GRIDSTART

GRIDSTART Technical Bulletin

Issue 1 • October 2002

Contents:

- Editorial -

Turning the GRID into Reality

We are now seeing a rapid convergence of computing and communication technologies at all levels. On one hand, computing systems are becoming highly connected. On the other, mobile phones are becoming more and more intelligent. A new generation of communicating, intelligent sensors and other devices are being developed. Larger and larger repositories of data are now moving on-line and increasing bandwidth means that these data can be accessed from anywhere at any time.

The GRID will be one of several mechanisms to exploit the highly connected sea of networked computers sensors and data repositories. It will be part of a rich landscape of pervasive computing, high-bandwidth communications and immediate access to data. It will reach into our professional lives by making data, computational power and scientific simulations widely available. It will impact our private lives by enabling improved healthcare and new ways of learning and by providing access to entertainment, to our cultural heritage and to new, as yet unimagined, services.
15 years ago the Internet was largely unknown outside scientific institutions. Now it provides a major public access point for information, services and entertainment. The GRID is widely seen as the next generation Internet. Science alone does not have the clout to take GRID developments to the high level of capability required by the scientific community. The challenge facing the GRID community is to move its technology rapidly to a point where commercial exploitation can take over and become the driver taking it forward as a mainstream and sustainable technology. This will quickly lead to significant benefits for scientists around the world because the “industrial strength” GRID will be a clear enabler of better science and of the most demanding scientific applications. This parallels what happened with the Internet and it will happen with the GRID if we get it right.

GRIDSTART has been funded by the European Commission to bring together current GRID activities across Europe so that they can share common developments and

In this issue:

This, the first issue of the GRIDSTART Technical Bulletin, features a selection of articles describing technical work in progress around the world. It is our main focus: to let the GRID development community know what is going on elsewhere and to encourage collaboration, sharing of ideas, the establishment of standards and, where appropriate, some healthy competition. Not surprisingly several of the articles feature work with Globus and with Unicore and deal with efforts to make the two operate together. Important operational requirements such as Quality of Service are also covered. Most significantly the implementation of real applications is given prominence. Some of these come from the world of big science. This is what one would expect because the requirements of big science have been the raison d’être for the GRID. However the GRID will be driven forward to commercial reality and sustainability because its functionality supports effective ways of doing business. It is therefore encouraging to see that, even at this early stage in GRID development, we can already report on work towards making the GRID usable by the wider world beyond science.

This publication will only make a worthwhile contribution to the long and challenging road that lies ahead if it gets the support of the GRID community. That support means the submission of articles, news items and any other material you feel is of interest to the rest of us involved in making the GRID a reality. Please send your contributions to bulletin@gridstart.org. We look forward to promoting your work and ideas in future issues.

technologies. This technical bulletin is part of an effort to make GRID developments widely known across the technical community. The central objective of GRIDSTART is to maximise the impact of these activities by encouraging coherent developments and a combined approach towards standards. Through the consolidation of GRID technical advances, the identification of synergies between projects, the stimulation of early take-up by industry and commerce and by playing a full role in the setting of international standards particularly through the Global GRID Forum, GRIDSTART aims to maximise the potential of the GRID for science and to speed its passage from the laboratory into wider commercial, business and private use.

No one region can act alone in the development of what will be a global resource for the use of all communities. In this, the first issue of the GRIDSTART technical bulletin, we are pleased to give pride of place in the foreword to an announcement by Alan Blatecky of the National Science Foundation about the start of a global collaboration aimed at bringing together GRID developments worldwide. I leave it to Alan to tell you about this exciting and important development.

Francis Wray, EPCC, f.wray@epcc.ed.ac.uk

- Foreword -

Closer cooperation between the US, EU and AP in the development of GRID technology being discussed

Alan Blatecky, NSF, ablateck@nsf.gov

Representatives from the US, EU and Asia-Pacific met at the GGF5 meeting in Edinburgh on July 24 to discuss global coordination of developments in GRID technologies which are being funded by a number of science agencies from these regions. The representatives felt that effective coordination would benefit the global science and research community, and established an international Task Force to develop a working paper based on the ideas presented in Edinburgh. This paper would be presented and discussed at the GGF6 meeting in Chicago.

The Task Force was charged to address three components. The first component was to develop a set of objectives to guide the activities. The second task was to outline several goals which could be implemented quickly, within 3-6 months. These goals would include activities such as collecting and synthesizing grid research activities worldwide, establishing effective communications channels or involving more international reviewers in proposal and project reviews. As soon as these first two components have been discussed and approved, the Task Force will be asked to expand their focus to address some longer term goals such as shared testbeds and links, share resources, development of a process for coordinated software releases and so forth.

The working document detailing progress of this initiative is available at www.gridstart.org

- Articles -

CROSSGRID - Project at a Glance

Piotr Nowakowski, Cyfronet, September 2002, ymnowako@cyf-kr.edu.pl

Summary
The CrossGrid Project belongs to a wide family of Grid projects funded by the European Union within the structures of the 5th Framework Programme. The central aim is to establish a Grid computing community for areas which – unlike earlier applications that do not require extensive control and frequent input, and hence yield easily to parallelization and distribution in Grid-type environments – are characterized by the presence of a human agent (called a *person in the loop*). Numerous software engineering tools (new releases of the Globus toolkit, XML tools, MPI interfaces and portal access solutions) are being tested as part of the Project. The Project unites researchers from various European countries (21 research and HPC institutions in 11 countries in all) and participates actively in the GridStart and GGF consortia, thus aiming for visibility on a global stage. The estimated duration of the Project is three years, with summer 2005 being the expected deadline for results.

Introduction
Grid-related research is currently on an upswing within the European Community and elsewhere. Numerous Grid projects are either being set up or already in progress. The computing requirements of modern industries and scientific communities necessitate international cooperation in providing the necessary solutions both to individual institutions and mutual enterprises. The CrossGrid Project aims to create a pan-European High Performance Computing (HPC) platform intended to serve applications which require frequent human interaction and real-time (or near-real-time) responses from the system.

The Project consists of the development of Grid-enabled applications which will demonstrate the achievements of CrossGrid and put them to practical use, a set of new Grid tools designed specifically for enabling the aforementioned “person in a loop”-type interaction, a collection of Grid services which elaborate on the basic functionality of Grid middleware (such as the Globus toolkit) to provide further optimization and extended monitoring capabilities for system administrators, and – finally – a European testbed which will utilize computing resources of various European partners (both academic and commercial) in testing CrossGrid solutions and applying them to real-life cases.

