

CERN-2005-002  
Volume 1  
4 April 2005

ORGANISATION EUROPÉENNE POUR LA RECHERCHE NUCLÉAIRE  
**CERN** EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

**COMPUTING IN HIGH ENERGY AND NUCLEAR PHYSICS 2004**

Congress Centre, Interlaken, Switzerland  
27 September – 1 October 2004

**PROCEEDINGS**

Editors: A. Aimar,  
J. Harvey, N. Knoors

**Volume 1**

GENEVA  
2005

842	BaBar Bookkeeping - a distributed meta-data catalog of the BaBar event store <i>D. Smith, T. Adye, A. Forti, A. Ceseracciu, D. Bukin, G. Dubois-Felsmann, D. Hutchcroft, P. Jackson, D. Kovalskyi, W. Roethel</i>	915
846	A Lightweight Monitoring and Accounting System for LHCb DC'04 Production <i>M. Sanchez Garcia, V. Garonne, R. Graciani Diaz, J. Saborido Silva, R. Vizcaya Carrillo</i>	921
850	Development and use of MonALISA high level monitoring services for the star unified Meta-Scheduler <i>E. Efsthadiadis, J. Lauret, I. Legrand, L. Hajdu</i>	924
854	DIRAC - The Distributed MC Production and Analysis for LHCb <i>A. Tsaregorodtsev, V. Garonne, I. Stokes-Rees, R. Graciani, M. Sanchez-Garcia, J. Closier, M. Frank, G. Kuznetsov, P. Charpentier, J. Blouw, N. Brook, U. Egede, M. Gandelman, I. Korolko, G. Patrick, A. Pickford, V. Romanovski, J. Saborido-Silva, A. Soroko, M. Tobin, V. Vagnoni, M. Witek, R. Bernet</i>	928
860		
864	<b>Track 5 - Distributed Computing Systems and Experiences</b>	<b>933</b>
868	The evolution of the distributed Event Reconstruction Control System in BaBar <i>A. Ceseracciu, T. Pulliam</i>	935
872	Production Management Software for the CMS Data Challenge <i>J. Andreeva, N. Darnenov, V. Lefebure, P. Garcia-Abia, A. Anzar, G. Graham, G. Guglielmo, N. Ratnikova, A. Fanfani, C. Grandi, T. Wildish</i>	939
876	ATLAS Production System in ATLAS Data Challenge 2 <i>L. Goossens, K. De</i>	943
879	A GRID approach for Gravitational Waves Signal Analysis with a Multi-Standard Farm Prototype <i>S. Pardi, F. Acernese, F. Barone, R. De Rosa, A. Eleuteri, R. Esposito, P. Mastroserio, L. Milano, G. Russo, L. Giordano, F. Taurino, G. Tortone</i>	947
883	The architecture of the AliEn system <i>P. Buncic, A. Peters, P. Saiz, J. Grosse-Oetringhaus</i>	951
887	A distributed, Grid-based analysis system for the MAGIC telescope <i>H. Kornmayer, C. Bigongiari, M. Hardt, M. Kunze, C. Bigoriari, M. Mazzucato, A. Deangelis, G. Cabras, A. Forti, M. Frailis, M. Piraccini, M. Delfino</i>	955
891		
895	HEP Applications Experience with the European DataGrid Middleware and Testbed <i>S. Burke, F. Harris, A. Tsaregorodtsev, K. Bos, D. Groep, W. Van Leeuwen, P. Capiluppi, A. Fanfani, C. Grandi, I. Augustin, R. Barbera, E. Luppi, G. Negri, L. Perini, S. Resconi, M. Reale, A. De Salvo, S. Bagnasco, P. Cerello, O. Smirnova, F. Carminati, O. Maroney, F. Brochu, D. Colling, I. Stokes-Rees, J. Closier, D. Boutigny, J. Blaising, V. Garonne</i>	959
	The Open Science Grid (OSG) <i>R. Pordes</i>	963
	Use of Condor and GLOW for CMS Simulation Production <i>S. Dasu, V. Puttabuddhi, S. Rader, D. Bradley, M. Livny, W. Smith</i>	967
899		
903	Application of the SAMGrid Test-Harness for Performance Evaluation and Tuning of a Distributed Cluster Implementation of Data Handling Services <i>A. Lyon, M. Leslie, M. Leslie, S. Stonjek, A. Baranovski, I. Terekhov, V. Bartsch, S. Belforte, L. Loebel-Carpenter, U. Kerzel, J. Trumbo, G. Garzoglio, A. Kreymer, R. Herber, R. Illingworth, R. Kennedy, L. Lueking, F. Ratnikov, R. St. Denis, S. White, W. Merritt, A. Sill, M. Burgon-Lyon</i>	970
907	Job Submission & Monitoring on the PHENIX Grid* <i>D. Morrison, R. Lacey, M. Reuter, I. Sourikova, T. Thomas, A. Withers, A. Shevel, B. Jacak</i>	974
911	The Condor based CDF CAF <i>F. Wuerthwein, E. Lipeles, M. Neubauer, I. Sfiligoi</i>	978

