Mem. S.A.It. Vol. 80, 454

Memorie della

Applications using the EGEE grid infrastructure

C. Loomis[1] and V. Floros[2]

[1] Laboratoire de l'Accélérateur Linéaire, Université Paris-Sud 11, Orsay, France
[2] GRNET, Athens, Greece

Abstract. Routine, heavy use of the EGEE grid infrastructure by a broad spectrum of different scientific disciplines is a reality. The infrastructure accommodates diverse applications with requirements ranging from massive, sophisticated data management to pseudo-interactivity with fast turnaround times. Most importantly, the grid infrastructure offers a collaborative environment where scientists can easily share their resources, data, and expertise. Grid technology has advanced scientific knowledge by speeding existing analyses and by motivating larger or more precise calculations. The EGEE project continues to provide this essential platform and works to ensure its long-term availability.

Key words. Grid – Distributed Computing – Applications – Collaborative Computing Platform

1. Introduction

A grid Foster and Kesselman (2000) provides uniform access to distributed computational resources. A grid *federates* resources from different administrative domains or institutions, requiring a coherent user authentication and authorization framework throughout. Moreover, the grid serves simultaneously users from many different scientific domains making it a generic infrastructure. Real grids consist of both the computing hardware and a grid "middleware" that glues the resources together into a powerful, coherent infrastructure.

In grid jargon, a scientific collaboration is called a Virtual Organization (VO) and consists of scientists that what to share their resources and expertise to achieve common (computational) goals. Scientists within a VO benefit from a grid infrastructure because it allows them to federate easily their resources, permitting individuals to perform larger or more precise calculations than they would be able to do alone. Similarly, the grid may also expose exotic or rare resources, allowing use by a larger community. The grid acts as a platform for sharing data and expertise (e.g. code or algorithms). Much like email and the web facilitate communication, the grid allows easy exchange of raw and processed data and the means to analyze them. In short, it provides an ideal platform to analyze data, to publish result datasets, and to combine previous results.

Resource providers and funding bodies also benefit from grid technology. Sharing via the grid allows better resource utilization by allowing "external" users to take advantage of periods of low "local" use and vice versa. Permitting a site to offload peak utilization to other sites potentially allows an institute to purchase resources corresponding to its average needs, rather than its peak needs, thus, reducing costs.

Send offprint requests to: C. Loomis

2. Enabling grids for E-sciencE

EGEE is a publicly-funded project that provides a service-oriented grid infrastructure. The infrastructure serves primarily the European scientific community, although the project has contacts with researchers throughout the world as well contacts with industry to disseminate the use of grid technologies outside academia. Although EGEE provides the overall operational policies and framework, the project does not own any of the resources on the grid. For authorization, Virtual Organizations decide their membership and resource providers decide which VOs can access their resources. As there is no central authority, resource allocation decisions are more difficult; however, the benefit is a more open, extensible infrastructure.

2.1. EGEE grid infrastructure

In nearly all measures, the EGEE grid infrastructure is the largest in the world with 260 sites from 45 countries participating. Around 80000 CPU "cores" and many petabytes of disk space are accessible via the grid infrastructure. The total size of the infrastructure can be somewhat misleading, though. Any particular VO will only have access to resources on sites that authorize the VO. Because of that, the resources available to any particular VO will be much smaller than the totals above.

More than 13000 users from nearly 130 registered VOs take advantage of these resources; the real number of users is undoubtedly higher if one includes the 100 additional, unregistered VOs. EGEE prefers that VOs using the grid infrastructure register with the project; registration allows the project to better understand the needs of the VO, its applications, and how EGEE is impacting the scientific community. Nonetheless, the large number of unregistered VOs attests to the openness of the infrastructure.

2.2. gLite middleware

The gLite middleware[1] provides the core functionality available of the EGEE grid infrastructure. Including software developed both by EGEE itself and by other projects, gLite provides a core software stack on which to run scientific applications.

