



SPACE, TIME AND THE GRID: THE SCIENCE BEHIND GRIDPP


- What are things made of? *What are the basic building blocks of matter?*
- How are things held together? *What forces allow particles to interact and change?*
- What are space and time? *What is the shape and size of the Universe?*
- Are there other Universes? *Are our laws of physics the only possible ones?*

These questions were familiar to the ancient Greek philosophers – yet after 2500 years' progress, they are still unanswered, and still relevant. Today's experimental physicists attack the problem by building apparatus to observe the interactions of fundamental particles, the most basic form of matter known. Analysis of data from these experiments is now so demanding that it requires a new approach to data-handling and computing: the Grid.

What we know (and what we don't know)

A Standard Model of particle interactions, representing matter and forces as excitations of many quantum fields, has been developed over the last fifty years. The model has been tested exhaustively in experiments, and no flaw has yet been found. However, it cannot be a truly fundamental description of nature – we know there are missing elements:

CGI illustration of CERN



● **A missing piece:** The model is mathematically inconsistent without a Higgs Boson, responsible for allowing the other particles mass – but it has not been observed experimentally, and may not exist in nature.

● **Hidden problems:** Our world may be a large-scale 'shadow' of a more complex multi-dimensional Universe at tiny scales. This means that the Standard Model should break down at higher energies than those accessible today, revealing a more fundamental picture.

To make progress, we must carry out experiments at high energies, where the Standard Model is known to fail. The newest particle accelerator at CERN, the Large Hadron Collider, will be completed in 2008, and will produce higher energy collisions than ever before.

The data deluge

For hundreds of years, microscopes have been used to examine objects below the human scale of existence. Modern particle physics carries on this approach; rather than using visible light waves, we rely upon the wave-like quantum nature of the particles themselves. Evidence of new physics – for instance, the production and decay of a Higgs Boson – will lie in an extremely rare type of collision between very high energy protons. We will carry out collisions at a rate of over 2GHz for thousands of hours per

● **Gravity:** Only three of the four known forces are described, with gravity playing no role, yet at very high energies (for instance, soon after the big bang), gravity plays an important role.

● **Symmetry:** There are many unexplained similarities between different types of particles; this implies that there are yet more underlying 'fundamental' objects which complete the symmetries.

Quarks	u	c	t
	d	s	b

Leptons	ν_e	ν_μ	ν_τ
	e	μ	τ

Gauge Bosons	γ	g	H
	W	Z	

Quantum fields in the Standard Model

year in order to distinguish the one hundred or so events that will contain the most interesting information. This is analogous to searching for a needle in 10,000 haystacks. Each collision record kept for later study will require storage capacity of around 2MB.


Why do we need the Grid?

The LHC will not only use high energies – it will also allow us to look for a range of rare phenomena that provide the clearest signs of new physics. We will need to work with more experimental data than ever before; around 15,000,000 Gigabytes per year.





Extracting scientific knowledge through the detailed analysis of this scale of data is an unprecedented challenge, beyond the capabilities of any single computer centre. The Grid provides a new way to harness very large international computing and data storage resources in a flexible way, such that they can be brought to bear on large problems.

Physicists around the world will use Grid tools to build a single worldwide computing system from resources at many institutes. It will be capable of moving, manipulating and analysing the gigantic LHC datasets on demand. Researchers will have access to all LHC data in their search for new phenomena – but they will not need to be distributed computing experts, nor know the location of their data, or where it is processed.




The LHCb experiment will carry out a detailed study of matter-antimatter symmetry, and to explain why the visible Universe consists only of matter, not antimatter.



The ALICE experiment will focus on observing a new liquid state of matter made when heavy nuclei collide at high energies (the quark-gluon plasma).

