The Grid2003 Production Grid: Principles and Practice

I. Foster, J. Gieraltowski, S. Gose, N. Maltsev, E. May, A. Rodriguez, D. Sulakhe, A. Vaniachine, *Argonne Natl. Lab.*

J. Shank, S. Youssef, *Boston U.*

D. Adams, R. Baker, W. Deng, J. Smith, D. Yu, *Brookhaven Natl. Lab.*

I. Legrand, S. Singh, C. Steenberg, Y. Xia, *Caltech*

A. Afaq, E. Berman, J. Annis, L.A.T. Bauerdick, M. Ernst, I. Fisk, L. Giacchetti, G. Graham, A. Heavey, J. Kaiser, N. Kuropatkin, R. Pordes, V. Sekhri, J. Weigand, Y. Wu, *Fermi Natl. Accelerator Lab.*

K. Baker, L. Sorrillo, *Hampton U.*

J. Huth, *Harvard U.*

M. Allen, L. Grundhoefer, J. Hicks, F. Luehring, S. Peck, R. Quick, S. Simms, *Indiana U.*

G. Fekete, J. vandenBerg, *Johns Hopkins U.*

K. Cho, K. Kwon, D. Son, H. Park, *Kyungpook Natl. University/KISTI*

S. Canon, K. Jackson, D.E. Konerding, J. Lee, D. Olson, I. Sakrejda, B. Tierney, *Lawrence Berkeley Natl. Lab.*

M. Green, R. Miller, *U. Buffalo*

J. Letts, T. Martin, *U.C. San Diego*

D. Bury, C. Dumitrescu, D. Engh, R. Gardner, M. Mambelli, Y. Smirnov, J. Voeckler, M. Wilde, Y. Zhao, X. Zhao, *U. Chicago*

P. Avery, R. Cavanaugh, B. Kim, C. Prescott, J. Rodriguez, A. Zahn, *U. Florida*

S. McKee, *U. Michigan*

C. Jordan, J. Prewett, T. Thomas, *U. New Mexico*

H. Severini, *U. Oklahoma*

B. Clifford, E. Deelman, L. Flon, C. Kesselman, G. Mehta, N. Olomu, K. Vahi, *U. Southern California*

K. De, P. McGuigan, M. Sosebee, *U. Texas Arlington*

D. Bradley, P. Couvares, A. De Smet, C. Kireyev, E. Paulson, A. Roy, *U. Wisconsin-Madison*

S. Koranda, B. Moe, *U. Wisconsin-Milwaukee*

B. Brown, P. Sheldon, *Vanderbilt U.*

Abstract

The Grid2003 Project has deployed a multi-virtual organization, application-driven grid laboratory ("Grid3") that has sustained for several months the production-level services required by physics experiments of the Large Hadron Collider at CERN (ATLAS and CMS), the Sloan Digital Sky Survey project, the gravitational wave search experiment LIGO, the BTeV experiment at Fermilab, as well as applications in molecular structure analysis and genome analysis, and computer science research projects in such areas as job and data scheduling. The deployed infrastructure has been operating since November 2003 with 27 sites, a peak of 2800 processors, work loads from 10 different applications exceeding 1300 simultaneous jobs, and data transfers among sites of greater than 2 TB/day. We describe the principles that have guided the development of this unique infrastructure and the practical experiences that have resulted from its creation and use. We discuss application requirements for grid services deployment and configuration, monitoring infrastructure, application performance, metrics, and operational experiences. We also summarize lessons learned.

1 Introduction

The Grid2003 Project [1] has deployed for the first time a persistent, shared, multi-virtual organization (VO) [2], multi-application grid laboratory capable of providing production level services for large-scale computation- and data-intensive science applications. The project was organized by representatives of the U.S. "Trillium" projects (the GriPhyN virtual data research project [3], Particle Physics Data Grid, PPDG [4], International Virtual Data Grid Laboratory, iVDGL [5]) and the U.S. ATLAS [6] and U.S. CMS [7] Software and Computing Projects of the Large Hadron Collider (LHC) [8] program at CERN [9]. The goal of Grid2003 was to build an application grid laboratory ("Grid3") that would provide:

- a platform for experimental computer science research by GriPhyN and other grid researchers;

- the infrastructure and services needed to demonstrate LHC production and analysis applications running at scale in a common grid environment;

- the ability to support multiple application groups, including the Sloan Digital Sky Survey (SDSS) [10]

and the Laser Interferometer Gravitational Wave Observatory (LIGO) [11, 12], core participants in GriPhyN and iVDGL.

A set of specific and quantitative goals defined for Grid2003 included performance targets and metrics. We used the SC2003 conference (Nov. 15-21, 2003) [13] to initiate sustained operations, and since that period have met or exceeded most performance targets. The deployed grid continues operations today. We view this demonstration of successful and sustained operations as a significant step forward in our ability to create and operate persistent, shared grid-based cyberinfrastructure.

In the rest of this paper, we present the overarching project requirements (Section 2), related work (Section 3), application requirements (Section 4), grid design (Section 5), application results (Section 6), milestones and metrics (Section 7), and lessons learned (Section 8).

2 Project Requirements

The ambitious goals of Grid2003 included providing production capabilities to many data-intensive applications while also maintaining a laboratory for computer scientists developing new grid systems. Many universities and national laboratories contributed to the project. Important considerations were to develop a simple architecture that could link many sites, provide software that could be easily installed, and run an operations center as a focal point for information gathering and dissemination for all aspects of the project. We refine the overall project goals further as follows.

Architecture: We needed a simple grid architecture that would link execution and storage sites and provide services for monitoring, information publication, and discovery. A centralized operations center was needed to provide services to several grid application frameworks.

