure that will enable specific science applications to more easily use both OSG and TeraGrid or broaden access to a capability to both TeraGrid and OSG users. It is organized around four technologies and their respective science users.

1. Lustre-WAN Distributed File System to be used by CMS and ATLAS in providing data to Tier 3 sites: The first project deliverables have been to set up both test and production hardware and Kerberos security infrastructure at the University of Florida, PSC, and FNAL for use by CMS. Testing of the UF installation and tuning the Lustre filesystem at PSC are in progress.
2. Virtual Machines for STAR and CMS: Initial deliverables have been enabling the CMS VO to run within a VM on Purdue's TeraGrid resources and the STAR VO to run its VM on Clemson's OSG cloud. STAR used the Clemson cloud for a full month, averaging 800 VMs and producing 12 billion PYTHIA events, their largest run to date. Hypervisors have been deployed at both Purdue and Clemson to create large clouds. Work is proceeding on verifying interoperability of the VMs on both the OSG and TeraGrid clouds.
3. Workflow and Client Tools for SCEC and Protein Folding applications: SCEC processing has been split into parts suitable for TeraGrid and OSG, both of which have been run successfully manually, and work continues on automating the decisions for task mapping, job scheduling, and data mapping. The protein folding applications have been successfully run on TeraGrid and have run on as many as 14 OSG sites. Elements of the automated decision making process for distributing the SCEC application will also be used for protein folding applications.
4. Job Submission Paradigms to execute Cactus applications on both TeraGrid and OSG for Coastal Modeling and for science projects using Ensemble Kalman Filters. The Simple API for Grid Applications (SAGA) job management system is being extended to use Condor-g and Condor glide-in to enable Cactus applications to effectively execute on both TG and OSG. The project start was delayed so only architectural work has begun.

#### **4.6 Corral WMS**

Under the NSF OCI-0943725 award, the University of Southern California, Fermi National Accelerator Laboratory and University of California San Diego have been working on the CorralWMS integration project, which provides an interface to resource provisioning across national as well as campus cyberinfrastructures. Software initially developed under the name Corral now extends the capabilities of glideinWMS. The resulting product, glideinWMS, is one of the main workload management systems on OSG, and enables several major user communities to efficiently use the available OSG resources.

Corral, a tool developed to complement the Pegasus Workflow Management System, was recently built to meet the needs of workflow-based applications running on the TeraGrid. It is being used today by the Southern California Earthquake Center (SCEC) CyberShake application. In a period of 10 days in May 2009, SCEC used Corral to provision a total of 33,600 cores and used them to execute 50 workflows, each containing approximately 800,000 application tasks, which corresponded to 852,120 individual jobs executed on the TeraGrid Ranger system. The 50-fold reduction from the number of workflow tasks to the number of jobs is due to job-clustering features within Pegasus designed to improve overall performance for workflows with short duration tasks.

GlideinWMS was initially developed to meet the needs of the CMS (Compact Muon Solenoid) experiment at the Large Hadron Collider (LHC) at CERN. It generalizes a Condor glidein system developed for CDF (The Collider Detector at Fermilab) first deployed for production in 2003. It has been in production across the Worldwide LHC Computing Grid (WLCG), with major contributions from the Open Science Grid (OSG) in support of CMS for the past two years, and has recently been adopted for user analysis. GlideinWMS also is currently being used by the CDF, D0, and MINOS experiments, and servicing the NEBioGrid and Holland Computing Center communities. GlideinWMS has been used in production with more than 12,000 concurrently running jobs; the CMS use alone totals over 45 million hours.

The integrated CorralWMS system, which will retain the glideinWMS product name, includes a new version of Corral as a frontend. It provides a robust and scalable resource provisioning service that supports a broad set of domain application workflow and workload execution environments. The system enables workflows to run across local and distributed computing resources, the major national cyberinfrastructure providers (Open Science Grid and TeraGrid), as well as emerging commercial and community cloud environments. The Corral frontend handles the end-user interface, the user credentials, and determines when new resources need to be provisioned. Corral then communicates the requirements with the glideinWMS factory, and the factory performs the actual provisioning.

The CorralWMS system also contributes to the glideinWMS factory development. During the year, many new features have been implemented. It is now possible to pass project information between the frontend and factory and for the frontend to specify that large groups of glideins should be submitted as one grid job. These features are mainly for the system to work better with TeraGrid, but are required to make workloads run across the infrastructures. Monitoring and scalability have been improved.

Communities using the system on OSG include SCEC and LIGO. The SCEC workflows described above start off with a couple of large earthquake simulation MPI jobs, and are followed by a large set of serial jobs to determine how different sites in the Los Angeles region would be affected by the simulated earthquake. SCEC has been a long-time Corral user and requirements driver, and has been using Corral in production for runs on TeraGrid. As a demonstration on how CorralWMS can be used across cyberinfrastructures, a SCEC workflow was planned to execute MPI jobs on TeraGrid and serial workloads on OSG using the glideinWMS system. Four such runs were completed successfully.

CorralWMS project supports the OSG-LIGO Taskforce effort whose mission is to enable LIGO workflows to perform better in the OSG environment. One problem LIGO had was that short tasks in the workflows were competing with a large amount of long-running LIGO Einstein at

Home jobs. When submitted to the same OSG site, the workflows were essentially starved. With the glideinWMS glideins, the workflow jobs are now on a more equal footing as the glideins retain their resources longer, and during the lifetime can service many workflow jobs. The glideinWMS glideins also helped when running multiple workflows at the same time, as job priorities were used to overlap data staging job and compute jobs. This was not possible when relying on the remote Grid site's batch system to handle scheduling rather than the glideinWMS scheduler.

The CorralWMS project has recently contributed sessions to the OSG Grid Schools in Madison and Sao Paulo, gave a presentation at Condor Week 2010, and presented a poster at TeraGrid '10.

#### **4.7 OSG Summer School**

As requested by DOE and NSF, OSG staff provides pro-active support in workshops and collaborative efforts to help define, improve, and evolve the US national cyber-infrastructure. As part of this support, we applied for an OSG satellite proposal to run a four-day school in high-throughput computing (HTC) in July 2010 at the University of Wisconsin-Madison. This was a joint proposal with TeraGrid, and all seventeen students who participated in the OSG-portion of this school also participated in the TeraGrid 2010 conference.

An important focus of the school was hands-on experience with HTC; the students learned how to use HTC in a campus environment as well as large-scale computations with OSG. In order to provide the students with the most knowledgeable instructors, we relied on OSG staff experienced in the use of these technologies to instruct at the school. The students received a variety of experiences including running basic computations, large workflows, scientific applications, and pilot-based infrastructures as well as how to use storage technologies to cope with large quantities of data. The students had access to two OSG sites so they could experience scaling of computations. They used exactly the same tools and techniques that OSG users currently use, but we emphasized the underlying principles so they can apply what they learned more generally.

The students came from a wide variety of scientific disciplines including physics, biology, GIS, computer science and more. Ten of the instructors were OSG staff; we were able to have such a strong staff representation because we co-located the school with the OSG staff retreat. Four of the instructors were invited from other projects: TeraGrid, Globus, the Middleware Security and Testing research project, and the Condor Project.

A highlight of the school was the HTC Showcase in which four local scientists gave lectures about their experience with HTC. They showcased how the use of HTC expanded not only the amount of science they could do, but the kinds of science they could do.

After the school concluded, we assigned OSG staff as mentors to the students and they made regular contact with the students so they could provide help and deepen students' participation in the broader HTC community. The students were paired with an OSG staff member based on factors such as common interests, organizational memberships, and geography. Many of the students made use of this mentorship program and have begun to use HTC in their research.

During and after the school, we asked the students to provide evaluations of their experiences. An independent researcher examined these evaluations. While this examination is too extensive to include here, an extract of the summary said:

*Overall, the majority of the participants were very happy with how the conference went, and raved about the accommodations and organization of it. The majority were also very happy with the instructors and the presentations and activities within the sessions.*

We are currently applying for funding to run the school again in the summer of 2011. The OSG Summer School 2010 Web Page which includes curricula and materials:

<https://twiki.grid.iu.edu/bin/viewauth/Education/OSGSummerSchool2010>

An article about the school in International Science Grid This Week:

<http://www.isgtw.org/?pid=1002696>

#### **4.8 Science enabled by HCC**


The Holland Computing Center (HCC) of the University of Nebraska (NU) began to use the OSG to support its researchers during Year 4. HCC's mission is to meet the computational needs of NU scientists, and does so by offering a variety of resources. In the past year, the OSG has served as one of the resources available to our scientists, leveraging the expertise from the OSG personnel and running the CMS Tier2 center at the site.

Support for running jobs on the OSG is provided by our general-purpose applications specialists and we have only one graduate student dedicated to using the OSG.

To distribute jobs, we primarily utilize the Condor-based GlideinWMS system originally developed by CDF and CMS (now independently supported by the NSF). This provides a user experience identical to using the Condor batch system, greatly lowering the barrier of entry for users. HCC's GlideinWMS install has been used as a submission point for other VOs, particularly the LSST collaboration with OSG Engage. The GlideinWMS has led a CS graduate student to start his master thesis on grid technologies for local campus grids and campus bridging.

For data management, we leverage the available high-speed network and Hadoop-based storage at Nebraska. For most workflows, data is moved to and from Hadoop, sometimes resulting in multiple terabytes of data movement a day.

In the last year, 6 teams of scientists have run over 11 million hours on the OSG. This is about 10% of our active research teams at HCC, but over 20% of HCC computing (only counting computing done opportunistically at remote sites). Less precise figures are available for data movement, but it is estimated to be about 50TB total. See the below figure.



**Figure 38: Monthly wall hours per facility**

Applications we have run on the OSG in the past year are:

- TreeSearch, a platform for the brute-force discovery of graphs with particular structures. This was written by a Mathematics doctoral candidate and accumulated 4.6 million wall hours of runtime. Without OSG, HCC would not have been able to provide a sufficient number of hours to this student to finish this work.
- DaliLite, a biochemistry application. This was a one-off processing needed due to reviewers of a paper requesting more statistics. This was moved from a small lab cluster to the OSG in a matter of days, and accumulated over 120 thousand runtime hours. Without OSG, the scientists would not have been able to complete their paper in time.
- OMSSA, an application used by a researcher at the medical school. This was another example where a researcher who had never used HCC clusters discovered he needed a huge amount of CPU time in order to make a paper deadline in less than 2 weeks. As with Dali-Lite, the researcher would have not made the deadline if he ran only on HCC resources.
- Interview Truth Table Generator for Boolean models, developed by a mathematical biology research group at the University of Nebraska-Omaha, is a Java-based tool to generate Boolean truth tables from user data supplied via a web portal. Depending on input size, as a single process the tool required several days to complete one table. After a small modification to the source code, the Condor glidein mechanism was utilized to deploy jobs to both HCC and external resources, enabling significant reduction in runtime. For one particular case with a table of 67.1 million entries, runtime was reduced from approximately 4 days to about 1 hour.

- CPASS, another biochemistry application. This was our first web application converted to the grid; Condor allows a user to submit a burst of jobs which first fill the local clusters, then migrate out to the grid if excess capacity is needed.
- AutoDock and CS-ROSETTA, two smaller-scale biochemistry applications brought to HCC and converted to the grid. The work done with these was essential for a UNL graduate student to finish his degree work.
- CMS, We have several local CMS students and professors whose work is enabled by the using the grid; this science is covered elsewhere.

HCC has gone from almost zero local usage of the OSG last year to millions of CPU hours this year. This has been done without local OSG funding for user support (there is other OSG personnel and one student who do share expertise). The OSG is an important part of our “toolbox” of solutions for Nebraska scientists. The OSG is not a curiosity or a toy for HCC, but something we depend on not only to offload jobs, but to support science which could not have been completed at HCC resources alone.

#### **4.9 Virtualization and Clouds**

In the past year, various teams have worked with OSG on virtualization and in a broader sense “clouds.” This work can be summarized with the following four points:

- Several grass roots efforts are happening. This was demonstrated at the OSG all hands meeting at Fermi Lab where a special technical session was organized around this theme. All stakeholders in OSG have initiated exploratory work on the use of clouds and virtualization. Project such as DOE Magellan and NSF Future Grid have received significant funding to provide cloud services and a higher level of abstraction to provision scientific computing services. ATLAS and CMS have shown interest in virtualization and demonstrated several workflows that can make efficient use of cloud resources but they have been worried about the performance impact of the virtualization overhead.
- Of all OSG stakeholder the STAR experiment has been the biggest advocate of Clouds, at the OSG All hands meeting they presented a comprehensive comparison of various workflows and technologies: Nimbus/EC2, Condor VM universe and the Clemson Kestrel cloud scheduler. All three showed successful runs, STAR documented that the Clemson Kestrel scheduler had great promise. Following on this work, STAR and the Clemson OSG group partnered over the summer and made use of a sustained cloud of approximately 1,000 virtual machine instances at Clemson. This work is now on-going through the OSG/TG ExtenCI satellite project. CMS is also working via this project to run CMS virtual machines on multiple TG and OSG sites.
- OSG staff at LBNL, Fermi and Clemson are participating in the HEPiX working group on virtualization which has created a policy for trusted virtual machines and is coming up with prototypes of virtual machines exchange mechanisms. These mechanisms leverage years of work in grid computing and its associated mechanisms. In broad terms, staff producing VM will have their VM endorsed by trusted people from VOs or sites, then each site will be allowed to approved VM images from the catalogue of the endorser building a site specific trusted VM image catalogue. CERN has already deployed such a scheme and Clemson is currently working on it.