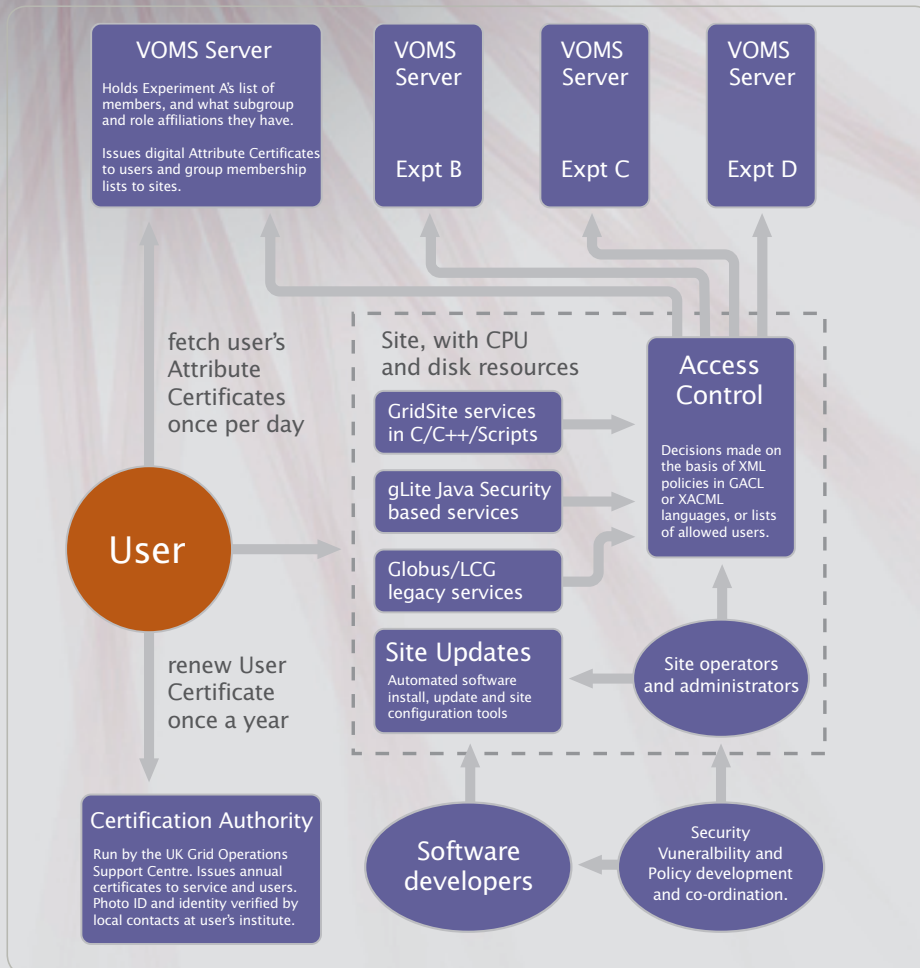


The ATLAS and CMS experiments are general purpose sets of apparatus designed to explore physics beyond the standard model – in particular, the detection of the Higgs boson, new 'superparticles', or evidence of extra spatial dimensions.





GRID SECURITY: BUILDING A SAFE AND SECURE GRID



The how and the why of Grid jobs

The GridPP and LCG infrastructures focus on the execution of batch jobs, processing particle physics data on CPU and storage farms with hundreds of nodes. The farms are presented to the Grid as sites, and the security model reflects this mode of working, with packages of work bundled up in jobs for remote execution at a suitable site selected by the Grid. X.509 Proxy Certificates are granted to jobs, to allow them to act on behalf of a user, either using the legacy Globus GSI protocol, or by the EGEE / GridSite Delegation portType originated by GridPP. When a job or a file access request arrives at a site, authorization policies are applied in the form of lists of users, or XML policies written in terms of VOMS memberships. Web Services written in Java or using the GridSite C/C++/Scripts hosting environment, and legacy Globus services, all use this common framework based on X.509 credentials, which provides a consistent authentication history for the purposes of auditing and accounting. GridPP has played a leading role in establishing the EUGridPMA and the worldwide International Grid Trust Federation, which co-ordinates policy and good practice amongst Certification Authorities used on the Grid.