Software: We opted for a middleware installation based on the Virtual Data Toolkit (VDT) [14], which provides services from the Globus Toolkit [15], Condor [16], GriPhyN, and PPDG, as well as components from other providers such as the European Data Grid Project (EDG) [17]. VDT allows grid facility administrators to configure their sites easily with simple and well-defined interfaces to existing facility configurations, information service providers, and storage elements. Additional services such as Replica Location Service (RLS) [18], Storage Resource Manager (SRM) [19], and dCache [20], can be provided by individual VOs if desired.

Policy management: Experiment groups should be able to run their applications effectively on non-dedicated resources, including resources not controlled by their VO and/or shared with local users. Automated application installation and publication is important so as to impose minimum requirements on grid facility managers.

Grid2003 is one of several large-scale grids in the U.S., Europe, and Asia. Many applications targeted by Grid2003 are also designed to run on other grids. Thus, efforts were made to ensure consistency with and "federate" with other Grid projects where possible, in particular the LHC Computing Grid Project (LCG) [21].

3 Related Work

The many successful grid projects worldwide encompass a variety of architectures, deployment approaches, and targeted application domains. For example, building on early experiences such as the NSF MetaCenter [22], I-WAY [23], and GUSTO [24], a number of U.S. grids link modest numbers of high-end systems: e.g., NASA's Information Power Grid [25], the NSF PACI grids [26] and TeraGrid [27] European efforts include the aforementioned LCG, the European Data Grid (EDG) and its follow-on (EGEE) [28], and DataTAG [29], which focused on transatlantic grid testbeds and high performance networks. NorduGrid [30] links computational centers in Scandinavia to deliver production services for high-energy physics applications. Also relevant is PlanetLab [31] which provides a uniform OS environment across PCs located at different sites to support experimentation with distributed system services.

Grid2003 extends these and other efforts in several respects. First, it is organized as a consortium among participating stakeholder grid and application software and computing organizations. This structure allows several project objectives to be met simultaneously, and a large scale production environment achieved with the aggregate of resources from the participating groups, while maintaining a development environment for computer science research. Second, the approach taken for construction was aimed to minimize site-specific requirements (e.g., for installation and configuration) while stressing site and VO autonomy. Like other grids, and unlike PlanetLab, Grid3 links high-value resources subject to often demanding local policies, and supports computation-intensive and data-intensive applications.

4 Application Requirements

Grid2003 was aligned with specific application milestones, in particular the LHC data challenges detailed below. Additional requirements were supplied by the milestones for the participating grid projects. This alignment with external project milestones helped to ensure strong participation in the project.

4.1 ATLAS Challenge Problems

The ATLAS application focused on Monte Carlo simulation of the physics processes that will occur in high energy proton-proton collisions at the LHC. Datasets recording the simulated response of the ATLAS detector

to these collisions were used as input to event reconstruction and analysis algorithms.

The application workflow comprises several steps and was implemented using Chimera and Pegasus virtual data tools [32-34] and other VDT services. The first step is to generate the physics processes. The Pythia Monte Carlo program [35] is used to simulate and record them into RLS. Next, the GEANT-based [36] core simulation package, built from the CERN software repository and packaged with grid-based installation scripts, creates datasets with an average size of about 2 GB. All datasets produced are archived at the Tier1 facility at Brookhaven National Laboratory (BNL). Finally, datasets are "reconstructed" either on Grid3 or at CERN, producing samples ready for physics analysis. The distributed analysis program DIAL [37] is used for creation and analysis of physics histograms.

4.2 CMS Challenge Problems

The CMS Collaboration was able to use Grid3 resources when they came online in October/November 2003 to produce events for their 2004 data challenge. Fifty million events with minimum bias pile-up at a beam luminosity of 2×10^{33} were needed in the final sample. CMS detector simulation consists of 3 steps: (1) event generation with Pythia, (2) event simulation with a GEANT-based simulation application, and finally (3) reconstruction and digitization with the additional pile-up events. The sample of simulated events was accumulated at CERN for primary reconstruction, and distributed in real time to Tier1 and Tier2 centers (some being Grid3 sites) for calibration and toy analysis. The software suite includes MCRunJob [38], a CMS tool for workflow configuration, and MOP [39], a CMS DAG writer, which were first grid-enabled during a previous "big n-tuple" production during the fall of 2002 [40]. CMS Production jobs are specified by reading input parameters from a control database and converting them to DAGs suitable for submission to Condor-G/DAGMan [41]. All datasets produced were archived through a Storage Element at the Tier1 facility at Fermi National Accelerator Laboratory (Fermilab).

4.3 Cluster finding in SDSS

SDSS contributed several challenge problems. A search for galaxy clusters in SDSS data resulted in workflows with several thousand processing steps organized by Chimera virtual data tools. A second application involved a pixel-level analysis of astronomical data, such as analysis of cutouts of images about galaxies with the aim of adding more information to existing catalogs. Other applications included a search for near earth asteroids, which calls for examining complete SDSS images in search of highly elongated objects.

4.4 Blind Gravitational Wave Searches

The LIGO challenge problem was an extensive, all-sky, blind search for continuous wave (pulsar) signals in the LIGO S2 data set. Each search required that a conventional binary short Fourier transform data file be accessible containing the frequency band that the target signal spans during the observation time. Additional data files containing the ephemeris data for the year are staged from LIGO facilities to Grid3 sites using GridFTP. The location of the staged data (on average 4 GB per job) is published in RLS so that its location is available to the job. The last job in the workflow stages the output results back to the LIGO facility and updates database entries. Each workflow instance runs for several hours on an average processor. The GriPhyN-LIGO working group developed the necessary infrastructure using Chimera and Pegasus to generate and execute the workflows.

