'Eighteen months ago only 8 per cent of companies were using Grid computing. Today that number has increased significantly.'

solutions

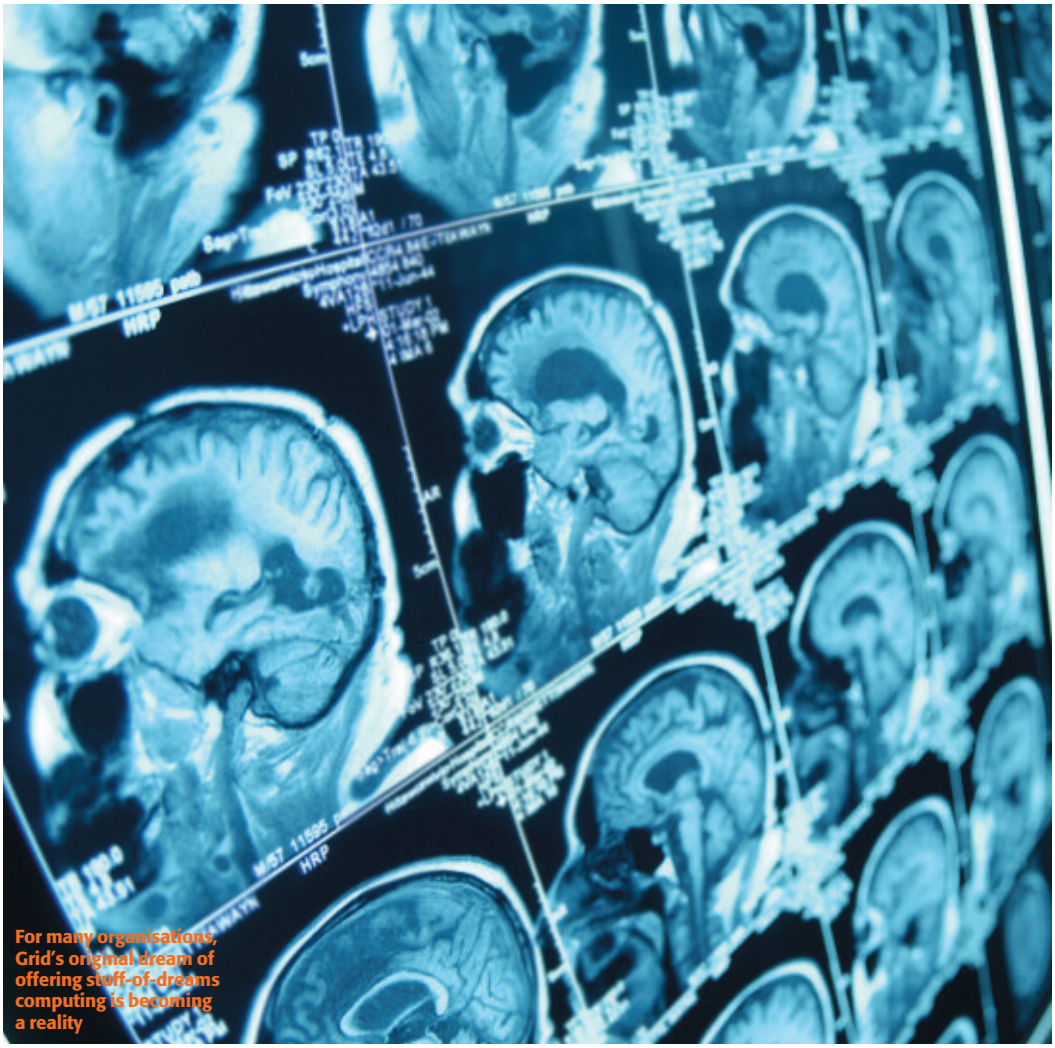

For many organisations, Grid's original dream of offering stuff-of-dreams computing is becoming a reality

BRINGING GRID INTO THE IT MAINSTREAM

As Grid technology starts to merge with open networks public and private, its viability for commercial applications is growing. **William Knight** charts its progress.

TRANSPORT PLANNERS argued for years that better roads eased congestion, oblivious to the fact, perhaps, that improving a product results in more customers, not fewer. The phenomenon is entirely consistent with the economic theory of induced demand: if supply grows, use will, perforce, increase.

Curiously, this theory also helps explain this year's rise in Grid computing. For the last 20 years Grids have attempted an escape from academia and into mainstream commerce, only to find the way blocked by concerns over where the data goes, who owns the infrastructure, and, not least, a 'what do I need that for?' attitude.

As long as Grid is perceived to be an 'exclusive' technology, operating in its own space, then perceptions of what it can deliver to the open market are skewed. But this is changing. Ian Osborne, project manager of the government-sponsored Knowledge Transfer Network, IECnet – promoting Grid concepts to UK organisations – thinks Grid technology is now merging with general operating infrastructures.

"It's becoming quite common to implement a system that allows you to register the resources you have, have them managed in some central way and have the work scheduled across all those resources," Osborne says.