Digital certificates

GridPP's security work is part of the worldwide LHC Computing Grid (LCG) and Enabling Grids for E-Science (EGEE) projects. X.509 digital certificates are used throughout to identify and authorize users, hosts, jobs and services, and we are developing middleware and operational policies to extend the scope and robustness of the Grid. Certificates for users and hosts are provided by the Certification Authority run by the UK Grid Operations Support Centre. Individuals who apply for certificates have their identity verified by local contacts at their institute, using their passport or another form of photographic ID.

Virtual organisations and being a member

Users are grouped into Virtual Organisations (VO), which correspond to experimental particle physics collaborations of hundreds of members, and other application areas such as biomedicine. These VOs are responsible for managing group and subgroup memberships, on which authorization decisions are made, by running a VO Membership Service (VOMS). These services publish lists of user certificate names and allow users to request a short-lived X.509 Attribute Certificate which can be used to prove VO and subgroup membership. GridPP operates VOMS servers for several UK-based particle physics collaborations, and supports access by members of other VOs based elsewhere in the LCG infrastructure.

Other uses for our certificates

GridPP also provides a central documentation website, using GridSite for the management of HTML and binary files such as PDF and Word documents; and GridSiteWiki for free-form documentation and notes, with easy cross referencing. Both of these systems use the same X.509 user certificates as the Grid itself to authenticate users, and to control write access and read access to internal documents.



The Real Time Monitor, developed by GridPP members at Imperial College, showing sites on the EGEE Grid.



Attendees at a 2007 meeting of the European Policy Management Authority for Grid Authentication (EUGridPMA) in the UK






VIRTUAL ORGANISATIONS: COMMUNITIES AND GRIDPP

The concept of the virtual organization (VO) is as inherent to the Grid as the sharing of computer resources. In practical terms, this means the ability to identify a Grid user with a well-defined and managed community via their Grid certificate, and to allocate and account for their use of computer resources based on that VO.

United Kingdom Cumulative Normalised CPU time by VO and DATE
(Excluded dteam & ops VOs) EGEE VOs. August 2006 - July 2007



Managing VOs

In particle physics, the primary VOs are established for each experiment, but it was quickly realised that more was needed; individuals need to be associated with groups within the VO, have well-defined roles and also have other attributes such as a regional identity. All of these attributes can affect the allocation of resources and priority to be given to a particular task, all of which is controlled by a combination of VO and local site policy.

To implement all of this, the Virtual Organization Management System (VOMS) was created, and the Grid middleware is increasingly making use of the VOMS information. The biggest and most pressing area of development in this regard is in access control and allocation of storage.


GridPP's computers have been helping find drugs against the avian flu virus, pictured left.

VOs for physics

GridPP supports the VOs of all experiments in the UK particle physics programme, as well as VOs for the phenomenology community (PhenoGrid), the MAGIC gamma-ray telescope, the PLANCK cosmic microwave background satellite and others such as fusion research. GridPP provides the infrastructure services for the management of many of these VOs, a role which it is now sharing with the National Grid Service (NGS).

VOs beyond physics

GridPP's supported and managed VOs are not restricted to particle physics. As part of our mission to disseminate the technologies developed, we have supported commercial communities such as Cambridge Ontology (attempting to automatically index and classify images in a way that allows them to be searched and retrieved). Other scientific communities have also benefited from GridPP, such as the BioMed biomedical VO, which has undertaken published studies on H5N1 (bird 'flu') and malaria using in part 'spare' computing capacity on the GridPP UK computing facilities. These latter academic communities are brought in through GridPP's association with the EU EGEE project; the intention is that reciprocal capacity will become available to particle physics, leading to efficiency gains on the computing investment for all partners, and possibly the establishment of an EU Grid infrastructure.



VOs for Grid development

Finally, GridPP supports and manages a GridPP VO and also a VO corresponding to each of the Tier 2s. These are to allow small-scale or initial usage of the facilities by small communities without the (initial) overhead of setting-up and managing their own VO infrastructure.

As a result of the VO system, UK particle physicists are accessing computer resources from the UK to Taiwan, and the UK is attracting involvement in research and other projects with the same world-wide distribution.