The gLite middleware uses a Public Key Infrastructure (PKI) to identify users and machines on the grid. The system allows delegation of rights which enables "single-sign" and ensures that tasks running on behalf of users can access other grid services. This Grid Security Infrastructure (GSI) comes from Globus Toolkit[2]. Basing the authentication and authorization framework on GSI allows gLite to be interoperable, at least at the security level, with other infrastructures based on the Globus Toolkit. The distribution of the Globus Toolkit used by EGEE comes from the Virtual Data Toolkit (VDT)[3].

At the resource level, it provides service interfaces for computing and storage resources. The interface to batch systems in widespread use was based on the Globus Toolkit gatekeeper and has been refined in previous projects to scale to the load on the EGEE infrastructure. A new, web-service-based gatekeeper, called CREAM, will be available in the near future. For storage, the project has adopted the Storage Resource Manager (SRM) interface as a standard. There are two implementations in widespread use on the grid—Disk Pool Manager (DPM) and dCache. Both have been developed outside of EGEE itself, but are distributed with the gLite middleware.

Higher-level services include a meta-scheduler and a File Transfer Service (FTS). The meta-scheduler frees users from worrying about detailed scheduling choices by automatically selecting a resource based on a job's resource requirements. Currently the infrastructure (and users) are moving from the older Resource Broker (RB) meta-scheduler to the newer Workload Management System (WMS). FTS reliably transfers files between

[1] `http://glite.org/`
[2] `http://www.globus.org/`
[3] `http://vdt.cs.wisc.edu/`

sites while limiting the total bandwidth consumed by them. Such services are critical for data-intensive science and also for maintaining the overall stability of the infrastructure.

A range of other services complete the core functionality of gLite. The complete list can be found on the gLite website.

2.3. Other infrastructures

There are a large number of grid projects and grid infrastructures throughout the world. EGEE has several peers: Open Science Grid (OSG)[4] in the United States, Nordic DataGrid Facility (NDGF)[5] in Northern Europe and NAREGI[6] in Japan, for example. These offer services similar to those on EGEE but have developed as separate projects largely because of political and financial constraints. Although technology choices have differed on occasion, these infrastructures are very similar. To support fully global scientific collaborations, work to achieve full interoperability is progressing.

All of the above infrastructures federate institutional resources. In contrast, some other grid projects federate different types of resources. The BOINC[7] software, for example, harnesses idle desktop computing resources and is best known for the SETI@home[8] initiative. For calculations that can be divided into short, independent tasks only requiring small amount of data input and output, BOINC-like infrastructures are unparalleled in the huge amount of free computing power they provide. BOINC-like infrastructures, however, are not well adapted to data-intensive applications.

At the other end of the spectrum are infrastructures like DEISA[9] that forms a grid between European supercomputing centers. DEISA uses a common authentication and authorization system that allows data sharing between sites and cross-submission of jobs. This infrastructure is ideal for huge, highly-parallel calculations.

Integration of these other types of resources (desktop computers or supercomputers) on the EGEE infrastructure faces no fundamental technical barrier. These could simply be specialized resources exposed via the same EGEE resource interfaces. In fact, the EDGeS project[10] aims to do exactly that for desktop grids. Unfortunately, there is not yet a similar project to link DEISA and EGEE.

3. Use of the grid

The EGEE grid infrastructure welcomes applications from all scientific disciplines, although the project itself focuses on six strategic scientific disciplines: astronomy, astrophysics & astro-particle physics; computational chemistry; earth sciences; fusion; high-energy physics; and life sciences.

In the first half of 2008, the total CPU utilization on the EGEE grid infrastructure was approximately equivalent to 28000 CPUs running continuously over that 6-month period. The largest users of, and largest contributors to the EGEE grid infrastructure are the four experiments of the Large Hadron Collider (LHC), the world's highest-energy particle accelerator. These four experiments account for 64% of the total use. Other high-energy physics collaborations account for 11% and collaborations outside of high-energy physics account for 25%.

