# EGEE

## "From Research to Business"

### 10 Successful EGEE Business Use Cases

www.business.eu-egee.org

# EGEE and Business

Enabling Grids for E-sciencE (EGEE), the world's largest grid computing infrastructure, has pioneered e-Infrastructures for eScience throughout Europe and beyond by providing researchers in academia and business with access to a production level grid infrastructure. Over the past six years, EGEE has attracted a wide range of new users and disciplines, enhancing technical performance to enable companies and researchers alike to do new things and save time on complex calculations.

EGEE has embarked upon a strategy to bring on board companies of different sizes and from diverse commercial sectors that have embraced grid technology and with an eye to leveraging the compute and IT resources needed to speed up time to take new products and services to market and gain efficiencies and cost savings.

As a result of our work together, a number of companies are now marketing solutions and EGEE developed middleware, gLite, such as CGGVeritas, Maat-g, NICE, and Constellation Technologies. Other companies have been able to create new niche markets in IT consultancy and training around EGEE technologies, such as Excelian, GridwiseTech, Imense, Linalis and Six2. Several of these companies, Imense and Constellation, have been formed thanks to the strategic alliances forged with EGEE and underpinned by the technologies developed over the past few years.

gLite serves as the "glue" that holds EGEE's grid infrastructure together. The middleware is released under an Open Source business-friendly licence, Apache 2, and is the next generation middleware for grid computing. gLite provides a framework for building grid applications tapping into the power of distributed computing and storage resources across the Internet. The gLite distribution is an integrated set of components designed to enable resource sharing. It encompasses a wide range of services, while ensuring ease of installation and configuration on chosen platforms such as Scientific Linux and Windows. It is currently deployed on hundreds of sites worldwide enabling global science and industrial research and development.


*"EGEE has successfully targeted the business community throughout the life of the project and is shown through the number of pilot business applications taking advantage of EGEE gLite open-source solutions."*

**EGEE Project Director, Dr. Bob Jones**

The gLite Open Collaboration has been established between EGEE partners involved in middleware activities as a new framework for the maintenance and future evolution of the middleware beyond the end of the EGEE series of projects to ensure the availability of the middleware, support and developments for years to come. This initiative is complemented by the European Middleware Initiative bringing together the major European middleware providers such as ARC, gLite and UNICORE to deliver a consolidated set of services that will provide the European and global research and business communities with unprecedented computing power.

These are exciting times as EGEE has transitioned to the European Grid Initiative, a pan-European Grid Infrastructure in collaboration with National Grid Initiatives (NGIs) guaranteeing the long-term availability of an e-Infrastructure for all European research communities and their international collaborators. From a business perspective, new

opportunities are emerging for commercial organisations to take part in the infrastructure and continue to reap the benefits that distributed computing delivers.

The ten successful EGEE business use cases described in this booklet showcase a variety of opportunities for companies to take advantage of the services developed by the EGEE project, such as adopting EGEE technology on a company's own infrastructure or integration into a new or existing solution, running proof-of-concept applications on the EGEE testbed infrastructure or working on areas of joint technical development.

# Philips Research: Scientific simulation, modeling and data mining supports healthcare

*"The EGEE grid was originally conceived to serve the academic world. However, with thousands of distributed assets, interconnected with unconventional technology, it has great potential to become a virtual hub for business; the role that Stanford University had played for what later became the Silicon Valley. Our solution, combined with LCG/gLite, can easily serve other industrial customers, thereby creating the foundations for a pan-European business grid."*

**Pawel Plaszczak, President of GridwiseTech**

Philips Research is one of the world's largest corporate research organisations that develops new technologies and investigates potential growth areas for research and development. Over the last few years, Philips Research ICT has been investigating the deployment of grid technology with the aim of enabling research that effectively draws together expertise, connecting networks for knowledge exchange and accessing external High Performance Compute (HPC) resources in areas such as medical research.

Developments in medical imaging and bioinformatics applications involve many high-resolution sensors that generate significant amounts of data, as well as data correlation from different sources. Open Innovation is aimed at addressing the need to have resources and/or data in the right place at the right time, in order to support the technology chain. The infrastructure has to be able to cope with 'data explosion' and 'large-scale data analysis' and ensure secure collaboration with partners. While large computations are needed, much can be parallelised or distributed. A key end objective is to support the shift from experiments that are model-driven to data-driven experiments, as well as ensuring active participation in the eScience project ecosystem and being able to both consume external resources and provide resources to externals.

Philips Research's adoption of gLite is connected with this Open Innovation concept, which has resulted strategic co-operation between companies, universities and research institutions. EGEE Business Associate, GridwiseTech, a vendor independent specialist in scalable solutions, played an instrumental role in the adoption of gLite by connecting Philips Research to some 30,000 CPU on the EGEE infrastructure and 15 Petabytes of storage, distributed among 195 data centres mainly in Europe. Built on WLCG/gLite grid middleware, EGEE is thought to be the first production grid system deployed worldwide.