The CrossGrid Project is closely related to another Grid undertaking, the European DataGrid (EDG), coordinated by CERN, which is currently in its second year of development and has already yielded interesting insight with regard to the software components it utilizes and the procedures it follows.

Detail
The CrossGrid Project consists of five Work Packages, each of which is further divided into tasks. The division is as follows:

1. **Work Package 1: Applications**

   The applications which have been selected for development within CrossGrid are at the same time compute- and data-intensive. Each of these applications can be put to practical use once the project concludes and additional applications may be developed for deployment on the CrossGrid testbed.

   - **Interactive simulation and visualization in a biomedical system**

   The aim of this task is to create an application which computes and visualizes the blood flow in vascular systems for pre-treatment and planning in vascular interventional procedures. This application will use CrossGrid for its computational and data management needs. A special module, called the Grid Visualization Kernel (GVK) is being developed for the manipulation of image files and automatic translation of visual data generated by dedicated hardware (i.e. PET scanners) into standardized formats (such as NetCDF).

- **Flooding crisis team support system**

This is a decision support system designed for predicting and visualizing potential flood hazards based on weather forecasts and observed rainfall in selected river basins. Advanced meteorological (ALLADIN/LACE, MM5), hydrological (HEC-1, TR-20, HSPF, NLC, HBV) and hydraulic (FESWMS, RMA2, NLN) models are being ported to the Grid environment. Models with unsatisfactory response time will be optimized using tools developed in Work Package 2.

- **Distributed data analysis in High Energy Physics**

This application will be a computing interface which uses large distributed databases for storage and manipulation of data generated by the new LHC collider which is currently under construction at CERN. The amount of data generated by the LHC will require new software tools and hardware infrastructures to sift through the millions of events (individual collisions) and find the handful that are of value for researchers. Hence, new data mining, storage and access optimization methods will be called for with regard to the extensive array of storage elements present in the Grid.

- **Weather forecasting and air pollution modelling**

The objectives of this task are to provide a representative collection of sample applications of Grid tools for use by the atmospheric/oceanographic community. A data mining system will thus be developed (in collaboration with the HEP application team mentioned earlier) for analysing of the vast geographically-distributed archive of operational data. Data mining techniques, including association rules, linear and non-linear correlation methods and neural networks (Self-Organized Maps) will be developed for the extraction of interesting patterns within the DBs.

2. Grid Application Programming Environment

This Work Package includes software which facilitates the development and tuning of parallel distributed high-performance and high-throughput computing applications within the Grid infrastructure. An MPI verification and debugging tool (to check whether user applications comply with MPI standards) will be developed, along with a set of semiautomatic performance evaluation tools. These tools will use the concept of *probes* – pieces of external code inserted into individual applications to gauge the performance of key function calls and rate it against a set of predefined metrics and benchmarks. This information can then be used by software developers to further optimize their applications, as well as by users who wish to know the estimated duration of computations for the problems they are submitting.

3. New Grid Services and Tools

This Work Package deals with new Grid functionality as required by the specifics of CrossGrid applications. Two distinct categories of services exist:

- **Application-Specific Services**

This category includes services tailored to the needs of individual applications, such as the Grid Visualization Kernel (GVK) for the biomedical application or the User Interaction Services which translate requests submitted by users of each application into generic jobs which can be handled by Grid scheduling agents.

- **Generic Services**

A layer of basic services which are the source of Grid functionality. Currently, CrossGrid is planning on using Globus Toolkit 2.0 (mostly to maintain compatibility with the DataGrid testbed); still, we are in the process of evaluating an OGSA (Open Grid Services Architecture) version of Globus (GT 3.0), which uses XML for communication. The Globus Replica Manager and Catalog are used, as well as GRAM, GIS/MDS and GridFTP. The GridFTP itself is being optimized for access to secondary and tertiary data storage elements (i.e. disk drives and tape drives). CrossGrid will also inherit the DataGrid scheduler and job submission service, which will be extended for the purposes of providing efficient job allocation for applications developed in WP1. Finally, a set of Web-based portals will be developed to enable CrossGrid users to easily manage their accounts and submit jobs at any site which has Internet access.

4. International Testbed Organization

This Work Package aims to provide a consistent international testbed spanning numerous sites across Europe. Computing elements have been set up in various European countries and we are currently in the process of linking them and acquiring the necessary security certificates.

5. **Project Management**
   The final Work Package devotes itself to managerial issues. ACC Cyfronet (Krakow) is the coordinator of the Project, with special teams assigned to work on issues of CrossGrid architecture, dissemination and supervision.

Conclusions

The CrossGrid Project is still in its early stages. The initial three months saw the development of Software Requirement Specification (SRS) documents for each software module. This was followed (in the summer of 2002) by another set of detailed design documents, describing the intended functionality and structure of each module in terms of UML component diagrams, class diagrams, interface specifications etc.
The goal for the next three months is to develop an initial software release, consisting of prototypes of all applications and several basic tools and services (some of them procured through cooperation with DataGrid). This release will be deployed on the development testbed prepared as part of WP4.
The Project will then enter a phase of incremental releases, with new functionality being added each time.

CrossGrid will continue to actively operate within the framework of European and global research communities in order to publish results and increase international awareness of the Project. A presentation for the IST conference in Copenhagen (November 2002) is currently being prepared.

References

1. http://www.eu-crossgrid.org

DAMIEN - Concept and Project Status

Edgar Gabriel, Michael M. Resch, High Performance Computing Center Stuttgart (HLRS), September 2002,
gabriel, resch@hlrs.de

Summary

DAMIEN is a project that started out in early 2001 to continue the successful work of the European pilot project for GRID-Computing METODIS. METODIS had successfully shown the feasibility of the GRID-approach in industry. The objective of DAMIEN is to develop further building blocks for a middleware environment for distributed industrial simulation and visualisation in the GRID. Besides the multi-protocol MPI-library PACX-MPI for heterogeneous networks this includes the handling of Quality of Service requirements in distributed simulations. The coupling code interface MpCCI will allow linking distributed applications. Tools for performance analysis (Vampir) and performance prediction (DIMEMAS) are extended to be made GRID-aware. Applications from industry serve as test cases for the developed software.

The project has recently released the prototypes of all tools presented above. They are currently tested by the industrial partner, which has the need to couple several of its resources on various sites distributed all over Europe.