# Job submission & monitoring on the PHENIX Grid\*

A. Shevel<sup>1†</sup>, B. Jacak<sup>2</sup>, R. Lacey<sup>1</sup> and M. Reuter<sup>2</sup>, D. Morrison<sup>3</sup>, I. Sourikova<sup>3</sup>, A. Withers<sup>3</sup> and T. Thomas<sup>4</sup>  
<sup>1</sup>Dept. Chemistry, State University of New York at Stony Brook, NY 11794, USA  
<sup>2</sup>Dept. Physics, State University of New York at Stony Brook, NY 11794, USA  
<sup>3</sup>Dept. Physics, Brookhaven National Laboratory, NY 11973, USA  
<sup>4</sup>Dept. Physics, State University of New Mexico, NM 87131, USA

## Abstract

The PHENIX collaboration currently records  $\sim 1/4$  PB of raw data per year for each experimental run. Efficient and timely analysis of this data benefits from a framework for distributed analysis via the PHENIX Grid. Initial experience with this newly implemented grid infrastructure is presented with particular emphasis on job monitoring and job submission.

## THE PHENIX GRID IN BRIEF

The PHENIX Grid is currently comprised of the RHIC Computing Facility (RCF), located at Brookhaven National Laboratory (BNL) and five moderately sized clusters located in five different geographical locations across the United States and Europe. A schematic layout of the grid is given in Fig. 1. The clusters have different computing power, independent administration, different job schedulers, different Internet connectivity, etc, which makes for many interesting challenges. In particular no multi-cluster administrative privileges are currently available to the grid administrator/s.

In its current configuration, the PHENIX Grid provides invaluable CPU cycles for data analysis and simulations important to physics interpretation and/or detector efficiency studies. In what follows, we give a brief description of our early experience with this Grid.

## EARLY EXPERIENCE WITH THE PHENIX GRID

### Infrastructure of the PHENIX Grid

The PHENIX Grid relies on existing Grid/Globus infrastructure. Specifically, for each cluster a dedicated or **Globus GateWay (GGW)** computer is assigned and all required components of the Globus server software (ie. GT-2.4 from <http://www.globus.org>) are deployed. That is,

- the GGW is a local machine in the cluster and is connected to the Internet.

<sup>†</sup> shevel@mail.chem.sunysb.edu

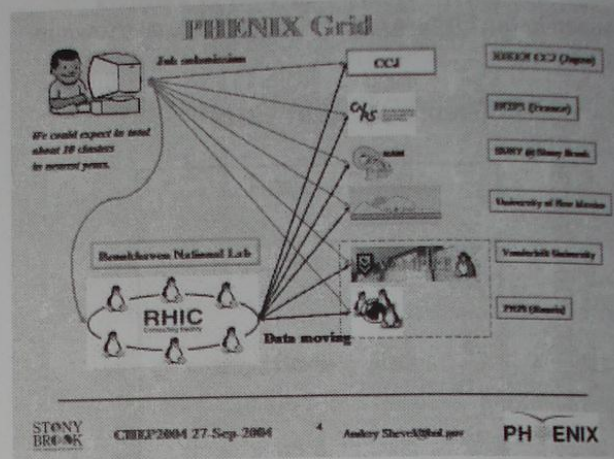


Figure 1: General scheme for PHENIX Grid computing resources

- All required Globus components are installed and running.
- The data and home directory of each user are accessible on the GGW.
- A batch scheduler software for the cluster is available on the GGW (i.e. it is possible to submit jobs from the GGW).

The procedure for achieving this implementation is relatively straightforward and works rather well. On the other hand, the flexible nature of the current management scheme for the clusters does not guarantee that all Grid enabled software components (see for example <http://www.globus.org>, <http://www.ppdg.net>, <http://www.griphyn.org> and <http://www.edg.org>) are uniformly available on all clusters.