The total CPU utilization has been doubling every 12 to 18 months. In the last quarter of 2008, the LHC will start producing scientific data and we expect a further large increase in the use from the LHC experiments as the focus shifts from simulation to analysis of real data. The fraction of use from non-LHC VOs has been steadily increasing; recent numbers show that more than a third of the use can be attributed to them.

4. Highlighted applications

A short paper such as this cannot summarize the hundreds of applications that are cur-

[4] http://www.opensciencegrid.org/
[5] http://www.ndgf.org/
[6] http://www.naregi.org/index_e.html
[7] http://boinc.berkeley.edu/
[8] http://setiathome.ssl.berkeley.edu/
[9] http://www.deisa.eu/
[10] http://www.edges-grid.eu/

rently running on the EGEE grid infrastructure. Broader surveys have been carried out by the project and are available elsewhere EGEE Project (2000). Here, three applications from different scientific disciplines are highlighted, showing the broad utility of grid technology and demonstrating its value for advancing scientific knowledge.

4.1. LHC applications

The four experiments[11] that will take data using the LHC[12] rely entirely on the EGEE grid infrastructure to provide the storage capacity for their data and the computing power to analyze it. Some characteristics of the accelerator and the detectors put the data management and analysis challenge into perspective. The accelerator will produce particle interactions at a rate of 40 MHz. Within those interactions, interesting ones (those to be saved) occur at a rate of 100 Hz, with each one producing a data record of 1-10 MB. The accelerator will operate continuously, producing petabytes of data to store every year. Moreover, the analysis of these data requires 1 to 10 times more simulated data. The real challenge for the LHC experiments is data management—effective storage, transfer, and retrieval of these data.

Each of the LHC collaborations consists of thousands of physicists from all over the world. Using the grid, they have created a tiered infrastructure[13] that duplicates all of the raw data at 11 "Tier-1" centers. These centers then serve the data to a larger number of "Tier-2" centers for analysis. With this system, the raw data flows from the accelerator ("Tier-0") where it is produced to all of the physicists throughout the world. All of the data are cataloged, allowing physicists to find the appropriate raw data based on meta-data queries.

Simulated data are vital for the scientific analysis of the raw data. The full simulated data set is as large or larger than the raw data set; in contrast to the raw data, these are not produced at a single site. Nevertheless, the grid must be able to handle large scale transfers of simulated data between the distributed producers of these data and the physicists who use them. This is a more difficult problem than the distribution of the raw data.

Over the past several years, various exercises have been used to test whether the EGEE grid infrastructure can respond to those challenges. Problems with the robustness and scalability of services have been found and fixed. Recent exercises have been successful and have given confidence that the grid can handle the LHC's data management requirements. The proof will come when data taking starts in the last quarter of 2008.

4.2. WISDOM

The WISDOM[14] collaboration aims to speed the development of new drugs for neglected and emerging diseases. To accomplish this, the collaboration uses a "virtual screening process" that evaluates the molecular interactions between a particular target protein and molecular compounds taken from a large database. This *in silico* screening can scan a larger number of compounds more quickly than traditional *in vitro* assays.

The initial calculation was done in the summer of 2005 and evaluated 42 million compounds against a protein of the malaria parasite. This calculation was completed in six weeks and produced a database of around 750 MB. The success of this compaign interested others working with the avian flu virus. Reusing the software for the initial calculation and the EGEE grid resources, a calculation targeting the avian flu virus was run in the spring of 2006 and took only three months to go from conception to conclusion. Subsequently, another calculation has been run against malaria and another is currently being run against avian flu.

All of these calculations have identified new potential drugs and those compounds are now being tested *in vitro* to determine their efficacity. Some have been shown to be as ac-

[11] http://lhc.web.cern.ch/lhc/LHC_Experiments.htm

[12] http://lhc.web.cern.ch/lhc/

[13] http://lcg.web.cern.ch/LCG/

[14] http://wisdom.healthgrid.org/

tive or more active than known compounds. The collaboration is now investigating using the grid to run detailed molecular dynamics calculations to determine if the lists of identified compounds can be further refined before initiating the expensive and time-consuming *in vitro* tests.