Through the installation of a local gLite/EGEE-based grid cluster connected to DutchGrid, Philips Research investigated and evaluated the technology with access via gLite and UNICORE grid middleware software, evaluating its benefits and potential challenges. Knowledge exchange and visibility have been ensured by GridwiseTech President, Pawel Plaszcak, offering a view on best practices in grid infrastructure testing and bringing industry-class reliability and quality of service, highlighting the benefits

of grid technology in the healthcare sector, including the advantages of adopting external sources like EGEE.

After several months of work, a solution was delivered and put in place. The engineers from GridwiseTech integrated job management environments providing a user-friendly, application-centric grid portal layer and enabled the application to run (both internally and externally) within the designated grid secure virtual organisation. Peak processing can now be scheduled in an on-demand fashion.

**http://www.gridwisetech.com**

**http://www.dutchgrid.nl/**

# WISDOM: Drug Discovery

*"Grid allows biologists and chemists to focus their experimental work on the most promising potential drugs, reducing the time required to develop commercial drugs for malaria and indeed, other diseases. Without grid, such large scale studies would be expensive and very time consuming. Grid could be the catalyst for drug development that brings together the actors - biochemists, physicians and computational chemists - and pushes them ahead in the same direction."*

**Dr. Martin Hofmann, Drug Discovery Application Supervisor at SCAI**

"Grid allows biologists and chemists to focus their experimental work on the most promising potential drugs, reducing the time required to develop commercial drugs for malaria and indeed, other diseases. Without grid, such large scale studies would be expensive and very time consuming. Grid could be the catalyst for drug development that brings together the various actors - biochemists, physicians and computational chemists - and pushes them ahead in the same direction." Dr. Martin Hofmann, Drug Discovery Application Supervisor at Fraunhofer SCAI

Malaria is a dreadful disease caused by a protozoan parasite, plasmodium, which kills a million people every year and affects 300 million more. The number of malaria victims has grown in many parts of the world as resistance to the most widely used drug (Chloroquine) has increased and as the Anopheles mosquitoes that carry malaria have become increasingly resilient to common insecticides. Molecular biology research has identified parasite proteins as potential targets for drugs against malaria. A new strategy to fight malaria aims at the haemoglobin metabolism, which is one of the key metabolic processes for the survival of the parasite.

The WISDOM initiative (Wide In Silico Docking On Malaria) demonstrates the relevance and impact of the grid approach to address drug discovery for neglected and emerging diseases. Partners committed to tackling these diseases are connected globally to use grid infrastructures with the aim of organising and accelerating their research, deploying production experiment of virtual screening at a large scale against diseases. These activities are called "data challenges". The ultimate goal is to identify which molecules could dock on the protein active sites inhibiting its action and therefore interfere with the molecular processes essential for the pathogen.

WISDOM was one of the front runners in leveraging grid enabled in silico docking to simulate the interaction of potential drugs with target proteins on the EGEE grid infrastructure. Researchers aim to increase the performance of the application and compute millions of potential drugs in just a few weeks – a real possibility with the EGEE grid. Working at this rate, the hope is to take a significant step towards finding a new drug to treat this widespread and devastating disease.

WISDOM also worked on the Avian Flu (H5N1), studying the impact of point mutation on drug resistance. Since the first H5N1 outbreak in 1987 there has been an increasing number of HPAI H5N1 bird-to-human

transmissions leading to clinically severe and fatal human infections, while scientists have showed that the N1 and N2 subtypes could evolve into variants under drug stress. WISDOM spearheaded an initiative to study the impact of point mutation on drug resistance. The goal is to screen a large set of compounds against the same target, the influenza A neuraminidase, with various structures predicted from homology methods thanks to grid infrastructure.

During the summer of 2008, WISDOM ran a test using gLite on Digital Ribbon resources (See Use Case – Digital Ribbon: Towards Computing as a Commodity in this Booklet). Data used for the test consisted of 750 chemical dockings - potential candidates for diabetes drugs. On a single computer, a study like this with 100,000 potential drugs would normally require six months to complete - but on the EGEE grid it was achieved in just two days. Such an achievement was hailed a success also in terms of gaining cost efficiencies in drug development, which is a very expensive process. The next step for the drug discovery application include sorting through the large amounts of data to identify potential drugs to treat a range of diseases and to reduce the gap between virtual screening and real world drug development.

The work carried out through the WISDOM experiment has resulted in more than 15 scientific publications and 5 patents. It clearly demonstrates that computing grids like EGEE are not only the cheapest and fastest ways to discover new drug leads but also that in silico approaches minimise the non-productive trial-and-error approach in the lab.