4.5 CP Violation in Heavy Quark Decay

The BTeV challenge problem was to simulate charge-parity (CP) violations in decays of heavy quarks produced in proton-antiproton collisions at the Fermilab collider. The clarity of the Chimera virtual data toolkit as a BTeV physics interface and the scalability of these tools for large Monte Carlo generation were goals to be tested with data challenges run at scale. The workflow processing time was about 15 seconds per event on a 2GHz machine, translating into a typical request for 2.5 million events generated with 1000 10-hour jobs across Grid3.

4.6 Computational Chemistry and Biology

SnB [42, 43], a computer program based on the *Shake-and-Bake* method, is the program of choice for structure determination in many of the 500 laboratories that have acquired it. The *SnB* program uses a dual-space direct-methods procedure for determining crystal structures from X-ray diffraction data. This program has been used in a routine fashion to solve difficult atomic resolution structures, containing as many as 1000 unique non-Hydrogen atoms, which could not be solved by traditional reciprocal-space routines. GADU [44] is a Genome Analysis and Databases Update Tool from the Mathematics and Computer Science division at Argonne National Laboratory, used to perform a variety of analyses of genome data. Both of these applications ran under the iVDGL VO.

4.7 Computer Science Challenge Problems

Computer science groups worked with experiment developers to provide the application middleware (e.g., Chimera and Pegasus, Globus client libraries, Condor-G, RLS) required by grid-based application frameworks. Various computer science groups also used Grid3 as a

vehicle for research studies. In addition, the following three demonstrators were provided.

A data transfer study was performed to evaluate whether we could perform large-scale reliable data transfers between Grid3 sites. A Java-based plug-in environment (Entrada) was used to generate simulated traffic between a matrix of sites in a periodic fashion [45].

NetLogger-instrumented GridFTP was used to monitor the Globus Toolkit GridFTP server and [46] URL copy program. NetLogger events were generated at program start, end, and on errors (the default) and for all significant I/O requests (by request) [46, 47].

An exerciser backfill application provided by the Condor group tested the status of the batch systems and operation characteristics of each Grid3 site. This application ran repeatedly with a low priority at 15 minute intervals.

5 Grid Design

We adopted a simple two-tier approach, in which each resource (compute, storage, application, site, user) was logically associated with a VO. At each site, a core set of grid middleware services with VO-specific configuration and additions were installed, with registration to a VO-level set of services such as index servers and grid certificate databases. Where appropriate, VO-level services were combined into top-layer services at the iVDGL Grid Operations Center (iGOC), which provided monitoring applications, display clients, and verification tasks and an aggregate view of the collective Grid3 resource and performance. Six VOs (U.S. ATLAS, U.S. CMS, SDSS, LIGO, BTeV, iVDGL) were configured. Appropriate policies were implemented at each local batch scheduler (OpenPBS, Condor, and LSF) and Unix group accounts were established at each site for each VO.

5.1 Site Installation Procedures

Procedures for installation, configuration, post-installation testing, and certification of the basic middleware services were devised and documented. The Pacman [48] packaging and configuration tool was used extensively to facilitate the process. A Pacman package encoded the basic VDT-based Grid3 installation, which included:

- The Globus Toolkit's Grid security infrastructure (GSI)[49], GRAM, and GridFTP services;

- Information service based on MDS, with registration scripts to VO-specific information index servers and VO-specific information providers;

- Cluster monitoring services based on Ganglia [50], with provisions for hierarchical grid views; and

- Server and client software for the MonALISA [51] agent-based monitoring framework.

Conventions were documented to provide grid facility administrators and operators with uniform instructions with the goal of obtaining a consistent Grid3 environment over the heterogeneous sites. In particular, information providers were developed for site configuration parameters such as application installation areas, temporary working directories, storage element locations, and VDT software installation locations. Only a few extensions to the GLUE [52] MDS schema were required.

5.2 Monitoring and Information Services

The software installed on Grid3 sites included components necessary to monitor the overall behavior and performance of the grid and its applications. Several packages sensed monitoring data and made it available to a distributed framework of services and client tools. The set of information providers deployed was determined by identifying and prioritizing desirable grid-level (such as overall resource availability and consumption) and VO-level (e.g., aggregate CPU usage) performance indicators. Other requirements derived from auditing, scheduling and debugging considerations.

The framework was built by integrating existing monitoring software tools into a simple architecture. Figure 1 shows the components of the framework. Producers provide monitored information, consumers use this information, and intermediaries have both roles, sometimes providing aggregation or filtering functions.
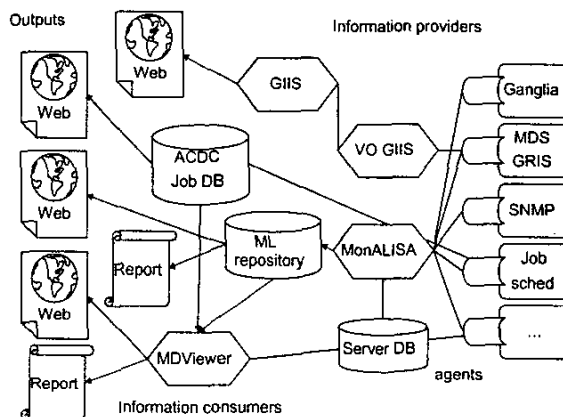

Figure 1 Grid3 monitoring architecture showing information providers and consumers, and the data flows between them.

Some monitoring components are located on Grid3 sites, some in central servers, and some are the clients of the users accessing the information. An aggregated data summary is available centrally, while more detailed data

239

and streams of updates are available from the sites. The main components of the monitoring framework are:

- The Globus Toolkit's *Monitoring and Discovery Service* (MDS) [53] is used to maintain site configuration and monitoring information. A schema extension, producers (MDS information providers), and intermediaries were developed to use this framework in Grid3.