- Recently Engage demonstrated that it could use its existing job submission mechanism to submit VM to site that had hypervisors deployed. A technical memo was submitted to the OSG council describing the mechanisms. The VM is packaged by Engage as input data for the job, transferred and unzipped at the site, then booted on the site's OSG cluster using the KVM hypervisors. Each VM appears as a regular Unix process owned by the OSG VO account at the site. This work from Clemson demonstrates that a highly interoperable solution between current job submission mechanisms and Cloud specific mode of operation is doable on OSG. OSG site coordinators are currently evaluating the workflow and documenting what is necessary for all sites to support such a service.

#### **4.10 Magellan**

The OSG is collaborating with the DOE Magellan project at the Argonne Leadership Computing Facility and the National Energy Research Scientific Computing Center. Magellan is a testbed to explore the effectiveness of scientific cloud computing. The OSG's goal is to use Magellan as a testbed for future grid deployment strategies.

In the course of enabling use of Magellan, the OSG client tools have been extended.

- Condor's Amazon EC2 interface was generalized to work with 3<sup>rd</sup> party clouds that run the EC2 specification including Magellan.
- GlideinWMS was extended to enable submission to Amazon and Magellan clouds using the generic Condor framework. This enables users to transparently use the same condor execution environment on the OSG and the Cloud.

The fixes for Condor and GlideinWMS were done by their respective developers, but motivated by the testing coordinated by the OSG. The OSG and Magellan hold monthly meetings to coordinate our efforts for cloud submission. The Magellan team at NERSC has installed a generic OSG Compute Element on Magellan that opportunistic OSG VO's have successfully utilized.

We are currently experimenting with cloud submission to the NERSC Magellan. Our next step in development is to adapt the workflow we currently use on Amazon EC2 to the Magellan testbed. We expect to run production OSG jobs on Magellan by the end of 2010.

#### **4.11 Internet2 Joint Activities**

Internet2 collaborates with OSG to develop and support a suite of tools and services that make it easier for OSG sites to support its widely distributed user community. Identifying and resolving performance problems continues to be a major challenge for OSG site administrators. A complication in resolving these problems is that lower than expected performance can be caused by problems in the network infrastructure, the host configuration, or the application behavior. Advanced tools can quickly isolate problem(s) and will go a long way toward improving the grid user experience and making grids more useful to science communities.

In the past year, Internet2 has worked with OSG software developers to update the advanced network diagnostic tools already included in the VDT software package. These client applications allow VDT users to verify the network performance between end site locations and per-SONAR-based servers deployed on the Internet2 and ESnet backbones by allowing on-demand diagnostic tests to be run. The tools enable OSG site administrators and end users to test any individual compute or storage element in the OSG environment thereby reducing the time it takes to diagnose performance problems. They allow site administrators to more quickly determine if a

performance problem is due to network specific problems, host configuration issues, or application behavior.

In addition to deploying client tools via the VDT, Internet2 staff, working with partners in the US and internationally, have continued to support and enhance a simple live-CD distribution mechanism for the server side of these tools (perfSONAR-Performance-Toolkit). This bootable CD allows an OSG site-admin to quickly stand up a perfSONAR-based server to support the OSG users. These perfSONAR hosts automatically register their existence in a distributed global database, making it easy to find new servers as they become available.

These servers provide two important functions for the OSG site-administrators. First they provide an end point for the client tools deployed via the VDT package. OSG users and site-administrators can run on-demand tests to begin troubleshooting performance problems. The second function they perform is to host regularly scheduled tests between peer sites. This allows a site to continuously monitor the network performance between itself and the peer sites of interest. The US-ATLAS community has deployed perfSONAR hosts and is currently using them to monitor network performance between the Tier-1 and Tier-2 sites. Internet2 attends weekly USATLAS calls to provide on-going support of these deployments, and has come out with regular bug fixes. Finally, on-demand testing and regular monitoring can be performed to both peer sites and the Internet2 or ESNet backbone network using either the client tools, or the perfSONAR servers. Internet2 will continue to interact with the OSG admin community to learn ways to improve this distribution mechanism.

Another key task for Internet2 is to provide training on the installation and use of these tools and services. In the past year Internet2 has participated in several OSG site-admin workshops, the annual OSG all-hands meeting, and interacted directly with the LHC community to determine how the tools are being used and what improvements are required. Internet2 has provided hands-on training in the use of the client tools, including the command syntax and interpreting the test results. Internet2 has also provided training in the setup and configuration of the perfSONAR server, allowing site-administrators to quickly bring up their server. Finally, Internet2 staff has participated in several troubleshooting exercises; this includes running tests, interpreting the test results and guiding the OSG site-admin through the troubleshooting process.

#### **4.12 ESNET Joint Activities**

OSG depends on ESnet for the network fabric over which data is transferred to and from the Laboratories and to/from LIGO Caltech (by specific MOU). ESnet is part of the collaboration delivering and supporting the perfSONAR tools that are now in the VDT distribution. OSG makes significant use of ESnet's collaborative tools with telephone and video meetings. And ESnet and OSG are planning the collaborative testing of the 100Gigabit network testbed as it becomes available in the future.

OSG is the major user of the ESnet DOE Grids Certificate Authority for the issuing of X509 digital identity certificates for most people and services participating in OSG (~67% of all certificates). Registration, renewal and revocation are done through the OSG Registration Authority (RA) and ESnet provided web interfaces. ESnet and OSG collaborate on the user interface tools needed by the OSG stakeholders for management and reporting of certificates.

OSG and ESnet are implementing features of COO and contingency plans to make certificates and CA/RA operations more robust and reliable by replication and monitoring. We also partner



as members of the identity management accreditation bodies in America (TAGPMA) and globally (International Grid Trust Federation, IGTF).

OSG and ESnet jointly organized a workshop on Identity Management in November 2009 with two complementary goals: (1) To look broadly at the identity management landscape and evolving trends regarding identity in the web arena and (2) to gather input and requirements from the OSG communities about their current issues and expected future needs. A main result of the analysis of web-based technologies is the ability to delegate responsibility, which is essential for grid computing, is just beginning to be a feature of web technologies and still just for interactive timescales. A significant result from gathering input from the users is that the communities tend to either be satisfied with the current identity management functionality or are dissatisfied with the present functionality and see a strong need to have more fully integrated identity handling across the range of collaborative services used for their scientific research. The results of this workshop and requirements gathering survey are being used to help plan the future directions for work in this area with two main thrusts being improvements to the registration process and closer integration between web and command line services. More details are included in the Security section of this report.

There is currently an effort underway lead by Mike Helm of ESnet to help and encourage DOE laboratories to use the Shibboleth identity federation technology and to join the InCommon Federation as a way to provide more efficient and secure network access to scientific facilities for the widely distributed user communities located at universities as well as laboratories. Technical discussions of issues particular to DOE laboratories are carried out on the Science Federation Google group as well as a demonstration collaborative web space at [confluence.scifed.org](http://confluence.scifed.org). This activity is of great interest to OSG as it leads to the next stage in the evolution of secure network identity credentials.

## 5. Training, Outreach and Dissemination

---

### 5.1 Training

The OSG Training program brings domain scientists and computer scientists together to provide a rich training ground for the engagement of students, faculty and researchers in learning the OSG infrastructure, applying it to their discipline and contributing to its development. During 2010, OSG sponsored and conducted training events. Training organized and delivered by OSG in the last year is identified in the following table:

<b>Workshop</b>	<b>Length</b>	<b>Location</b>	<b>Month</b>
Grid Colombia Workshop	5 days	Bucaramanga, Colombia	Mar., 2010
OSG Summer School	4 days	Madison, Wisconsin	July, 2010
Site Administrators Workshop	2 days	Nashville, TN	Aug., 2010
South American Grid Workshop	5 days	Sao Paulo, Brazil	Dec., 2010

The Grid Colombia workshop in March 2010 was supported by OSG core staff and provided contributions from the OSG community at large to help train Colombian site administrators. This was a follow-up to a two-week workshop held the previous October that aimed to help them launch a National Grid Infrastructure (NGI). The Grid Colombia project aims to connect 11 public and private universities (representing more than 100,000 students and 5,000 faculty members). The workshops provided technical training and hands-on experience in setting up and managing Grid sites, central Grid services (information, VO management, and accounting) as well contributed presentations from OSG experts in workload management, security and operations.

Other training activities include the organization of the Site Administrators workshop held in Nashville TN in August 2010; support for VO-focused events such as the US CMS Tier 3 workshop, and grid schools like the OSG Summer School and the Sao Paulo OSG school. For the Site Administrators workshop, experts from around the Consortium contributed to course-ware tutorials that participants used during the event. Some forty site administrators, several brand new to OSG, participated and were impressed by the many experts on hand which included not only the instructors but also seasoned OSG site administrators all who were eager to help.

OSG staff also participated as keynote speakers, instructors, and/or presenters at a venue this year as detailed in the following table:

<b>Venue</b>	<b>Length</b>	<b>Location</b>	<b>Month</b>
IX DOSAR Workshop	3 days	Pilanesburg, South Africa	April, 2010
Oklahoma Supercomputing Symposium	2 days	Norman, Oklahoma	Oct., 2010

## 5.2 Outreach Activities

We present a selection of the activities in the past year:

- The NSF Task Force on Campus Bridging
- HPC Best Practices Workshop
- Workshops on Distributed Computing, Multidisciplinary Science, and the NSF's Scientific Software Innovation Institutes Program
- Chair of the Network for Earthquake Engineering Simulation Cyberinfrastructure sub-committee of the Project Advisory Committee
- Member of the DOE Knowledge Base requirements group
- Continued co-editorship of the highly successful International Science Grid This Week newsletter, [www.isgtw.org](http://www.isgtw.org) (see next section)

## 5.3 Internet dissemination

OSG co-sponsors the weekly electronic newsletter called International Science Grid This Week (<http://www.isgtw.org/>) in collaboration with the European project e-Science Talk. The newsletter has been very well received, having published 203 issues as of November 2010 with subscribers totaling approximately 6,838, an increase of over 13% in the last year. With support from NSF and DOE, a full-time US editor has been hired to develop and improve the publication. This in turn has improved our ability to showcase US contributions.

Articles are drawn from a wide range of locales including the United States, Europe, the Americas, the Pacific Rim and the Far East. Effort is made to insure that readers are treated to a wide range of topics in the area of computing-enabled scientific research done via various flavors of computing—advanced, distributed, grid, cloud, and other cyber infrastructure techniques.

In 1Q-2011 International Science Grid This Week will be renamed The Digital Scientist <http://www.thedigitalscientist.org/> and the web site re-launched with the intent that the clearer name and more modern web site features will bring even more subscribers.

In addition, the OSG maintains a web site, <http://www.opensciencegrid.org>, intended to inform and guide stakeholders and new users of the OSG. Sections show various accomplishments of the OSG including the large number of scientific publications that have resulted from the OSG.

## **6. Cooperative Agreement Performance**

---

OSG has put in place processes and activities that meet the terms of the Cooperative Agreement and Management Plan:

- The Joint Oversight Team meetings are conducted, as scheduled by DOE and NSF, via phone to hear about OSG progress, status, and concerns. Follow-up items are reviewed and addressed by OSG, as needed.
- Two intermediate progress reports were submitted to NSF in February and June of 2007.
- In February 2008, a DOE annual report was submitted.
- In July 2008, an annual report was submitted to NSF.
- In December 2008, a DOE annual report was submitted.
- In June 2009, an annual report was submitted to NSF.
- In January 2010, a DOE annual report was submitted.
- In June 2010, an annual report was submitted to NSF.
- In December 2010, this annual report submitted to DOE.



As requested by DOE and NSF, OSG staff provides pro-active support in workshops and collaborative efforts to help define, improve, and evolve the US national cyberinfrastructure.


# **Production on Open Science Grid**

Year Ending December 2010

1. OSG Facility Summary.....	2
2. ATLAS on the OSG .....	4
3. CMS on the OSG .....	5
4. CDF Usage .....	6
5. D0 Usage .....	7
6. LIGO on the OSG .....	8
7. ENGAGE Usage.....	9
8. SBGrid Usage.....	10
9. HCC Usage .....	11
10. Non-HEP Usage .....	12


## 1. OSG Facility Summary







### Job Count by Facility (Sum: 157,043,909)

51 Weeks from Week 48 of 2009 to Week 48 of 2010





## 2. ATLAS on the OSG









### 3. CMS on the OSG





## 4. CDF Usage





## 5. D0 Usage





## 6. LIGO on the OSG





## 7. ENGAGE Usage




## 8. SBGrid Usage



## 9. HCC Usage



## 10. Non-HEP Usage





**Publications**  
**Enabled by**  
**Open Science Grid**  
**January 2010 - December 2010**

## Contents

ATLAS Virtual Organization .....	3
CDF Virtual Organization.....	9
CDMS Virtual Organization .....	15
CIGI Virtual Organization .....	15
CMS Virtual Organization.....	15
ENGAGE Virtual Organization .....	30
GLOW Virtual Organization .....	32
HCC Virtual Organization .....	34
LIGO Collaboration .....	34
MiniBooNE Virtual Organization.....	35
MINOS Virtual Organization .....	35
NanoHUB Virtual Organization .....	36
NYSgrid Virtual Organization .....	36
SBGrid Virtual Organization .....	36
STAR Virtual Organization .....	37