For most enterprises, this merger has little to do with promises of massive parallel computing but with the perennial business concerns of cost and efficiency.

"The technologies and principles of [Grid] management and the principles of resource sharing are now widely adopted within virtualisation software," explains Osborne. "You can basically increase your servers loading from 7 per cent, on average, to 80 per cent or 90 per cent – people are taking their installed base down by a factor of eight or ten."

GREEN CROSSROADS

Such efficiency improvements are timely. With the seriousness of climate change, IT departments are feeling the pinch of green fingers, and the game changes further when emissions trading – or 'cap and trade' – comes into effect, says Osborne, referring to the EU's flagship emissions trading strategy.

At the moment, only one in ten data centre managers even views the energy bills, Osborne believes, but carbon costs will start to bite in two to three years: "IT managers will have to be ready to argue with the retail manager and the logistics manager for their share of the energy budget."

And yet, for many companies, Grid's original dream of offering stuff-of-dreams computing power is also a reality. ▶

'Thousands of simulations every day are run to forecast, control, to get statistics and simulate different scenarios – it costs about $100m to drill a well, so we have to get it right'
John Hybertsen, StatoilHydro

A year ago the Norwegian oil company StatoilHydro embarked on a project to link five of its Linux clusters – four in Norway and one in Houston, Texas – into a multi-cluster Grid using Platform Computing's LSF workload management software. The result was an 850 CPU Grid capable of running thousands of 'reservoir simulations' each day.

Each simulation models an underground reservoir by splitting the rock strata into thousands of 3D blocks or cells each with their own properties of permeability, porosity, oil-saturation temperature and structure, among many others. It answers fundamental questions like, "if I drill here, and pump in water there, how much oil will come out?"

Such an application sucks up immense quantities of computer power, but by running hundreds or even thousands of such simulations StatoilHydro can identify the likely optimum strategies.

John Hybertsen, StatoilHydro principle engineer, explains: "We have thousands of simulations every day. They are run to forecast, control, to get statistics and simulate different scenarios. It costs about $100m to drill a well, so we have to get it right, and you have to control how you drain the reservoir."

Hybertsen foresees the requirement for more simulation needing ever greater computing power and he expects improvements to be ongoing: "We are looking at more nodes, more CPUs. We have refined the granularity – the blocks are smaller – and we need to have more memory to run it in a fair amount of time."

The investment is well worth it, as Hybertsen reports, "There is a huge effort in getting more from the reservoirs. In some reservoirs they started out by saying that current production technology would get only about 45 per cent of the oil. But we have aims of 55 per cent or up to 60 per cent, and you don't do that without controlling the production and knowing what is happening in the reservoirs."

Terje S. Knudsen/StatoilHydro

DRIVING SIMULATIONS

Platform Computing's vice president for high performance computing (HPC), Christoph Reichert, reports sharp growth in use of Grid simulations in automotive, electronics design, animation as well as oil and gas exploration.

"Everywhere you have huge calculations, in particular simulation jobs – whether you are simulating electronic circuits on a chip, whether you are doing the wingtip aero-dynamics on an Airbus A380, or whether you are looking at the chassis of a formula one race car

Osborne: IT demands more and more

– all of these software modules and applications are a target for our kinds of technology," Reichert avers.

But in financial services Grid technology has more than proven itself. Driven by success with the massive calculations in the so called Monte Carlo simulation, "Grid has become the *de facto* mechanism within most of the finance houses, certainly in America," explains Frank Falcon, senior project manager at Colt, the city based telecoms enterprise. "Every Wall Street firm does Grid now. It's become a standard mantra from the CIOs that they want to get the same benefits for the trans-actional-intense applications."

Falcon believes some transaction-heavy applications, if enabled for Grids, could easily see a 45 per cent drop in hardware requirement and perhaps a 30 per cent drop in operational costs. Colt is positioning itself as a Grid service provider offering utility computing and other Grid-centric services, since Falcon's analysis sees finance houses ravenous for more computing power, a hunger too great and too expensive to satisfy from in house resources.

This is largely as IECnet's Ian Osborne also predicts: "Businesses are all running on IT; all decision making and all business processes, are maintained by IT. IT has become so fundamental to the way we do business, to the productivity of the knowledge economy, that we will simply continue to demand more and more."

The double benefit of squeezing out inefficiencies – and undreamt of computing power – is seeing Grid computing racing, almost unseen, into new markets. And just like a good road system attracts users, now Grid infrastructure is common place, and its benefits apparent, there will be no end of applications vying to take advantage.

GRID CONSEQUENCES FOR IT PROFESSIONALS

In the three months to October 2007, ITjobswatch.com reported a nearly 300 per cent rise in demand for Grid computing over the same period the year before. Such growth is predicted to continue.

"Eighteen months ago only 8 per cent of companies were using Grid computing. Today that number has significantly increased, with many businesses using Grid computing in some guise as they look to improve utilisation of IT assets by pooling resources," says Peter Critchley, strategy director at IT consultancy Morse.