**http://wisdom.eu-egee.fr/**

# Health-e-Child: Enabling Grid for personalised and preventative Healthcare for Children

*"I am touched when I see the dedication of clinicians in the hospitals we work with. If we can make their job a little easier, it will be the most rewarding thing I can think of. I know that if we can save them a little time, they will spend it with the patient"*

**Martin Huber, Siemens Corporate Technology and Technical Coordinator of Health-e-Child.**

The European project Health-e-Child brings together clinicians, researchers and technology innovators with an eye to developing and delivering personalised healthcare for children by connecting grid technology to paediatric medicine by providing clinicians with integrated, multi-level information. The Health-e-Child Gateway is a platform underpinned by the EGEE gLite grid middleware that has been enhanced by the company Maat Gknowledge, together with other European partners, to store, retrieve and manipulate medical records. "*This grid-enabled platform for European paediatrics provides the seamless integration of traditional and emerging sources of biomedical information; the long-term goal being to deliver uninhibited access to universal biomedical knowledge repositories for personalised and preventive healthcare, large-scale biomedical research and training, and informed policy making*", explained Jörg Freund, Siemens and Project Co-ordinator.

Health-e-Child has placed much emphasis on some of the most challenging areas in medicine and paediatrics, from medical imaging to the treatment of rare diseases. In the case of paediatrics, several challenges are complicated by the fact that physicians are dealing with smaller human beings, as Martin Huber pointed out: "*We have the additional problem that a child's heart is smaller and beats faster than an adult's. Imagine trying to take a picture of a hummingbird's wings in flight, that's roughly what we're trying to do. This is a huge challenge for medical imaging, and current scanners, Magnetic Resonance Imaging (MRI) and computer tomography provide images of spectacular quality. Our task is to transform these into graphs and present them to the doctor.*"

Important headway has been made with regard to cardiac modelling, which is normally a lengthy, manual process. To improve this state of play, Health-e-Child has developed an application that creates individual-specific heart models for patients with enlarged right ventricles, making modelling near-automatic, saving time and giving more accurate measurements.

Another successful application is CaseReasoner for rare diseases. While physicians are very familiar with common diseases and life-threatening conditions, this is not the case when it comes to diseases that are rarely seen by the average physician. In these cases, they are dependent mainly on textbook knowledge. When a doctor has a patient with unusual or conflicting symptoms, CaseReasoner allows him or her to use a grid to search for similar conditions from anonymous patient records in any participating hospital. The doctor can compare recorded diagnoses, treatments and outcomes—helping prescribe treatment more confidently. In the words of Andrew Taylor, Great Ormond Street Hospital in London, UK, "*Health-e-Child's application, the tools running on its grid platform and the models that the project has developed will allow doctors to plan hopefully not only the treatment, but also whether they need that treatment, or*

*when it is the best time for them to get that treatment*".

Such important work in the field has been widely recognised at a European level, culminating in one of the most prestigious and sought after awards for aspiring European research groups - the Exhibit Grand Prize at the ICT2008 conference and exhibition, which presented priorities for ICT R&D and the funding agenda in Europe for the years ahead. David Manset, Maat Gknowledge, said "*This prize is truly a welcome acknowledgement by the European Commission and fellow colleagues throughout Europe of Health-e-Child's work and achievements, illustrating the potential to tackle the technical challenges involved and encompass other medical domains.*"

One of the strengths of Heath-e-Child in addressing key areas of paediatric healthcare lies in the ability to bring together the right mix of experts and practitioners. This is why, in addition to the Consortium members from academia and enterprise, Health-e-Child includes European hospitals, research facilities and labs.

**http://www.health-e-child.org/**

# Imense - from Lab to Market

*"Having access to the grid was critical to launch imense.com. Analysing so many images would have been unfeasible for us as a small start-up, but the grid allowed us to test our software, impress our investors and secure funding for our company."*

**David Sinclair, Imense Co-founder and CEO.**

Imense, a small, high-tech UK company based in Cambridge, has developed the Imense® Web 3.0 product to significantly improve the accuracy and depth of search to attract new users. The company, which targets image agencies; publishing groups; digital asset management providers; search engines; photo-sharing and social networking sites, was founded by David Sinclair and Chris Town and today has seven employees. The company's value-add was acknowledged in 2007, when the BBC described it as one of the UK's "hottest emerging companies".

Images and video make up over 70% of the digital data available on the Internet, an estimated 15 billion images, but traditional search engines cannot index this information directly, relying instead on text descriptions entered by hand. Imense's key innovation is a new form of image retrieval that automatically analyses images in terms of their content, without the need for human generated captions. They have also developed a powerful query language that lets people search for the images they need.