## ATLAS Virtual Organization

- 1) [Readiness of the ATLAS tile calorimeter for LHC collisions](#)  
By: The ATLAS Collaboration  
[[arXiv:1007.5423v2](#)] (19 Nov 2010)
- 2) [Search for New Particles in Two-Jet Final States in 7 TeV Proton-Proton Collisions with the ATLAS Detector at the LHC](#)  
By: The ATLAS Collaboration  
[[arXiv:1008.2461v2](#)] [Phys. Rev. Lett. 105, 161801](#) (14 Aug 2010)
- 3) [Search for Quark Contact Interactions in Dijet Angular Distributions in 7 TeV Proton-Proton Collisions with the ATLAS Detector at the LHC](#)  
By: The ATLAS Collaboration  
[[arXiv:1009.5069v1](#)] (26 Sep 2010)
- 4) [Measurement of inclusive jet and dijet cross sections in proton-proton collisions at 7 TeV centre-of-mass energy with the ATLAS detector](#)  
By: The ATLAS Collaboration  
[[arXiv:1009.5908v2](#)] (30 Sep 2010)
- 5) [Measurement of the  \$W \rightarrow l\nu\$  and  \$Z/\gamma^\* \rightarrow ll\$  production cross sections in proton-proton collisions at  \$\sqrt{s} = 7\$  TeV with the ATLAS detector](#)  
By: The ATLAS Collaboration  
[[arXiv:1010.2130v1](#)] 11 Oct. 2010)
- 6) [Observation of a centrality-dependent dijet asymmetry in lead-lead collisions at  \$\sqrt{s\_{NN}} = 2.76\$  TeV with the ATLAS detector at the LHC](#)  
By: The ATLAS Collaboration  
[[arXiv:1011.6182v2](#)] [Phys. Rev. Lett.](#) (25 Nov 2010)
- 7) [Studies of the performance of the ATLAS detector using cosmic-ray muons](#)  
By: The ATLAS Collaboration  
[[arXiv:1011.6665v1](#)] [CERN-PH-EP-2010-062](#) (1 Dec 2010)
- 8) [Measurement of underlying event characteristics using charged particles in pp collisions at  \$\sqrt{s} = 900\$  GeV and 7 TeV with the ATLAS detector](#)

By: The ATLAS Collaboration  
[\[arXiv:1012.0791v1\]](#) (3 Dec 2010)

- 9) [Relative luminosity measurement of the LHC with the ATLAS forward calorimeter.](#)  
By A. Afonin, et al.,  
JINST 5:P05005,2010,. [\[arXiv:1005.1784\]](#)
- 10) [The ATLAS discovery reach for SUSY models with early data.](#)  
By Janet Dietrich & for the ATLAS collaboration.  
[\[arXiv:1005.2034\]](#) (May 2010) 4p.
- 11) [Commissioning of the ATLAS Liquid Argon Calorimeter.](#)  
By S. Laplace.  
Nucl.Instrum.Meth. A617:30-34,2010,. [\[arXiv:1005.2935\]](#)
- 12) [Collective flow in \(anti\)proton-proton collision at Tevatron and LHC.](#)  
By Tanguy Pierog, Sarah Porteboeuf, Iurii Karpenko, Klaus Werner.  
[\[arXiv:1005.4526\]](#) (May 2010) 4p.
- 13) [The ATLAS Simulation Infrastructure.](#)  
By The ATLAS Collaboration.  
[\[arXiv:1005.4568\]](#) (May 2010)
- 14) [Performance of the ATLAS Detector using First Collision Data.](#)  
By The ATLAS Collaboration.  
[\[arXiv:1005.5254\]](#) (May 2010)
- 15) [Readiness of the ATLAS detector: Performance with the first beam and cosmic data.](#)  
By Francesca Pastore.  
Nucl.Instrum.Meth.A617:48-51,2010,.
- 16) [Prospects for Observing the Standard Model Higgs Boson Decaying into  \$b\bar{b}\$  Final States Produced in Weak Boson Fusion with an Associated Photon at the LHC.](#)  
By D.M. Asner, M. Cunningham, S. Dejong, K. Randrianarivony, C. Santamarina, M. Schram.  
[\[arXiv:1004.0535\]](#) (Apr 2010) 15p.
- 17) [The ATLAS Inner Detector commissioning and calibration.](#)  
By ATLAS Collaboration  
[\[arXiv:1004.5293\]](#) (Apr 2010) 34p.
- 18) [Identification of muons in ATLAS.](#)  
By Zdenko van Kesteren.  
(Mar 2010) 159p.

- 19) [LHC Detectors and Early Physics.](#)  
By Guenther Dissertori.  
[arXiv:1003.2222] (Mar 2010)
- 20) [Charged-particle multiplicities in pp interactions at  \$\sqrt{s} = 900\$  GeV measured with the ATLAS detector at the LHC.](#)  
By ATLAS Collaboration  
Phys Lett B 688, 2010, Issue 1, 21-42, Phys.Lett.B688:21-42,2010,. [arXiv:1003.3124]
- 21) [Higgs Boson Search Sensitivity in the  \$SH \rightarrow WW\$  Dilepton Decay Mode at  \$\sqrt{s} = 7\$  and 10 TeV.](#)  
By Edmond L. Berger, Qing-Hong Cao, C.B. Jackson, Tao Liu, Gabe Shaughnessy.  
[arXiv:1003.3875] ANL-HEP-PR-10-8 (Mar 2010)
- 22) [Topics in the measurement of top quark events with ATLAS: Pixel detector optoelectronics, track impact parameter calibration, acceptance correction methods.](#)  
By Stephan A. Sandvoss.  
(Mar 2010) 170p.
- 23) [Commissioning and early physics analysis with the ATLAS and CMS experiments.](#)  
By Andreas Hoecker.  
[arXiv:1002.2891] (Feb 2010) 84p.
- 24) [Drift Time Measurement in the ATLAS Liquid Argon Electromagnetic Calorimeter using Cosmic Muons.](#)  
By The ATLAS Collaboration.  
[arXiv:1002.4189] (Feb 2010) 30p.
- 25) [Search for Supersymmetry in Tripleton Final States with the ATLAS Detector and the Alignment of the ATLAS Silicon Tracker.](#)  
By Oleg Brandt.  
[arXiv:1001.1365] (Jan 2010) 236p.
- 26) [The Hunt for New Physics at the Large Hadron Collider.](#)  
By P. Nath, et al.,  
Nucl.Phys.Proc.Suppl.200-202:185-417,2010,. [arXiv:1001.2693]
- 27) [B Physics and Quarkonia studies with early ATLAS data.](#)  
By for the ATLAS Collaboration  
[arXiv:1001.3806] ATL-PHYS-PROC-2010-006 (Jan 2010)
- 28) [NLO QCD corrections to top anti-top bottom anti-bottom production at the LHC: 2. full hadronic results.](#)  
By Axel Bredenstein, Ansgar Denner, Stefan Dittmaier, Stefano Pozzorini.  
JHEP 1003:021,2010,. [arXiv:1001.4006]

- 29) [The construction and testing of the silicon position sensitive modules for the LHCf experiment at CERN.](#)  
By O. Adriani, et al.,  
JINST 5:P01012,2010,.
- 30) [Micromegas study for the sLHC environment.](#)  
By T. Alexopoulos, et al.,  
JINST 5:P02003,2010,.
- 31) [Atlas TRT: Research & design B-type module mass production.](#)  
By Yu.V. Gusakov, et al.,  
Phys.Part.Nucl.41:1-26,2010,.
- 32) [New developments in data-driven background determinations for SUSY searches in ATLAS.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:297-300,2010,.
- 33) [Prospects for the discovery of supersymmetry based on inclusive searches in ATLAS.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:301-304,2010,.
- 34) [Multi-lepton SUSY searches with the ATLAS detector.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:305-308,2010,.
- 35) [Search for dilepton and lepton + E\(T\)\\*\\*miss resonances at high mass with ATLAS.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:309-312,2010,.
- 36) [Reconstruction of tau leptons and prospects for SUSY in ATLAS.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:341-344,2010,.
- 37) [Exclusive measurements for SUSY events with the ATLAS detector at the LHC.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:345-348,2010,.
- 38) [R-parity violation with jet signatures at the ATLAS detector.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:354-357,2010,.
- 39) [Signatures of GMSB with non-pointing photons at the ATLAS detector.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:362-365,2010,.

- 40) [Associated standard model Higgs boson search with ATLAS.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:378-381,2010,.
- 41) [Inclusive standard model Higgs boson searches in ATLAS.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:382-385,2010,.
- 42) [Beyond the standard model Higgs searches at ATLAS.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:394-397,2010,.
- 43) [SM Higgs properties measurement at ATLAS.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:402-405,2010,.
- 44) [Search for resonances decaying into top quark pairs with ATLAS.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:678-681,2010,.
- 45) [Electron reconstruction and identification with the ATLAS detector.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:693-696,2010,.
- 46) [Photon reconstruction and identification with the ATLAS detector.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:697-700,2010,.
- 47) [In-situ determination of the ATLAS muon performance.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:705-708,2010,.
- 48) [Search for new physics with long-lived particles.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:742-745,2010,.
- 49) [Reconstruction and identification of heavy long-lived particles at the ATLAS detector at the LHC.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:762-765,2010,.
- 50) [Prospects for early discoveries at the LHC with dileptons, jets and no missing energy with the ATLAS detector.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:802-805,2010,.

- 51) [Study of a micromegas chamber in a neutron beam.](#)  
By T. Alexopoulos, et al.,  
JINST 5:P02005,2010,.
- 52) [Measurement of the response of the ATLAS liquid argon barrel calorimeter to electrons at the 2004 combined test-beam.](#)  
By M. Aharrouche, et al.,  
Nucl.Instrum.Meth.A614:400-432,2010,.
- 53) [The ATLAS muon spectrometer.](#)  
By Giora Mikenberg.  
Mod.Phys.Lett.A25:649-667,2010,.
- 54) [How physics defines the LHC environment and detectors.](#)  
By D. Green.  
Int.J.Mod.Phys.A25:1279-1313,2010,.
- 55) [Higgs physics with ATLAS.](#)  
By ATLAS Collaboration  
Nucl.Phys.Proc.Suppl.198:149-156,2010,.
- 56) [The hybrid tracking system of ATLAS.](#)  
By Leonardo Rossi.  
Int.J.Mod.Phys.A25:1519-1540,2010,.
- 57) [The ATLAS electromagnetic calorimeters: Features and performance.](#)  
By Luciano Mandelli.  
Int.J.Mod.Phys.A25:1739-1760,2010,.
- 58) [ATLAS electronics: An overview.](#)  
By Philippe Farthouat.  
Int.J.Mod.Phys.A25:1761-1784,2010,.
- 59) [TileCal: The hadronic section of the central ATLAS calorimeter.](#)  
By K. Anderson, T. Del Prete, E. Fullana, J. Huston, C. Roda, R. Stanek.  
Int.J.Mod.Phys.A25:1981-2003,2010,.
- 60) [Readiness of the ATLAS Liquid Argon Calorimeter for LHC Collisions](#)  
By ATLAS Collaboration  
[arXiv:0912.2642] (Dec 2009) 31p.
- 61) [A Detailed Far-Ultraviolet Spectral Atlas of Main Sequence B Stars.](#)  
By Myron A. Smith.  
Astrophys.J.Suppl.186:175-190,2010,. [arXiv:0911.4747]
- 62) [Searches for SUSY with the ATLAS detector.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:32-40,2010,. [arXiv:0910.0559]



- 63) [Performance of the ATLAS Detector on First Single Beam and Cosmic Ray Data.](#)  
By Martin Aleksa.  
AIP Conf.Proc.1200:689-692,2010,. [arXiv:0910.2198]
- 64) [ScotGrid: Providing an Effective Distributed Tier-2 in the LHC Era.](#)  
By Sam Skipsey, David Ambrose-Griffith, Greig Cowan, Mike Kenyon, Orlando Richards, Phil Roffe, Graeme Stewart.  
J.Phys.Conf.Ser.219:052014,2010,. [arXiv:0910.4507]
- 65) [Discovery Potential of R-hadrons with the ATLAS Detector at the LHC.](#)  
By ATLAS Collaboration  
AIP Conf.Proc.1200:750-753,2010,. [arXiv:0909.1911]
- 66) [Determination of QCD Backgrounds in ATLAS: A Challenge for SUSY searches.](#)  
By Bernhard Meirose.  
AIP Conf.Proc.1200:717-720,2010,. [arXiv:0909.4427]
- 67) [The Software of the ATLAS beam pick-up based LHC monitoring system.](#)  
By C. Ohm & T. Pauly.  
J.Phys.Conf.Ser.219:022040,2010,. [arXiv:0909.5378]

## **CDF Virtual Organization**

- 1) Measurement of the WW + WZ Production Cross Section using a Matrix Element Technique in Lepton + Jets Events, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D82, 112001 (2010). arXiv: 008.4404. Abstract, PostScript (673KiB). 10275: Added: Thursday, 2010 December 02 16:15:26 CST
- 2) Updated Search for the Flavor-Changing Neutral-Current Decay  $D^0 \rightarrow \mu^+\mu^-$  T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D82, 091105 (2010). Fermilab-Pub-10-337-E. arXiv: 1008.5077. Abstract, PDF (415KiB). 9899: Added: Tuesday, 2010 November 16 13:44:46 CST
- 3) Search for Supersymmetric Partner of the Top Quark in p anti-p Collisions at  $\sqrt{s} = 1.96$  TeV, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D82, 092001 (2010). arXiv: 009.0266. Abstract, PDF (1509KiB). 10164: Added: Tuesday, 2010 November 02 11:52:56 CDT
- 4) Search for R-Parity Violating Decays of tau Sneutrinos to e  $\mu$ ,  $\mu$  tau, and e tau Pairs in p anti-p Collisions at  $\sqrt{s} = 1.96$  TeV T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 105, 191801 (2010). arXiv: 1004.3042. Abstract, PDF (251KiB). 10006: Added: Tuesday, 2010 November 16 11:09:20 CST

