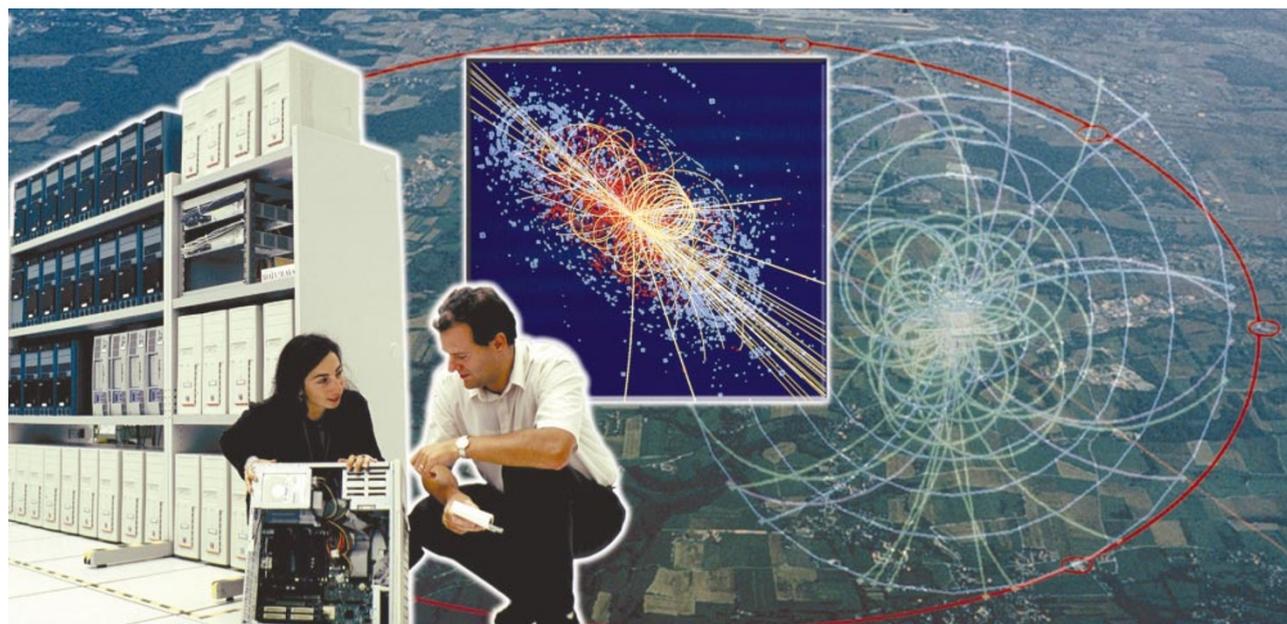


Tomorrow's computing today

The physics lab that gave us the World Wide Web is now gearing up to make the first practical deployment of the Internet's next big thing: the Grid. Declan Butler takes a test drive.



P. LOEZZE, MALANDAIN/CERN

So this is it ... the future of scientific computing. I'm at the headquarters near Geneva of CERN, the European Laboratory for Particle Physics, getting a sneak preview of a prototype Grid. The advanced publicity has led me to expect something special. Grids, according to the hype, will soon allow scientists to harness the power of a global supercomputing network. Just think about it: greater number-crunching power than in your wildest dreams, and seamless access to mountains of raw data stored in colleagues' labs worldwide.

From an ordinary-looking Linux PC, you submit a supercomputing job in high-energy physics. A few hours later, the Grid tells you the results are ready. To be frank, I'm underwhelmed. But that's just what David Williams, formerly head of CERN's computing and networks division, now responsible for the lab's relations with the European Union (EU), wanted to hear. "Grids will have succeeded when they become invisible to the user; when we stop talking about them," he says.

It is only when you peer under the hood that you see why those in the know are getting excited about Grids, and why they may revolutionize the way in which research is conducted. Grids should be able to give scientists

unprecedented computing speed, data access and functionality thanks to a suite of software tools known as 'Globus middleware', which was developed by Grid godfathers Ian Foster of the Argonne National Laboratory near Chicago, and Carl Kesselman of University of Southern California's Information Sciences Institute in Los Angeles.

Live link-up

The smartest component is the 'resource broker'. When you give the Grid a job, this goes off and negotiates with computers signed up to its network to book the computing power you need, which might include a supercomputer in Boston and a 'farm' of thousands of PCs in Madrid. It also scours data centres worldwide for the files and data that you want, and, if needs be, converts them into compatible formats. It will even fire up the software needed to perform a particular analysis.

To pull this off, the resource broker must chatter constantly with the other pillars of the Grid middleware. The 'information service' keeps its finger on the pulse of the network, tracking up-to-the-minute information on where and how much computing power is available, and how fast various data links are running. The 'replica catalogue' is

Combined effort: physicists hoping to identify the signatures of rare subatomic particles (inset) plan to link up computers around the world to give them the data-crunching power they need.

the Grid's Yellow Pages, keeping a note of the type and whereabouts of all data and files. The 'logging and bookkeeping' server, meanwhile, acts as an office manager, tracking the status and history of all submitted jobs.

Many labs are working on prototype Grids, but at CERN the concept is about to be given its first proper outing. Here, the geeks are working with a gun against their heads. They have been set a deadline of July to prove that the Grid can deliver to the lab's scientists the awesome computing power that will be needed when the US\$2-billion Large Hadron Collider (LHC) comes online in 2007.