Introduction

The Grid is generally seen as a concept for `coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organization'. The original idea came from scientists, who were mainly interested in the scientific solution of their problems. Their jobs can be executed on any machine, respectively on any set of machines, to which they have access. Like in Metacomputing, which was a popular concept in the mid-90s, the idea of distributing jobs onto several machines is an important part of the Grid concept. The need for doing so is mainly coming from applications, which have a high demand for computational resources, for increasing throughput on the machines and reducing turnaround times.
Distributing a parallel job onto several machines imposes however many problems to the user and the application. Most of the problems stem from the fact, that the Grid is heterogeneous in several senses. First problems arise with different data representations on different machines, which require data conversion either in the communication layer or in the application. Various processor speeds, differences in the available memory and the usage of shared resources require the application to have

a smart initial load distribution and dynamic load balancing. The differences in the communication characteristics for communication between processes located on potentially different sites require distinct programming techniques for hiding the wide area latency and dealing with the low bandwidth. Another level of heterogeneity is introduced by the different methods of access to the different machines in a Grid, e.g. ssh, UNICORE or globus.

Detail

Software life-cycle of an application in the Grid

A scientific application has to pass two different stages during its lifetime. In the development phase the application is modified to enable the solution of certain problems. The second stage is the production phase, where the code is used for solving the problems which it was originally designed for.

The boundaries between these two stages are not well defined. It is quite common to modify programs which have already been used for production runs. Therefore, during its lifetime, the code will pass both phases several times. Starting point for a Grid-enabled application is either a sequential code or a parallel (MPI-) code, which has already been used on HPC platforms. Although rarely done, a program may also be developed from scratch targeting from the very beginning for distributed Grid-environments. Furthermore, applications which combine several scientific areas are getting more and more popular among scientists. These applications, known as multi-disciplinary or coupled applications, are a promising approach for Grid-environments, since with careful handling by the end-user the communication pattern of the application may reflect quite well the communication characteristics of the Grid-computer.

In the second phase, the user will already have a working code. The end-user has now to figure out, how many processes and which machine(s) he should use for optimal performance for the given problem. This decision is usually driven by the need for minimizing the turn-around time of his simulation. This, on one hand is depending on the optimal performance of the application on the used machines and networks, and on the other hand by the current load of the machines. After determining these parameters, he can finally start the job on the cluster of machines he decided to use. The execution of the application has again different problems, like co-scheduling of all used resources, including the network, and authentication and authorization of the user on all resources.

The DAMIEN toolset

The goal of the DAMIEN project is to support the end-user during both steps presented above by developing a tool-set for Grid-computing, based mainly on existing and widely accepted tools from the area of high performance computing.

For the development phase, the tool-set consists of the following elements:

- The parallelization paradigm supported in the DAMIEN project is based on the Message Passing Standard MPI. The communication library used is called **PACX-MPI** and is an optimized implementation of MPI for Grid-computing environments.
- For coupling of several applications, the Code Coupling Interface **MpCCI** is extended to support Grid-environments. MpCCI provides the end-user data exchange functions on a higher level by providing all required functionality (e.g. interpolation between different meshes).
- For performance analysis of a Grid-enabled scientific application, **MetaVampir** and its tracing library **MetaVampirtrace** are included in the tool-set.

For the production phase, the following tools are included in the DAMIEN tool-set:

- **DIMEMAS** is a performance prediction tool which can be used for estimating the behaviour of the application when certain parameters of the run, e.g. latency or bandwidth, are modified. This can be used than for determining which machine (or combination of machines) the user should use for achieving the best performance.
- A **Configuration Manager** is developed in the frame of the project for launching and monitoring the application.
- A critical step for achieving the required performance for a production run is to ensure a certain level of quality of the network between the machines. Therefore, a **Quality of Service Manager** (QoS Manager) is developed in the frame of the DAMIEN project, having the goal to ease the modification of certain QoS parameters whenever required.

Conclusions

Scientific applications have to pass several steps before they can be used in production. Each of these steps requires several tools and libraries to support the end-user. The DAMIEN project develops a tool-set which is based on widely accepted tools and libraries from the area of high performance computing, and supports the end-user both during the development of the application and the execution of production runs.

The project has recently released the prototypes of all tools presented above. They are currently tested by the industrial partner in the project. The participating company has the need to couple several of its resources on various sites distributed all over Europe for solving problems, which are too big to be solved on one of their available resources. The final release of the tool-set is expected for spring 2003.

EUROGRID - European Testbed for GRID Applications

Hans–Christian Hoppe, Pallas GmbH Daniel Mallmann, Forschungszentrum Jülich, September 2002,
hans-christian.hoppe@pallas.com, d.mallmann@fz-juelich.de

Summary

The EUROGRID project (IST–1999–20247) is a Research and Technology Development project funded as part of the IST Programme by the European Commission with a project term from November 1, 2000 through October 31, 2003. EUROGRID demonstrates the use of GRIDs in important scientific and industrial communities, addresses the specific requirements of these communities in a European GRID middleware, and highlights the benefits of using trans–European GRIDs.

The objectives of the EUROGRID project are:

- Establish a trans–European GRID of leading High Performance Computing centres.
- Operate and support the EUROGRID software infrastructure.
- Develop important GRID software components and integrate them into EUROGRID (dynamic resource broker, accounting and billing, interface for coupled applications and interactive access).
- Demonstrate distributed simulation codes from significant application areas (biomolecular simulations, meteorology, coupled CAE simulations).
- Make the EUROGRID middleware available to other GRID projects and prepare productisation.

Introduction

Computational GRIDs are rapidly becoming a cornerstone of future high–performance computing research and infrastructure. In the world–wide computer science research community, the benefits of GRID technology have been amply demonstrated, and GRID systems are beginning to be phased into daily use at universities and research centres, for instance for remote experiments in biology and astrophysics.

However, the uptake of GRID technology in the broad community of scientists that depend on computational simulation techniques for their work has yet to happen, and industrial end–users are lagging even farther behind. The EUROGRID project demonstrates the use of GRIDs in selected scientific and industrial communities, addresses their specific GRIDs. The choice of communities ensures that the results can be transferred to a broad range of potential users.

Within the EUROGRID project four application–specific GRIDs are being operated:

- Bio–GRID for biomolecular simulations
- Meteo–GRID for localized weather prediction
- CAE–GRID for coupled engineering applications
- HPC–GRID for general HPC end–users

The EUROGRID software is based on the UNICORE system developed and used operationally by the leading German HPC centers. Functional extensions are being developed in the following areas:

requirements, and highlights the benefits of using GRIDs. The choice of communities ensures that the results can be transferred to a broad range of potential users.
Within the EUROGRID project four application–specific GRIDs are being operated:

- Bio–GRID for biomolecular simulations
- Meteo–GRID for localized weather prediction
- CAE–GRID for coupled engineering applications
- HPC–GRID for general HPC end–users

The EUROGRID software is based on the UNICORE system developed and used operationally by the leading German HPC centers. Functional extensions are being developed in the following areas:

- High–performance data transfer
- Dynamic resource discovery and brokering
- Accounting and billing
- Application coupling
- Interactive access and application steering

These extensions are being integrated into the operational EUROGRID software, and evaluated in operational use by the domain–specific GRIDs.