And yet there is a possible staffing crunch in the offing. "One of the key issues the technology industry faces today is ensuring software developers are able to build new applications to take advantage of distributed, scalable computing," says Ian Osborne.

He believes the recently announced IBM and Google joint initiative to help students better understand net-scale projects is of major importance. "This initiative could be key to helping the next-generation of Internet software engineers to understand distributed computing and the software challenge we face," says Osborne.

Understanding Grid computing and its associated technologies is essential for information professionals: not just for the short-to-medium terms, but for ten years' time, when 50 core processors will be commonplace.

Many of the advertised Grid computing positions (ITjobswatch.com) are in the finance sector (40 per cent) and for developers (Java 54 per cent), but the importance of understanding distributed and Grid technologies, for all IT professionals from developer to CTO, is growing daily.

Further, environmental focus will add greater impetus to consolidation and virtualisation efforts. A recent poll of 335 companies by Storage Expo 2007 found that 95 per cent of organisations are making IT infrastructure investments to help reduce their carbon footprint.

"Companies are more aware than ever of the importance of reducing energy consumption and costs through better management of IT assets," says Morse's Critchley. "It is key that IT professionals have the knowledge on where and how resources can be effectively pooled in order to successfully reduce costs and energy consumption.

"Whether this is referred to as Grid computing, utility computing, on demand or virtualisation, it is a core competency in today's pressurised green IT environment," Critchley adds. ■

Further information

www.platform.com
www.statoil.com
www.gridcomputing.com
www.ibm.com/grid
www.gridforum.org
www.gridcomputingplanet.com
www.gridcomputing.org

CASE STUDY: STATOIL ASA

Reservoir engineers in the Sub-Surface Division of Statoil ASA, a Norway-based oil and gas company and one of the world's largest crude oil traders, use 3D simulation applications to search for potential oil-bearing structures. These applications involve vast amounts of data, complex calculations, and require that thousands of iterations be run.

Local computing clusters at each of Statoil's four Norway locations were not meeting engineers' needs for increased ease of use in job submission and for additional computing power to run more model iterations.

Using LSF workload management software – from Canadian software company Platform Computing – with its MultiCluster capability to interconnect the four clusters, in concert with EnginFrame, a Web portal for job submission, Statoil implemented a single, enterprise-wide computing Grid, giving engineers access to all computing resources in the division.

Another key aspect in Statoil's project is the ongoing work carried out by reservoir engineers, who use computers and sophisticated 3-D modelling software to simulate the flow of fluids in oil and gas reservoirs. Their primary tools in this pursuit are the ECLIPSE Simulators from Schlumberger, which provide numerical simulation techniques for reservoirs and all degrees of extraction complexity. ECLIPSE can be very compute-intensive, involving huge amounts of data, large numbers of highly complex calculations and often, hundreds of simulations for every field analysis. With ever increasing demands for more oil and gas output from existing fields, Statoil engineers are performing more simulations than ever before.

Initially, each of Statoil's four Norway locations independently operated its own local computing infrastructure consisting of engineering workstations and server clusters that ranged in size from 64 to 400 CPUs. This approach was throwing up a number of problematic issues.

Platform Computing's LSF solution lets Statoil grow its Grid by harvesting unused CPU cycles

Statoil ASA

For example, each location was investing in its own computing equipment, and the clusters on each site needed to be oversized to accommodate peak periods of activity. In addition, home-grown interfaces based on old scripting technology were used for job submission and to integrate the ECLIPSE application with the Platform LSF software used to control the computing clusters.

These scripts were old, difficult to maintain and did not provide adequate visibility into the job execution activity going on within the clusters.

The disparity in the relative number of CPUs available to each location also caused problems in terms of differences in engineering process consistency, performance and reservoir simulation accuracy.

To better meet current needs, and position the company for continued growth, Statoil's IT department turned to Platform's Grid technology to create a computing environment which would foster a consistent methodology for job submission among Statoil reservoir engineers worldwide.

Platform LSF MultiCluster software was used to tie the four local server clusters into a single computing grid, giving each engineer user, irrespective of their location, access to the division's entire complement of computing resources, which is fast approaching 1,000 CPUs.

Platform worked with Schlumberger's software development team to integrate ECLIPSE, Platform LSF and EnginFrame. This allows ECLIPSE users to perform significantly more iterations for multi-realisation simulations on cost-effective server clusters, which in turn yields more accurate results – when you're drilling for oil.

Platform LSF allows Statoil to dynamically grow the size of its Grid by 'harvesting' unused CPU cycles as needed from any of the 185 user workstations across the division that might not be running at full capacity. This accounts for nearly 25 per cent of Statoil's overall computing capacity, and is totally transparent to the workstation users.

Next, a Web portal was implemented based on EnginFrame from NICE, giving users an interface that simplified the submission of jobs to the Grid. EnginFrame's ability to support the submission and monitoring of jobs by any skill level of user from anywhere in the world is a key component of Statoil's grid solution.

As a result of giving users access to greater computing power, Statoil is seeing a significant increase in the amount of simulation work taking place throughout the grid – without having to add more computer hardware.