A primary focus of the PHENIX Grid is to provide each user and administrator with easy access to:

- Information about data location. This is accomplished via a file catalog.
- Seamless job submission across clusters.
- Seamless job monitoring across clusters.
- Grid maintenance with minimal down-time.

## Job submission and monitoring on the PHENIX Grid

Multi-cluster/multi-site job processing requires a reliable set of job monitoring and control functions. For example, it is often the case that a user wants to;

- test the ability of a remote cluster to run a job over the Globus infrastructure.
- Query which of his jobs are currently running and where they are running (on which site or cluster);
- Query whether or not there is output from his job and find its location (directory file, cluster).
- Query where and how results will be collected (on which cluster, by which tool, script, etc.);
- Kill all or a specific selection of his submitted jobs on all clusters.

In addition to the points raised above, it is easy to envisage that rather detailed information about each job (in the multi-cluster environment) could be stored in a database for later access via a web-based interface. A similar (or the same) interface could provide primary control functions to submit, kill, and resubmit jobs. Such control functions would also provide fine granularity selections for each job such as job name, output length, execution time, input file names, etc.

Several administrative and technical considerations are important to the realization of seamless job submission and job monitoring on the PHENIX Grid. We list a few of these below

**Administrative issues** In addition to the requirement that each user has to be given the privilege to use clusters across sites, special consideration had to be given to various issues related to the fact that;

- each cluster is independently administered;
- all clusters are not singularly devoted to PHENIX activity;
- each cluster does not necessarily run the the same version of the Linux operating system;
- each cluster has a different local batch systems;

Given this reality, we found that it was particularly useful to implement a common cluster status board (see <http://ram3.chem.sunysb.edu/phenix-grid>) which allowed a user to interrogate several parameters for each cluster.

Other issues included;

- Reliable porting and testing of PHENIX software on each cluster. Typically the porting step was achieved via the Packman packaging scheme or by direct mirroring over AFS. After porting, each software of interest was tested to ensure that each cluster gave the same result.

A number of interesting challenges were encountered.

– *Problems associated with hardware (CPU architecture).* Porting is obviously straightforward between clusters having the same hardware and can be achieved by simply moving binary programs to new clusters. If the hardware on a new cluster is different then the porting process is much more involved.

– *Operating system.*

This issue is similar to the one above. If one has exactly the same operating system on two clusters, one can copy most of ones scripts and configuration files and be pretty sure that most of them will work as intended. However if one cluster is running under different type of operating system the port of the software will require more effort.

– *Versions of system software in operating system.*

It is also important to know the version of libraries of type 'glibc', the version of the kernel, libtool, and related utilities. Serious difference in different versions of these components serve to make the port of the software more difficult.

– *Available resources.*

Issues concerning the availability of resources on each cluster is of course important. For example, in many clusters a user might not have access to enough disk space to accommodate the files generated by his job.

– *Batch scheduler.*

Usually the user has no choice and have to use on new remote cluster standard for remote cluster batch scheduler. From the experience we know it might be: PBS, LSF, BQS, PBSPro, Condor, Sun Grid Engine, or something else. That means we have no possibility to copy user scripts for job submission as they are in the case when the user scripts use specific command names or/and parameters for batch job scheduler. From the real experience we know that batch scheduler may have some peculiar configuration. For example, at the university U the scheduler ProPBS has default queue the name 'default'. The permitted number of running jobs per the user is limited to 10. That means if user will submit 15 jobs first 10 jobs will be really submitted and will run but other 5 jobs will wait till the moment when at least one running user jobs will finish. It will happen in spite of the number of free machines in the cluster and the activity of other users in the cluster.

**Job monitoring** Two types of monitoring have been investigated for the PHENIX Grid; service monitoring and

job/task monitoring. A wide range of service monitoring tools are readily available from the Internet. A short list which proved to be quite useful follows:

- Monalisa
- Ganglia
- Zabbix
- Big Brother

These tools provide seamless access to a variety of services such as:

- network status;
- computer status (CPU load, memory usage, swap usage, etc.);
- web service status;
- other types of service,

which proved to be immensely useful for monitoring the current status or past history of the activities of each cluster ie. CPU usage, memory usage, IO channel usage, etc.