Dr David Sinclair, CEO and co-founder: "We built a prototype of our new image analysis and search technology, but simply were not able to test our software on sufficiently large numbers of photos. We knew we could search tens of thousands of pictures, but could not afford to try it on hundreds of thousands or millions of images. This made it difficult for Imense to get the investment we needed to develop a commercially viable product. That is where our partnership with the particle physics grid came in."

Spread across 17 sites, the UK Particle Physics Grid (GridPP) has been built to analyse the peta-bytes of data expected from Europe's particle accelerator, the Large Hadron Collider, located at the European Organisation for Nuclear Research (CERN) in Switzerland. Its now 20,000 computers have been shared with other researchers, from geophysicists to biologists. Participation at an EGEE-II Industry Day in Cambridge in 2006 and subsequently a meeting arranged by the Science and Technology Facilities Council (STFC) on opportunities for industry revealed how grid technology could be the answer to Imense's problem. Image analysis is a naturally parallel process that fits perfectly with the capabilities of the grid used by STFC scientists to process data in particle physics.

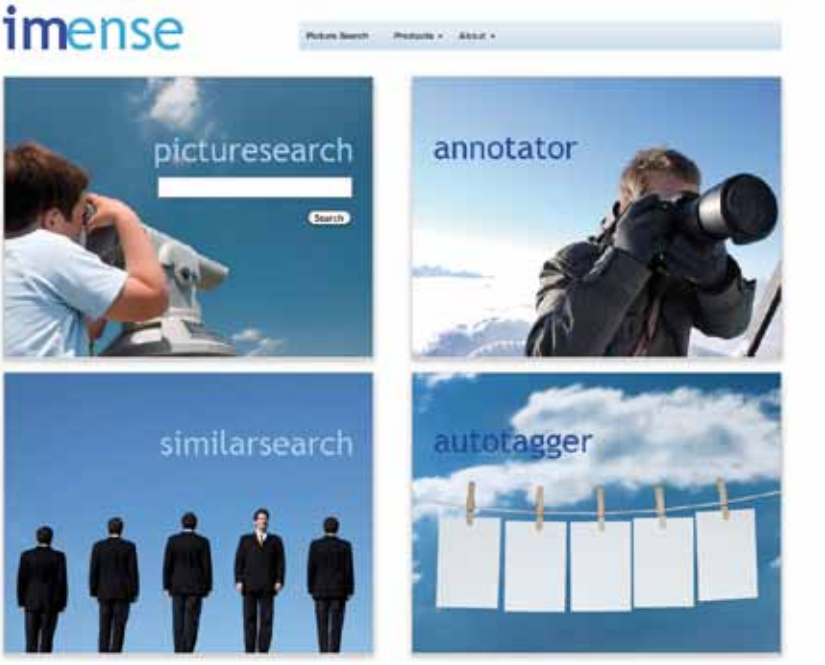


Imense's work with the Cambridge e-Science Centre was funded through an STFC mini-PIPSS award. PIPSS is a knowledge transfer scheme that supports the development of effective, long-term collaborations between UK Universities, research bodies and industry. An investment of more than £500,000 helped the company bring a new product to market, illustrating the positive outcomes of this particular grid experience.

One of the key success factors from both a funding and technology transfer perspective is explained by Professor Keith Mason, Chief Executive of the STFC: "*We actively encourage the researchers we fund to consider the wider applications of the work they do.*

*In this case, computing problems that had to be addressed for particle physics can also be used to solve other challenges with large amounts of data. The Council's Knowledge Exchange Service put the two teams together and provided modest funding to start them off – the new investment attracted by Imense represents a ten-fold return on the initial development funds."*

The creation of Imense is one of the finest examples of how EGEE and grid technology has helped turn an idea into a business.

**http://www.imense.com**

# Total, UK: In-house or out?

*"The availability of a test-bed service offered by EGEE's GILDA pre-production infrastructure allowed the Geoscience Resource Centre to run test applications on an external grid in order to assess the usefulness and compare the benefits of in-house vs. external run applications. This has proved invaluable through a unique opportunity to analyse business options in terms of both cost and resource optimisation."*

**Hannah Cumming, Total UK External Grid Project Leader**

Today, there are currently large production grid Infrastructures in both the commercial and research domains. However these two communities are quite separate and different. In the commercial domain the majority of grid Infrastructures are internal (company-wide only) or are on-demand type services (one provider, multiple users) and use mainly commercial "grid" software. On the other hand, the grids in the Research Domain are mainly External (multiple administrative domains sharing resources), and use Open Source software.

Engineers and geoscientists use seismic, geo-models and reservoir simulation models to improve their knowledge of how petroleum reservoirs behave during production and their potential to deliver. Dynamic flow simulation and seismic processing and modelling are essential for this and are computationally demanding, generating large amounts of data.