- *Ganglia* is used to collect cluster monitoring information such as CPU and network load and memory and disk usage. Ganglia-collected information is available through web pages served at the sites and a summary [54] a central server at iGOC. Intermediaries have been developed for it too.

- *MonALISA* [55], Monitoring Agents in a Large Integrated Services Architecture, provides access to monitoring data provided by a variety of information providers, including agents which monitored the GRAM logfiles, job queues, and Ganglia metrics. The MonALISA client allowed access to both the central repository as well as site servers through a graphical interface. Custom agents were developed to collect VO-specific activity at sites such as jobs run, compute element usage, and I/O.

- The *MonALISA central repository* collects its information in a central server at the iGOC, storing it in a round robin-like database, and makes it available through the web [55].

- The *ACDC Job Monitor* [56] from the Advanced Computational Data Center (ACDC) at the University of Buffalo collects information from local job managers using a typical pull-based model. Statistics and job metrics are collected and stored in a web-visible database, available for aggregated queries and browsing.

- The *Site Status Catalog* [57] periodically tests all sites and stores some critical information centrally. A web interfaces provides a list of all Grid3 sites, their location on a map, their status, and other important information.

- The *Metrics Data Viewer* (MDViewer) [58] allows for the analysis and display of collected metrics information. It provides an API for manipulating, comparing and viewing information and a set of predefined plots, parametric in arbitrary time intervals, sites and VOs, tailored to Grid2003 needs.

The Grid3 monitoring and analysis system allows similar information to be collected by different paths. This redundancy might appear unnecessary, but we have found that it has the advantage of permitting crosschecks on the data collected. A coordinated system has been deployed that adapts and combines the different monitoring tools.

Information producers collect information close to its source, a common intermediary defines a uniform representation and access methods, and information is centrally collected to produce aggregated information, statistics and documents. Client consumers can access centrally stored data, or more detailed data from participating sites, in a uniform manner.

5.3 Virtual Organization Management

To simplify user access to Grid3 resources and reduce the burden on grid facility administrators, we deployed EDG's Virtual Organization Management System (VOMS) [59]. We also used group accounts at sites, with a naming convention for each VO. We generated the local grid-map files that map user identities presented in X509 certificates to local accounts by calling an EDG script to contact each VO's VOMS server.

5.4 Support and Operations

The deployment and operation of the Grid3 environment required a number of centralized support activities. The iGOC hosted centralized services, including the Pacman cache, the top-level MDS index server, the Site Status Catalog, the MonALISA central repositories, and web services for Ganglia. A simple trouble ticket system was used intermittently during the project. An acceptable use policy modeled after that used by the LCG was adopted.

Ongoing support for Grid3 sites and applications is distributed according to responsibility. Site administrators provide for the operation and support of their sites. The VO central support organizations provide the organization and effort for the support and maintenance of their applications and virtual facilities.

6 Results

An important strategic goal for Grid2003 was to "Provide the infrastructure and services needed to demonstrate LHC production and analysis applications running at scale in a common grid environment." Figure 2 shows the integrated and Figure 3 the differential Grid3 usage during a 30 stretch beginning October 25, 20003. Both U.S. ATLAS and U.S. CMS ran production systems at scale during this period using shared facilities. Note that the experiments continue to exercise production on Grid3 with an average of 700 CPUs in daily use in April 2004.

6.1 U.S. ATLAS GCE and DIAL

ATLAS deployed its grid-enabled application package GCE-Server on 22 Grid3 sites. Automated user-level installation tools based on Pacman used the Grid3 MDS information schema extensions for application installation attributes. Client hosts (GCE-Client) were installed outside Grid3 for job submission. More than 5000 jobs

(Geant3-based simulation followed by reconstruction) were processed at 18 sites, with total data I/O of about 1.1 TB. A dataset catalog was created for produced samples, making them available to the DIAL distributed analysis package. Output datasets were stored at BNL by the grid jobs, and continue to be analyzed by DIAL developers and the SUSY physics working group.
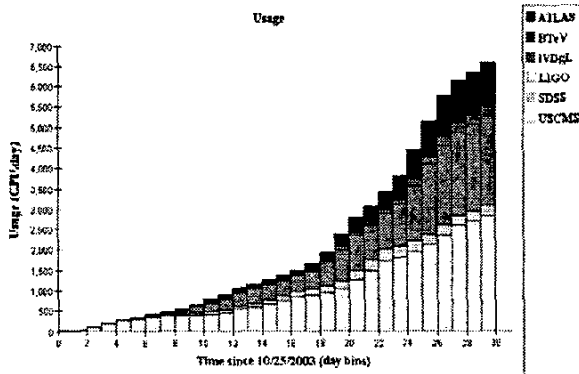

Figure 2: Integrated CPU usage (CPU-days) during the 30 day running for SC2003, by VO.

We observed a failure rate of approximately 30%, where failures are defined as jobs experiencing errors in any processing step that prevented perfect completion (pre-stage, job execution producing the output files, post-stage to the final storage element at BNL, and registration to RLS). Approximately 90% of failures were due to site problems: disk filling errors, gatekeeper overloading, or network interruptions. For example, we did not handle ACDC's nightly roll over of worker nodes gracefully, and so jobs still running had to be re-processed.
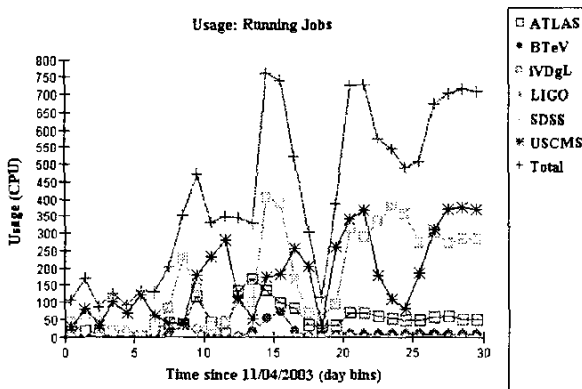

Figure 3: Differential CPU usage (measured in time-averaged number of CPUs used) during the 30 day running period for SC2003, organized by VO.

