


SOFTWARE

ETICS assures quality on the Grid

On 1 January 2006, an EU-funded project was launched that will improve the coherence and quality of Grid software. The e-Infrastructure for Testing, Integration and Configuration of Software project (ETICS) will integrate existing procedures, tools and resources to create a facility where distributed research projects can integrate their code, libraries and applications; validate the code against standard guidelines; run extensive automated tests and benchmarks; produce reports; and improve the overall quality and interoperability of the software.

Scientific software today is often the product of large distributed collaborations and increasingly uses new technologies like the Grid to solve complex, computer-intensive problems. Owing to a large variety of available tools, programming languages, platforms and so on, these software stacks risk suffering from a lack of coherence and quality. Limited timescales, manpower and funding often prohibit the creation of a dedicated build and test infrastructure for each new project.

ETICS will provide such a capability for software configuration, integration, testing and benchmarking for the scientific community, including software engineering tools and support infrastructures developed by other projects (Enabling Grids for E-science



The defect/kSLOC chart is used to monitor the general quality of the code. The pre-release quality index tends to settle on a constant value that can be taken as measure of the quality of the software development process, whereas the post-release quality index tends to zero as the software reaches maturity. ETICS launched on the first day of 2006.

[EGEE], the LHC Computing Grid, the NSF Middleware Initiative) and open-source or industrial entities. To promote international collaboration and facilitate interoperability of analysis already at early stages of development and implementation, ETICS will collect, organize and publish middleware and application configuration information as well as common quality guidelines and principles.

The ETICS project is a spin-off from EGEE.

That project's Grid middleware, gLite, is based on middleware from many projects worldwide and inherits code from various sources and languages (*CERN Courier* September 2005 p21). EGEE's software quality assurance has developed ways to monitor code development and adherence to standards. ETICS will make these available to developers of middleware and application software worldwide, as well as a dedicated testing infrastructure.

AUTHENTICATION

Grid researchers get one online identity

A new Grid federation is helping scientists around the world to access computers and information in more than 50 countries and regions simply, securely and easily. The

International Grid Trust Federation (IGTF) was established at the 15th Global Grid Forum meeting held in Boston in October 2005, and will bring together organizations representing Asia, the Americas and Europe – from Canada to China and from Portugal to Pakistan. Members provide systems allowing scientists to identify themselves to any Grid resource in the world with just a single online identity in the form of a digital certificate.

"Reliably identifying Grid users and their work is critical to the success of Grid computing across the world," explained Neil Geddes, director of e-science at the CCLRC Rutherford Appleton Laboratory. "The formation of the IGTF is a huge step forward in agreeing mechanisms whereby users of one Grid can seamlessly become users of all Grids, thereby unlocking the full potential of Grid computing worldwide." ▸

Les gros titres de l'actualité informatique

ETICS veille à la qualité de la Grille	13	Une réunion sur le potentiel des grilles	16
Les utilisateurs de la grille seront identifiables	13	Deux nouveaux intergiciels pour le programme NMI	16
Les développeurs de Python échangent leurs idées	14	Le W3C publie ses recommandations pour XML	17
SC 05 : une conférence sur l'informatique haute performance	15	Produits	17
Le bulletin CNL du CERN a 40 ans	15	Événements	17

IGTF is a federation of three “policy management authorities”, one covering the EU and beyond (www.eugridpma.org), one for the Americas (www.tagpma.org) and one for Asia-Pacific (www.apgridpma.org). Individual members are national certification authorities

which issue digital certificates to scientists to enable them to use the Grid, international Grid collaborations which rely on the authorities for authenticating their scientists, and major infrastructure providers which rely on certificates for protecting their resources.

These resources include more than 40 000 computer processors and several petabytes of storage – equivalent to millions of DVDs. The federation today has 61 members and covers 50 countries and regions.

● For more on authentication see p21.

CONFERENCE

Python developers swap ideas

The European Python community came together in 2005 for the annual EuroPython conference, held as in 2004 at the campus of the School of Architecture at the Chalmers University of Technology in Göteborg, Sweden. EuroPython brings together developers, designers and business people from the Python and Zope communities, where Zope is a Python-based open-source application server for building content-management systems, intranets, portals and custom applications. The conference attracts a good mix of people, from those who have just started to learn about Python to more experienced users. Also, the “benevolent dictator” and inventor of Python, Guido van Rossum, was again present at the conference and, as usual, always available for a chat.

The 100 talks presented during the three days were categorized into 12 “tracks” and covered a wide range of activities, including success stories of using Python in companies as well as educational talks and tutorials. Frameworks written in Python that help to ease the life both of developers and users were also presented to the attendees, who numbered around 270.

A common theme in the tracks was the methodology used for the development of projects. The clear favourite here was the “agile process”, used in virtually all large projects. The main feature of the agile process is that it is a lightweight yet disciplined process aimed at making software development more predictable and efficient. Some, like the Zope and PyPy projects, have even established special teams to analyse and improve the development process. Agile methods emphasize real-time communication, preferably face-to-face, over written documents, and they attempt to minimize risk by developing software in short slices of time, called iterations, which



EuroPython delegates at the pub event.



Guido van Rossum (far right), “benevolent dictator” and inventor of Python, was at the conference and chatted with delegates.

typically last one to four weeks.

One such suggestion is “sprint-driven” development, which makes the distributed development effort more efficient, significantly increasing productivity. A “sprint” is a two- or three-day focused development session, where a maximum of 10 developers get together in a room and focus on building a particular subsystem. Sprints are lively bursts of development activity and have proven to be very effective and successful. A panel discussion on methodologies wrapped up this theme, addressing how tools and artefacts can be used in various types of “sprints” (coding, design, requirements, tutorial sprints

and so on) to improve the overall product.