In contrast to service monitoring tools, a typical user also has a need for actual job monitoring. For example, when a job is submitted to the PHENIX Grid it can run on any of the six clusters which comprise the Grid. A user needs to know his;

- job ID (probably simple unique number which does not depend of hostname, site name, etc);
- job name;
- job submission time;
- host name where job was submitted;
- job start execution time;
- other suitable parameters.
- etc, etc

Job monitoring tools are not pervasive. This is especially true for monitoring tools which can provide detailed information about a users job. One monitoring tool that we have found to be useful is the Batch Object Submission System **BOSS** (see <http://www.bo.infn.it/cms/computing/BOSS/>).

**BOSS** consists of a binary interpreter which performs Command Line Interface (CLI) commands, a database which stores job status information, and a daemon, which is launched for each user job, that sends job information data to the database. A web interface (**BOSS** Database Explorer or **BODE** (see <http://filine.home.cern.ch/filine/bode/bode-2.1.1.tgz>)) provides easy access to the database. An overview of the job flow with **BOSS** is illustrated in Fig. 2. More than often, it is desirable to keep the history for each job for an extended period of time. We found this

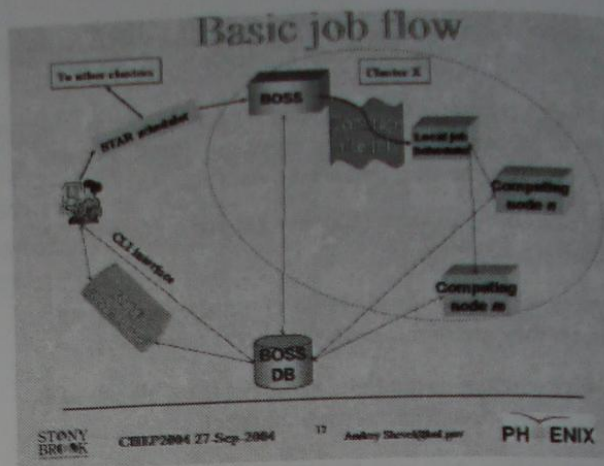


Figure 2: Basic job flow with use of **BOSS** and **STAR** scheduler.

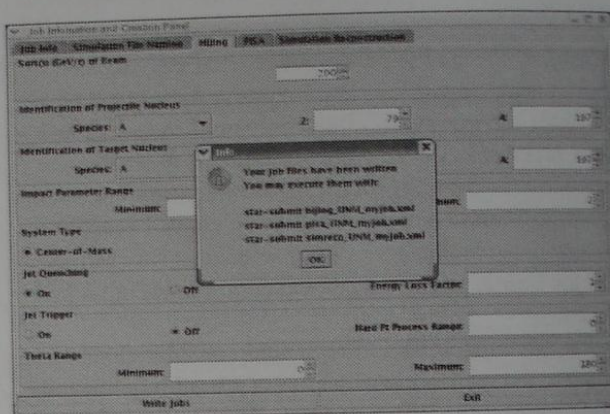


Figure 3: The panel to generate the production jobs for simulation. The user can choose a range of physics parameters for the simulation.

to be particularly important for jobs executed on multiple clusters. For this, it was necessary to implement additional databases so that a user can store information about a current job for up to a year or more.

As stated earlier, it is important to have the possibility to submit and control jobs on multiple clusters. In order to do this, a tool kit **GSUNY** (see <ftp://ram3.chem.sunysb.edu/pub/suny-gt-2/gsuny.tar.gz>) was developed. Thus if a user wants to cancel all all jobs which are running on different clusters he would use the command **G KillAllRun**. Currently, all commands in the **GSUNY** toolkit interacts with the **BOSS** database.

Job submission and control is possible from the command line or via a web interface. Fig. 3 shows a sample panel which is used for job submission. In this panel the user is able to enter several physics parameters for concrete simulation and after that start the simulation production (many jobs through **STAR** scheduler).

Despite the success of the **PHENIX** Grid, a number of challenges still remain. A brief list of some of the more important ones follow.

- Administration service (maintenance).
- More sophisticated job control and accounting across clusters.
- More complete technology for automatic deployment of the PHENIX software.
- More checking tools for regular testing (once an hour?) to be sure that all components of multi cluster multi site PHENIX Grid is running properly.
- A portal to integrate all PHENIX Grid tools in one user window.

## SUMMARY

Initial investigations for implementation of a PHENIX Grid across a diverse range of institutions have proved to be quite successful. This success is reflected in the ability of many users to simultaneously marshal computational resources from several participating institutions for data analysis.

It is anticipated that that our current strategy of using existing subsystems as bricks to build a robust PHENIX distributed computing environment will continue.