6.2 USCMS MOP Production

U.S. CMS has used Grid3 resources to produce simulated events for the upcoming CMS data challenge. U.S. CMS ran a GEANT3-based, statically linked FORTRAN application called CMSIM and a GEANT4-based, dynamically linked, C++ application called OSCAR. Since SC2003, U.S. CMS has used Grid3 resources on 11 sites to simulate more than 14 million GEANT4 full detector simulation events. Figure 4 shows usage since mid-November. Efficiency on Grid3 resources is roughly as high as on the original U.S. CMS production grid, once sites are fully validated. The official OSCAR production jobs are long (some more than 30 hours) and not all sites have been able to accommodate running them. The effort required to run the application has been about 2 FTEs, split between the application administrator and site operations support.

Approximately 70% of CMSIM and OSCAR jobs completed successfully, which is consistent with US-ATLAS estimates. Jobs often failed due to site configuration problems, or in groups from site service failures. We saw few random job losses: more frequently a disk would fill up or a service would fail and all jobs submitted to a site would die. Service level monitoring needs to be improved and some services probably need to be replaced. For example, storage reservation (e.g., as provided by SRM) would have prevented various storage-related service failures.
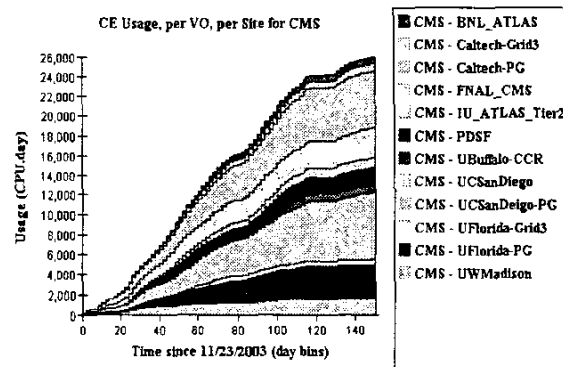

Figure 4 CMS cumulative use of Grid2003. The chart plots the distribution of usage (in CPU-days) by site in Grid2003 over a 150 day period beginning in November 2003.

6.3 GridFTP Data Transfer Demonstrator

We met our goal of transferring 2 TB across Grid3 per day, and long-running data transfers ran reliably. Issues of account privileges, ports, and firewalls caused the main problems in deployment and configuration. Figure 5 shows data "consumed" by Grid3 sites according to the VO responsible.
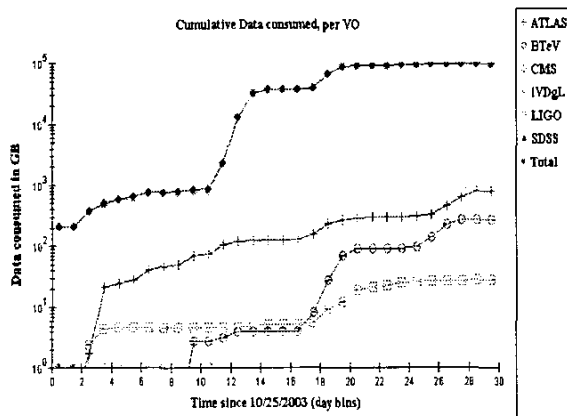
Figure 5 Data consumed by Grid3 sites, by VO. Nearly 100 TB was transferred during 30 days before and after SC2003 (top curve is total from all sources). The GridFTP demonstrator accounted for most data transferred on Grid3.

6.4 Analysis of Grid Usage

In summary, Grid3 users could be classified into seven application demonstrator classes corresponding to their VO, as shown in Table 1. Each class contained its own set of users which in turn evaluated their applications on the Grid3 production resources. Several basic application requirements drove how users selected sites:

1. Internet connectivity of compute nodes: some applications needed outbound internet connectivity to databases located outside of privately addressed production nodes.

2. Availability of required disk space: a given Grid3 resource may not have had sufficient disk space available for the proposed task.

3. Maximum allowable runtime: queue managed Grid3 resources required every computational job to specify the runtime requested which may not have been long enough for the proposed task.

4. Gatekeeper network bandwidth capacity: applications requiring large quantities of application data or that produced a large number of output files would select only those Grid3 resources having the highest bandwidths.

We analyzed a portion of the monitoring data logged during the last seven months. Using a sample of 291052 job records, each application demonstrator completed a widely varying number of jobs with average job runtimes varying from minutes to days. The total CPU consumption of an application class did not directly correspond to the total number of jobs completed.

The gatekeeper load created by scheduling and managing grid-enabled resource computational jobs was quite different depending on the frequency and duration of the submitted jobs. In general, a typical gatekeeper using a queue manager will experience a sustained one minute load of ~225 when managing ~1000 computational jobs. This load can sharply increase when the job submission frequency is high, thus short duration high frequency computational jobs tend to sharply increase the gatekeeper loading. For computational jobs that only require a minimal amount of production node file staging, a factor of two can be applied to the sustained load; on the other hand computational jobs requiring a substantial amount of file staging the factor can increase to three or four.