Detail

Bio–GRID
Scientists in biomolecular research regularly use a set of established simulation and visualization packages and molecule databases. Incompatibilities between packages, substantial differences in user interfaces and the need to become familiar with the high–performance execution systems complicates the daily work. Non–experts have to face a particularly steep learning curve before they can start to become productive. In the Bio–GRID, intuitive user interfaces for biomolecular packages and compatibility interfaces for their databases are being developed, resulting in an integrated biomolecular toolkit that allows streamlined work processes, and access to all systems in the BIO–GRID with a uniform and intuitive user interface.

Meteo–GRID
For many scenarios, exact meso– or microscale weather predictions are required: agriculture, pollution prediction, traffic and public event planning all depend on accurate localized weather data.
To accommodate these, a flexible framework for on–demand localized weather prediction is being developed in EUROGRID: from a request issued from a workstation somewhere on the Internet, a run of a special local weather model is scheduled based on the regular forecast data, and the result data is transferred to the customer. The EUROGRID middleware takes care of transporting requests and results, and runs the local forecast on a Meteo–GRID system.

CAE–GRID
Use of CAE packages is standard practice in many industrial sectors: automotive, aerospace, electronics etc., but it often require more computing power than is available locally. Thus, jobs need to be run at external computing centres, or internal computing resources have to be combined. EUROGRID demonstrates the use of GRID technology in two key areas: coupling of applications that simulate different aspects of a complex system, and providing billed access to computing resources and standard engineering applications (ASP).

HPC–GRID
The participating HPC centres have established a trans–European computational GRID, making available part of their local resources for the EUROGRID users in science and industry. This GRID serves as a testbed for the development of distributed applications, the integration of full production systems in a GRID environment, and for the close cooperation between sites operating a wide range of different HPC platforms. Focus is on the application domains not represented in the domain–specific activities discussed above.

GRID Infrastructure
EUROGRID relies on a proven GRID system originally developed in the German UNICORE Plus project. It consists of three distinct software tiers: a *client* the user interacts with to construct, submit and control the execution of computational jobs, a *gateway* acting as single point–of–entry into the protected domains of the supercomputing centres, and a *server* that schedules and executes the locally. All components and protocols are written or specified in Java.

The system emphasizes strong authentication and data security and relies on X.509 certificates and SSL. The server uses platform–specific *incarnation databases* to translate abstract jobs to concrete command sequences for the execution system.

Within EUROGRID, functional extensions as required by the end–users are being developed and integrated:

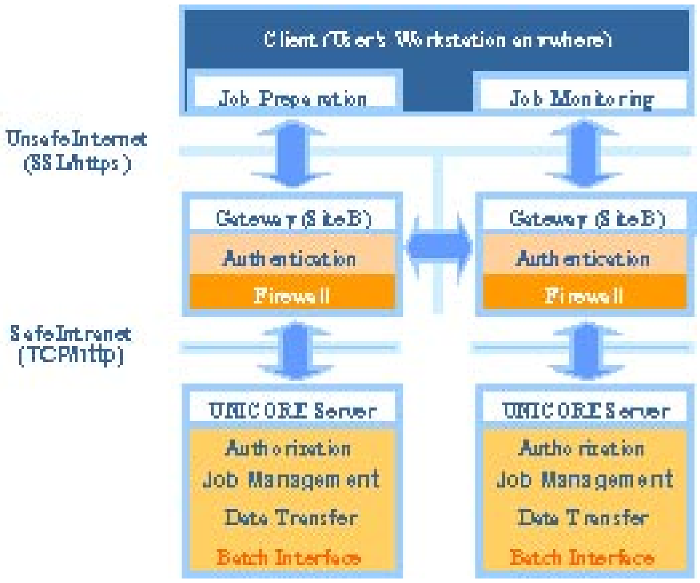

- High–performance data transfer: integrate third–party transfer mechanisms and allow for overlapping of data transfer and processing.
- Dynamic resource discovery and brokering: identify available resources that match criteria specified by the end–user and broker for best–case execution according to a variety of metrics (cost, time–to–completion, reliability).
- Accounting and billing services: in an ASP setting, keep track of resource usage and handle invoices/payment in a secure manner.
- Application coupling: accommodate proven techniques for the coupling of simulation applications (Corba, MpCCI, ...).
- Interactive access and application steering: allow for interactive use of resources (shell commands) and demonstrate the integration of application steering mechanisms.

Conclusions

In the first two years, the EUROGRID project has clearly demonstrated the viability of trans–national GRIDs and the benefits to end–users: access to a broader range of computing resources, ability to run larger jobs, and better reliability due to increased flexibility. Development of domain–specific extensions has progressed well, with *plugins* for important applications (Gaussian, CPMD, Molpro, Gromos etc.) emerging as the key element for increasing acceptance of GRIDs and user productivity. A plugin presents a friendly interface for running an application to a scientific end–user, hides system and GRID details and works at an appropriate level of abstraction. An advanced plugin for driving the on–demand localized weather prediction is nearing completion, as is a plugin for seamless integration of Corba–based application coupling.

In the current version of EUROGRID, data transfer has been enhanced with the integration of GidFTP. A resource broker component will become operational end of 2002, and the first implementation of interactive access allows the execution of selected shell commands and the steering of applications based on input streams.

The source code of the base EUROGRID system and of many extensions is available under an Open Source license from www.unicore.org/downloads.htm

---

GrangeNet - GRID and Advanced Computing Services in Australia

David Abramson, Monash University, Paul Davis, GrangeNet, davida@mail.csse.monash.edu.au, paul.davis@grangenet.net

Summary

Grangenet is an incubator for the development and deployment of innovative Grid and Advanced Communications services. To facilitate this GrangeNet will install, operate and develop a multi-gigabit network servicing the East Coast of Australia. It will be part of the global research and education network through AARNet's capacity across the Pacific. It will seek interconnection and collaboration with other research and education networks both nationally and internationally.

As part of the GrangeNet, a number of projects are aimed at building Grid aware applications. One of these is the Australian Nimrod Testbed (ANT), which aims to deploy the Nirmrod/G package onto a range of GrangeNet connected machines, and then to mount a number of national demonstrators.