The Total Group Geoscience Research Centre (GRC), located in Aberdeen, Scotland, plays an important role in research and development for Total's exploration and production worldwide. Their geologists, geophysicists and reservoir engineers work together to develop and test new techniques in areas such as history matching, fault and sedimentary modelling, and seismic for reservoir and pore network modelling. To carry out the research that the Geoscience Resource Centre needs to get the most from its available computing resources.

In order to evaluate the value-add of grid computing, the Geoscience Resource Centre became an early EGEE adopter, initiating the External Grid Computing project to experiment with grid computing and assess the potential of an external approach. Potential benefits include providing computational power and data storage on demand, by linking together distributed computing resources and implementing open source grid in the oil and gas industry. Hannah Cumming, leader of the external grid project stated, "*The external grid community is at the forefront of driving standards and developing new ideas within the grid community*".

The Geoscience Resource Centre operated a hands-on approach, porting an existing in-house application to the EGEE external grid using gLite middleware. The EGEE grid was chosen due to its international status, the many different institutions taking part, the variety of disciplines involved, including geosciences organisations like CGGVeritas, and the availability of a production service for researchers. The Geoscience Resource Centre leveraged EGEE's GILDA pre-production test-bed service to run its test applications, which also demonstrated an effective way of gaining experience of grid computing.

Though the outcomes remain internal, the Total UK case study offers an excellent example of how companies can make more informed decisions when it comes to building in-house or opting to outsource.

**http://www.uk.total.com/activities/geoscience_research_centre.asp**

# Seismic imaging & reservoir simulation: Solution for Oil & Gas Market

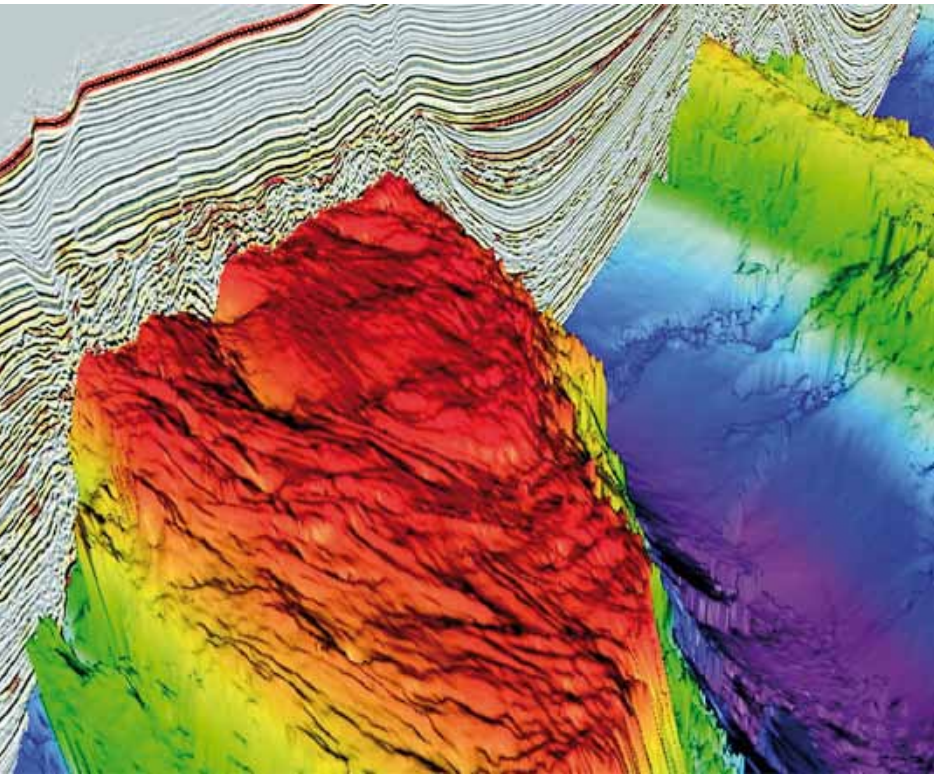

Finding new oilfields, monitoring resource recovery and practising sustainable development are all complex tasks. Worldwide collaboration, huge computing power, fostering innovation and cutting-edge geosciences software all require large financial investments. Building huge processing centres is an industrial challenge in itself and small laboratories face the most difficult challenges when it comes to making the investments needed. Competitive edge and overall benefits for the Geoscience community come from access to state-of-the-art geosciences software, shorter processing times for seismic data and the ability to collaborate with colleagues working on the same project irrespective of their geographic location.