Accounting and Grid use

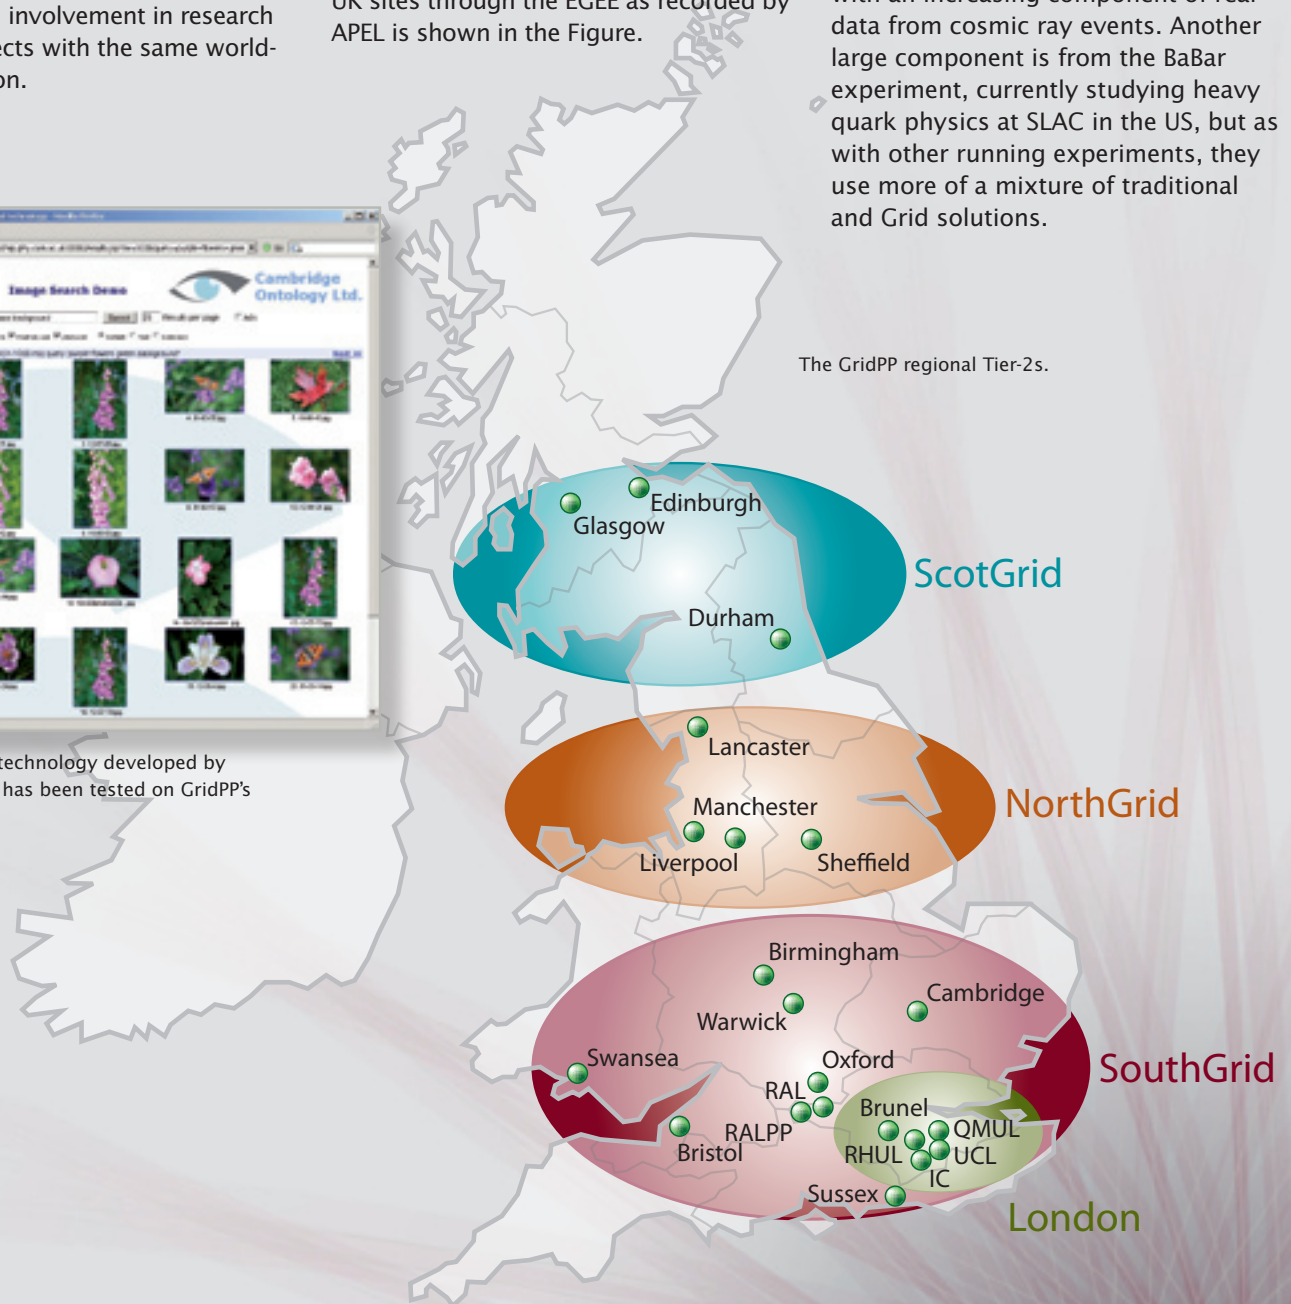
In all of this, accounting for the usage is vital in order to correctly implement policies, to allow the correct attribution of the contributed resources and perhaps in time to allow appropriate cost sharing. The UK has made major contributions to the EGEE accounting structure, including the APEL system and the underlying R-GMA information system technologies. The usage of CPU by the major VOs on UK sites through the EGEE as recorded by APEL is shown in the Figure.