- 5) Measurement of the Top Pair Production Cross Section in the Dilepton Decay Channel in p anti-p Collisions at  $\sqrt{s}$  1.96 TeV, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D82, 052002 (2010). arXiv: 1002.2919. Abstract, PDF (333KiB). 9991:  
Added: Thursday, 2010 September 23 15:27:46 CDT
- 6) Search for Anomalous Production of Events with Two Photons and Additional Energetic Objects at CDF, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D82, 052005 (2010). arXiv: 910.5170 Abstract, PostScript (1106KiB). 9886:  
Added: Wednesday, 2010 September 29 13:23:55 CDT
- 7) Search for New Physics with a Dijet Plus Missing E(t) Signature in p anti-p Collisions at  $\sqrt{s}$  1.96 TeV, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 105, 131801 (2010). arXiv:0903.2850. Abstract, PostScript (364KiB). 9329:  
Added: Thursday, 2010 September 23 16:08:29 CDT
- 8) Measurement of  $Z_\gamma$  Production in p anti-p Collisions at  $\sqrt{s}$  1.96 TeV T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D82, 031103 (2010). arXiv: 004.1140. Abstract, PDF (1081KiB). 9904: Added: Friday, 2010 August 27 16:24:22 CDT
- 9) Studying the Underlying Event in Drell-Yan and High Transverse Momentum Jet Production at the Tevatron, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D82, 034001 (2010). arXiv:1003.3146. Abstract, PDF (4872KiB). 9567:  
Added: Wednesday, 2010 August 18 13:49:10 CDT
- 10) Search for the Production of Scalar Bottom Quarks in p anti-p Collisions at  $\sqrt{s}$  1.96 TeV, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 105, 081802 (2010). arXiv:1005.3600. Abstract, PDF (173KiB). 10091:  
Added: Friday, 2010 August 20 11:54:02 CDT
- 11) Exclusion of an Exotic Top Quark with  $-4/3$  Electric Charge Using Soft Lepton Tagging T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 105, 101801 (2010). arXiv: 1006.4597. Abstract, PDF (256KiB). 10028:  
Added: Thursday, 2010 September 23 15:45:23 CDT
- 12) Measurement of W-Boson Polarization in Top-quark Decay in p anti-p Collisions at  $\sqrt{s}$  1.96 TeV, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 105, 042002 (2010). Fermilab-Pub-10-041-E. arXiv: 1003.0224. Abstract, PostScript (244KiB). 9995:  
Added: Monday, 2010 July 26 16:21:08 CDT
- 13) Search for a New Color-Octet Vector Particle Decaying to t anti-t Collisions at  $\sqrt{s}$  1.96 TeV, Phys. Lett. B, 691 (2010). Fermilab-Pub-09-578-E. arXiv: 0911.3112. Abstract, PDF (236KiB). 9636: Added: Monday, 2010 July 26 14:40:35 CDT

- 14) Search for WW and WZ Resonances Decaying to Electron, Missing E(t), and Two Jets in p anti-p Collisions at  $s^{1/2}$  1.96 TeV T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 241801 (2010). Fermilab-Pub-10-096-E. arXiv: 1004.4946. Abstract, PDF (362KiB). 10141: Added: Monday, 2010 July 26 14:16:09 CDT
- 15) Combined Tevatron Upper Limit on  $gg \rightarrow H \rightarrow W+W-$  and Constraints on the Higgs Boson Mass in Fourth-Generation Fermion Models T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D82, 011102 (2010). arXiv: 1005.3216. Abstract, PDF (267KiB). 10101: Changed: Monday, 2010 July 26 13:48:55 CDT
- 16) Measurement of the  $W+W-$  Production Cross Section and Search for Anomalous WWg and WWZ Couplings, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 201801 (2010). arXiv: 0912.4691. Abstract, PostScript (414KiB). 9902: Added: Tuesday, 2010 July 27 16:13:41 CDT
- 17) Search for Pair Production of Supersymmetric Top Quarks in Dilepton Events from p anti-p Collisions at  $s^{1/2}$  1.96 TeV, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 251801 (2010). arXiv: 0912.1308. Abstract, PDF (257KiB). 9823: Added: Wednesday, 2010 July 28 14:54:19 CDT
- 18) First Measurement of the Ratio  $\sigma_{t\bar{t}}/\sigma_{Z/\gamma^*} \rightarrow e^+e^-$  and Precise Extraction of the  $t\bar{t}$  Cross Section, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 105, 012001 (2010). arXiv: 1004.3224. Abstract, PDF (244KiB). 10019: Added: Wednesday, 2010 July 28 15:16:38 CDT
- 19) Search for Single Top Quark Production in p anti-p Collisions at 1.96 TeV in the Missing Transverse Energy Plus Jets Topology T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D81, 072003 (2010). arXiv: 1001.4577. Abstract, PDF (1213KiB). 9933: Added: Thursday, 2010 April 22 11:06:33 CDT
- 20) A Search for the Higgs Boson Using Neural Networks in Events with Missing Energy and b-quark Jets in p anti-p Collisions at  $s^{1/2}$  1.96 TeV T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 141801 (2010). arXiv: 0911.3935. Abstract, PDF (303KiB). 9838: Added: Monday, 2010 April 19 10:53:25 CDT
- 21) Measurement of the b-jet Cross Section in Events with a W Boson in p anti-p Collisions at  $s^{1/2}$  1.96 TeV, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 131801 (2010). arXiv: 0909.1505. Abstract, PDF (264KiB). 9649: Added: Monday, 2010 April 05 15:56:31 CDT
- 22) Measurement of the Top Quark Mass and p anti-p  $\rightarrow t\bar{t}$  Cross Section in the All-Hadronic Mode with the CDF II Detector T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D81, 052001 (2010). arXiv: 1002.0365. Abstract, PDF (413KiB). 9976: Added: Monday, 2010 April 05 10:05:03 CDT

- 23) Measurement of the  $WW+WZ$  Production Cross Section using Lepton+Jets Final State at CDF II, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 101801 (2010). arXiv: 0911.4449. Abstract, PDF (181KiB). 9964:  
Changed: Thursday, 2010 March 18 11:58:12 CDT
- 24) Search for New Bottomlike Quark Pair Decays  $QQ \rightarrow (t W^{+/-})(\bar{t} W^{+/-})$  in Same-Charge Dilepton Events, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 091801 (2010). arXiv: 0912.1057. Abstract, PDF (251KiB). 9919:  
Added: Monday, 2010 March 08 11:42:12 CST
- 25) Search for Technicolor Particles Produced in Association with a W Boson at CDF T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 111802 (2010). arXiv: 0912.2059. Abstract, PDF (189KiB). 9907:  
Added: Monday, 2010 March 22 16:28:03 CDT
- 26) Study of the Associated Production of Photons and b-quark Jets in p anti-p Collisions at  $\sqrt{s} = 1.96$  TeV, T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D81, 052006 (2010). arXiv: 0912.3453. Abstract, PostScript (1069KiB). 9803:  
Added: Thursday, 2010 March 18 11:56:11 CDT
- 27) Measurement of the  $\Lambda_b$  Lifetime in  $\Lambda_b \rightarrow \Lambda_c \pi$  Decays in p anti-p Collisions at  $\sqrt{s} = 1.96$  TeV T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 102002 (2010). arXiv: 0912.3566. Abstract, PDF (163KiB). 9763:  
Added: Thursday, 2010 March 18 11:43:45 CDT
- 28) Measurements of the Branching Fraction Ratios and CP Asymmetries in  $B \rightarrow D_0 CP$   $K^{+/-}$  Decays in Hadron Collisions T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D81, 031105 (2010). arXiv: 0911.0425. Abstract, PDF (218KiB). 9708:  
Added: Wednesday, 2010 March 03 11:33:24 CST
- 29) Search for Single Top Quark Production in p anti-p Collisions at 1.96 TeV in the Missing Transverse Energy Plus Jets Topology. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D81, 072003 (2010). arXiv: 1001.4577.
- 30) Measurement of  $Z_\gamma$  Production in p anti-p Collisions at  $\sqrt{s} = 1.96$  TeV. T. Aaltonen et al., The CDF Collaboration, submitted to Phys. Rev. D April 7, 2010. Fermilab-Pub-10-073-E. arXiv: 1004.1140.
- 31) A Search for the Higgs Boson Using Neural Networks in Events with Missing Energy and b-quark Jets in p anti-p Collisions at  $\sqrt{s} = 1.96$  TeV. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 141801 (2010). arXiv: 0911.3935.

- 32) Measurement of the B- Lifetime Using a Simulation Free Approach for Trigger Bias Correction. T. Aaltonen et al., The CDF Collaboration, submitted to Phys. Rev. D April 27, 2010. Fermilab-Pub-10-095-E. arXiv: 1044.4855.
- 33) Measurement of the b-jet Cross Section in Events with a W Boson in p anti-p Collisions at  $\sqrt{s}=1.96$  TeV. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 131801 (2010). arXiv: 0909.1505.
- 34) Observation of Single Top Quark Production and Measurement of  $|V_{tb}|$  with CDF. T. Aaltonen et al., The CDF Collaboration, submitted to Phys. Rev. D April 7, 2010. Fermilab-Pub-10-063-E. arXiv: 1004.1181.
- 35) First Measurement of the Ratio  $\sigma_{t\bar{t}}/\sigma_Z/\gamma^* \rightarrow e^+e^-$  and Precise Extraction of the  $t\bar{t}$  Cross Section. T. Aaltonen et al., The CDF Collaboration, submitted to Phys. Rev. Lett. April 19, 2010. Fermilab-Pub-10-084-E. arXiv: 1004.3224.
- 36) Search for R-parity Violating Decays of tau Sneutrinos to  $e\mu$ ,  $\mu\tau$ , and  $e\tau$  Pairs in p anti-p Collisions at  $\sqrt{s}=1.96$  TeV. T. Aaltonen et al., The CDF Collaboration, submitted to Phys. Rev. Lett. April 19, 2010. Fermilab-Pub-10-077-E. arXiv: 1004.3042.
- 37) Measurement of W-Boson Polarization in Top-quark Decay in p anti-p Collisions at  $\sqrt{s}=1.96$  TeV. T. Aaltonen et al., The CDF Collaboration, submitted to Phys. Rev. Lett. March 1, 2010. Fermilab-Pub-10-041-E. arXiv: 1003.0224.
- 38) Measurement of the Top Quark Mass and p anti-p  $\rightarrow t\bar{t}$  Cross Section in the All-Hadronic Mode with the CDF II Detector. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D81, 052001 (2010). arXiv: 1002.0365.
- 39) Measurement of the WW+WZ Production Cross Section using Lepton+Jets Final State at CDF II. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 101801 (2010). arXiv: 0911.4449.
- 40) Search for New Bottomlike Quark Pair Decays  $QQ \rightarrow (t W^-)(\bar{t} W^+)$  in Same-Charge Dilepton Events. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 091801 (2010). arXiv: 0912.1057.
- 41) Search for Technicolor Particles Produced in Association with a W Boson at CDF. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 111802 (2010). arXiv: 0912.2059.
- 42) Study of the Associated Production of Photons and b-quark Jets in p anti-p Collisions at  $\sqrt{s}=1.96$  TeV. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D81, 052006 (2010). arXiv: 0912.3453.

- 43) Measurement of the  $\Lambda^0(b)$  Lifetime in  $\Lambda^0(b) \rightarrow \Lambda^+_c \pi^-$  Decays in p anti-p Collisions at  $\sqrt{s}=1.96$  TeV. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 102002 (2010). arXiv: 0912.3566.
- 44) Measurements of the Branching Fraction Ratios and CP Asymmetries in  $B \rightarrow D^0_{CP} K^\pm$  Decays in Hadron Collisions. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D81, 031105 (2010). arXiv: 0911.0425.
- 45) Studying the Underlying Event in Drell-Yan and High Transverse Momentum Jet Production at the Tevatron. T. Aaltonen et al., The CDF Collaboration, submitted to Phys. Rev. D March 17, 2010. Fermilab-Pub-10-053-E. arXiv: 1003.3146.
- 46) Measurement of  $d\sigma/dy$  of Drell-Yan  $e^+e^-$  Pairs in the Z Mass Region from p anti-p Collisions at  $\sqrt{s}=1.96$  TeV. T. Aaltonen et al., The CDF Collaboration, submitted to Physics Letters B March 29, 2010. Fermilab-Pub-10-057-E. arXiv: 0908.3914.
- 47) Measurement of the Top Pair Production Cross Section in the Dilepton Decay Channel in p anti-p Collisions at  $\sqrt{s}=1.96$  TeV. T. Aaltonen et al., The CDF Collaboration, submitted to Phys. Rev. D February 15, 2010. Fermilab-Pub-10-032-E. arXiv: 1002.2919.
- 48) Measurement of the Top Quark Mass in the Dilepton Channel Using  $m_{t\bar{t}}$  at CDF. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D81, 031102 (2010). arXiv: 0911.2956.
- 49) Inclusive Search for Standard Model Higgs Boson Production in the WW Decay Channel Using the CDF II Detector. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 061803 (2010). arXiv: 1001.4468.
- 50) Measurement of the  $t\bar{t}$  Production Cross Section in p anti-p Collisions at  $\sqrt{s}=1.96$  TeV using Soft Electron b-Tagging. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. D81, 092002 (2010). arXiv: 1002.3783.
- 51) Search for Supersymmetry with Gauge-Mediated Breaking in Diphoton Events with Missing Transverse Energy at CDF II. T. Aaltonen et al., The CDF Collaboration, Phys. Rev. Lett. 104, 011801 (2010). arXiv: 0910.3606.
- 52) Combination of Tevatron Searches for the Standard Model Higgs Boson in the  $W^+W^-$  Decay Mode. T. Aaltonen et al., The CDF Collaboration, submitted to Phys. Rev. Lett. January 25, 2010. Fermilab-Pub-10-017-E. arXiv: 1001.4162.