Modern seismic data processing and geophysical simulations require ever greater amounts of computing power, data storage and increasingly sophisticated software. However, small and medium sized research centres encounter difficulties in exploiting their innovative algorithms due to an inability, at research level, to keep pace of enhancements in the current state of the art. Grid computing offers an opportunity to foster the sharing of computer resources and gain access to large computing power for a limited period of time at an affordable cost, as well as to share data and sophisticated software. The opportunity to solve complex problems and validate innovative algorithms on real scale problems is also a way to attract and keep the brightest researchers for the benefit of both the academic and industrial R&D geosciences communities.

Back in 2007, CGGVeritas, a French company leading the development of natural resources through geophysics, embarked on a mission to evaluate the benefits of grid computing through the Expanding GEOsciences on DEmand (EGEODE) Virtual Organisation. The approach adopted included access to EGEE grid resources coupled with a proposed solution to geophysicists working in research laboratories that needed access to a generic platform for seismic processing.

A year on, CGGVeritas developed software called Geocluster to process seismic data from a variety of sources, including trucks that mechanically send sound-waves into the earth with the use of large, hydraulics devices, and ship-borne "airguns" that fire at regular intervals just below the water's surface as the vessel moves along pre-determined survey lines. Regardless of the source of the sound-waves, or whether they take place on land or sea, the principle is the same: the reflected compression wave created is detected by a network of sensors, forming a signal. The challenge for the developers of Geocluster was to process this signal, distinguish it from background noise, store and interpret the information, and put it all into a format easily understood by humans.

Geocluster is able to accomplish this by creating three-dimensional underground maps that outline properties of the subsurface, along with the locations of oil and gas reserves. A sister application, known as Reservoir Simulation, models how these reserves evolve during the drilling process, enabling more efficient extraction. The data are interpreted by a geophysicist who instructs the exploratory company on where to drill. As with Geocluster, Reservoir Simulation was developed to operate in a grid environment using gLite through the EGEODE Virtual Organisation, using EGEE's infrastructure as a research environment and testbed. As an outcome, CGGVeritas has marketed its Geocluster application, now called Geovation. This includes the new generation of massive high-density wide-azimuth datasets and the latest compute-intensive imaging algorithms with a full range of advanced applications for interactive processing and quality control.

**http://www.cggveritas.com**

# Digital Ribbon: Towards Computing as a Commodity

*"Computing is moving towards becoming a commodity; however it is much more complex than the oft-cited example of electricity. Key requirements for computing consumers include not just power, but also bandwidth, connectivity, that is, the speed at which processing cores can talk to each other, and memory; having enough resources available at the right time is key."*

**Erik Weaver, Digital Ribbon CEO.**

Digital Ribbon, which is a US-based service registry company covering a multitude of platforms, processors and architectures providing resources to both the private and public sectors, aims to act as a clearing house for computational resources. The company uses its service registry model to connect resource consumers with the right resource providers. This approach has the potential to transform the way users are currently running jobs on clouds or grids.

Having more options could be a great advantage to users. "*We are always looking for ways to link user communities, applications and resource providers,*" said Bob Jones, Project Director of EGEE. In the summer of 2008, WISDOM, a global initiative for discovering new medicines for neglected and emerging diseases, ran a test using EGEE's gLite middleware on Digital Ribbon resources. The trial run was aimed at paving the way for an effective EGEE and Digital Ribbon collaboration, with an eye to pinpointing additional areas of co-operation.

WISDOM harnesses the power of grid computing to look for drug candidates to treat specific diseases. Drug development is a very expensive process, but grid technology and collaboration between partners in academia and industry can significantly cut costs for its initial step. Data used for the test consisted of 750 chemical docking potential candidates for drugs. In about 12 hours they ran calculations corresponding to 55 days on a single processing unit.

While the trial run used a relatively small set of calculations, it helped test new ways of sending jobs, pointing to a new window of opportunity for EGEE. Jean Salzemann, computing researcher with the WISDOM collaboration, explained, "*We test millions of potential drugs in our data challenges. Ideally we would have access to one core per potential drug or ligand. With this we would have the complete results in about an hour. When we are preparing future data challenges, we could consider splitting the data load between EGEE and Digital Ribbon resources.*"

"*This successful test with WISDOM shows that applications within EGEE can run well on Digital Ribbon resources", says Bob Jones. "Now we can see if other user communities might benefit.*"

**http://www.digitalribbon.com/**

# Finance – Stock Analysis Application

*"The finance sector is one of the most advanced within the realm of distributed computing. An application of this kind and the tools developed could prove valuable for the EGEE grid community at large, showcasing also the potential exploitation of gLite in Finance research.""*

**Ezio Corso, Excelian.**


Athens University of Economics and Business and the International Centre for Theoretical Physics (ICTP), which forms part of the EU-IndiaGrid project, have developed a Stock Analysis Application for intensive statistical analysis of financial research data related to securities (stocks, bonds, options) able to run on different grid infrastructures including EGEE and underpinned by its gLite middleware.