This calculation highlights a couple important aspects concerning the grid. First, the grid is an excellent collaborative platform. WISDOM initially started as a limited endeavor involving only a few researchers. This has since grown to a global collaboration. The easy access to computing resources and opportunity to store the results offered by the grid has been an important factor in this transition from a "local" to "global" program.

Second, it shows the willingness of the resource providers to support high-quality science. Most of the resources used in these calculations were not owned by biomedical institutes; nonetheless, they provided access to their resources for these calculations. Being able to share resources between disciplines to make better overall use of the infrastructures is one of the fundamental tenets of grid technology.

4.3. GEOSCOPE

GEOSCOPE[15] is a global network of seismic centers. This network works to localize earthquakes, to determine the rupture mode of faults, and to provide a 3D tomography of the interior of the planet. The initial calculation to localize an earthquake and calculate the rupture mode of the fault takes several days on a single machine. Because it takes so long, this prerequisite calculation delayed other, more detailed analyses. The initial motivation for using the grid was to reduce the delay from a few days to a few hours.

Given the unpredictable nature of earthquakes, it is not possible to reserve the necessary computing power in advance and it is wasteful to leave a large number of computers idle waiting for the next earthquake. Sharing resources via the grid allows this calculation to use a large number of CPUs with short notice without having to waste resources. Currently, the results from the initial calculation are now obtained within 24 hours with only 3-4 hours of that from the calcuation itself. The remaining time is from the data transfers from the 28 participating seismic stations.

Because of this positive experience with the grid, this group decided to analyze the noise from the 28 seismic stations over the last 25 years. The goal of this calculation was to understand better the noise from each center and to design eventually better seismic stations. This ambitious calculation was successfully completed using the grid.

This is a good example of what typically happens with use of grid technology. Because the computational aspects are no longer the limiting constraint, one finds larger, more ambitious analyses to do.

5. Evolution

EGEE is actually a series of projects funded in part by the European Commission. Its roots originate in 2001 with the European DataGrid project, which has been followed by the EGEE, EGEE-II, and EGEE-III projects. The latest, EGEE-III started in May 2008 and will continue to April 2010, to be followed, hopefully, by a perennial European Grid Infrastructure. Over this time, the project has changed from an R&D program to a grid service provider. In parallel, the application activities have changed from running a few example applications to providing support services to a large and growing user community. In EGEE-III, the program of work has been redefined to reflect this, focusing on user support, community building, and advanced functionality.

5.1. User support

Users of the grid infrastructure must receive prompt, accurate responses when problems arise on the grid. In the past, most of this help was provided by the operations activity within EGEE. Naturally, people involved were best suited to handle issues related to transient operational problems with grid services. The application activity in EGEE-III will now com-

[15] `http://geoscope.ipgp.jussieu.fr/`

plement this with a team focused on how to use the grid and its services. This team will survey the available documentation, review it, and request (or produce) new documentation where there are holes. This should help resolve a long-standing issue concerning documentation.

Help with porting applications to the grid is another important area. Often people come to the grid with vague ideas how the grid can improve their scientific work. For people not experienced with grid technology it can be extremely difficult to discover the best services and techniques to use when running an application on the grid. An application support team will act as consultants to discuss with scientists their applications and to develop a specific porting plan. This helps people better understand the EGEE grid services and avoids initial frustration with the significant learning curve associated with grid technology.

5.2. Community building

Given the size of the user community (more than 13000 people) relative to the number of people in the application activity of EGEE-III (around 300), the project cannot hope to provide comprehensive services to the full user community. Instead the project must work to make the user community as self-supporting as possible. First, the project sponsors User Forums each year where the full user community can come together to discuss their achievements, to exchange techniques, and to "steal" useful generic services created by others. Second, the project sponsors similar meetings for individual scientific disciplines to disseminate information about grid technology and to focus on topics relevant only to that discipline, for example, access to particular databases or software from the grid.