## CDMS Virtual Organization

- 1) Analysis of the low-energy electron-recoil spectrum of the CDMS experiment.

By CDMS Collaboration ([Z. Ahmed et al.](#)). FERMILAB-PUB-09-340-E, Jul 2009.  
(Published Feb 15, 2010). 4pp.  
Published in Phys.Rev.D81:042002, 2010.  
e-Print: arXiv:0907.1438 [astro-ph.GA]

- 2) Dark Matter Search Results from the CDMSII Experiment

*Science* The CDMS II Collaboration 327 (5973): 1619 (published March 2010)

## CIGI Virtual Organization

- 1) Padmanabhan, A., Ghosh, S., and Wang, S. 2010. “A Self-Organized Grouping (SOG) Method for Efficient Grid Resource Discovery.” *Journal of Grid Computing*, In press. DOI 10.1007/s10723-009-9145-0.
- 2) Padmanabhan, A., and Wang, S. “A Distributed Resource Broker for Spatial Middleware Using Adaptive Space-Filling Curve”. In: *Proceedings of the ACM SIGSPATIAL International Workshop on High Performance and Distributed Geographic Information Systems (ACM HPDGIS 2010)*, November 2, 2010, San Jose, CA, USA

## CMS Virtual Organization

- 1) Measurement of the Isolated Prompt Photon Production Cross Section in  $pp$  Collisions at  $\sqrt{s} = 7$  TeV / CMS Collaboration

The differential cross section for the inclusive production of isolated prompt photons has been measured as a function of the photon transverse energy  $E_T$ -gamma in  $pp$  collisions at  $\sqrt{s}=7$  TeV using data recorded by the CMS detector at the LHC. [...]

[arXiv:1012.0799](#) ; [CERN-PH-EP-2010-053](#) ; CMS-QCD-10-019. - 2010. - 26 p.



- 2) Search for Stopped Gluinos in pp collisions at  $\sqrt{s} = 7$  TeV / CMS Collaboration The results of the first search for long-lived gluinos produced in 7 TeV pp collisions at the CERN Large Hadron Collider are presented. [...]

[arXiv:1011.5861](#) ; [CERN-PH-EP-2010-049](#) ; CMS-EXO-10-003. - 2010. - 26 p.

- 3) Charged particle multiplicities in pp interactions at  $\sqrt{s} = 0.9, 2.36, \text{ and } 7$  TeV / CMS collaboration Measurements of primary charged hadron multiplicity distributions are presented for non-single-diffractive events in proton-proton collisions at centre-of-mass energies of  $\sqrt{s} = 0.9, 2.36, \text{ and } 7$  TeV, in five pseudorapidity ranges from  $|\eta| < 0.5$  to  $|\eta| < 2.4$  [...]

[arXiv:1011.5531](#) ; [CERN-PH-EP-2010-048](#) ; CMS-QCD-10-004. - 2010. - 36 p.

- 4) Prompt and non-prompt  $J/\psi$  production in pp collisions at  $\sqrt{s} = 7$  TeV / CMS Collaboration

The production of  $J/\psi$  mesons is studied in pp collisions at  $\sqrt{s} = 7$  TeV with the CMS experiment at the LHC. [...]

[arXiv:1011.4193](#) ; [CERN-PH-EP-2010-046](#) ; CMS-BPH-10-002. - 2010. - 38 p.

- 5) First Measurement of the Cross Section for Top-Quark Pair Production in Proton-Proton Collisions at  $\sqrt{s} = 7$  TeV / CMS Collaboration

The first measurement of the cross section for top-quark pair production in pp collisions at the LHC at center-of-mass energy  $\sqrt{s} = 7$  TeV has been performed using 3.1  $\sqrt{\text{fb}}$  0.3 inverse pb of data recorded by the CMS detector. This result utilizes the final state with two isolated, highly energetic charged leptons, large missing transverse energy, and two or more jets [...]

[arXiv:1010.5994](#); [CERN-PH-EP-2010-039](#); [CMS-TOP-10-001](#).- Geneva : CERN, 2010 - 26 p. [Preprint](#) ; [Elsevier Open Access article](#)



- 6) Search for Quark Compositeness with the Dijet Centrality Ratio in pp Collisions at  $\sqrt{s} = 7$  TeV / CMS Collaboration

A search for quark compositeness in the form of quark contact interactions, based on hadronic jet pairs (dijets) produced in proton-proton collisions at  $\sqrt{s}=7$  TeV, is described [...]

[arXiv:1010.4439](#) ; [CERN-PH-EP-2010-038](#) ; CMS-EXO-10-002. - 2010. - 24 p.

- 7) Search for Dijet Resonances in 7 TeV pp Collisions at CMS / CMS Collaboration

A search for narrow resonances in the dijet mass spectrum is performed using data corresponding to an integrated luminosity of 2.9 inverse pb collected by the CMS experiment at the LHC.

Upper limits at the 95% confidence level (CL) are presented on the product of the resonance cross section, branching fraction into dijets, and acceptance, separately for decays into quark-quark, quark-gluon, or gluon-gluon pairs.

[arXiv:1010.0203](#); [CERN-PH-EP-2010-035](#); [CMS-EXO-10-010](#).- Geneva : CERN, 2010 - 26 p. - Published in : [Phys. Rev. Lett. 105 \(2010\) 211801](#) [preprint](#) ; [APS Open Access article](#)

- 8) Observation of Long-Range, Near-Side Angular Correlations in Proton-Proton Collisions at the LHC / CMS Collaboration

Results on two-particle angular correlations for charged particles emitted in proton-proton collisions at center-of-mass energies of 0.9, 2.36, and 7 TeV are presented, using data collected with the CMS detector over a broad range of pseudorapidity ( $\eta$ ) and azimuthal angle ( $\phi$ ). Short-range correlations in  $\Delta\eta$ , which are studied in minimum bias events, are characterized using a simple "independent cluster" parametrization in order to quantify their strength (cluster size) and their extent in  $\eta$ (cluster decay width). [...]

[arXiv:1009.4122](#); [CERN-PH-EP-2010-031](#); [CMS-QCD-10-002](#).- Geneva : CERN, 2010 - Published in : [J. High Energy Phys. 09 \(2010\) 091](#) [Fulltext](#) ; [Springer Open Access article](#)

9) CMS Tracking Performance Results from Early LHC Operation / CMS Collaboration

The first LHC pp collisions at centre-of-mass energies of 0.9 and 2.36 TeV were recorded by the CMS detector in December 2009.

[arXiv:1007.1988](#) ; [CERN-PH-EP-2010-019](#) ; CMS-TRK-10-001. - 2010. - 29 p. [Preprint](#)

10) Transverse momentum and pseudorapidity distributions of charged hadrons in pp collisions at  $\sqrt{s} = 0.9$  and 2.36 TeV

Measurements of inclusive charged-hadron transverse-momentum and pseudorapidity distributions are presented for proton-proton collisions at  $\sqrt{s} = 0.9$  and 2.36 TeV. The data were collected with the CMS detector during the LHC commissioning in December 2009. [...]

[arXiv:1002.0621](#); [CMS-QCD-09-010](#); [CERN-PH-EP-2010-003](#).- Geneva : CERN, 2010 - 33 p. - Published in : *J. High Energy Phys.* 02 (2010) 041 [Preprint](#) ; [Supplemental material](#) ; [Figure\\_001a](#) ; [Figure\\_001b](#) ; [Figure\\_002a](#) ; [Figure\\_002b](#) ; [Figure\\_003a](#) ; [Figure\\_003b](#) ; [Figure\\_004a](#) ; [Figure\\_004b](#) ; [Figure\\_005a](#) ; [Figure\\_005b](#) ; [Figure\\_006a](#) ; [Figure\\_006b](#) ; [Figure\\_007a](#) ; [Figure\\_007b](#) ; [Springer Open Access article](#) ;

11) Commissioning and Performance of the CMS Pixel Tracker with Cosmic Ray Muons / CMS Collaboration

The pixel detector of the Compact Muon Solenoid experiment consists of three barrel layers and two disks for each endcap. The detector was installed in summer 2008, commissioned with charge injections, and operated in the 3.8 T magnetic field during cosmic ray data taking. [...]

[arXiv:0911.5434](#); [CMS-CFT-09-001](#).- Geneva : CERN, 2010 - 37 p. - Published in : *J. Instrum.* 5 (2010) T03007 [Preprint](#) ; [SISSA/IOP Open Access article](#)

## 12) Measurement of the Muon Stopping Power in Lead Tungstate / CMS Collaboration

A large sample of cosmic ray events collected by the CMS detector is exploited to measure the specific energy loss of muons in the lead tungstate of the electromagnetic calorimeter. The measurement spans a momentum range from 5 GeV/c to 1 TeV/c. [...]

[arXiv:0911.5397](#); [CMS-CFT-09-005](#).- 2010 - 31 p. - Published in : *J. Instrum.* 5 (2010) P03007 Preprint ; [SISSA/IOP Open Access article](#)

## 13) Performance of the CMS Level-1 Trigger during Commissioning with Cosmic Ray Muons and LHC beams / CMS Collaboration

The CMS Level-1 trigger was used to select cosmic ray muons and LHC beam events during data-taking runs in 2008, and to estimate the level of detector noise. This paper describes the trigger components used, the algorithms that were executed, and the trigger synchronisation. [...]

[arXiv:0911.5422](#); [CMS-CFT-09-013](#).- 2010 - 49 p. - Published in : *J. Instrum.* 5 (2010) T03002 Preprint ; [SISSA/IOP Open Access article](#)

## 14) Performance of CMS Muon Reconstruction in Cosmic-Ray Events / CMS Collaboration

The performance of muon reconstruction in CMS is evaluated using a large data sample of cosmic-ray muons recorded in 2008. Efficiencies of various high-level trigger, identification, and reconstruction algorithms have been measured for a broad range of muon momenta, and were found to be in good agreement with expectations from Monte Carlo simulation. [...]

[arXiv:0911.4994](#); [CMS-CFT-09-014](#).- 2010 - 47 p. - Published in : *J. Instrum.* 5 (2010) T03022 Preprint ; [SISSA/IOP Open Access article](#)

15) Commissioning and Performance of the CMS Silicon Strip Tracker with Cosmic Ray Muons / CMS Collaboration

During autumn 2008, the Silicon Strip Tracker was operated with the full CMS experiment in a comprehensive test, in the presence of the 3.8 T magnetic field produced by the CMS superconducting solenoid. Cosmic ray muons were detected in the muon chambers and used to trigger the readout of all CMS sub-detectors. [...]

[arXiv:0911.4996](#); [CMS-CFT-09-002](#).- 2010 - 45 p. - Published in : *J. Instrum.* 5 (2010) T03008 Preprint ; [SISSA/IOP Open Access article](#)

16) Performance of the CMS Cathode Strip Chambers with Cosmic Rays / CMS Collaboration

The Cathode Strip Chambers (CSCs) constitute the primary muon tracking device in the CMS endcaps. Their performance has been evaluated using data taken during a cosmic ray run in fall 2008. [...]

[arXiv:0911.4992](#); [CMS-CFT-09-011](#).- 2010 - 39 p. - Published in : *J. Instrum.* 5 (2010) T03018 Preprint ; [SISSA/IOP Open Access article](#)

17) Performance of the CMS Hadron Calorimeter with Cosmic Ray Muons and LHC Beam Data / CMS Collaboration

The CMS Hadron Calorimeter in the barrel, endcap and forward regions is fully commissioned. Cosmic ray data were taken with and without magnetic field at the surface hall and after installation in the experimental hall, hundred meters underground. [...]

[arXiv:0911.4991](#); [CMS-CFT-09-009](#).- 2010 - 35 p. - Published in : *J. Instrum.* 5 (2010) T03012 Preprint ; [SISSA/IOP Open Access article](#)

18) CMS Data Processing Workflows during an Extended Cosmic Ray Run / CMS Collaboration

The CMS Collaboration conducted a month-long data taking exercise, the Cosmic Run At Four Tesla, during October-November 2008, with the goal of commissioning the experiment for extended operation. With all installed detector systems participating, CMS recorded 270 million cosmic ray events with the solenoid at a magnetic field strength of 3.8 T. [...]

[arXiv:0911.4842](#); CMS-CFT-09-007.- 2010 - 43 p. - Published in : *J. Instrum.* 5 (2010) T03006 Preprint ; [SISSA/IOP Open Access article](#)

19) Performance of the CMS Drift Tube Chambers with Cosmic Rays / CMS Collaboration

Studies of the performance of the CMS drift tube barrel muon system are described, with results based on data collected during the CMS Cosmic Run at Four Tesla. For most of these data, the solenoidal magnet was operated with a central field of 3.8 T. [...]

[arXiv:0911.4855](#); CMS-CFT-09-012.- 2010 - 47 p. - Published in : *J. Instrum.* 5 (2010) T03015 Preprint ; [SISSA/IOP Open Access article](#)

20) Calibration of the CMS Drift Tube Chambers and Measurement of the Drift Velocity with Cosmic Rays / CMS Collaboration

This paper describes the calibration procedure for the drift tubes of the CMS barrel muon system and reports the main results obtained with data collected during a high statistics cosmic ray data-taking period. The main goal of the calibration is to determine, for each drift cell, the minimum time delay for signals relative to the trigger, accounting for the drift velocity within the cell. [...]

[arXiv:0911.4895](#); CMS-CFT-09-023.- 2010 - 39 p. - Published in : *J. Instrum.* 5 (2010) T03016 Preprint ; [SISSA/IOP Open Access article](#)

21) Commissioning of the CMS Experiment and the Cosmic Run at Four Tesla / CMS Collaboration

The CMS Collaboration conducted a month-long data-taking exercise known as the Cosmic Run At Four Tesla in late 2008 in order to complete the commissioning of the experiment for extended operation. The operational lessons resulting from this exercise were addressed in the subsequent shutdown to better prepare CMS for LHC beams in 2009. [...]