In one user example a team of only three people from the R&D department of a pharmaceutical company developed a flexible tool (with around 35 000 lines of Python code) in less than three years. This is now deployed and used by the 1000 targeted users. The success of the project was attributed to good communication, frequent and intensive testing, an excellent architecture and the use of Python as the language of choice for the implementation. The importance of the close interaction between the scientists who use the tool and the software developers who wrote it was emphasized, especially the fact that the developers often had to follow frequently changing requirements.

A definite highlight in the tutorial track for newcomers was the “batteries included” tour of the standard modules library that comes with Python. Also very well received was the tutorial on “Acceptance Testing Using TextTest”, a tool that can help developers in various ways. TextTest can, for example, measure CPU time and other system information, and notifies the developer if one of the values is outside a tolerance margin set by the user. In addition, this tool can be used to test a graphical user interface (GUI) by recording which methods and functions the GUI software layer calls then “playing them back” in a later testing session and comparing the results. These tests can be run in batch mode (ideal for test runs during the night) or in parallel mode using various Grid engines (SunGridEngine, LSF, etc).

As every year, one of the two keynote talks was given by van Rossum, from the Python Software Foundation, who presented new features that will be included in future Python releases, as well as giving reasons for rejecting other requests. The second keynote

talk was by Steven Pemberton of CWI Amsterdam, who is best known for his efforts in the World Wide Web Consortium (W3C), and is the chair of the working groups responsible for the XHTML and XForms standards. He was once involved in the development of the ABC language, the forerunner of Python, where he worked together with van Rossum. Pemberton gave an inspiring talk entitled "Hypothesis: Programmers are Humans Too!", arguing that computing languages – the interfaces that the programmers see – have not evolved significantly over the past 30 years. He concluded that it is better (and easier) to change programs than programmers.

The conference concluded as usual with the "lightning talks" session of short talks – a time limit of five minutes was strictly controlled, but in a humorous way. These were in general very illuminating and ranged

from specific technical issues to a very lively demonstration on how to give a presentation as well as a call to collect sonnets from the Python community. Other "lightnings" covered new developments of tools and modules for a personal financial management system, for managing WiFi cards under Linux, for an implementation of the XPath 1.0 subset for ElementTree in pure Python, and a library for morphing splines.

Some of the short demos showed how to use Python on a Nokia phone to communicate with a GPS personal digital assistant (PDA) via bluetooth. A short demo from van Rossum showed what you can do with the new generators, which are Python's mechanism for lazy evaluation of a function that would otherwise return a space-prohibitive or computationally intensive list.

● The next conference will take place at CERN on 3–7 July 2006.

CNL ARCHIVE

CNL is 40 years old

After CERN's 50th anniversary in 2004, this year there is another (admittedly smaller) anniversary: the first *CERN Computer Newsletter* (CNL) was circulated in 1966. As CNL celebrates its 40th year, we will take a look back at some of the highlights – and historical curiosities – of the early years of computing at CERN, seen through its pages. You will find a brief look at previous issues of CNL here; for an extended retrospection read the January–March 2006 issue of CNL, available online at www.cerncourier.com.

Why CNL?

"As computing becomes a more and more widespread and complex activity in the laboratory, the need will increase for a means to have a wider general circulation of background information about different aspects of computing activities than is possible with the present system of Computer Notices." (February 1966, from the introduction to the first issue by G R Macleod, the leader of the Data Handling division from 1964 to 1975)

CERN's first computer expires

"The Ferranti Mercury, installed at CERN in 1957, has now been dismantled [...]. The



The Ferranti Mercury in building 2 in 1958.

computer was the first to be installed at CERN and remained as the only machine until the arrival of the IBM 709 in 1961." (April 1966)

The end of Fortran?

"FORTRAN: FORmula TRANslating. The most commonly used program language exists in many versions e.g. CERN Fortran. ALGOL: from Algorithm. A more advanced language now available on many computers but rarely efficient for production work. PL1: The new programming language specified by the IBM users' organization (SHARE) and IBM to supplant Fortran." (April 1966)

SUPERCOMPUTING

SC|05 showcases high-energy computing

As Grid computing projects gear up to serve data from the Large Hadron Collider (LHC) to thousands of physicists worldwide, scientists and engineers developing the LHC Computing Grid (LCG), the Enabling Grids for E-science (EGEE) and the Open Science Grid travelled to SC|05, the 2005 International Conference for High Performance Computing, Networking, Storage and Analysis held in Seattle on 12–18 November. Members of the Grid collaborations presented talks and posters at the Grid 2005 Workshop and the SC|05 technical programme, and particle-physics applications demonstrating interoperability across multiple Grid infrastructures featured in several SC|05 exhibits.

For the third consecutive year, the high-energy physics team captured first prize in the SC|05 Bandwidth Challenge. Led by the California Institute of Technology, Fermilab, SLAC and the University of Michigan, with participation from CERN and 10 other institutions, the team of physicists, computer scientists and network engineers transferred physics data at a rate of over 150 Gbit/s. The entry provided a showcase for the worldwide Grid system being built for the LHC experiments, using production storage systems and file servers to transfer 475 TB of high-energy physics data in 24 h.



David Foster, left, accepts the HPC Public Awareness Award on behalf of CERN.

CERN's commitment to educating the public about high-performance computing was recognized during SC|05 by HPCWire, the online magazine of high-performance computing. The laboratory was awarded the HPC Public Awareness Editor's Choice Award.