5.3. Advanced functionality

The gLite middleware provides the core grid functionality. Real applications, however, usually need to use quite a lot of software on top of gLite. Although some of this additional software really is specific to the application, much of it can take the form of generic functionality or services.

The user community and the application activity in EGEE have been very active in developing some high-level functionality and services. A meta-data catalog (AMGA), medical data management services (MDM), and job management services (Ganga, Diane) have been developed within the application activity. In addition, the user community has been very active in working groups to find technical solutions to issues like low-latency scheduling, MPI on the grid, and job prioritization.

5.4. Commercial software

As the user community has expanded, there has been more demand for the use of commercial software on the grid. The challenge to using commercial software on the grid is not technical. The issue is to find a licensing model that works well with the distributed nature of the grid but still guarantees a revenue stream for the vendor. Two different models have been tried: a VO-based scheme and a client/server scheme.

For Gaussian, a popular computational chemistry program, a separate VO was created. Institutes that have a site license for Gaussian and also provide resources for EGEE can deploy Gaussian for use by members of this VO. The VO itself is only open to members of institutes that have site licenses. This arrangement allows members of this VO to access more resources without having to deploy those resources themselves.

However, this arrangement has several disadvantages. First, it puts a strong legal burden on the VO manager to check that all members come from institutes with a valid Gaussian license. Second, it makes it difficult for users to use other commercial software if a separate VO is needed for each package. This arises because the grid does not easily permit sharing of data between different VOs. Third, making the software on a site only accessible to a single VO takes additional effort from the system administrators.

A different licensing model is being used for MATLAB. Here there are separate licenses for the server and client. This allows a site to setup and advertise the service side of the "Parallel Computing Toolkit" without having to worry about limiting access to a particular set of users. The server itself will check if a particular client has a valid license or not. On the client side, the user need only worry about purchasing his own license. This model, currently being tested, fits better with the service-oriented nature of the grid infrastructure.

Because the user community has grown and because users are demanding access to commercial software on the grid, software vendors are starting to take notice. EGEE is becoming are target platform for them and hopefully they will be responsive to our experiences concerning licensing models on the grid.

6. Conclusions

Routine, heavy use of the EGEE grid infrastructure by a broad spectrum of different scientific disciplines is a reality. The grid infrastructure can accommodate diverse applications with needs ranging from the intense data management requirements for the LHC experiments to fast turnaround times from the GEOSCOPE application. Most importantly, the grid serves as an ideal collaborative plaform allowing scientists to analyze, to publish, and to combine previous results all within the same environment. Examples, such as WISDOM that has grown from a small group into a global collaboration, best illustrate the potential of grid technology to advance scientific knowledge.

EGEE and grid techology have evolved over the years and will continue to evolve to better meet the needs of the user community by expanding the types of resources available on the grid, by introducing new services, and by incorporating promising technologies. The hope is that the EGEE projects will be followed by an initiative to guarantee the availability of this critical computational platform into the future.

Acknowledgements. The authors gratefully acknowledge the support of the European Commission and of the various national governments participating in the EGEE project and their collaborators within the application activity (NA4) within the project. This work is co-funded by the European Commission through the EGEE-III project, contract number INFSO-RI-222667.

References

Ian Foster and Carl Kesselman, "The Grid: Blueprint for a New Computing Infrastructure", Morgan Kaufmann Publishers, 1999 and "The Grid 2: Blueprint for a New Computing Infrastructure", Elsevier, Inc., 2004.

EGEE Project. "Application Deployment Plan", https://edms.cern.ch/document/722131/2, 2006 and "Summary and Evaluation of Application Usage of EGEE-II Services and Updated Application Deployment Plan", https://edms.cern.ch/document/722132/3, 2007.