The purpose of the application was to augment the grid's core middleware services with application-level infrastructure services. The application aimed to identify the building blocks in order to combine or provide solutions to other research challenges. The application automatically analyses a large mass of financial data focusing on the latest trade volume, best-buy price, best-sell price, and so on at the time that a particular event occurs, for example the arrival of a new order. The application facilitates processes that could also be achieved by grid scripting.

So why grid? Grid is well suited for applications dealing with large amounts of data. This finance application handles more than 4 terabytes of input from 700 securities such as stocks, bonds and options, each equalling several years' worth of data. The resources provided by EGEE make it a logical choice for the primary objective of analysing massive financial databases on an instrument-by-instrument basis (one instrument's data is analysed at each node) but may have many other application domains.

The application comprises three distinct parts: Business logic - code that encapsulates the finance/statistics knowledge; Analysis interface - accepts analysis requests and manages grid jobs; and Local environment manager - encapsulates grid knowledge when interacting with Business Logic. The result is the automatic processing of large numbers of files, providing the basis for a full Grid Web Service which can be exploited by user-friendly web applications. Other web services for more elaborate operations are also possible using simple command-line clients.

The application transfers and unzips large quantities of data from secure storage to each node, performs identical computationally intensive statistical analysis of the data at each node and then zips and securely stores the voluminous results of this analysis.

The analysis of the data collected is then broken down into two steps: characterising the behaviour of a security (e.g. trade prices, returns, volatility, waiting time between trades, etc.) and then the time intervals of interest that can change during open-end research, depending on the properties of the data.

Overall, this application has served as an attempt use an architectural approach to building a solution through co-operating services and has the potential to serve as a building block for both research and industry.

**https://euindia.ictp.it/stock-analysis-application**

# Multimedia: GridVideo

*"A grid-based multimedia application for the distributed tailoring and streaming of media files demonstrates how grid technologies can be used for the development of non-scientific applications."*

**Giuseppe Iellamo, GridVideo Developer, University of Messina**

GridVideo, a grid-based multimedia application using gLite for the distributed tailoring and streaming of media files lays the foundations for a Video On-Demand (VOD) service widely used in both educational and commercial environments.

The development of GridVideo was carried out using a grid testbed that has been built as part of the PI2S2 Project spearheaded by the University of Messina. For any commitment on grid environments, the middleware is always a crucial component. Combining components from the best middleware projects currently available (Condor, Globus Toolkit) with components developed for the WLCG project, the gLite middleware facilitates the implementation of grid applications offering an opportunity to use the basic services in an easy and efficient manner.

As the number of Internet-based multimedia applications increases, the provision of tailored content has emerged as a key issue, in particular in environments, such as mobile computing. GridVideo enables the adaptation of media files to the specific features of client devices with enormous benefits in terms of download and processing time and without overloading user devices.

The main idea behind GridVideo is to use grid resources both to store high-quality videos and to perform on-the-fly parallel tailoring to enable the provision of user adapted multimedia contents. Thus, it represents a powerful way to implement a Video- On-Demand Service overcoming challenges associated with tailoring techniques. Services of this kind can have two types of tailoring techniques: the Client Side Adaptation, and the Server Side Adaptation. Client Side Adaptation, where the processing is done by the consumer is subject to resource overload restrictions and channel overload caused by the transmission of original high-quality media. On the other hand, scalability issues due to the potential high number of users affect the Server Side Adaptation, where processing is done at source. In a nutshell, the value-add of a grid-based solution lies in a robust and scalable system that does not cause the overloading of user devices.

The implementation of GridVideo requires constant interaction between Storage Elements, where media files are stored, and Worker Nodes, where the transcoding step is carried out. Thanks to gLite's Data Management System, this requirement is met quickly and easily.

Overall, performance meets the stringent quality-of-service requirements of real-time applications and a reservation based architecture. Users see a fast, responsive service; providers are able to run an efficient system. Architecture of this type is able to manage the grid resource allocation, thus enabling the provisioning of advanced services with high quality of service.

**http://www.italiangrid.org/appdb/appdetails/138**

# S-Sicilia – a Grid-based infrastructure for business

*"The benefits offered through the infrastructure directly impact costs and productivity of companies. Business customers have the opportunity to experience an improvement on their business processes thereby affecting their performance and have the potential to accelerate time to market, improve product quality or increase Return on Investment in terms of cost reduction and increased revenue and profit."*

**Antonio Puliafito, COMETA**

Today, most business processes are carried out over the Internet. The challenge for companies is to make sure their online processes run securely, reliable and efficiently. While different technologies are used to achieve these aims, two approaches have proved better suited in terms of advanced services with reduced costs: Web services along with the Service Oriented Architecture (SOA) paradigm and the grid. Service Oriented Architectures use web services to perform business processes, assuring the right level of security, reliability and performance in a business setting. Grid provides a solid infrastructure with built-in mechanisms such as accounting, load balancing and security. The combination of those paradigms leverages the emerging concept of a service oriented market, where companies can move their businesses, in part or in full, from in-house to online, improving the level of their services and reducing costs.