[arXiv:0911.4845](#); [CMS-CFT-09-008](#).- 2010 - 37 p. - Published in : *J. Instrum.* 5 (2010) T03001 Preprint ; [SISSA/IOP Open Access article](#)

22) Identification and Filtering of Uncharacteristic Noise in the CMS Hadron Calorimeter / CMS Collaboration

Commissioning studies of the CMS hadron calorimeter have identified sporadic uncharacteristic noise and a small number of malfunctioning calorimeter channels. Algorithms have been developed to identify and address these problems in the data. [...]

[arXiv:0911.4881](#); [CMS-CFT-09-019](#).- 2010 - 31 p. - Published in : *J. Instrum.* 5 (2010) T03014 Preprint ; [SISSA/IOP Open Access article](#)

23) Commissioning of the CMS High-Level Trigger with Cosmic Rays / CMS Collaboration

The CMS High-Level Trigger (HLT) is responsible for ensuring that data samples with potentially interesting events are recorded with high efficiency and good quality. This paper gives an overview of the HLT and focuses on its commissioning using cosmic rays. [...]

[arXiv:0911.4889](#); [CMS-CFT-09-020](#).- 2010 - 31 p. - Published in : *J. Instrum.* 5 (2010) T03005 Preprint ; [SISSA/IOP Open Access article](#)

24) Aligning the CMS Muon Chambers with the Muon Alignment System during an Extended Cosmic Ray Run / CMS Collaboration

The alignment system for the muon spectrometer of the CMS detector comprises three independent subsystems of optical and analog position sensors. It aligns muon chambers with respect to each other and to the central silicon tracker. [...]

[arXiv:0911.4770](#); CMS-CFT-09-017.- 2010 - 35 p. - Published in : *J. Instrum.* 5 (2010) T03019 Preprint ; [SISSA/IOP Open Access article](#)

25) Performance of the CMS drift-tube chamber local trigger with cosmic rays / CMS Collaboration

The performance of the Local Trigger based on the drift-tube system of the CMS experiment has been studied using muons from cosmic ray events collected during the commissioning of the detector in 2008. The properties of the system are extensively tested and compared with the simulation. [...]

[arXiv:0911.4893](#); CMS-CFT-09-022.- 2010 - 33 p. - Published in : *J. Instrum.* 5 (2010) T03003 Preprint ; [SISSA/IOP Open Access article](#)

26) Fine Synchronization of the CMS Muon Drift-Tube Local Trigger using Cosmic Rays / CMS Collaboration

The CMS experiment uses self-triggering arrays of drift tubes in the barrel muon trigger to perform the identification of the correct bunch crossing. The identification is unique only if the trigger chain is correctly synchronized. [...]

[arXiv:0911.4904](#); CMS-CFT-09-025.- 2010 - 33 p. - Published in : *J. Instrum.* 5 (2010) T03004 Preprint ; [SISSA/IOP Open Access article](#)

27) Performance of CMS hadron calorimeter timing and synchronization using test beam, cosmic ray, and LHC beam data / CMS Collaboration

This paper discusses the design and performance of the time measurement technique and of the synchronization systems of the CMS hadron calorimeter. Timing performance results are presented from the Cosmic Run At Four Tesla and LHC beam runs taken in the Autumn of 2008. [...]

[arXiv:0911.4877](#); [CMS-CFT-09-018](#).- 2010 - 33 p. - Published in : *J. Instrum.* 5 (2010) T03013 Preprint ; [SISSA/IOP Open Access article](#)

28) Alignment of the CMS Muon System with Cosmic-Ray and Beam-Halo Muons / CMS Collaboration

The CMS muon system has been aligned using cosmic-ray muons collected in 2008 and beam-halo muons from the 2008 LHC circulating beam tests. After alignment, the resolution of the most sensitive coordinate is 80 microns for the relative positions of superlayers in the same barrel chamber and 270 microns for the relative positions of endcap chambers in the same ring structure. [...]

[arXiv:0911.4022](#); [CMS-CFT-09-016](#).- 2010 - 41 p. - Published in : *J. Instrum.* 5 (2010) T03020 Preprint ; [SISSA/IOP Open Access article](#)

29) Performance Study of the CMS Barrel Resistive Plate Chambers with Cosmic Rays / CMS Collaboration

In October and November 2008, the CMS collaboration conducted a programme of cosmic ray data taking, which has recorded about 270 million events. The Resistive Plate Chamber system, which is part of the CMS muon detection system, was successfully operated in the full barrel. [...]

[arXiv:0911.4045](#); [CMS-CFT-09-010](#).- 2010 - 33 p. - Published in : *J. Instrum.* 5 (2010) T03017 Preprint ; [SISSA/IOP Open Access article](#)



30) Time Reconstruction and Performance of the CMS Electromagnetic Calorimeter / CMS Collaboration

The resolution and the linearity of time measurements made with the CMS electromagnetic calorimeter are studied with samples of data from test beam electrons, cosmic rays, and beam-produced muons. The resulting time resolution measured by lead tungstate crystals is better than 100 ps for energy deposits larger than 10 GeV. [...]

[arXiv:0911.4044](#); CMS-CFT-09-006.- 2010 - 27 p. - Published in : *J. Instrum.* 5 (2010) T03011 Preprint ; [SISSA/IOP Open Access article](#)

31) Precise Mapping of the Magnetic Field in the CMS Barrel Yoke using Cosmic Rays / CMS Collaboration Task physics.acc-ph

The CMS detector is designed around a large 4 T superconducting solenoid, enclosed in a 12000-tonne steel return yoke. A detailed map of the magnetic field is required for the accurate simulation and reconstruction of physics events in the CMS detector, not only in the inner tracking region inside the solenoid but also in the large and complex structure of the steel yoke, which is instrumented with muon chambers. [...]

[arXiv:0910.5530](#); CMS-CFT-09-015.- 2010 - 37 p. - Published in : *J. Instrum.* 5 (2010) T03021 Preprint ; [SISSA/IOP Open Access article](#)

32) Performance and Operation of the CMS Electromagnetic Calorimeter / CMS Collaboration Task physics.acc-ph

The operation and general performance of the CMS electromagnetic calorimeter using cosmic-ray muons are described. These muons were recorded after the closure of the CMS detector in late 2008. [...]

[arXiv:0910.3423](#); CMS-CFT-09-004.- 2010 - 39 p. - Published in : *J. Instrum.* 5 (2010) T03010 Preprint ; [SISSA/IOP Open Access article](#)

33) Alignment of the CMS Silicon Tracker during Commissioning with Cosmic Rays  
/ CMS Collaboration Task physics.acc-ph

The CMS silicon tracker, consisting of 1440 silicon pixel and 15148 silicon strip detector modules, has been aligned using more than three million cosmic ray charged particles, with additional information from optical surveys. The positions of the modules were determined with respect to cosmic ray trajectories to a precision of 3-4 microns RMS in the barrel and 3-14 microns RMS in the endcap in the most sensitive coordinate. [...]

[arXiv:0910.2505](#); [CMS-CFT-09-003](#).- 2010 - 41 p. - Published in : *J. Instrum.* 5 (2010) T03009 Preprint ; [SISSA/IOP Open Access article](#)

## DZero Virtual Organization

### Submitted:

- 1) [Search for WH production in  \$5.3 \text{ fb}^{-1}\$  of pp Collisions at the Fermilab Tevatron](#)  
Submitted 12/7/10: Phys. Lett. B, [arXiv.org:1012.0874](#) [plots](#)
- 2) [Measurement of the W Boson Helicity in Top Quark Decays using  \$5.4 \text{ fb}^{-1}\$  of pp Collision Data](#)  
Submitted 11/30/10: Phys. Rev. D, [arXiv.org:1011.6549](#) [plots](#)
- 3) [Search for Resonant WW and WZ Production in pp Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Submitted 11/30/10: Phys. Rev. Lett., [arXiv.org:1011.6278](#) [plots](#)
- 4) [Search for Neutral Higgs Bosons in the Multi-b-Jet Topology in  \$5.2 \text{ fb}^{-1}\$  of pp Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Submitted 11/8/10: Phys. Lett. B, [arXiv.org:1011.1931](#) [plots](#)
- 5) [A Measurement of the Ratio of Inclusive Cross Sections  \$\sigma\(\text{pp} \rightarrow Z + \text{b jet}\) / \sigma\(\text{pp} \rightarrow Z + \text{jet}\)\$  at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Submitted 10/29/10: Phys. Rev. Lett., [arXiv.org:1010.6203](#) [plots](#)
- 6) [Search for Single Vector-Like Quarks in pp Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Submitted 10/8/10: Phys. Rev. Lett., [arXiv.org:1010.1466](#) [plots](#)
- 7) [Precise Study of the Z/  \$\gamma^\*\$  Boson Transverse Momentum Distribution in pp Collisions Using a Novel Technique](#)  
Submitted 10/1/10: Phys. Rev. Lett., [arXiv.org:1010.0262](#) [plots](#)

- 8) [Search for Pair Production of the Scalar Top Quark in the Electron+Muon Final State](#)  
Submitted 9/27/10: Phys. Lett. B, [arXiv.org:1009.5950](#) [plots](#)
- 9) [High Mass Exclusive Diffractive Dijet Production in pp Collisions at  \$\sqrt{s} = 1.96\$  TeV](#)  
Submitted 9/13/10: Phys. Lett. B, [arXiv.org:1009.2444](#) [plots](#)

#### Accepted:

- 1) [Determination of the Width of the Top Quark](#)  
Accepted 12/9/10: Phys. Rev. Lett., [arXiv.org:1009.5686](#) [plots](#)
- 2) [The DØ Silicon Microstrip Tracker](#)  
Accepted 11/29/10: Nucl. Instrum. Methods in Phys. Res. Sect. A, (limited author- list paper), [arXiv.org:1005.0801](#) [photos and plots](#)
- 3) [Search for  \$ZH \rightarrow l^+l^-bb\$  Production in  \$4.2 \text{ fb}^{-1}\$  of pp Collisions at  \$\sqrt{s} = 1.96\$  TeV](#)  
Accepted 11/22/10: Phys. Rev. Lett., [arXiv.org:1008.3564](#) [plots](#)
- 4) [Search for a Heavy Neutral Gauge Boson in the Dielectron Channel with  \$5.4 \text{ fb}^{-1}\$  of pp Collisions at  \$\sqrt{s} = 1.96\$  TeV](#)  
Accepted 10/27/10: Phys. Lett. B, [arXiv.org:1008.2023](#) [plots](#)
- 5) [Measurement of the  \$WZ \rightarrow l\nu ll\$  Cross Section and Limits on Anomalous Triple Gauge Boson Couplings in pp Collisions at  \$\sqrt{s} = 1.96\$  TeV](#)  
Accepted 10/20/10: Phys. Lett. B, [arXiv.org:1006.0761](#) [plots](#)

#### Published:

- 1) [Search for Diphoton Events with Large Missing Transverse Energy in  \$6.3 \text{ fb}^{-1}\$  of pp Collisions at  \$\sqrt{s} = 1.96\$  TeV](#)  
Published 11/24/10: Phys. Rev. Lett. **105**, 221802 (2010), [arXiv.org:1008.2133](#) [plots](#)
- 2) [Search for New Fermions \("Quirks"\) at the Fermilab Tevatron Collider](#)  
Published 11/19/10: Phys. Rev. Lett. **105**, 211803 (2010), [arXiv.org:1008.3547](#) [plots](#)
- 3) [Search for Events with Leptonic Jets and Missing Transverse Energy in pp Collisions at  \$\sqrt{s} = 1.96\$  TeV](#)  
Published 11/19/10: Phys. Rev. Lett. **105**, 211802 (2010), [arXiv.org:1008.3356](#) [plots](#)

- 4) [Search for Sneutrino Production in  \$e\mu\$  Final States in  \$5.3 \text{ fb}^{-1}\$  of pp Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Published 11/5/10: Phys. Rev. Lett. **105**, 191802 (2010), [arXiv.org:1007.4835](#) [plots](#)
- 5) [Measurement of  \$t\bar{t}\$  Production in the  \$\tau\$ +Jets Topology using pp Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Published 10/27/10: Phys. Rev. D **82**, 071102 (2010), [arXiv.org:1008.4284](#) [plots](#)
- 6) [Search for the Rare Decay  \$B\_s^0 \rightarrow \mu^+ \mu^-\$](#)   
Published 9/18/10: Phys. Lett. B **693**, 539 (2010), [arXiv.org:1006.3469](#) [plots](#)
- 7) [Measurement of the Dijet Invariant Mass Cross Section in pp Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Published 9/17/10: Phys. Lett. B **693**, 531 (2010), [arXiv.org:1002.4594](#) [plots](#)
- 8) [Measurement of the Normalized  \$Z/\gamma^\* \rightarrow \mu^+ \mu^-\$  Transverse Momentum Distribution in pp Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Published 9/16/10: Phys. Lett. B **693**, 522 (2010), [arXiv.org:1006.0618](#) [plots](#)
- 9) [Dependence of the  \$t\bar{t}\$  Production Cross Section on the Transverse Momentum of the Top Quark](#)  
Published 9/15/10: Phys. Lett. B **693**, 515 (2010), [arXiv.org:1001.1900](#) [plots](#)
- 10) [Search for Scalar Bottom Quarks and Third-Generation Leptoquarks in pp Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Published 8/19/10: Phys. Lett. B **693**, 95 (2010), [arXiv.org:1005.2222](#) [plots](#)
- 11) [Search for Flavor Changing Neutral Currents via Quark-Gluon Couplings in Single Top Quark Production Using  \$2.3 \text{ fb}^{-1}\$  of pp Collisions](#)  
Published 8/11/10: Phys. Lett. B **693**, 81 (2010), [arXiv.org:1006.3575](#) [plots](#)
- 12) [Measurement of the  \$t\bar{t}\$  Cross Section using High-Multiplicity Jet Events](#)  
Published 8/20/10: Phys. Rev. D **82**, 032002 (2010), [arXiv.org:0911.4286](#) [plots](#)
- 13) [Evidence for an Anomalous Like-Sign Dimuon Charge Asymmetry](#) **PRL Editors' Suggestion** ["Physics" Viewpoint article](#)  
Published 8/16/10: Phys. Rev. Lett. **105**, 081801 (2010), [arXiv.org:1007.0395](#) [plots](#) [press coverage, seminars](#)
- 14) [Evidence for an Anomalous Like-Sign Dimuon](#) ["Physics" Viewpoint article](#)  
Published 8/16/10: Phys. Rev. D **82**, 032001, (2010), [arXiv.org:1005.2757](#) [plots](#) [press coverage, seminars](#)
- 15) [Search for CP Violation in  \$B\_s^0 \rightarrow \mu^+ D\_s^- X\$  Decays in pp Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Published 7/26/10: Phys. Rev. D **82**, 012003 (2010), [arXiv.org:0904.3907](#) [plots](#)