Each application class showed a fairly wide usage of Grid3 sites during the peak months (Fall 2003) but the general trend is that applications tend to favor the resources provided within their VO. There are many factors that contribute to this observed behavior (including VO ownership of certain sites, site policies, and production cycles. Each application class performs differently on each individual resource and some resources are better suited for processing low frequency long running jobs whereas other resource may not be able to process long running jobs at all. Additionally, application demonstrators tended to have "favorite" Grid3 resources and submitted more computational jobs to them. In any case it is evident that the peak production months for each application class did not account for a substantial percentage of the total CPU days. Thus, a substantial amount of the computational jobs are processed on a continual basis and not just during intensive submission periods. This would indicate that a persistent production grid would indeed increase the overall production rate of all application classes. This is also illustrated by Figure 6, where the obvious ramp up of computational production jobs appears in 2003 and a more sustained production rate appears in 2004.

7 Milestones and Metrics

At the outset of Grid2003, we defined milestones for use in tracking progress and evaluating success. We have met and even surpassed most of these milestones. Here we summarize some highlights.

- **Number of CPUs (target = 400, actual = 2163).** The number of processors in Grid3 fluctuates over time as sites introduce and withdraw resources. A peak of over 2800 processors occurred during SC2003. More than 60% of CPU resources are drawn from non-dedicated facilities that are both shared among Grid3 participants and available to local users.

- **Number of users (target = 10, actual = 102).** About 10% of users are application administrators who perform most job submissions. However, more than 102 users are authorized to use Grid3 resources through their respective VOMS services.

- **Number of applications (target > 4, actual = 10).** Seven scientific applications, including at least one from each of the five participating experiments, continue to run on Grid3. In addition, the three computer science demonstrators are run periodically.
- **Number of sites running concurrent applications (target > 10, actual = 17).** The number of sites capable of running applications from multiple VOs.
- **Data transferred per day (target = 2-3 TB, actual = 4 TB).** This metric was met with the aid of the GridFTP demo that was run concurrently with the scientific applications. Plots of statistics collected may be found at the project website [45].
- **Percentage of resources used (target = 90%, actual = 40-70%).** The maximum number of CPUs on Grid3 exceeds 2500 most of the time. On Nov. 20, 2003 there were sustained periods when over 1300 jobs ran simultaneously (the metrics plots are averages over specific time bins, which can report less that the peak depending on chosen bin size).
- **Efficiency of job completion (target = 75%; actual: varies).** The value of this metric varies depending on the application and on the definition of failure. Generally speaking, for well-run Grid3 sites and stable applications, this figure exceeds 90%. Work is under way to collect more detailed statistics.
- **Peak number of concurrent jobs (target = 1000, actual = 1300).** Achieved on 11/20/03.
- **Rate of faults/crashes: (target < 1/hour, status: varies).** We have not started to measure this metric quantitatively, but have begun to collect summaries from the application groups.
- **Operations support load: (target < 2 FTEs, status: typically 10 part-time).** We added applications and sites continuously throughout SC2003, and this process continues today. Once a site becomes stable, it usually remains so except for hardware problems. Several sites replaced disks and/or nodes without perturbation to overall system operation. The infrastructure has been stable since November with a small support load of less than 2 FTEs. The number of jobs from different applications ramps up and down without impacting overall stability.

8 Project Lessons

We learned that we can indeed build, sustain, and operate a fairly large common grid from many autonomous organizations, and with reasonable effort and efficiency. We also learned that we can provide ongoing science benefit to stakeholders. Our experiences to date suggest several areas where improvements are needed, including the following.

- **Automated configuration, testing, and tuning** scripts are needed to give immediate feedback

regarding potential software installation issues, and to further reduce the cost of operating Grid3.

- **API for accessing troubleshooting and accounting** information are needed, particularly for the GRAM job submission and GridFTP file transfer systems. These APIs should provide *direct* information without the necessity of parsing log files.
- **Contact and support model.** We identified the need to revise the contact, operations and support model. Factorization of responsibilities, perhaps at the service level, is being explored.
- **Efficiency metrics.** Grid2003 efficiency targets were not met. Understanding why will require increased analysis of end-to-end applications.
- **Job Execution Policies:** Tools should be deployed and analyses done to check that the current Grid3 job policies are being properly enforced.
- **Job Resource Requirements:** Sites should publish more information about job execution and resource usage policies, such as maximum CPU time allowed. This information will aid in efficient job scheduling.
- **Storage Services and Data Management:** Grid3's current data management model is based on GridFTP and RLS. Additional infrastructure services are needed to support managed persistent and transient storage.
- **Troubleshooting:** Additional tools are necessary for troubleshooting, specifically tools for analyzing and querying log files, the ability to link a job ID on the execution side with a job ID at the submit (VO) side.
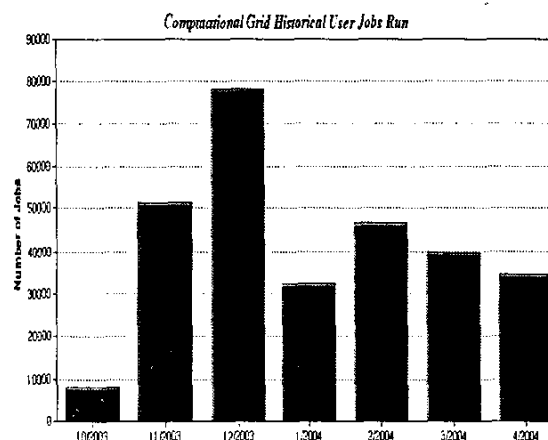

Figure 6 Distribution of the number of jobs run on Grid3 by month starting from October 2003.

Table 1: Grid3 computational job statistics based on completed production jobs from the period of October 23, 2003 to April 23, 2004 (source ACDC University at Buffalo).