The aim of the S-Sicilia project, a collaboration between Oracle and the COMETA consortium, is to set up a grid-based business infrastructure capable of making the above scenario a reality. The infrastructure has been divided into three difference parts: Service Level Agreement (SLA) and Quality of Service (QoS) Engines and Virtualisation.

The SLA engine provides the mechanisms to handle business contracts for B2B or B2C. On one side, it acts as the interface with customers, providing access to the whole infrastructure. On the other side, it interfaces with the grid infrastructure, which deals with the underlying resources. It receives customer requests, queries the underlying

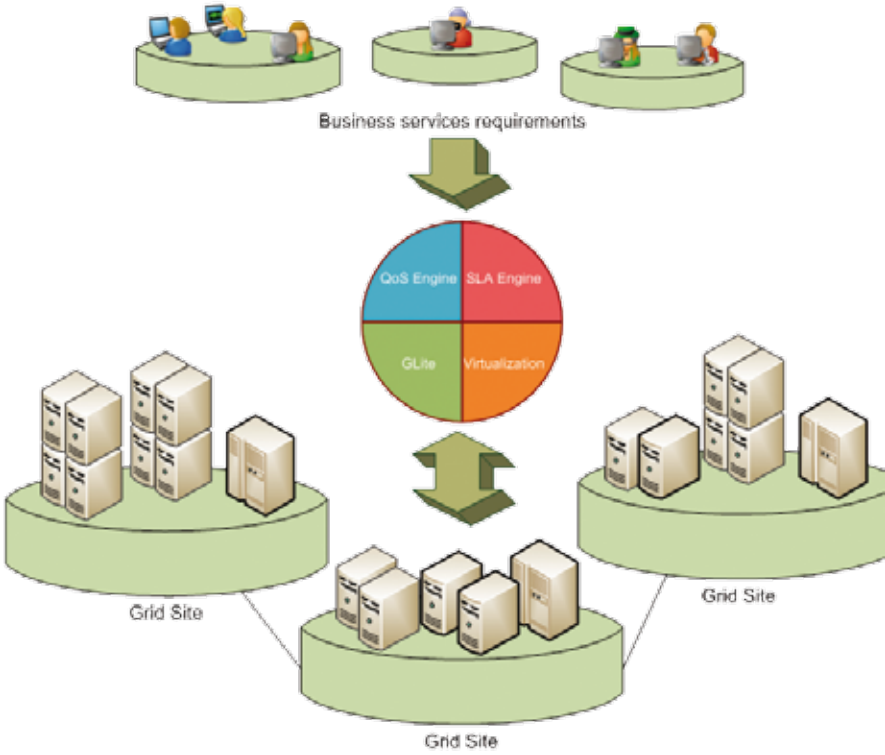


layer for availability and sets up an SLA offer. The customer can then choose whether to accept the SLA. From that point, the SLA Engine which has been designed as a set of interconnected web services that implement these functionalities.

This Quality of Service module works between the SLA Engine and the grid layer, providing the functionalities that the grid infrastructure does not provide on its own. Decisions are taken by considering computation resources (CPU and RAM), storage resources and also networking resources. It also provides a prediction service to the SLA Engine, providing an estimate on the particular service that has been queried.

This Virtualisation module makes users to run applications with particular requirements, offering isolation and security mechanisms complementary to the operating system, customisation and encapsulation of entire application environments, while support for legacy applications. Using virtual machines in the context of grid and large-scale distributed system offers several advantages from providing a confined environment where non-trusted applications can be run; limited hardware resource access and usage; and a

runtime environment that is adapted to the application instead of the other way round; and the management of applications and processes running as a whole.

The S-Sicilia collaboration and its business infrastructure for commercial use, underpinned by EGEE grid technologies, illustrates the value-add of an effective regional infrastructure made available through commercial agreements.

**http://s-sicilia.unime.it**

**http://www.consorzio-cometa.it/**

**http://www.oracle.com/technologies/soa/soa-suite.html**

# For more information

Visit the EGEE and Business pages at: http://business.eu-egee.org

Download the EGEE and Business Brochure for an overview of EGEE's Programme for Business, current Market Offers and EGEE Business Associates: http://business.eu-egee.org/index.php?id=120

Contact: Trust-IT Services – info@trust-itservices.com (www.trust-itservices.com) or Gabriel Zaquine, Chair of the Business Task Force gabriel.zaquine@bt.com

e-infrastructure

EGEE-III INFSO-RI-222667