- 16) [Combined Tevatron Upper Limit on  \$gg \rightarrow H \rightarrow W^+W^-\$  and Constraints on the Higgs Boson Mass in Fourth-Generation Fermion Models](#)  
Co-authors: DØ and CDF collaborations. Published 7/15/10: Phys. Rev. D **82**, 011102 (2010), [arXiv.org:1005.3216](#) [plots](#)
- 17) [Search for Randall-Sundrum Gravitons in the Dielectron and Diphoton Final States with  \$5.4 \text{ fb}^{-1}\$  of Data from pp Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Published 6/17/10: Phys. Rev. Lett. **104**, 241802 (2010), [arXiv.org:1004.1826](#) [plots](#)
- 18) [Measurement of Direct Photon Pair Production Cross Sections in pp Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Published 5/12/10: Phys. Lett. B **690**, 108 (2010), [arXiv.org:1002.4917](#) [plots](#)
- 19) [Search for Single Top Quarks in the Tau+Jets Channel using  \$4.8 \text{ fb}^{-1}\$  of pp Collision Data](#)  
Published 5/6/10: Phys. Lett. B **690**, 5 (2010), [arXiv.org:0912.1066](#) [plots](#)
- 20) [Search for the Associated Production of a b Quark and a Neutral Supersymmetric Higgs Boson that Decays into  \$\tau\$  Pairs](#), Phys. Rev. Lett. **104**, 151801 (2010)
- 21) [Double Parton Interactions in  \$\gamma\$ +3 Jet Events in p̄p Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Phys. Rev. D **81**, 052012 (2010)
- 22) [Search for the Standard Model Higgs Boson in the  \$ZH \rightarrow \nu\bar{\nu}bb\$  Channel in  \$5.2 \text{ fb}^{-1}\$  of p̄p Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#) , Phys. Rev. Lett. **104**, 071801 (2010)
- 23) [Search for Higgs Boson Production in Dilepton and Missing Energy Final States with  \$5.4 \text{ fb}^{-1}\$  of p̄p Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#) , Phys. Rev. Lett. **104**, 061804 (2010)
- 24) [Combination of Tevatron Searches for the Standard Model Higgs Boson in the  \$W^+W^-\$  Decay Mode](#), Phys. Rev. Lett. **104**, 061802 (2010)
- 25) [Search for a Resonance Decaying into WZ Boson Pairs in pp Collisions](#)  
Phys. Rev. Lett. **104**, 061801 (2010)
- 26) [Measurement of the t-Channel Single Top Quark Production Cross Section](#)  
Phys. Lett. B **682**, 363 (2010)
- 27) [Measurement of  \$Z/\gamma^\* + \text{jet} + X\$  Angular Distributions in p̄p Collisions at  \$\sqrt{s} = 1.96 \text{ TeV}\$](#)   
Phys. Lett. B **682**, 370 (2010)

## ENGAGE Virtual Organization

- 1) D. Chahal, S. Stuart, S. Goasguen and C. J. Trout “Automated, Parallel Optimization of Stochastic Functions using a Modified Simplex Algorithm” 2010 e-science workshop on parallel optimization and parameter optimization. Australia, December 2010.
- 2) R. Wartel, T. Cass, B. Moreira, E. Roche, U. Schwickerath and S. Goasguen “Image Distribution in Large Scale Cloud Providers” 2nd IEEE Cloud Computing Conference, Indianapolis IN, October 30<sup>th</sup> - November 3<sup>rd</sup> 2010.
- 3) M. Murphy, L. Stout and S. Goasguen "Involuntary Computing: Hacking the Cloud" Emerging Research Paper, 2nd IEEE Cloud Computing Conference, Indianapolis IN, October 30<sup>th</sup> - November 3<sup>rd</sup> 2010.
- 4) P. Ruth, J. Ree, D. Xu and S. Goasguen “Autonomic Live Adaptation of Virtual Networked Environments in a Multi-Domain Infrastructure” Submitted to Journal of Internet Service and Applications, Springer, November 2010.
- 5) L. Stout, M. Walker, J. Lauret and S. Goasguen “Using Kestrel and XMPP to Run Scientific Applications in the Cloud” Submitted to the journal of Scientific Programming, November 2010.
- 6) J. Lauret, M. Walker, S. Goasguen and L. Hadju “From Grids to Cloud the STAR Experience” Submitted to Journal of Physics, October 2010
- 7) S. Goasguen “Building IaaS Clouds for Scientific Applications”, Petascale Data Analytics on Clouds: Trends, Challenges and Opportunities, ACM/IEEE Supercomputing conference, November 14th, 2010, New Orleans, LA, USA.
- 8) J. Lauret, M. Walker, S. Goasguen “When STAR meets the Clouds: Virtualization and Cloudified Grids Experience” Computing in High Energy and Nuclear Physics, CHEP 2010, Taipei, Taiwan
- 9) S. Goasguen, M. Guijaro, B. Moreira, E. Roche, U. Schwickerath, R. Wartel “Integration of Virtual Machines in the Batch System at CERN, CHEP 2010, Taipei, Taiwan4.
- 10) M. Fenn, J. Lauret and S. Goasguen “Contextualization in Practice: The Clemson Experience” ACAT, Jaipur, India February 2010.

- 11) M. A. Murphy, L. Abraham, M. Fenn and S. Goasguen “Autonomic Clouds on the Grid”  
Journal of Grid Computing Volume 8, Number 1 (March 2010), pages 1-18.
- 12) V. Patel, J. McGregor and S. Goasguen “SysFlow: A Workflow Engine for Service  
Based Infrastructures” IEEE Sysconf 2010, May 2010, CA
- 13) D. Rosato, O. Dybenko, V. Ratnaswamy, D. Horntrop, Nathan Andryscio, Xavier  
Tricoche, Lou Kondic, "Microstructure Evolution in Density Relaxation by Tapping",  
Frontiers in Applied and Computational Mathematics, New Jersey Institute of  
Technology, Newark, NJ, USA, May 2010.
- 14) L. Stout, M. Fenn, M. Murphy and S. Goasguen "Scaling Virtual Organization Clusters  
over a Wide Area Network using the Kestrel Workload Management System" VTDC,  
June 2010, Chicago
- 15) L. Abraham, M. Fenn, M. Murphy and S. Goasguen “Self-Provisioned Hybrid Clouds”  
7th IEEE International Conference on Autonomic Computing, June 7-11, 2010, D.C,  
USA
- 16) D. Rosato, O. Dybenko, V. Ratnaswamy, D. Horntrop, Nathan Andryscio,  
Xavier Tricoche, Lou Kondic, "Microstructure Evolution in Density Relaxation by  
Tapping", Gordon Research Conference: Granular & Granular Fluid Flow, Colby  
College, Waterville, Maine, USA, June 2010
- 17) D. Rosato, O. Dybenko, D. J. Horntrop, V. Ratnaswamy, L. Kondic, “Microstructure  
Evolution in Density Relaxation by Tapping”, to appear in Phys. Rev. E
- 18) Tao Wu, X. Sheldon Wang, Barry Cohen and Hongya Ge. Molecular Modeling of  
Normal and Sick Hemoglobins, IJMCE. Accepted.
- 19) Feng Zhao, Jian Peng and Jinbo Xu. Fragment-free Approach to Protein Folding Using  
Conditional Neural Fields. Bioinformatics (Proceedings of ISMB 2010), 2010.
- 20) Feng Zhao, Jian Peng, Joe DeBartolo, Karl F. Freed, Tobin R. Sosnick and Jinbo Xu. A  
probabilistic and continuous model of protein conformational space for templatefree  
modeling. Journal of Computational Biology, 2010. In Press.

## GLOW Virtual Organization

- 1) Methods for point-source analysis with temporal structure in high-energy neutrino telescopes, J. Braun, et al., *Astropart.Phys.*33:175-181,2010. e-Print: arXiv:0912.1572
- 2) Teague, B., Waterman, M.S., Goldstein, S., Potamouisis, K., Zhou, S., Reslewic, S., Sarkar, D., Valouev, A., Churas, C., Kidd, J., Kohn, S., Runnhein, R., Lamers, C., Forrest, D., Newton, M.A., Eichler, E.E., Kent-First, M., Surti, U., Livny, M., and Schwartz, D.C. High-resolution human genome structure by single molecule analysis. *Proc. Natl. Acad. Sci. USA*, epub May 31, 2010 doi:10.1073/pnas.0914638107.
- 3) Ma, L., van der Doe, H. C., Borkovich, K.A., Coleman, J.C., Daboussi, M.-J., Di Pietro, A., Dufresne, M., Freitag, M., Grabherr, M., Henrissat, B., Houterman, P.M., Kang, S., Shim, W.-B., Woloshuk, C., Xie, X., Xu, J.-R., Antoniw, J., Baker, S.E., Bluhm, B.H., Breakspear, A., Brown, D.W., Butchko, R.A.E., Chapman, S., Coulson, R., Coutinho, P.M., Danchin, E.G.J., Diener, A., Gale, L.R., Gardiner, D.M., Goff, S., Hammond-Kosack, K.E., Hilburn, K., Hua-Van, A., Jonkers, W., Kazan, K., Kodira, C.D., Michael Koehrsen, M., Kumar, L., Lee, Y.-H., Li, L. Manners, J.M., Miranda-Saavedra, D., Mukherjee, M., Park, G., Park, J., Park, S.-Y., Robert H. Proctor, Regev, A., Ruiz-Roldan, M.C., Sain, D., Sakthikumar, S., Sykes, S., Schwartz, D.C., Turgeon, B.G., Wapinski, I., Yoder, O., Young, S., Zeng, Q., Zhou, S., Galagan, J., Cuomo, C.A., Kistler, H.C and Rep, M.. Comparative genomics reveals mobile pathogenicity chromosomes in *Fusarium*. *Nature*, epubl March 18, 2010.
- 4) Interpolation in the Directed Assembly of Block Copolymers on Nanopatterned Substrates: Simulation and Experiments. Francois A. Detcheverry, Guoliang Liu, Paul F. Nealey, Juan J. de Pablo. *Macromolecules* 2010 43 (7), 3446-3454
- 5) Remediation of Line Edge Roughness in Chemical Nanopatterns by the Directed Assembly of Overlying Block Copolymer Films. Mark P. Stoykovich, Kostas Ch. Daoulas, Marcus Miller, Huiman Kang, Juan J. de Pablo, Paul F. Nealey. *Macromolecules* 2010 43 (5), 2334-2342
- 6) 2D IR Line Shapes Probe Ovispirin Peptide Conformation and Depth in Lipid Bilayers. Ann Marie Woys, Yu-Shan Lin, Allam S. Reddy, Wei Xiong, Juan J. de Pablo, James L. Skinner, Martin T. Zanni. *Journal of the American Chemical Society* 2010 132 (8), 2832-2838



- 7) "Heterogeneous dynamics during deformation of a polymer glass". Robert A. Riggleman, Hau-Nan Lee, Mark D. Ediger, and Juan J. de Pablo. *Soft Matter* 6(2):287-291, January 2010. [doi:10.1039/b912288e ]
  
- 8) "Antiplastization and the elastic properties of glass-forming polymer liquids". Robert A. Riggleman, Jack F. Douglas, and Juan J. de Pablo. *Soft Matter* 6(2):292-304, January 2010. [doi:10.1039/b915592a ]
  
- 9) "Human embryonic stem cell-derived keratinocytes exhibit an epidermal transcription program and undergo epithelial morphogenesis in engineered tissue contacts". Christian M. Metallo, Samira M. Azarin, Laurel E. Moses, Lin Ji, Juan J. de Pablo, and Sean P. Palecek. *Tissue Engineering Part A* 16(1):213-223, January 2010. [doi:10.1089/ten.tea.2009.0325 ]
  
- 10) "Simulations of theoretically informed coarse grain models polymeric systems". Francois A. Detcheverry, Darin Q. Pike, Paul F. Nealey, Marcus Meuller, and Juan J. de Pablo. *Faraday Discussions* 144():111-125, January 2010. [doi:10.1039/b902283j ]
  
- 11) GNNQQNY - Investigation of Early Steps during Amyloid Formation. Allam S. Reddy, Manan Chopra, Juan J. de Pablo. *Biophysical Journal* - 17 March 2010 (Vol. 98, Issue 6, pp. 1038-1045) [doi:10.1016/j.bpj.2009.10.057]
  