Description	Grid3 User Classification (VO)						
	BTEV	iVDGL	LIGO	SDSS	USATLAS	USCMS	Exerciser
Number of Users	1	24	7	9	25	26	3
Grid3 Sites Used	8	19	1	13	18	18	14
Number of Jobs	2598	58145	3	5410	7455	19354	198272
Avg. Runtime (hr)	1.77	1.22	0.01	1.46	8.81	41.85	0.13
Max. Runtime (hr)	118.27	291.74	0.02	152.90	292.40	1238.93	36.45
Total CPU (days)	191.88	2945.79	0.01	329.44	2736.05	33750.14	1034.28
Peak Production Rate (jobs/month)	2377	25722	3	1564	3198	8834	72224
Number of Peak Prod. Resources	7	15	1	4	17	17	7
Max. Prod. from Single Resource (jobs/month) [%]	1421 [59.8]	22671 [88.1]	3 [100]	1120 [71.6]	901 [28.2]	4820 [48.4]	38512 [53.4]
Peak Production Month-Year	11-2003	11-2003	12-2003	02-2004	11-2003	11-2003	12-2003
Peak Production CPU (days)	129.46	1244.97	0.01	65.91	696.48	1981.95	51.78

9 Summary

We have discussed the deployment and use of a persistent, shared, multi-virtual organization, multi-application grid, the first of its kind. The infrastructure remains in place and is currently undergoing upgrades for future application demonstrators. Grid3 is giving us practical experience that will enable us to better define, plan, and achieve the additional scale, technologies and efforts needed for ubiquitous common grids providing long term production quality services to stakeholders. As well as serving as a valuable proving ground for grid operation techniques, Grid3 continues to deliver new scientific results and benefits for its application communities and, in addition, to attract new users from a range of disciplines, including computer science.

References

[1] The Grid2003 Project, http://www.ivdgl.org/grid2003/.

[2] I. Foster, Kesselman, C. and Tuecke, S., "The Anatomy of the Grid: Enabling Scalable Virtual Organizations," Intl. J. Supercomputer Applications, vol. 15 (3), pp. 200-222, 2001.

[3] P. Avery, I. Foster, Towards Petascale Virtual Data Grids (GriPhyN Project), http://www.griphyn.org/.

[4] Particle Physics Data Grid, http://www.ppdg.org/.

[5] P. Avery , I. Foster, R. Gardner, H. Newman, A. Szalay, "An International Virtual-Data Grid Laboratory for Data Intensive Science," Technical Report: GriPhyN-2001-2, 2001, http://www.griphyn.org/.

[6] U.S. ATLAS Software and Computing Project, http://www.usatlas.bnl.gov/atlas_psc/.

[7] U.S. CMS Software and Computing Project, http://www.uscms.org/scpages/sc.html.

[8] The Large Hadron Collider Project at CERN, http://lhc-new-homepage.web.cern.ch/lhc-new-homepage/.

[9] CERN, the European Laboratory for Particle Physics, http://cern.ch/public/.

[10] The Sloan Digital Sky Survey Project (SDSS), http://www.sdss.org/sdss.html.

[11] Laser Interferometer Gravitational Wave Observatory, http://www.ligo.caltech.edu/.

[12] B.C. Barish and R. Weiss, "LIGO and the Detection of Gravitational Waves," Physics Today, vol. 52, pp. 44, 1999.

[13] SC2003, Phoenix, Arizona, November 15-21, http://www.sc-conference.org/sc2003/.

[14] The Virtual Data Toolkit (VDT), http://www.lsc-group.phys.uwm.edu/vdt/.

[15] I. Foster, Kesselman, C., "Globus: A Metacomputing Infrastructure Toolkit," International Journal of Supercomputer Applications, vol. 11(2), pp. 115-129, 1998.

[16] M.J. Litzkow, Livny, M. and Mutka, M.W., "Condor - A Hunter of Idle Workstations," 8th International Conference on Distributed Computing Systems, pp. 104-111, 1988.

[17] The European Data Grid Project (EDG), http://eu-datagrid.web.cern.ch/eu-datagrid/.

[18] A. Chervenak, E. Deelman, et al., "Giggle: A Framework for Constructing Sclable Replica Location Services," presented at SC'02: High Performance Networking and Computing., 2002.

[19] A. Shoshani, Sim, A. and Gu, J., "Storage Resource Managers: Essential Components for the Grid," in

Resource Management for Grid Computing, J. Nabrzyski, Schopf, J. and Weglarz, J., Ed., 2003.

[20] *The dCache Project*, http://www.dcache.org/.

[21] *The LHC Computing Grid Project (LCG)*, http://lcg.web.cern.ch/LCG/.

[22] L. Smarr C. Catlett, "Metacomputing," *Communications of the ACM*, vol. 35, pp. 44-52, 1992.

[23] T. DeFanti, Foster, I., Papka, M., Stevens, R. and Kuhfuss, T., "Overview of the I-WAY: Wide Area Visual Supercomputing," *International Journal of Supercomputer Applications*, vol. 10 (2), pp. 123-130, 1996.

[24] S. Brunett, Czajkowski, K., Fitzgerald, S., Foster, I., Johnson, A., Kesselman, C., Leigh, J. and Tuecke, S., "Application Experiences with the Globus Toolkit," presented at 7th IEEE International Symposium on High Performance Distributed Computing, 1998.

[25] W.E. Johnston, Gannon, D. and Nitzberg, B., "Grids as Production Computing Environments: The Engineering Aspects of NASA's Information Power Grid," presented at In 8th IEEE International Symposium on High Performance Distributed Computing, 1999.

[26] R. Stevens, Woodward, P., DeFanti, T. and Catlett, C., "From the I-WAY to the National Technology Grid," *Communications of the ACM*, vol. 40 (11), pp. 50-61, 1997.

[27] C. Catlett, *The TeraGrid: A Primer*, http://www.teragrid.org/.

[28] *EGEE: Enabling Grids for E-Science in Europe*, http://public.eu-egee.org/.

[29] *Research and Technological Development for a Data TransAtlantic Grid*, http://datatag.web.cern.ch/datatag/project.html.