Introduction

GrangeNet (GRid And Next GEneration Network) is a joint venture between the Australian Academic Research Network (AARNet), the Australian Partnership for Advanced Computing (APAC), the Distributed Systems Technology Centre (DSTC), CISCO and PowerTel. Participants include AARNet members, APAC partners and several research and commercial organisations. In addition seven organisations responsible for advanced computing and communications programs in Asia and North America will be collaborating with GrangeNet.

The program, supported by the Australian Government through the Building on IT Strengths (BITS) Advanced Networks Program (ANP) of the Department of Communications, Information Technology and the Arts, is an incubator for the development and deployment of innovative Grid and Advanced Communications services. To facilitate this GrangeNet will install, operate and develop a multi-gigabit network.

GrangeNet will be driven by demanding applications from the Australian research community, from users of advanced computing systems and from specific communities of interest. It will support a number of user communities that have focussed interests. Some of these communities are expected to be in computational physics, bio-informatics, astronomy, computational engineering, on-line health, environmental modelling, distance education and media services.

One such body of work is the Australian Nimrod Testbed – ANT. Nimrod/G is a specialized parametric modeling system that employs a simple declarative parametric modeling language to express a parametric experiment. Nimrod/G is grid aware and uses the Globus toolkit for high level grid services. The aim of ANT is to deploy Nimrod/G and utilise it to perform a number of national demonstrators that both highlight the power of the computational grid for this type of work, but also the capabilities of high speed low latency networks of the type implemented in GrangeNet.

Detail

The GrangeNet Network

The GrangeNet backbone linking Melbourne, Canberra, Sydney and Brisbane consists of dual 2.5 Gbps DWDM PoS services from PowerTel using Cisco active regeration. Initially eight organisations will be connected using gigabit Ethernet (GbE); peering between AARNet and GrangeNet is facilitated by the collocation of the AARNet and GrangeNet POPs (Points-of-Presence).

GrangeNet will be part of the global research and education network through access to the dual STM-1 capacity that AARNet has acquired on Southern Cross Cable Network (SCCN). It will seek interconnection and collaboration with other research and education networks both nationally and internationally.

ANT ~ The Australian Nimrod Testbed

Over the past several years as part of a DSTC funded project, we have developed a specialized parametric modeling system, called Nimrod, to support large parameter sweep applications. Nimrod uses a simple declarative parametric modeling language to express a parametric experiment and provides machinery that automates the task of formulating, running, monitoring, and collating the results from the multiple individual experiments. Equally important, Nimrod incorporates a distributed scheduling component that can manage the scheduling of individual experiments to the processors of a cluster or idle computers in a local area network. Together, these features mean that even complex parametric experiments can be defined and run with little programmer effort. In many cases it is possible to establish a new experiment in minutes.

Over the years, Nimrod has been applied to a range of application areas, including Bioinformatics, Operations Research, Network Simulation, Electronic CAD, Ecological Modelling and Business Process Simulation. To date, all of these experiments have been confined to single clusters of the type used by Kopp.

Nimrod/G is a version of Nimrod that works over a global computational Grid. It uses novel resource management and scheduling algorithms based on economic principles to distribute the computational load. Specifically, it supports user-defined deadline and budget constraints for schedule optimisations and manages the supply and demand of resources in the Grid using a set of resource trading services called GRACE (Grid Architecture for Computational Economy).

Because GrangeNet will link both high-end computational resources and more significant clusters in Australia, it provides an ideal platform for a much more aggressive role out of the Nimrod technology than has been possible to date. In combination, GrangeNet and Nimrod/G will provide the infrastructure required to perform routine parametric computational experiments of the type discussed above. The service will be known as "The Australian Nimrod Testbed".

Conclusions

The overall GrangeNet strategy is to operate an experimental network with a professional level of performance, recognising the limitations of the network design and its implementation with leading edge equipment and advanced services. Service levels will be commensurate with the experimental nature of the network.

The backbone network and five of the eight GbE tails became operational during September this year.

The ANT project will develop in three stages, namely:

1. The deployment of existing technology (Nimrod/G Version 2.0)
2. The enhancement of Nimrod/G V 2.0 for GrangeNet
3. Development of National Demonstrators

Stages (1) and (2) will take the existing Nimrod/G prototype (version 2.0) and deploy it on a number of machines that are linked to GrangeNet. This will involve a fair degree of customization and modification to accommodate the different local environment on those machines. Most significantly, a number of the components have only been tried as research prototypes, and these need to be scaled to work in a network of the size of GrangeNet. For example, the experimental computational economy will require the development and deployment of a unified accounting system across GrangeNet.

Stage (3) will involve taking a number of important applications from the APAC expertise program and using these as case studies in parametric computation. Ideally this will involve applications from at least Melbourne, Canberra, Sydney and Brisbane.

In a future paper we will outline some of the other application projects being undertaken as part of GrangeNet.

References

1. Nimrod/G http://www.csse.monash.edu.au/~davida/nimrod
2. DSTC http://www.dstc.edu.au

GRIA - GRID Resources for Industrial Applications

Mike Surridge, IT Innovation, September 2002, ms@it-innovation.soton.ac.uk

Summary

The GRIA project will take Grid technology into the real world, enabling industrial users to trade computational resources on a commercial basis to meet their needs more cost effectively.

We are developing business models, processes and semantics to allow resource owners and users to discover each other and negotiate terms for access to high-value resources. GRIA will implement an overall business process to find, procure and utilise resources capable of carrying out these high-value, expert-assisted computations.

The project will focus on two application sectors: structural and hydrological engineering, and digital TV and Movie post-production. These have very different technological requirements, but share a need for large amounts of computational power to meet peak loads.

By focusing on business processes and the associated semantics, GRIA will enable users to provision for their computational needs more cost effectively, and develop new business models for some of their services.

Introduction

Today, the main driver for using the Grid is the need for higher availability or peak computational performance. GRIA will create a Grid testbed that allows end-users to investigate and achieve the following:

- increased utilisation of in-house resources for running critical applications,;
- increased peak capacity for running these applications by exploiting external resources run by third parties;
- increased accessibility of applications, by converting key components into Grid-based services that can be incorporated into higher-level business processes; and
- improved overall management of these processes, through greater reliability and accountability of resourcing.

The GRIA consortium is small and highly focused on achieving these goals by using off the shelf software and standards as far as possible. We will take the best features from existing Grid systems and combine them with more conventional e-commerce and web-based software and standards, to support a range of applications and out-sourcing or in-sourcing scenarios.