By smashing protons into one another at close to the speed of light, the LHC will spew out petabytes (10^{15} bytes) of data. That's about the same quantity of information involved if everyone on Earth were each talking into 20 telephones simultaneously. Clever algorithms that filter out 'noise', events that physicists are reasonably confident are not interesting, will eventually stem the flow somewhat. But if researchers across more than 40 countries are to detect the signatures of exotic subatomic



L. GUIRAUD/CERN

Information overload: the Large Hadron Collider will generate huge amounts of data.

particles such as the Higgs boson, which is hypothesized to give other particles their mass, only the Grid will do. "It's not an exaggeration to say: 'no Grid, no LHC,'" says Williams.

Given these high stakes, high-energy physicists want to do a few 'dry runs' — which is where the July deadline comes in. CERN is now gearing up for a groundbreaking project in which physicists will simulate the creation and analysis of the giant sets of data that they will have to handle, and so 'tune' their computing models and software. "This is not yet another Grid technology project, it is a Grid deployment project," says Les Robertson, who is heading the effort.

Until now, CERN has been operating the EU-funded European DataGrid testbed. This involves some 1,000 networked machines in 15 countries, 5 terabytes (5×10^{12} bytes) of disk space and 350 test users. July is about making this network available to all LHC physicists. From there on comes a series of deadlines by which the network must be ratcheted up to a system that will eventually require the equivalent of 200,000 of today's fastest PCs, store up to 8 petabytes of data annually — the same as 16 million CD-ROMs — and serve 8,000 physicists across the world.

The challenge facing CERN's Grid specialists is the problem of 'scalability', the catch-all term used by engineers to explain why something that worked in the prototype doesn't work in the real world. Take the resource

broker. When physicists run operations such as Monte Carlo simulations — a common statistical technique that can create many computing tasks at once — it can get hammered with requests and often packs up. To get round this, information technology (IT) staff at CERN are rewriting their software to include multiple resource brokers in the Grid working in concert to distribute the load.

Although Grids are designed to let you weave together myriad computers or operating systems into one seamless whole, and shuffle data around irrespective of location, in practice incompatibilities crop up all of the time. "There is a gap between the hype and technical performance," admits Ben Segal, who formally retired last year from his post as manager of the European DataGrid, but still spends some of his time working on the project.

Language barriers

Although scientists worldwide know how to collaborate, their computers do not. Compilers, for instance, are key components of any computer that convert the various languages used to write software into machine code that a computer can understand and execute. But the compilers used in LHC experiments often can't communicate with those on machines around CERN's Grid, which likewise are often unable to talk to one another. Quality of service is another major problem: one university in the Grid may have an excellent IT infrastructure, but another's systems may crash all the time.

Add these problems together, and it becomes clear why 20% of all jobs submitted to the prototype LHC Grid currently fail to complete. At present, keeping the network online also requires a great degree of human intervention. "You've got to turn all the knobs simultaneously to the right place and hold them there if it is to work," says Ian Bird,



Ian Bird is confident the compatibility problems that may limit the Grid can be overcome.

deployment manager for the LHC Grid.

But if anyone can succeed in ironing out the glitches, CERN's computing specialists would be among the favourites. To avoid the Grid descending into gridlock, a key step over the coming months will be to marry the best technologies from CERN's own European DataGrid with three parallel research efforts in the United States: the Particle Physics Data Grid, the Grid Physics Network and the International Virtual Data Grid Laboratory. It won't be easy, but Bird is quietly confident. "Of course we will make it," he chuckles.

In addition to the technical obstacles, getting the LHC Grid online requires unprecedented formal collaboration between the IT service departments of the institutions taking part, to agree on joint standards and best practices. Any institution taking part in the Grid has, for example, the right to carry out on-site inspections to verify that another's security and system performance are up to scratch. Thankfully, says Roger Jones, a physicist at Lancaster University, UK, it has proved "surprisingly easy" to get agreement among institutions, probably because of the high stakes involved.

Indeed, one could argue that the human element of Grids is just as important as the fancy technology. Increasingly, argue Grid enthusiasts, scientists will see themselves less as belonging to individual bricks-and-mortar institutions, and more as members of 'virtual organizations', communities of researchers in defined research areas or associated with particular experiments, who together decide what computing and data resources they will share over the Grid. Gone will be the logjams caused by limitations in computing power and data storage at one institution, and the need to rack up endless frequent-flyer miles to participate effectively in a project.

The primary goal of the LHC Computing Grid is to prove that physicists will be able to handle the avalanche of data that they will be presented with when CERN's new collider comes online, and to make all this information seamlessly available to individual researchers as if it were all on their own hard disk. But the project is also an experiment in a new sociology of e-science. CERN hopes to show the way, and encourage other disciplines to embrace the Grid's possibilities.

That may sound like a grandiose goal, but remember that CERN has already once changed the way in which the world interacts with computers, through the work of Tim Berners-Lee, who invented the technology that became the World Wide Web as a means of allowing physicists to share documents. Turning the Grid from testbed to reality would be a fitting encore. ■

Declan Butler is Nature's European correspondent.

LHC Computing Grid

♦ <http://lcg.web.cern.ch/LCG>

European DataGrid

♦ <http://eu-datagrid.web.cern.ch/eu-datagrid>

CERN