[30] *NorduGrid: Nordic Testbed for Wide Area Computing and Data Handling*, http://www.nordugrid.org/.

[31] M. Bowman, A. Bavier, B. Chun, D. Culler, S. Karlin, S. Muir, L. Peterson, T. Roscoe, T. Spalink, M. Wawrzoniak, "Operating System Support for Planetary-Scale Services.," *Proceedings of the First Symposium on Network Systems Design and Implementation (NSDI)*, 2004.

[32] I. Foster, J. Voeckler, et al., "Chimera: A Virtual Data System for Representing, Querying, and Automating Data Derivation," presented at 14th Intl. Conf. on Scientific and Statistical Database Management, Edinburgh, Scotland., 2002.

[33] E. Deelman, J. Blythe, Y. Gil, C. Kesselman, G. Mehta, S. Patil, M.-H. Su, K. Vahi, and M. Livny, "Pegasus : Mapping Scientific Workflows onto the Grid," presented at 2nd EUROPEAN ACROSS GRIDS CONFERENCE, Nicosia, Cyprus, 2004.

[34] James Blythe Ewa Deelman, Yolanda Gil, Carl Kesselman., "Pegasus: Planning for Execution in Grids," *GriPhyN Technical Reports*, vol. 2002-20, 2002.

[35] P. Edén T. Sjöstrand, C. Friberg, L. Lönnblad, G. Miu, S. Mrenna and E. Norrbin, "PYTHIA 6.154," *Computer Phys. Commun.*, vol. 135, pp. 238, 2001.

[36] *GEANT - Detector Description and Simulation Tool*, http://wwwasd.web.cern.ch/wwwasd/geant/index.html.

[37] *Distributed Interactive Analysis of Large Datasets (DIAL)*, http://www.usatlas.bnl.gov/~dladams/dial/).

[38] Dave Evans Gregory E. Graham, Iain Bertram, "McRunjob: A High Energy Physics Workflow Planner for Grid Production Processing," presented at CHEP 2003, La Jolla, California, 2003.

[39] *The MOP Project*, http://www.uscms.org/s&c/MOP.

[40] G.E. Graham, Bauerdick, L.A.T., Cavanaugh, R., Couvares, P., Livny, M., *Distributed Data Analysis: Federated Computing for High Energy Physics, (Chapter 10 in The Grid 2: Blueprint for a New Computing Infrastructure)*: Morgan Kaufman, 2003.

[41] J. Frey, Tannenbaum, T., Foster, I., Livny, M. and Tuecke, "Condor-G: A Computation Management Agent for Multi-Institutional Grids. Cluster Computing," *Cluster Computing*, vol. 5 (3), pp. 237-246, 2002.

[42] *The SnB Program*, http://www.hwi.buffalo.edu/SnB/.

[43] G.T. DeTitta C.M. Weeks, R. Miller, & H.A. Hauptman, "Applications of the minimal principle to peptide structures," *Acta Cryst.*, vol. D49, pp. 179-181, 1993.

[44] D. Sulakhe A. Rodriguez, E. Marland, V. Nefedova, G. X. Yu, and N. Maltsev, "GADU - Genome Analysis and Database Update Pipeline," *Preprint ANL/MCS-P1029-0203*, 2003.

[45] Scott Gose, *Entrada, a Lightweight Application Hosting Environment*, http://www-unix.mcs.anl.gov/~gose/entrada/.

[46] K. Jackson, "pyGlobus: A Python Interface to the Globus Toolkit," *Concurrency and Computation: Practice and Experience*, vol. 14, pp. 1075-1083, 2002.

[47] *Netlogger-Instrumented GridFTP Data Archive*, http://netlogger.lbl.gov:8080/grid3.rpy).

[48] S. Youssef, *Pacman, a Package Manger*, http://physics.bu.edu/~youssef/pacman/.

[49] I. Foster, Kesselman, C., Tsudik, G. and Tuecke, S., "A Security Architecture for Computational Grids," presented at 5th ACM Conference on Computer and Communications Security, 1998.

[50] Mason J. Katz Federico D. Sacerdoti, Matthew L. Massie, David E Culler, "Wide Area Cluster Monitoring with Ganglia," presented at IEEE Cluster 2003 Conference, Hong Kong, 2003.

[51] I.C. Legrand H.B. Newman, P.Galvez, R. Voicu, C. Cirstoiu, "MonALISA: A Distributed Monitoring Service Architecture," presented at CHEP 2003, La Jola, California, 2003.

[52] DataTAG and iVDGL Interoperability Working Group, *Grid Laboratory Uniform Environment*, http://grid.infn.it/datatag/wp4/doc/glue-v0.1.2.pdf.

[53] K. Czajkowski, Fitzgerald, S., Foster, I. and Kesselman, C., "Grid Information Services for Distributed Resource Sharing," *10th IEEE International Symposium on High Performance Distributed Computing*, pp. 181-184, 2001.

[54] *Grid3 Ganglia frontend*, http://gocmon.uits.iupui.edu/ganglia-webfrontend.

[55] *Grid3 MonALISA frontend, 2004*, http://gocmon.uits.iupui.edu:8080/index.html.

[56] *ACDC job monitor, 2004*, http://acdc.ccr.buffalo.edu/statistics/acdc/fullsizeindexqueue.php.

[57] *Grid2003 site catalog, 2004*, http://www.ivdgl.org/grid2003/catalog.

[58] M. Mambelli and D. Bury, *MDViewer: a Metrics Data Viewer of the Grid*, http://grid.uchicago.edu/metrics/.

[59] EU DataGrid Java Security Working Group, *VOMS Architecture v1.1*, http://grid-auth.infn.it/docs/VOMS-v1_1.pdf.