- 12) Solution Structures of Rat Amylin Peptide: Simulation, Theory, and Experiment. Allam S. Reddy, Lu Wang, Yu-Shan Lin, Yun Ling, Manan Chopra, Martin T. Zanni, James L. Skinner, Juan J. De Pablo. *Biophysical Journal* - 3 February 2010 (Vol. 98, Issue 3, pp. 443-451) [doi:10.1016/j.bpj.2009.10.029]
  
- 13) Disruption and formation of surface salt bridges are coupled to DNA binding in integration host factor (IHF): acomputational analysis, L. Ma, M. T. Record, Jr., N. Sundlass, R. T. Raines and Q. Cui, {it J. Mol. Biol.}, Submitted
  
- 14) An implicit solvent model for SCC-DFTB with Charge-Dependent Radii, G. Hou, X. Zhu and Q. Cui, {it J. Chem. Theo. Comp.}, Submitted
  
- 15) Sequence-dependent interaction of  $\beta$ -peptides with membranes, J. Mondal, X. Zhu, Q. Cui and A. Yethiraj, {it J. Am. Chem. Soc.}, Submitted

- 16) A new coarse-grained model for water: The importance of electrostatic interactions, Z. Wu, Q. Cui and A. Yethiraj, *J. Phys. Chem. B* Submitted
- 17) How does bone sialoprotein promote the nucleation of hydroxyapatite? A molecular dynamics study using model peptides of different conformations, Y. Yang, Q. Cui, and N. Sahai, *J. Langmuir*, Submitted
- 18) Preferential interactions between small solutes and the protein backbone: A computational analysis, L. Ma, L. Pegram, M. T. Record, Jr., Q. Cui, *J. Biochem.*, 49, 1954-1962 (2010)
- 19) Establishing effective simulation protocols for  $\beta$ - and  $\alpha/\beta$ -peptides. III. Molecular Mechanical (MM) model for a non-cyclic  $\beta$ -residue, X. Zhu, P. K\"onig, M. Hoffman, A. Yethiraj and Q. Cui, *J. Comp. Chem.*, In press (DOI: 10.1002/jcc.21493)

### HCC Virtual Organization

- 1) "PROFESS: A PROtein Function, Evolution, Structure and Sequence database" Thomas Triplet, Matthew D. Shortridge, Mark A. Griep, Jaime L. Stark, Robert Powers and Peter Revesz. It has been accepted with revisions to the journal Database.

### LIGO Collaboration

- 1) S5Bursts LSC submitted for publication [arXiv:1011.1357](https://arxiv.org/abs/1011.1357) 201011 *A search for gravitational waves associated with the August 2006 timing glitch of the Vela pulsar*
- 2) S5/VSR1/A5Bursts LSC, Virgo, external submitted for publication [arXiv:1011.4079](https://arxiv.org/abs/1011.4079) 201011 *Search for Gravitational Wave Bursts from Six Magnetars*
- 3) S5/VSR1 Bursts LSC, Virgo *Astrophys. J.* 715 (2010) 1438  
arXiv:0908.3824, 200908. Search for gravitational-wave bursts associated with gamma-ray bursts using data from LIGO Science Run 5 and Virgo Science Run 1.
- 4) S5/VSR1 CW LSC, Virgo *Astrophys. J.* 713 (2010) 671 arXiv:0909.3583, 200909. Searches for gravitational waves from known pulsars with S5 LIGO data.

- 5) S5/VSRI CBC LSC, Virgo Astrophys. J. 715 (2010) 1453  
arXiv:1001.0165, 201001. Search for gravitational-wave inspiral signals associated with short Gamma-Ray Bursts during LIGO's fifth and Virgo's first science run.
- 6) S5/VSRI Bursts LSC, Virgo Phys. Rev. D 81 (2010) 102001  
arXiv:1002.1036, 201001. All-sky search for gravitational-wave bursts in the first joint LIGO-GEO-Virgo run.

### MiniBooNE Virtual Organization

- 1) A.A. Aguilar-Arevalo et al., [First Measurement of the Muon Neutrino Charged Current Quasielastic Double Differential Cross Section](#), arXiv:1002.2680 [hep-ex], Phys. Rev. D81, 092005 (2010), [Data release](#).
- 2) A.A. Aguilar-Arevalo et al., ["Measurement of  \$\nu\_\mu\$  and  \$\bar{\nu}\_\mu\$  induced neutral current single  \$\pi^0\$  production cross sections on mineral oil at  \$E\_\nu \sim \mathcal{O}\(1 \text{ GeV}\)\$ "](#), arXiv:0911.2063 [hep-ex], Phys. Rev. D81, 013005 (2010) [Data release](#).
- 3) A.A. Aguilar-Arevalo et al., ["A Search for Core-Collapse Supernovae using the MiniBooNE Neutrino Detector"](#), arXiv:0910.3182 [hep-ex], Phys. Rev. D81, 032001 (2010)

### MINOS Virtual Organization

- 1) P. Adamson et al.(152 authors), "Neutrino and antineutrino inclusive charged-current cross section measurements with the MINOS near detector," Fermilab-Pub-09-468-E, Phys.Rev.D81:072002(2010) (Vol.81, No.7, 8 April 2010); in arXiv:0910.2201 and MINOS-doc-6252.
- 2) E.W. Grashorn, J.K. de Jong, M.C. Goodman, A. Habig, M.L. Marshak, S. Mufson, S. Osprey, P.Schreiner, "The atmospheric charged kaon/pion ratio using seasonal variation methods," Fermilab-Pub-10-058-E, Astroparticle Physics 33,140-145 (2010) (Issue 3, 6 March 2010); in arXiv:0909.5382 and MINOS-doc-6236.
- 3) P. Adamson et al.(149 authors), "Search for sterile neutrino mixing in the MINOS long-baseline experiment," Fermilab-Pub-09-650-E, Phys.Rev.D81:052004 (2010)(11 March 2010); also in arXiv:1001.0336 and MINOS-doc-6028.
- 4) P. Adamson et al. (154 authors), "Observation of muon intensity variations by season with the MINOS far detector," Fermilab-Pub-09-427-E, Phys.Rev.D81:012001 (2010), 6 January 2010; in arXiv:0909.4012 and MINOS-doc-3719.

## **NanoHUB Virtual Organization**

- 1) Michael McLennan, Rick Kennell, "HUBzero: A Platform for Dissemination and Collaboration in Computational Science and Engineering," Computing in Science and Engineering, vol. 12, no. 2, pp. 48-53, Mar./Apr. 2010, doi:10.1109/MCSE.2010.41 (URL: <http://www.computer.org/portal/web/csdl/doi/10.1109/MCSE.2010.41>)
- 2) Alejandro Strachan, Gerhard Klimeck, Mark Lundstrom, "Cyber-Enabled Simulations in Nanoscale Science and Engineering," Computing in Science and Engineering, vol. 12, no. 2, pp. 12-17, Mar./Apr. 2010, doi:10.1109/MCSE.2010.38 (URL: <http://www.computer.org/portal/web/csdl/doi/10.1109/MCSE.2010.38>)
- 3) Gerhard Klimeck, Mathieu Luisier, "Atomistic Modeling of Realistically Extended Semiconductor Devices with NEMO and OMEN," Computing in Science and Engineering, vol. 12, no. 2, pp. 28-35, Mar./Apr. 2010, doi:10.1109/MCSE.2010.32 (URL: <http://www.computer.org/portal/web/csdl/doi/10.1109/MCSE.2010.32>)

## **NYSgrid Virtual Organization**

- 1) A.J. Schultz and D.A. Kofke, "Virial coefficients of model alkanes", J. Chem. Phys (submitted).

## **SBGrid Virtual Organization**

- 1) Stokes-Rees, I and Sliz, P. Protein structure determination by exhaustive search of Protein Data Bank derived databases. Proc Natl Acad Sci USA (2010) doi:10.1073/pnas.1012095107.
- 2) Stokes-Rees, I and Sliz, P. Compute and data management strategies for grid deployment of high throughput protein structure studies. 3rd IEEE workshop on Many-Task Computing on Grids and Supercomputers. (2010)
- 3) Lazarus, M.B, Nam, Y., Jiang, J., Sliz, P., Walker, S. Structure of human O-GlcNAc transferase and its complex with a peptide substrate. Nature (2010) in press.

## STAR Virtual Organization

- 1) [Measurement of the Bottom contribution to non-photon electron production in p+p collisions at  \$\sqrt{s}=200\$  GeV](#)  
*Submitted Jul. 7, 2010 , published Nov. 12, 2010*  
Phys. Rev. Lett. 105 (2010) 202301  
e-Print Archives (1007.1200) : [Abstract](#) | [PS](#) | [PDF](#)  
SLAC-Spires HEP: [Entry](#)|[Cited by](#)|[Citebase](#)  
Journal article: [Phys. Rev. Lett. server](#)
- 2) [Balance Functions from Au+Au, d+Au, and p+p Collisions at  \$\sqrt{s\_{NN}} = 200\$  GeV](#)  
*Submitted May. 13, 2010*  
e-Print Archives (1005.2307) : [Abstract](#) | [PS](#) | [PDF](#)
- 3) [Higher Moments of Net-proton Multiplicity Distributions at RHIC](#)  
*Submitted Apr. 29, 2010*  
e-Print Archives (1004.4959) : [Abstract](#) | [PS](#) | [PDF](#)  
Data and figures: [click here](#)
- 4) [Azimuthal di-hadron correlations in d+Au and Au+Au collisions at  \$\sqrt{s\_{NN}}=200\$  GeV from STAR](#)  
*Submitted Apr. 14, 2010*  
e-Print Archives (1004.2377) : [Abstract](#) | [PS](#) | [PDF](#)  
Data and figures: [click here](#)
- 5) [Pion femtoscopy in p+p collisions at  \$\sqrt{s}=200\$  GeV](#)  
*Submitted Apr. 6, 2010*  
e-Print Archives (1004.0925) : [Abstract](#) | [PS](#) | [PDF](#)  
Data and figures: [click here](#)
- 6) [Longitudinal scaling property of the charge balance function in Au + Au collisions at 200 GeV](#)  
*Submitted Feb. 12, 2010*  
e-Print Archives (arXiv:1002.1641) : [Abstract](#) | [PS](#) | [PDF](#)  
Data and figures: [click here](#)
- 7) [Charged and strange hadron elliptic flow in Cu+Cu collisions at  \$\sqrt{s\_{NN}} = 62.4\$  and 200 GeV](#)

*Submitted Jan. 28, 2010 , published Apr. 9, 2010*

Phys. Rev. C 81 (2010) 44902

e-Print Archives (arXiv:1001.5052) : [Abstract](#) | [PS](#) | [PDF](#)

SLAC-Spires HEP: [Entry](#)|[Cited by](#)|[Citebase](#)

Journal article: [Phys. Rev. C server](#)

Data and figures: [click here](#)

8) [Upsilon cross section in p+p collisions at  \$\sqrt{s} = 200\$  GeV](#)

*Submitted Jan. 15, 2010*

e-Print Archives (1001.2745) : [Abstract](#) | [PS](#) | [PDF](#)

Data and figures: [click here](#)

9) [Three-particle coincidence of the long range pseudorapidity correlation in high energy nucleus-nucleus collisions](#)

*Submitted Jan. 8, 2010*

e-Print Archives (0912.3977) : [Abstract](#) | [PS](#) | [PDF](#)

Data and figures: [click here](#)

10) [Inclusive  \$\pi^0\$ , eta, and direct photon production in p+p and d+Au collisions at  \$\sqrt{s\_{NN}} = 200\$  GeV](#)

*Submitted Dec. 18, 2009*

e-Print Archives (0912.3838) : [Abstract](#) | [PS](#) | [PDF](#)

11) [Studying Parton Energy Loss in Heavy-Ion Collisions via Direct-Photon and Charged-Particle Azimuthal Correlations](#)

*Submitted Dec. 9, 2009*

e-Print Archives (0912.1871) : [Abstract](#) | [PS](#) | [PDF](#)

Data and figures: [click here](#)

12) [Observation of  \$\pi^+\pi^-\pi^+\pi^-\$  photoproduction in ultraperipheral heavy-ion collisions at  \$\sqrt{s\_{NN}} = 200\$  GeV at the STAR detector](#)

*Submitted Dec. 4, 2009 , published Apr. 2, 2010*

Phys. Rev. C 81 (2010) 44901

e-Print Archives (0912.0604) : [Abstract](#) | [PS](#) | [PDF](#)

SLAC-Spires HEP: [Entry](#)|[Cited by](#)|[Citebase](#)

Journal article: [Phys. Rev. C server](#)

Data and figures: [click here](#)

- 13) [Observation of an Antimatter Hypernucleus](#)  
*Submitted Oct. 29, 2009 , published Mar. 4, 2010*  
Science 328 (2010) 58  
e-Print Archives (1003.2030) : [Abstract](#) | [PS](#) | [PDF](#)  
SLAC-Spires HEP: [Entry](#)|[Cited by](#)|[Citebase](#)  
Journal article: [Science server](#)  
Data and figures: [click here](#)
- 14) [Identified particle production, azimuthal anisotropy, and interferometry measurements in Au+Au collisions at  \$\sqrt{s\_{NN}} = 9.2\$  GeV](#)  
*Submitted Sep. 23, 2009 , published Feb. 26, 2010*  
Phys. Rev. C 81 (2010) 24911  
e-Print Archives (0909.4131) : [Abstract](#) | [PS](#) | [PDF](#)  
SLAC-Spires HEP: [Entry](#)|[Cited by](#)|[Citebase](#)  
Journal article: [Phys. Rev. C server](#)  
Data and figures: [click here](#)