"This is a significant honour for CERN," said David Foster, head of CERN's Network and Communications Group. "All our institutional and industrial partners in LCG, EGEE and CERN openlab deserve to share in the credit. The Grid technology that is being deployed for the LHC is inevitably something that spans many institutions, all of whom are contributing to the broader public awareness concerning this new approach to high-performance computing."

GRID EVENTS

Meeting highlights Grid potential

At the end of September 2005, the UNESCO Regional Bureau for Science in Europe (ROSTE) and CERN organized a joint event on Grid computing, which was sponsored by Hewlett-Packard (HP). Current and future Grid participants, universities, and research institutions – with a special focus on south-eastern European countries – were invited to CERN to learn about the capabilities and opportunities of the Grid. The event included talks by representatives of existing Grids such as the Enabling Grids for E-science project, but best practices were also presented by regional research institutions and universities on establishing new – and connecting to existing – e-infrastructures.

“Grid computing is as important a shift as the Web has been,” said Howard Moore, director of UNESCO-ROSTE. “It will change the face of how science resources are shared and management of results will be handled.” The more tangible benefits of Grid computing – access to computing resources and cutting-edge technologies outside the local region – help universities and small research groups to attract students, thus maintaining talent in the area and alleviating the brain drain. Bringing people together to talk about Grid computing can increase researchers’ motivation, by giving them the feeling of being part of a bigger whole.

Over the past decade, countries in south-



Current and future Grid participants attended the meeting held by CERN and UNESCO.

east Europe suffered emigration of up to 70% of skilled professionals. In 2003, UNESCO launched the Alleviating Brain Drain project with HP to provide Grid computing technology to universities and financial support to encourage young scientists to remain in the region and cooperate with their fellow-nationals living abroad.

The event at CERN was an important step, but the road does not end there, as Bernard

Meric from HP pointed out: “Grids will move from scientific research to the citizen. It is just a question of time. A meeting like this can accelerate this movement from scientific research to beyond.”

The event led to the decision to hold a training event in Split, Croatia, later this year. This training, to be organized by CERN and SEE-GRID, will act as a pilot event for a CERN School of Computing in Croatia in 2007.

GRID MIDDLEWARE

Middleware has two important additions

Working with research communities to provide development and access management tools for Grid and other research environments, the eighth release of the National Science Foundation Middleware Initiative (NMI-R8) helps to facilitate the complex resource management and security required in a shared cyber-infrastructure.

NMI-R8 marks two important “firsts” for the NMI programme: the addition and integration of Ninfg-G, the first non-US developed

component included in the GRIDS Center software suite; and GridShib, the first software enabling interoperability between the Globus Toolkit and Shibboleth federating software.

Ninfg-G is a GridRPC reference implementation developed at the National Institute of Advanced Industrial Science and Technology in Japan. “With the integration of Ninfg-G into the GRIDS Center stack, we’re learning about the practical issues of sharing and supporting software across borders and languages, and see this activity as an important pilot study,” said Philip Papadopoulos, program director, Grid and Cluster Computing at San Diego Supercomputer Center.

GridShib enables interoperability between

the Globus Toolkit and Shibboleth federating software. Researchers who are members of both brick-and-mortar institutions and Globus-enabled virtual organizations can use their local campus credentials to access their distributed Grid-based resources.

NMI-R8 comprises several updated tools, software packages, practice documents, and schema that support institutional and federated identity management environments. The stack is flexible because the components are bundled together as one easy-to-deploy package. Once installed, only the parts of interest need to be configured. NMI-R8 is available to the public for downloading under open-source licences at www.nsf-middleware.org.

W3C RECOMMENDATION

W3C improves features for transforming and querying XML

The World Wide Web Consortium (W3C) has published eight documents in the XML family as Candidate Recommendations, sending a signal to the developer community that powerful new features for transforming and querying XML are ready for implementation.

XSLT 2.0 is a major revision to the very successful XSL Transformations language, which transforms XML content into other formats. XSLT 2.0 standardizes many features and includes stronger support for internationalization and richer tools for the programmer. In addition to new functionality,

XSLT 2.0 introduces strong typing and supports the optional use of W3C XML schema. Strong typing is a feature of enterprise-strength programming languages such as Java, C++ and C#, and is designed to reduce errors in programs, slashing the cost of developing and maintaining large systems.

XML Query brings the power of database search and select to XML. With XML Query, one can run cross-vendor cross-database joins between multiple forms of data, including XML documents, XML-native stores, relational database tables and more.

Connections between large applications, databases, operating systems, Web services and Web servers have traditionally used special software that, on demand, converts and manipulates data between the formats used by applications. With a standard way to integrate tools, a standard set of data formats and standard ways to query and manipulate those formats, users will be able to focus on their higher-level business logic and integrate new sources of data much more quickly.

● For a list of all recommendations see www.w3.org/TR/.

PRODUCT INFORMATION

The latest version of **COMSOL Multiphysics**, a software package for scientific simulations and formerly called FEMLAB, has been made public. The name change reflects the growing functionality of the product, which started off as software for finite-element modeling (FEM). The new software can include the whole spectrum of basic CAD formats, and has a new independent script language, COMSOL

Script, which allows simulations via the command line. The graphical user interface supports physical units and the software provides new algorithms for the simulation of moving particles and geometrical parameterization. For documentation see www.femlab.ch and www.comsol.com.

