

CERN COMPUTER NEWSLETTER

Volume 43, issue 1 January–March 2008

Contents

Editorial

- Computing models pass first test of readiness 1
ETICS software project keeps enforcing quality assurance 3
Les Robertson looks back at six eventful years at the helm of the LHC Computing Grid 4

Announcements and news

- Mail servers must register to accept e-mail from the Internet 6
New spam detection for all CERN mailboxes 6
School of Computing 2008 to be held in Norway 6
Changes to support structure for engineering tools 6
Collaboration just gets better as EVO takes over from VRVS 7
User Support services welcomes a new face 7

Grid news

- GridTalk: spreading the word 8
DILIGENT: from digital libraries to virtual research 8
The business of Grid computing: what can Grids do for you? 9
New Rome Tier2 centre prepares for CMS data 10

Technical brief

- A precious stream of data: from Point1 to analysis sites 11
New web-application tool boosts usability of JavaScript and CSS 12

Conference and event reports

- Genome centre hosts HEPiX event 13
EGI workshop in Rome: defining the European Grid infrastructure of the future 14
EGEE User Forum looks to build bridges in the science community 15

Information corner

- A new editor for the CNL 16
Computer security: think before you click! 16
Calendar 16

Editor Natalie Pocock, CERN IT Department, 1211 Geneva 23, Switzerland. E-mail cnl.editor@cern.ch. Fax +41 (22) 766 9204. Web cerncourier.com/articles/cnl.

Advisory board Wolfgang von Rueden (head of IT Department), François Grey (IT Communication team leader), Christine Sutton (*CERN Courier* editor), Tim Smith (group leader, User and Document Services).

Produced for CERN by IOP Publishing Dirac House, Temple Back, Bristol BS1 6BE, UK. Tel +44 (0)117 929 7481. E-mail jo.nicholas@iop.org. Fax +44 (0)117 930 0733. Web iop.org.

Published by CERN IT Department

©2008 CERN

The contents of this newsletter do not necessarily represent the views of CERN management.

IOP Publishing



Computing models pass first test of readiness

This article is a high-level overview of the Common Computing Readiness Challenge – known as CCRC'08. For those who have been involved, the most important message is undoubtedly “thanks for all the hard work – and save some energy for May!”

For several years, the four LHC experiments (ALICE, ATLAS, CMS and LHCb) have been running periodic stress tests of their planned computing operations when LHC data arrives. They have written design documents that describe their models for how they will acquire, distribute and process their data. They differ in the details but have a common base in a model of levels or “Tiers” of resources.

Master copies of the raw data are kept on tape at CERN, the Tier0 site. CERN also does a first reconstruction pass on the raw data using approximate calibration constants. The final calibrations are not ready for weeks or even months, but will be used for a second reconstruction that is done entirely at the next level – the Tier1 sites. (One experiment model – LHCb – differs here in that first pass and later reconstruction is done entirely at their Tier1, but CERN is considered part of that.) A complete copy of the raw data is exported to the Tier1 sites, where it is also put on tape. There are 11 such sites but they do not all support all the experiments.

A total of 10 Tier1 sites support ATLAS, the largest experiment in data rates and volumes: seven support CMS, seven support ALICE and six are behind LHCb (the smallest

experiment). The Tier1 resources are not evenly distributed, e.g. BNL (near New York) will provide 28% of ATLAS Tier1 resources, while CNAF in Bologna will provide only 5%.

Raw data is sent from the experiment data acquisition system to disks in the CERN computer centre. There it is migrated to tape under the Castor system, reconstructed by batch jobs at the Tier0, and sent (both the raw and reconstructed data) to the experiment Tier1 sites using the Grid file transfer service. At the Tier1 site it is copied to their local mass storage systems (the most common is called dCache) and when the calibration is ready the raw data is recalled from tape and reconstructed using Grid submitted batch jobs. The output data is then ready for further physics analysis.

There is then a third level – the Tier2 sites – which get a copy of the analysis files for their physicists. These comprise about 250 smaller computer centres, mostly at universities. They also generate simulated (Monte Carlo) events that are needed by the experiments and which are stored on tape at Tier1 sites and, usually, at CERN.

Over the years the experiments, as part of the Worldwide LHC Computing Grid project, and together with CERN-IT, have exercised the various components of their computing models. For example, last year over several days ATLAS succeeded in sending simulated raw data to the computer centre at the required 320 MB/s, processing it with several thousand




Fig. 1. Graph showing the rate of data transfer between CERN and all the Tier1 sites.