Our approach is based on the Semantic Grid model proposed by our colleagues at the University of Southampton (see www.semanticgrid.org/). GRIA is still in its early stages, but it is already clear that Semantic Grid ideas provide a rich environment in which to develop flexible and dynamic business processes in terms of Web Ontology and Web Service standards.

GRIA will promote take-up of the Grid by European industry, and contribute to the standardisation of the global Grid infrastructure through the Global Grid Forum and W3C.

Detail

Consortium

The GRIA project is led by the IT Innovation Centre, an autonomous institute of the University of Southampton Department of Electronics and Computer Science, specialising in promoting the take-up of advanced information technology by industry through collaborative projects and commercial consultancy.

The other research and technology partners are the Intelligence, Agents and Multimedia group (IAM), also from the University of Southampton, the Institute of Communication and Computer Systems at the National Technical University of Athens (NTUA), and Dolphin Interconnect Solutions, an industrial manufacturer of low-latency, high-performance networking technology and clustering solutions. The end users are ENEL Hydro's Hydraulic and Structural Engineering Centre, who specialise in the design, installation and operation of hydroelectric and thermal generating plant, and KINO who are one of Europe's most innovative producers of television commercials.

Applications

ENEL Hydro are experts in structural identification, which is of great value when dealing with structures that may be degraded like dams and ancient bridges in earthquake zones. Physical measurements are compared with a calculated response and the model adjusted until it agrees with the measurements. The model is then used to predict how the structure will respond to future stresses, and to design structural improvements if required. However, this approach isn't used every day so ENEL can't justify owning enough computers to run the computations quickly. GRIA will allow ENEL to outsource peak loads to a Grid service provider, or to set up their own Grid services to sell their skills and any spare computing power to other users.

KINO has similar requirements. Many of their commercials can be made cost effectively in a virtual studio using 3D digital rendering, but to meet production deadlines would require a very large computing facility which would often be idle. Moreover, KINO likes to involve their clients in the artistic side of developing a commercial, and want to provide secure access to the digital studio for this purpose. GRIA will allow KINO to outsource their peak loads for rendering to a Grid service provider. GRIA will also allow KINO to set up collaborative Grid services giving clients secure access to preliminary compositions and results, enabling participation in the creative process.

Approach

A typical GRIA user will exploit the Grid to solve a business problem by running a reasonably complex sequence of calculations. This "workflow" may involve people since our applications involve skilled structural engineers, and artistically creative digital directors and animators. GRIA will build upon existing Web and Grid infrastructure by adding services in three key areas:

i) quality of service: end-to-end performance and availability estimation, with efficient mapping of loads onto resources;

ii) interoperability: through open standards for describing resources, services and interactions;

iii) business processes: software support for Grid business models and processes via secure interactions between users and suppliers.

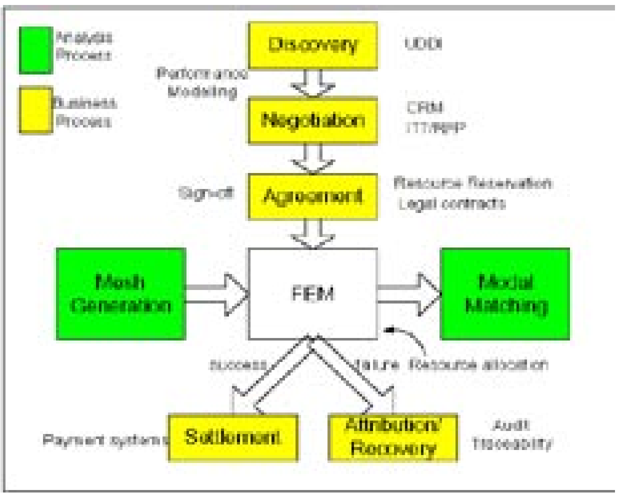

At each step, GRIA will support an orthogonal workflow between end-user, supplier(s) and trusted third parties that make up the business process of Grid service procurement, including:

- discovering suppliers of the required services;
- negotiating terms for access (including price and quality of service);
- exchanging signed agreements to prevent repudiation by either party;
- executing the Grid service including data exchange; and
- settling according to the agreed terms, including attribution and recovery in the event of failure.

The key to GRIA is the use of a Semantic Grid approach. This is based on the use of machine-readable schema for describing data, services, and processes that use them. The project is at an early stage, but already we see that the Semantic Grid is an appropriate paradigm for describing and implementing Grid-based business processes.

Conclusions

The GRIA project will develop, apply and evaluate a Grid testbed, based on an existing open-source infrastructure, which will incorporate services supporting Grid business models, business processes, and quality of service. This will enable the Grid to be used for commercial outsourcing computational services;

We will show how companies like KINO can operate a "virtual digital studio" over the Grid in which they and their clients can collaborate to produce fantastic video sequences using state of the art digital composition and rendering methods exploiting 3rd party computing resources.

We will show how companies like ENEL Hydro can exploit their specialist engineering skills to provide services to their colleagues and to external customers using a combination of outsourcing and insourcing business models over the Grid.

Since our applications come from two widely different sectors (structural engineering and TV/Movie production), our work should be widely applicable across European industry.

Above all, GRIA will demonstrate how the Semantic Grid, and the merger of Grid with broader Web and e-Commerce technologies through the Open Grid Services Architecture, will support a new class of dynamic and responsive yet fully commercial virtual businesses.

References

1. www.semanticgrid.org/

GRIP - Work in the Global Grid Forum

John Brooke, Manchester University, September 2002, j.m.brooke@man.ac.uk

Summary

We describe the work of participants from the Grid Interoperability Project (GRIP) towards development of Global Grid standards in three important areas.

1. Open Grid Services Architecture (OGSI Working Group in Architecture area)
2. Minimal Grid Functions (GPA-WG in Architecture area)
3. Grid Resource Allocation Agreement Protocol (GRAAP-WG in the Scheduling area).

The GRIP project [1] aims to make two major Grid systems, UNICORE and Globus interoperate and as part of this to establish standards in the Global Grid Forum that take account of middleware technology developed in Europe, e.g. UNICORE. As a result of this work, three representatives from GRIP are now co-chairs of the above groups, respectively Dave Snelling of Fujitsu European Laboratories and John Brooke and Jon MacLaren of the University of Manchester (UK).