Stylus Studio, a provider of XML development tools and components for advanced data integration, has announced the availability of

Stylus Studio 2006 XML Enterprise Edition and **Stylus Studio 2006 XML Deployment Adapters**. The new release of XML IDE provides a significant upgrade in usability and XML standards compliance, and many new features further simplify XML data integration. The adapters provide high-performance Java components that offer bidirectional, programmatic access to virtually any data source as XML. For more information see www.stylusstudio.com.

Calendar of events

February

13-17 15th International Conference on Computing in High Energy and Nuclear Physics (CHEP06) Mumbai, India, www.tifr.res.in/chep06/

13-16 GGF16 Athens, Greece, www.ggf.org/GGF16/ggf_events_ggf16.htm

March

1-3 TridentCom 2006 Barcelona, Spain, www.tridentcom.org

1-3 EGEE User Forum Geneva, Switzerland,

<http://egee-intranet.web.cern.ch/egee-intranet/User-Forum/index.html>

April

2-6 High Performance Computing Symposium (HPC 2006) Huntsville, Alabama, www.caip.rutgers.edu/hpc2006/

3-7 HEPiX meeting CASPUR in Rome, Italy, www.hepix.org/

25-29 20th IEEE International Parallel & Distributed Processing Symposium Rhodes Island, Greece, www.ipdps.org

June

27-30 ISC2006 Dresden, Germany, www.supercomp.de/, paper deadline 20 February

August

29-1 September Euro-Par 2006 Dresden, Germany, www.europar2006.de/, paper deadline 31 January

September

13-15 HPCC-06 Munich, Germany, <http://hpcc06.lrr.in.tum.de/>, paper deadline 13 March

www.heinzinger.de

HIGHCURRENT HIGHVOLTAGE HIGHPOWER
POWER SUPPLIES & CAPACITOR CHARGERS

HIGHPRECISION
up to <0.001%

DC Power Supplies & more

Heinzinger®
power supplies
supplies your world

Heinzinger electronic GmbH
The Power Supply Company
Anton-Jakob-Straße 4
83026 Rosenheim - Germany
Tel. +49-8031-2458-0
Fax: +49-8031-2458-58
info@heinzinger.de
www.heinzinger.de

Intelligent networking – a vital part of today's IT infrastructure

By Wenceslao Lada
Vice President and General
Manager, ProCurve Networking,
Europe, Middle East & Africa



Whether you consider a computer in its traditional sense to be a means of calculating large numbers quickly or your understanding is closer to the French word for a computer – ordinateur – which implies a tool for organising and managing work, it is clear in this day and age, such a device has a limited use unless connected to other computers via a reliable and intelligent network.

This is as true for a nuclear research establishment as it is for a commercial office although the amount of computing power available to each may be vastly different. Each organisation needs its employees to be able to communicate quickly and easily with each other and the outside world; it needs to be able to add and remove users from the network with the minimum of fuss and effort, and it needs to be able to guarantee its security and integrity.

As a key supplier of networking equipment to CERN, ProCurve Networking by HP has had to ensure that the networks we supply are able to deliver on all the requirements of a world-class organisation. That we are able to do so owes a lot to our own commitment to research and development and careful attention to the most important part of a modern network – the ability to manage it so that control is firmly in the hands of management but the ability of the individual users to operate it is not impaired.

This conundrum is not as easy to solve as it may appear. Traditionally, network security could be realised quite simply by concentrating all decision making in the hands of those controlling the central core of the network and imposing great restrictions on access, bandwidth and flexibility to users throughout the organisation.

An alternative, to have loose central control and devolve responsibility for access control, security and bandwidth control to departments at the edge of the network gave great freedom to users but left the way open for all manner of abuse – ranging from unauthorised access to company data to inefficient use of network infrastructure.

Adaptive EDGE Architecture™

ProCurve Networking by HP has met this challenge with a strategy we call the Adaptive EDGE Architecture™ or AEA. This

underpins our whole product design and the intention is to allow control of policies regarding security and resource allocation to be decided upon at the centre of the network while allowing implementation to be pushed out to PC users at the edge. To do this we have had to build in significant intelligence to all our products ranging from the large complex switches and routers which sit in the centre of the network to intelligent switches which are to be found at the edge.

All ProCurve's networking products and solutions are based on proven open industry standards such as Ethernet switching, 802.11 wireless networking and other open standards for delivering power to network devices via the network connection itself, the so-called Power over Ethernet standard. Adhering to open industry standards means we are able to offer our customers best in class value as well as maximum choice and flexibility.

Furthermore, the majority of ProCurve Networking products and solutions are based on our own custom designed ASICs (application specific integrated circuits) which capture our design expertise in mass-produced silicon chips. By cutting down the number of components needed, these ASICs allow us to offer life-time guarantees on many of our products and to


offer free software upgrades to enable our installed base of equipment to take advantage of new networking protocols as they become available.

Today's networks are no longer just about connecting computers in a server room to computers on a desktop. As employees become more mobile, they want to access their corporate networks from notebook computers, PDAs or even the latest generation of mobile phone. They will be calling in from wherever they happen to be, which may be from a hotel room, an airport or an Internet café. The modern network has to be able to handle these requests from a variety of media and to respond without sacrificing ease of use or security.

As such, the need to devise policies centrally is as important as ever, so that standard operating procedures can be developed and maintained, but the necessity to devolve the implementation of such policies to the edge of a network is also vital. Otherwise, traffic to the centre would become so congested as to be unworkable and the likelihood of employees attempting to circumvent such policies would increase. Far better for employees to have a workable system that they can trust than to have a well-meaning but autocratic and totally unworkable system that they will lose faith in.