The Tier 1 capacity was funded by PPARC/STFC, while the regional Tier 2 sites were funded by the GridPP partner institutions, but with effort provided by GridPP. In both cases, the dominant usage is from the LHC experiment communities as part of their ramp-up ready for the large data volumes to be produced by the Large Hadron Collider. At present, the data is mainly from simulation studies, but with an increasing component of real data from cosmic ray events. Another large component is from the BaBar experiment, currently studying heavy quark physics at SLAC in the US, but as with other running experiments, they use more of a mixture of traditional and Grid solutions.



The image search technology developed by Imense Ltd, which has been tested on GridPP's infrastructure.

The GridPP regional Tier-2s.





SUSTAINABILITY: BUILDING A LONG-TERM GRID INFRASTRUCTURE

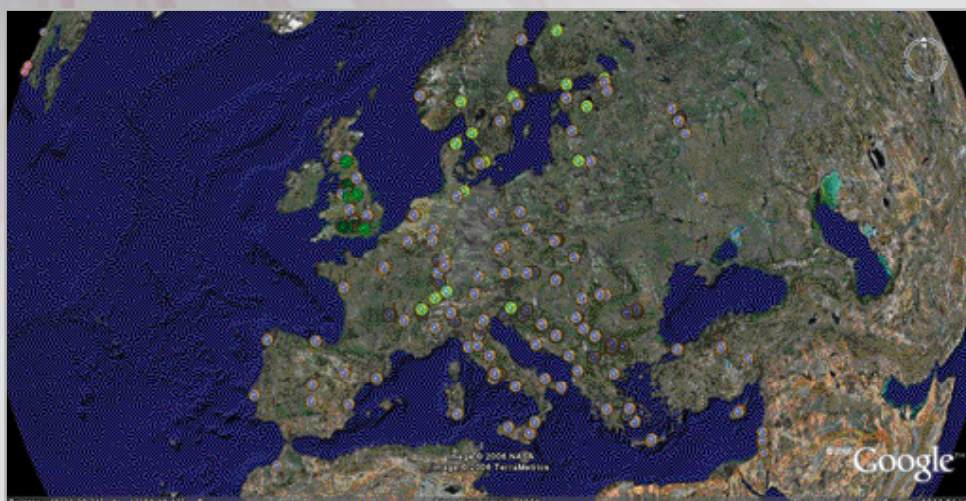


Image from a monitor running on Google Earth that shows nine Grid infrastructures

Sustainability

From the outset GridPP was to create a UK Particle Physics Grid and the computing technologies required for the Large Hadron Collider (LHC) at CERN. This would help place the UK in a leadership position in the international development of an overall EU Grid infrastructure, which could become multi-disciplinary.