Introduction

The development of standards is crucial for a working Global Grid. The success of the http protocol in ensuring the widespread adoption of the World Wide Web is well known. The Globus Project [2] from the US was specifically set up to drive towards adoption of such standards by providing a reference implementation (Globus toolkit) and working in the Global Grid Forum for standards adoption. Consequently Globus is now a de facto standard. However, technology developed in Europe, in the UNICORE project [3] and extended to a European Grid in EuroGrid [4], complements Globus and provides a well-tested and reliable European technology for Grid Computing. The aim of GRIP is to develop working interoperability of Globus and UNICORE and to work towards developing standards in the GGF that capture the strengths of these projects and exploit synergy between them..

A key feature of UNICORE, not replicated in Globus, is the concept of an Abstract Job Object which can be “incarnated” on differing platforms using information from an incarnation data base. This concept allows UNICORE to be readily adapted to the OGSA architecture and Dave Snelling posted the first working OGSA compliant application to the GGF in April 2002. The AJO also enshrines the concept of a hosting environment which is crucial for heterogeneous multi-site Grids and John Brooke is currently working with Bill Johnston of the NASA Information Power Grid to develop this concept. Experience in the US of running jobs across 5 different Grids proves the necessity of the AJO concept. Finally the construction of a working European Grid based on UNICORE has revealed the importance of advanced reservation and co-allocation of resources and Jon MacLaren is developing this work in the GRAAP-WG.

Strong European participation contributes to the health of the GGF and it is clear that European basic research into fundamentals of Grid Architectures has an important role to play. The GRIP project is interested in hearing of the experiences of other EU projects developing basic Grid architecture and feeding this into the relevant GGF groups. GRIP is also working in the Applications RG alongside projects such as GridLab. An early fruit of the interoperability work, is that jobs launched on UNICORE can run on resources controlled by Globus and UNICORE and Globus generated certificates can interoperate. This was proved by interoperability of certificates from the EuroGrid CA (Certificate Authority) and the UK eScience CA demonstrated by a team at the University of Manchester.

Detail

Globus operates on a toolkit approach. It provides the building blocks of the common services layer of the Grid, the neck of the “hour-glass” between Grid applications and the schedulers and networking tools that control the Grid Fabric (e.g. PBS, LSF, NQE, ftp etc…). APIs are provided with C language bindings and the toolkit can be utilised from programmes or via scripts or even from the command line. On this basic functionality have grown applications and portals that provide an application level interface. An example is Cactus [5] and various portals such as the

UK HPC Portal [6]. Globus also provides functionality for querying Grid resources via the MDS and GIIS. The Globus team have also developed interfaces via Java and Python, the Globus CoG kits.
However site policy is a key issue that is not part of the Globus abstraction. To take one example, consider the problem of filestore. User A submits to two different sites. At site X she has a user account to which her certificate is mapped and as part of this she has permanent file space for staging and holding files before and after the job runs. At site Y she is coming in onto a temporary guest account and site Y has a policy of deleting all files associated with this account when the job finishes. Clearly any assumption of pre-staging files in preparation for a job at Y will fail as will any attempt to hold files at the site for later processing or as part of a job chain. Again site X may have a virtual shared memory model where all processors have access to files and memory whereas site Y may be a workstation cluster with local disk and memory only and no common model of filestore. Thus the environment in which Grid jobs execute may be heterogeneous in both architecture and policy. It is, of course, possible to interrogate each site individually about these matters and tailor the job scripts accordingly. However this is not a scalable approach and is contrary to an important element of Grid computing namely seamless access to resources.

The importance of the abstractions of the UNICORE architecture and especially the AJO, is that they allow the architecture and environment to be described in an abstract way and then this can be actually implemented in scripts tailored to the site by the UNICORE software itself rather than by the user. Essentially this is a three-tier architecture. The UNICORE client allows the user to build the job constrained by the structure of the UNICORE model via a GUI based job preparation tool. At each UNICORE site, a Network Job Supervisor is responsible for receiving the UNICORE AJOs as Java objects and uses the UNICORE Incarnation Database to specialise the jobs for the target site, which are finally executed by the third part of the architecture, namely the Target System Interface (TSI). The AJOs may contain sub-AJOs which are sent along to other sites in the chain. At present the user has to run on specified sites but the recent development of a Resource Broker in EuroGrid allows UNICORE to search for suitable resources and report to the user. Currently, components of an AJO are ordered using dependences, represented by a DAG (Directed Acyclic Graph); in UNICORE 4 (currently in Alpha release), this system has been extended to include branches and loops. All components of the AJO are signed by the users permanent certificate. Security is an important feature of the system, but one which we do not address here. UNICORE is a vertically structured system encompassing all levels of the Grid from user applications to scheduling and resource management systems.

The current GRIP technical concept is to create a TSI for a Globus Grid, handing over the UNICORE certificates and AJO to the Globus toolkit rather than directly to the low level Grid fabric. The advent of OGSA is causing a re-examination of the strategy in that it may be advantageous to go more directly the the OGSA architecture as is being targeted by Globus Toolkit 3. The work in the OGSI WG is aimed at exploring these issues. Joint work with the RealityGrid [7] project from the UK eScience programme has developed a computational steering application and work done by IDRIS in EuroGrid exploits multi-site processing and visualization. This highlights a weakness at the resource allocation level of the Grid, namely the need to co-allocate resources at multiple sites simultaneously. In particular, most scheduling systems do not permit negotiation, so co-allocation of multiple resources by Grid middleware is difficult. This inadequacy has led directly to the work being pursued in the GRAAP-WG of the Scheduling area in GGF.

Conclusions

The development of well-constructed and robust abstractions to describe multi-site workflows on the Grid is an important component of standards and protocols in Grid Architectures. It is also important to have mechanisms for automatic “incarnation” of such abstractions at particular sites and to be able to update information on site policy via a database approach. The UNICORE technology developed in Europe is a leading implementation of this approach. Funding of the EuroGrid and GRIP projects by the EU is now enabling such technology to play an important role in the GGF. The work of the three named co-chairs rests on the technical achievements and support of members of the and the above mentioned EU projects and the UNICORE-Plus project funded by the German government .

The experience of US sites in running multiple and heterogeneous Grids confirms this European experience and there is mutual feedback between the work being done in the US on Production Grids with the European Grid built by EuroGrid. Achieving interoperability of UNICORE and Globus will allow Europe to create very flexible Grids uniting the work done in various projects, many of which use Globus.

Such work provides valuable pointers for OGSA, by providing working demonstrators for the OGSA concepts such as that provided by Snelling, the theoretical concepts of the architecture can be tested on real Grid testbeds. The Reality Grid steering demonstration of the mixing of a two component fluid is now an application that can be productively applied to research in the structure of soft condensed matter. Proposals for extending such an approach to modelling in the life and environmental sciences will be discussed at a meeting on Realistic Modelling of Complex Systems in Life and Environmental Sciences at the ERA meeting in Brussels from 12:00 to 14:00 on November 13th 2002.