ProCurve Networking by HP has over 25 years experience in the networking industry and has been responsible for many innovations. Its customers include some of the world's leading companies such as SAP, Hilton Hotels and Kodak. Its products are also used in leading universities and research establishments throughout Europe including Sahlgrenska in Sweden, the British Library and the University of York in the UK.

Further information on ProCurve Networking solutions and products is available at:
www.hp.com/eur/procurve



Tackling the challenge of lattice QCD

Don Holmgren describes Pion, Fermilab's new commodity solution for lattice quantum chromodynamics computing, which went online at the end of last year.

In December 2005, Fermilab brought a new cluster online devoted to lattice quantum chromodynamics (lattice QCD) calculations. "Pion" consists of 520 single-processor Intel Pentium computers connected with Infiniband, and is the latest lattice QCD cluster funded by the US Department of Energy (DOE).

Lattice QCD is the numerical technique used to study QCD, the theory that describes the strong force. Because lattice QCD calculations require such enormous computing power, simplifying assumptions (such as the quenched approximation) have been required in the past to make progress on the supercomputers available (*CERN Courier* June 2004 p23). In recent years, improvements in algorithms and a steady increase in the capabilities of computers have led to more complete lattice QCD simulations. This has enabled lattice theorists to make a number of predictions of physical quantities that were matched by new experimental results with equal precision (Aubin *et al.* 2005a and 2005b, Allison *et al.* 2005, Kronfeld *et al.* 2005 and *CERN Courier* July/August 2005 p13).

In lattice QCD computations, a finite volume of the space-time continuum is represented within a supercomputer by a four-dimensional lattice of sites. At each site, one or more SU(3) vectors, i.e. 1×3 arrays of complex numbers, represent the field strengths of quarks on the lattice. SU(3) matrices, i.e. 3×3 complex arrays, reside on the links between sites and represent the gluon fields that interact with the quark fields. Performing a Monte Carlo calculation to generate representative configurations of the QCD vacuum involves repeated sweeps through all sites, multiplying the link matrices by neighbouring site vectors and updating their values.

In current lattice QCD problems, the representation of the lattice requires around 10 GB of memory and calculation rates of around 1 Tflops (10^{12} floating-point operations per second) to make reasonable progress. A typical desktop PC is capable of about 1 Gflops.

In recent years, two approaches have been taken in the design of dedicated supercomputers for lattice QCD. The first uses custom




The Pion cluster at Fermilab consists of 520 Intel Pentium PCs.

processors that include on-chip connections to the neighbouring processors in three or more dimensions. Each processor chip also includes several megabytes or more of on-chip memory. The QCDOC (*CERN Courier* September 2005 p17) and apeNEXT (*CERN Courier* September 2005 p18) machines are examples of this system-on-a-chip approach. The second approach, used on Pion, relies on commodity computers connected via a high-performance network. Although similar to the large Linux farms used for reconstruction processing for experiments at Fermilab and CERN, it is the high-performance network that distinguishes the computer-cluster approach. Also, while the computers in reconstruction farms operate independently, those on a lattice QCD cluster are tightly coupled, constantly exchanging data during the computation.

Purpose-built systems like QCDOC and apeNEXT require several years of design, development and fabrication before they are ready for production. When new, these supercomputers provide the largest capability, with the fastest calculation for a single problem. By contrast, commodity clusters take advantage of each year's newest processors, memory technologies and networks; lattice QCD supercomputers based on clusters can be built relatively quickly and frequently at a lower cost. In terms of performance, the purpose-built US QCDOC machine sustains 5 Tflops, compared with about 1 Tflops for Fermilab's Pion cluster. By the end of next summer, the successor to Pion at Fermilab will sustain about 2.25 Tflops.

The design of lattice QCD clusters requires careful attention to the balance between calculation speed, memory bandwidth and network performance to achieve the most cost-effective designs. Lattice QCD codes require high memory bandwidth and strong floating-point performance. A double-precision SU(3) matrix-vector multiplication, for example, consumes 192 B of operands, produces 48 B of results and uses 66 floating-point operations. This bytes-to-flops ratio of more than 3.6 stresses memory subsystems. Typical Intel and AMD microprocessors, for example, are capable of more than 10 Gflops >




The lattice QCD prediction of the mass of the B_c meson is now approaching the precision of the value measured by the CDF collaboration at Fermilab (Allison et al. 2005).

when their parallel floating-point instructions (Streaming SIMD Extensions, or SSEs) are used. However, the typical peak memory bandwidths of these processors of about 6 GB/s fall far short of the 40 GB/s that would be required to feed and consume the data for SU(3) matrix-vector multiplications at the peak speed of the SSE floating-point unit. The calculations also stress the networks used to communicate data between the nodes of a cluster, requiring bandwidths and latencies superior to those provided by TCP/IP running over Gigabit Ethernet.

The networking requirements of lattice QCD on clusters have been met by using high-performance switched networks, such as Myrinet, Quadrics or recently Infiniband, or by using multiple Gigabit Ethernet networks running a specialized non-TCP/IP communications stack. The latter approach was used on two clusters built at the DOE's Jefferson Lab. The 2003 machine, "3G", uses six Gigabit Ethernet interfaces connected as a 3D toroidal mesh in which each computer can communicate directly only with neighbours in the positive and negative x, y and z directions. The 2004 machine, "4G", uses a 5D toroidal mesh. The toroidal Gigabit Ethernet mesh machines, pioneered by Zoltan Fodor of Eötvös University, Budapest, and colleagues were very cost-effective because of their use of commodity network interfaces (Z Fodor et al. 2003). However, in the past year the newer Infiniband network technology has become more favourable, offering greater performance and flexibility at a low cost. These factors led Fermilab to choose Infiniband for the Pion cluster. One advantage the switched Infiniband network provides for Pion is that any computer in the cluster may communicate directly with any other computer on the network.