GridPP

GridPP has fulfilled its early promise and is in its third full year as a Production Grid, has more than 5,000 CPUs and more than 1/2 Petabyte of disk storage. This has led the UK to be the largest CPU provider on the EGEE Grid, with total CPU used of 15 GS12k-hours in 2006. Now nearing the end of its second phase the project has met 69% of its original targets with 92% of the metrics within specification. The initial LCG service is now underway and has been run for the first 6 months of 2007. The aim is to continue to improve reliability and performance ready for start-up of the full Grid service.

EGEE

Enabling Grids for E-science (EGEE) is the EU wide multi-science Grid involving more than 100 institutions worldwide. The aim is for this Grid to have by the end of 2008 almost 100MS12k CPUs and close to 80PB of storage. Starting in 2003 and now in its second phase it currently has 177 sites, 29,266 CPUs, and 13,815 TB of storage. The main vision of EGEE's third phase (April 2008 to March 2010) is to make a strong move towards a sustainable world-wide production quality Grid infrastructure by appropriate technical and organisational evolutions.

NGIs, EGI and The Future

While the second phase of GridPP finishes in April 2008 a third phase, GridPP3, has been accepted by STFC (£30m) to extend the project to March 2011, securing the future of the UK's Grid till then. However EGEE and GridPP are still "just" projects and there is a need to prepare a permanent, common Grid infrastructure across Europe and the rest of the world. It is for this reason that the European Grid Initiative (EGI) which will ensure the long-term sustainability of the European e-infrastructure independent of short project funding cycles is being plotted out.


As it stands there are many Grid infrastructures and sometimes more than one in each country. EGI aims to encourage the creation of National Grid Infrastructures (NGIs) and coordinate the integration and interaction between NGIs while also operating the European level of the production Grid infrastructure which links the NGIs.

Interoperability

The future of the Grid is very dependent on EGEE and other grid infrastructures co-existing and being able to interoperate. The EGI/NGI era will rely heavily on interoperability and coexistence. As mentioned there are many Grid infrastructures and interoperability has to be considered at many levels – campus, local, national, regional, international. When EGI comes on stream these will likely be large production systems and they will have inertia, changing quickly is not easy. Introducing new software and standards is slow, there is a need to maintain backward compatibility. To prepare for this we need to work on this problem as soon as possible. It is for this reason that EGEE and GridPP use gLite, a choice based on interoperability needs and impact assessment on the infrastructure.



Computers in the GridPP cluster at Queen Mary University of London




Work being done


Grid Interoperability Now (GIN) has already put in effort in several areas key for EGEE. It works very closely with the Open Science Grid in U.S. to ensure that the two Grids can work together as much as possible. The weekly EGEE operations meetings are attended by OSG staff, processes set up with OSG for operations and user support workflows and OPS VO defined to support joint operations for testing/monitoring use and collaboration on monitoring tools and procedures. This has already led to CMS being able to submit production work to OSG from EGEE sites. EGEE and GIN are also working with other grid projects on specific interoperability at the level of middleware e.g. NAREGI, Unicore, NDGF(ARC).

Moving ahead to convergence

To ensure interoperability between infrastructures in the short term and a smooth transition to the EGI/NGI era it is important that decisions and solutions are driven by users and providers. Appropriate and workable standards for both sides are essential. Dialogue between both parties will help actively push convergence for the most pressing needs, leading to best practices which in turn should drive standardization efforts. Care needs to be taken not to fix standards too soon as Grids are not a mature technology.



UK ATLAS user job success



User experience getting better