References

1. Grid Interoperability Project http://www.grid-interoperability.org
2. Globus Project http://www.globus.org
3. UNICORE Project http://www.unicore.org
4. EuroGrid http://www.eurogrid.org
5. Cactus Computational Toolkit http://www.cactuscode.org
6. UK HPC Portal http://esc.dl.ac.uk/HPCPortal/
7. RealityGrid http://www.realitygrid.org

UNICORE - GRID Computing in Germany

UNICORE software available

Dietmar Erwin, FZJ, July 2002, D.Erwin@fz-juelich.de

Summary

Starting in 1997, German Supercomputer Centres initiated and created the grid infrastructure UNICORE (Uniform Interfaces to Computing Resources) to allow an easy-to-use, secure and browser-based access to the high-end machines from everywhere. In the meantime the development reached a production stage. Thus the UNICORE Forum decided to announce an „open source" licence of the UNICORE Software for research purposes.

The grid system allows the researcher to start jobs from their client everywhere without knowing the job control language as well as the file system and the compiler parameters of the target computer. In the next step the developers work on portals for specific application packages like computational chemistry and computer aided engineering and simulation.

The UNICORE development was funded by the German Ministry of Research (BMBF). The software is used in several computer centres. Additionally it is the base in several European projects like EUROGRID. In a specific project, GRIP, the interoperability of UNICORE and the US-based grid-system Globus, and the compatibility with the new grid standard OGSA will be realised.

Introduction

The UNICORE Forum

In the beginning of 2000 users and developers founded the UNICORE Forum as an autonomous, open and a non-profit organisation, which pushes and supports the development, distribution and the use of UNICORE. It controls the specifications and the application programming interfaces of the UNICORE system and guarantees the very important independence. The participation is open for grid developers, grid users, hard- and software vendors. Actually it has 25 members, including all the vendors interested and involved in grid technology. Additionally international computer centres in academia and research are members of the forum.

UNICORE Forum offers sources of UNICORE software

The availability of the first production of UNICORE inspired the Forum in its last session in April. It decided to make the software available for all interested partners for research purposes. The base is an community source licence model. It can be downloaded from http://www.unicore.org. For production environments Pallas GmbH, Brühl, Germany, offers a commercial UNICORE version. It can be used in industrial environments like Intranet solution or as a technology for the Application Service Provider model.

UNICORE Test Site

A UNICORE test system is available for people not involved in the project. It includes the current client and a set of target sites to send jobs to. The Central Institute for Applied Mathematics (ZAM) of the Research Centre Jülich installed a dedicated test grid for functionally testing. The installation and initialisation as well as the 30 day usage restriction are described on the web site. The Test Grid simulates there UNICORE Sites with up to two target systems. An access to real supercomputers is not possible. Thus interested users can check the user interface and the functionality even if they are not involved in the project.

User Experiences

Professor Geerd-Rüdiger Hoffmann, Head of the Department Systems and Operations at the German Weather Services (Deutscher Wetterdienst, DWD) and chairman of Unicore Forum introduces UNICORE at DWD. „In the near future we allow and control the access to our machines and supercomputer by UNICORE. The networked computer capacity incorporates a huge saving potential and will improve the efficiency for all participating partners. Thus we realise an DWD grid computing."

The hardware vendors are extremely interested in the project and support it heavily. „UNICORE realised since 1997 networked computing - long before the Grid term was used", mentioned Dr. Ulla Thiel, Director Scientific and Technical Computing, IBM EMEA, representing the view of the hardware vendors.

References

1. http://www.unicore.org (UNICORE Forum)
2. http://www.unicore.de (UNICORE Starting Page)
3. http://www.fz-juelich.de/unicore plus (BMBF-Projekt: 1.1.2000 - 31.12.2002)
4. http://www.fz-juelich.de/unicore-test (UNICORE Testgrid)
5. http://www.eurogrid.org (Eurogrid-Projekt)
6. http://www.grid-interoperability.org (GRIP-Projekt)
7. http://www.dwd.de (German Weather Service)

About the GRIDSTART Technical Bulletin

This technical bulletin is published every three months to all interested parties. To receive it, please register on the GRIDSTART website, www.gridstart.org. The emphasis is very much on work in progress and the stimulation of technical discussion amongst the target audience of those leading GRID developments. This bulletin aims to promote technical developments in GRID technology to all interested parties and to stimulate feedback on what is being developed. Contributions in the form of articles, letters and announcements are particularly welcome and should be sent to bulletin@gridstart.org. For the submission of articles, you should mail the above address for details of the format for articles.

Your attention is also drawn to the companion to the Technical Bulletin, the GRIDSTART Technical Newsletter, which focuses on the application of GRID technology to science, industry and commerce. This newsletter is also published every three months to all interested parties and can be received by registering on the GRIDSTART website. Contributions to the newsletter are particularly welcome and should be sent to newsletter@gridstart.org.

The GRIDSTART website

The GRIDSTART website, www.gridstart.org, aims to provide a central source of information about grid activities, events and other information sites. It references not only technical developments, but also practical applications of GRID technologies. A clear goal is to increase awareness of what is being developed not only in the technical GRID community, but also amongst potential early adopters from research, industry and commerce. If you have material suitable for inclusion in this website, would like to announce an event or would like your site linked, please contact webmaster@gridstart.org.

About GRIDSTART

The GRID is widely seen as a step beyond the Internet, incorporating pervasive high-bandwidth, high-speed computing, intelligent sensors and large-scale databases into a seamless pool of managed and brokered resources, available to industry, scientists and the man in the street. The name, GRID, itself draws the analogy between the pervasive availability of electrical power and that of computing and data, coupled with mechanisms for their effective use.

GRIDSTART is an initiative sponsored by the European Commission with the specific objective of consolidating technical advances in Europe, encouraging interaction amongst similar activities both in Europe and the rest of the world and stimulating the early take-up by industry and research of GRID-enabled applications. The initiative brings together technologists, scientists and industry in multi-disciplinary approach to developing the GRID infrastructure. The clear goal is to develop sustainable, effective and universal solutions addressing the needs of science, industry and the public. The GRIDSTART website at www.gridstart.org provides a portal to GRID initiatives taking place across the world.

For further information about GRIDSTART contact:

Maureen Wilkinson
GRIDSTART project manager
email: m.wilkinson@epcc.ed.ac.uk
Tel: +44 131 650 5027