Pion uses Intel Pentium 640 processors, each running at 3.2 GHz. With 800 MHz memory buses, these machines have one of the highest memory bandwidths per processor available. Pion was built in two halves, with the first 260 nodes last May and the second 260 in November, and the price per node dropped from \$1970 in the spring to \$1550 in autumn, with approximately \$840 for the computer and \$650 for the Infiniband. Pion exceeds 1.7 Gflops per processor on lattice QCD codes, or roughly \$0.90/Mflops on the second half.



The Jefferson Laboratory's 3G cluster dates from 2003 and uses six Gigabit Ethernet interfaces connected in a 3D mesh.

On 1 October 2005 the DOE launched a four-year project to support lattice QCD. This follows five years of support for lattice QCD by the DOE Office of Science via the SciDAC Lattice Gauge Theory Computational Infrastructure project, as well as grants used to purchase the US QCDOC in 2004 and 2005. The project will operate the 5 Tflops US QCDOC machine, which resides at Brookhaven National Laboratory, as well as the clusters at Jefferson Lab (3G and 4G, which total 0.55 Tflops) and Fermilab ("QCD", a 0.15 Tflops, 128-node Myrinet cluster built in 2004, and the 1 Tflops Pion).

The project will also support the construction of new machines. In the first year, two clusters will be constructed: a 0.5 Tflops Infiniband cluster based on 256 dual-core Intel processors at Jefferson Lab will come online at the end of March, and a 2.25 Tflops Infiniband cluster based on 1000 dual-core Intel processors at Fermilab will come online by the end of September. Other systems planned for 2007–2009 will provide an additional 11 Tflops of computing power.

Further reading

C Aubin et al. 2005 *Phys. Rev. Lett.* **94** 011601.
C Aubin et al. 2005 *Phys. Rev. Lett.* **95** 122002.
I F Allison et al. 2005 *Phys. Rev. Lett.* **94** 172001.
A S Kronfeld et al. 2005 *Proc. Sci.* **LAT2005** 206.
Z Fodor et al. 2003 *Comput. Phys. Commun.* **152** 121.

Résumé


Des grappes d'ordinateurs pour la CDQ sur réseau

En décembre 2005, le Laboratoire Fermi a mis en service la nouvelle grappe d'ordinateurs "Pion" pour les calculs de chromodynamique quantique (CDQ) sur réseau. Cette grappe, qui comprend 520 ordinateurs Intel Pentium à processeur unique interconnectés par un réseau haute performance, Infiniband, tire parti des processeurs, technologies de mémoire et réseaux les plus récents. Les superordinateurs de CDQ sur réseau s'appuyant sur cette technologie peuvent être produits relativement rapidement et à moindre coût que les systèmes spécialisés, dont la conception, le développement et la fabrication s'étendent sur des années.

Don Holmgren, Fermilab.

Towards a read–write Web

Security efforts led by the LCG and EGEE projects are uniting the Web and the Grid, and helping to fulfil the vision of a worldwide collaborative Web, says **Andrew McNab**.



Thousands of Internet users edit Wikipedia pages every month.


When Tim Berners-Lee invented the World Wide Web 15 years ago at CERN, he always intended that it should be easy for people to write to it, not just read from it. But if websites are opened up to additions from everyone, they often get vandalized or “spammed”. As well as direct disruption, this can lead to bad publicity, as shown by recent mainstream press coverage of a hoax story about US politician John Seigenthaler in the collaborative encyclopaedia Wikipedia. These problems have put security centre-stage in the development of a true read–write Web. Fortunately, solutions are emerging from large high-energy physics Grid projects.

“Wikis” like Wikipedia are websites where users can add, edit and cross-reference content without learning mark-up languages like HTML. Along with diary-like weblogs and the reader comments they often allow, wikis are a major part of the current trend away from a pure publishing model (“the Web as an electronic book”) to a more collaborative model in which users also produce content.

High-profile wikis have been plagued with problems of trust and identity. Although some wikis require a username and password, people running the sites still have no idea who their contributors really are, and the better known wikis have to be constantly on the look-out for offensive or simply irrelevant additions.

Stronger methods of authenticating users are needed, so that they can be held accountable for the additions they make by their peers, the website owners or the authorities, and so that sensitive sections can have restricted access. However, if each website continued to maintain a database of usernames and passwords, then each site would have to bear greater administrative costs compared with giving accounts freely to everyone.

Furthermore, the existing username and password systems also put an increasing burden on the users, who have to maintain more and more accounts. When account details change, such as the user's preferred e-mail address for password reminders, the user is



GridPP uses GridSite to authenticate contributors to websites.

forced to spend time updating this information on all the sites they can remember – or risk losing access when their old e-mail address is no longer active. Users are therefore tempted to use the same passwords at multiple sites, and to store passwords in their Web browser, where they can readily be stolen by computer viruses and Trojans. What is really needed are ways of reusing the same authentication procedure for multiple websites.

Large Grid projects, such as the LHC Computing Grid (LCG) led by CERN, have faced a similar problem: how to give thousands of users access to hundreds of computing facilities while ensuring proper control and logging, without users having to arrange access at each site one by one.


The LCG, the Enabling Grids for E-sciencE (EGEE) and the Open Science Grid (OSG) projects, which have been prominent in high-energy physics, have all settled on using X.509 digital certificates for user and service authentication. X.509 specifies how to store and distribute public and private cryptographic key pairs (which take the form of large numbers), and how to associate them with unique names for users and services. X.509 was originally a standard of the International Telecommunication Union, although a third revision has now been defined by the Internet Engineering Task Force.

Using X.509, users can establish their identity by proving they have the private key associated with it. Since X.509 is based on asymmetric cryptography, it is possible to prove possession by revealing only the public key. In turn, the public key is “signed” by a trusted third-party organization or “Certification Authority”, and this digital signature states that it is indeed the public key of the user in question. Grid services accepting X.509-based authentication need to establish trust for only a small number of certification authorities, even though they serve thousands of users.

To simplify the process whereby users and service-owners decide to trust a certification authority, consortia of certification authori- ▷


integrated

ELECTROMAGNETIC SIMULATION SOFTWARE



Photomultiplier tube modeled in LORENTZ 6.2, a program for analyzing particle trajectories, electron and ion beam optics. LORENTZ's calculations include full secondary emissions, emittance, spot size and radius. Image courtesy of ADIT.

ES' easy-to-use field solver software gives you fast and accurate answers.



3D magnetic position sensor modeled in AMPERES 6.1, a program for analyzing magnetic components incorporating linear, non-linear and permanent magnet materials. AMPERES' calculations include force, torque, flux linkage and inductance.

Choose from a variety of programs to help you design products better and faster than your competitors.

**REQUEST YOUR
FULL-VERSION
SOFTWARE EVALUATION
PACKAGE TODAY!**

Call 204.632.5636,
email emsimulation@integratedsoft.com, or
visit www.evaluation.integratedsoft.com.



INTEGRATED
ENGINEERING SOFTWARE

WEB SECURITY

ties have signed agreements and established criteria that users must meet to receive a signed public key. During the EU DataGrid project (the predecessor of EGEE), a group of European certification authorities formed the EU Grid Policy Management Authority (EUGridPMA) to coordinate their efforts, and published details of approved certification authorities. This year the International Grid Trust Federation has been formed by the EUGridPMA and equivalent PMAs in Asia and the Americas. Preparation for LHC Grid computing has been a major driving force in this area, and is having wider benefits for worldwide Grid interoperability.

The Grid software using X.509 keys to authenticate users is largely adapted from World Wide Web e-commerce systems, implementing HTTPS (HTTP transported over an encrypted channel). X.509 signed keys given to users of the Grid can thus also be used with standard Web servers and Web browsers.

The GridPP project in the UK decided to base authentication for its collaboration website on X.509 Grid credentials, to exploit this crossover. GridPP members at the University of Manchester developed the GridSite security system, a software toolkit and set of extensions to the Apache Web server. As well as support for X.509, GridSite accepts "Distinguished Name" and "Virtual Organization" credentials. Together, these components allowed GridPP members to gain write access to their group's section of the GridPP website, to maintain Web pages or upload documents. However, users have had to write in HTML, providing only the security side of the convenient read-write Web that is beginning to emerge.

To bridge this gap, GridPP has merged GridSite and Wiki technology to give the convenience of a wiki coupled with the security of X.509 credentials. Since the hard work of establishing a user's identity is done once by the certification authority or local contacts, users can install their X.509 signed key in the Web browser then move between GridSite- or Wiki-enabled websites without having to present (and maintain) their username and password again and again.

Although this does not solve the problem faced by public wikis – where users in repressive states or expressing unwelcome professional opinions may have legitimate privacy concerns – it does work well in an academic environment where an article's provenance and author names are already the accepted currency.

● For more information see www.gridsite.org.

Résumé

Vers un Web "lu-écrit"

En inventant le World Wide Web il y a 15 ans, Tim Berners-Lee avait non seulement l'intention de permettre aux utilisateurs de lire les pages des sites, mais également de pouvoir y écrire. Mais si les sites Web sont ouverts aux contributions de tous, ils sont souvent vandalisés ou pollués. La sécurité est par conséquent devenue une préoccupation centrale pour le développement d'un Web réellement "lu-écrit". Des solutions commencent à émerger des projets à grande échelle de Grille de calcul pour la physique des hautes énergies. Ces projets font en effet face à des difficultés similaires en donnant à des milliers d'utilisateurs l'accès à des centaines d'installations informatiques.

Andrew McNab, University of Manchester.

Globus Toolkit upgrade aids the Grid community

Ian Foster reports on the latest release of the Globus open-source software toolkit.


Globus software has long played a significant role in the LHC Computing Grid (LCG) and related Grid deployments and applications. In 2005 the Globus community achieved the following significant milestones: the release of Globus Toolkit version 4; the establishment of a new community development process; major new funding from the US National Science Foundation and the UK Engineering and Physical Science Research Council; and the establishment of companies focused on supporting commercial use of the software.

Globus software has been developed since the late 1990s to support service-oriented distributed-computing applications and infrastructures. Core components of the Globus Toolkit (GT) address, within a common framework, basic issues relating to security, resource access, resource management, data movement, resource discovery and so forth. These components enable a broader “Globus ecosystem” of tools and components that build on, or operate with, core GT functionality to provide a wide range of useful application-level functions. These tools have in turn been used to develop a wide range of Grid infrastructures and applications.

A major upgrade

GT4, released on 30 April 2005, provides significant improvements in terms of robustness, performance, usability, documentation, standards compliance and functionality (I Foster 2005). For example, the GRAM job submission service now easily manages 32 000 jobs, up from a few hundred in GT3. A UK e-Science evaluation speaks favourably of the quality, usability and performance of the GT4 code and its documentation (T Harmer *et al.* 2005). It notes, for example, that “GT4 installation was straightforward”, “GT4 services demonstrated significant improvements in performance and reliability over their GT3 versions” and “GT4 package descriptions were of a high quality, well structured, and accurate”.

GT4 makes extensive use of Web Services mechanisms to define its



GT4 architecture schematic, showing many but not all of the components. The violet boxes denote GT4 code, while the white boxes are user code.

interfaces and structure its components. Web Services provide flexible, extensible and widely adopted XML-based mechanisms used for describing, discovering and invoking network services. Furthermore, its document-oriented protocols very much suit the loosely coupled interactions that many argue are preferable for robust distributed systems. These help the development of service-oriented architectures: systems and applications structured as communicating services, in which service interfaces

are described, operations are invoked, access is secured and so on – all happening in uniform ways.

The figure illustrates the principal GT4 components. A set of service implementations (the bottom half of the figure) implement useful infrastructure services. These address such concerns as execution management (GRAM), data access and movement (GridFTP for data movement, reliable file transfer, OGSA-DAI for database access), replica management (replica location service, data replication service), monitoring and discovery (Index, Trigger, WebMDS), credential management (MyProxy, Delegation, SimpleCA) and instrument management (grid telecontrol protocol). Most are Java Web Services but some (indicated at the bottom right) are implemented in other languages and use other protocols. A powerful authorization framework enables the implementation of diverse authorization policies.

In addition, three containers can be used to host user-developed services in Java, Python and C. These containers provide implementations of security, management, discovery, state management and other mechanisms frequently required when building services. They extend open-source service-hosting environments with support for a range of useful Web Service specifications, including WS Resource Framework (WSRF), WS-Notification and WS-Security.

Finally, client libraries allow client programs in Java, C and Python to invoke operations on both GT4 and user-developed services (top of the figure). In many cases, multiple interfaces provide different >

GRID SOFTWARE

levels of control: for example, in the case of GridFTP, there is not only a simple command-line client (globus-url-copy) but also control and data channel libraries for use in programs – and the XIO library allowing for the integration of alternative transports.

GT4's use of uniform abstractions and mechanisms across these different components means that clients can interact with different services in similar ways, which facilitates the construction of complex, interoperable systems, and encourages code reuse.

Globus is more than software: it is also a community of users and developers, and the infrastructure that supports this community – code repositories, e-mail lists, a problem-tracking system and so forth. To facilitate external contributions, the Globus team recently introduced a community governance and development process, GlobDev, hosted at <http://dev.globus.org>. A growing number of Grid software projects are hosted and developed there.

Further reading

I Foster 2005 *IFIP International Conference on Network and Parallel Computing LNCS 3779* 2.

T Harmer et al. 2005 *UK e-Science Technical Report Series UkeS-2005-03*.

Acknowledgments

This report describes the work of many talented colleagues and collaborators (see www.globus.org). Work on Globus has been supported in part by the Mathematical, Information, and Com-

putational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, US Department of Energy, by the National Science Foundation (NSF) under its NSF Middleware Initiative and other programmes, and by IBM, DARPA, NASA, Microsoft, the UK Engineering and Physical Sciences Research Council and Department of Trade and Industry, and the Swedish Research Council. The author is also co-founder and chief open-source strategist at Univa Corporation.

Résumé


La boîte à outils Globus s'améliore

Globus Toolkit est développé depuis la fin des années 90 pour fournir des infrastructures et des applications de calcul décentralisées, axées sur le service. Ce progiciel joue un rôle clé pour la mise en œuvre et les applications de la grille de calcul pour le LHC (LCG) et de grilles associées. Dans cet article, Ian Foster commente la version 4 de Globus Toolkit (GT4), le dernier né de ce progiciel de Globus en libre accès pour le calcul sur grille. Disponible depuis avril 2005, GT4 présente d'importantes améliorations, que ce soit en matière de robustesse, performance, convivialité, documentation, conformité aux normes ou fonctionnalités.

Ian Foster, Math and Computer Science Division, Argonne National Laboratory.

X-Ray Detector

XR-100CR at **149 eV FWHM** Resolution



No Liquid Nitrogen
Solid State Design
Easy to Use
Low Cost

APPLICATIONS

- Nuclear Physics
- Synchrotron Radiation
- High Energy Physics
- Neutron Experiments
- Astrophysics
- Research & Teaching
- Nuclear Medicine
- X-Ray Fluorescence



XR100CR X-Ray Detector with PX4 Digital Pulse Processor, Power Supply, Shaping Amplifier & MCA



XR100CR fitted for vacuum applications

Visit Us Now
www.amptek.com



AMPTEK Inc. 14 Deangelo Drive, Bedford, MA 01730-2204 USA
Tel: +1 (781) 275-2242 Fax: +1 (781) 275-3470
E-mail: sales@amptek.com www.amptek.com

NEW!!! Complete X-Ray Spectrometer

INCLUDES

- 1 X-Ray Detector and Preamplifier
- 2 Digital Pulse Processor
- 3 Power Supply



2.7 x 3.9 x 1 in. (7 x 10 x 2.5 cm)

The X-123 is a complete X-Ray Detector System in one small box that fits in your hand.

Features of the X-123:

- Simple to Operate
- Low Power (1.2 Watts)
- Small Size
- USB and RS232 Communication
- Accommodates all Types of Amptek Detectors