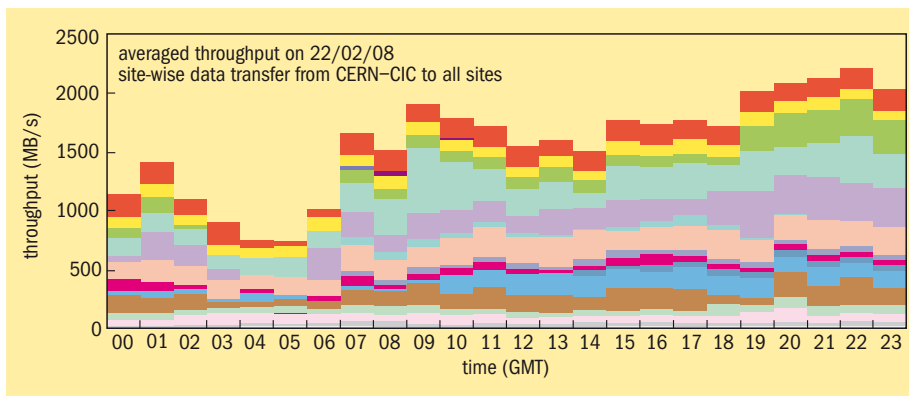


Fig. 2. The rate of data distribution from CERN to the Tier1 sites over a single day.

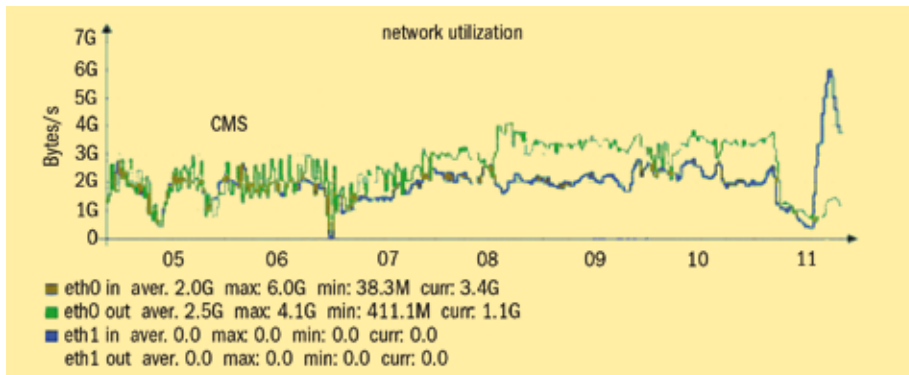


Fig. 3. Chart showing the data rates in and out of the Castor2 system at CERN.

concurrent batch jobs and exporting the raw and processed data to their 10 Tier1 sites at about 1000MB/s. Some of the Tier1 sites put the required data to tape, but not all did. When all experiments take data together, the aggregate rate out of CERN will be 1600MB/s and we will be storing about 40TB of data (about 100 full tape cartridges) into Castor each day. The current LHC plan is to run for 100 days in 2008, from July to December.

Stress testing

A weakness of all the tests done so far is that they have never overlapped all of the experiments together, and so have not reached the full rates expected at CERN and the Tier sites. Also, they have never had to cope with the lack of homogeneity of having the four different experiment data-management and batch-processing models being busy at the same time. CCRC'08 addresses these issues by exercising all the experiments' computing at the target rates and at the same time.

The plan is to run in two phases:

- The first is to test all the functionality involved and attempt to build up to the target rates as far as resources allow (the full 2008 resources at the remote sites will not be available until 1 April and some sites will be late). This was run from 4–28 February. As expected we found many problems in this period, including software, hardware and organizational problems, and are correcting them during March and April.

- In this way we will be ready for a second run during the whole of May, which is planned to reach the full 2008 data rates rapidly over the whole distributed processing and then run stably for many weeks. The May exercise is therefore a full-scale dress rehearsal for the accelerator run, which is due to start in July.

The accelerator should then be in action continuously for many months, with periodic stoppages of several days each. The current operational plans mean that most physics data from the experiments will arrive during a 12 hour period overnight, but the batch processing – both local and over the Grid – and data export will be continuous. The experiment and site support teams will play a vital role in ensuring these continuous operations.

A Twiki providing many details of CCRC'08 is available at <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGCommonComputingReadinessChallenges>, and an overview of the whole LHC Computing Grid (LCG) project is at <http://lcg.web.cern.ch/LCG/>.

What happened next?

As described above, the CCRC'08 was designed to bring together all four experiments and to exercise all the computing models from data acquisition through to data analysis at the Tier2 sites. Here we report on progress in the first phase and preparation for the second.

In terms of data transfer, several significant goals were achieved. The total

rates transferred from CERN to the Tier1 sites were higher than those achieved in earlier tests, and have been sustained over several weeks. All four experiments have shown sustained rates in excess of their requirements for 2008. Rates of greater than 2.1 GB/s were achieved in aggregate between all experiments from CERN to all 11 Tier1 sites (see figures 1 and 2). Figure 1 shows the distribution to sites by experiment over several weeks and figure 2 shows that on one day. The experiments have found the testing sufficiently useful and so are continuing.

The performance of the Castor 2 system at CERN had also been of concern, but performed reliably at rates well in excess of those needed for data taking. CMS in particular were able to demonstrate aggregate rates in and out of Castor of 3–4 GB/s (figure 3), and sustained rates to tape of 1.3 GB/s. Unfortunately this level of use with several experiments together was not demonstrated, since ATLAS and ALICE were later in starting the challenge. In total during the one-month challenge CMS moved more than 4.5 PB of data between all participating sites. All of their Tier1 sites achieved the targets to receive data from CERN and migrate to tape, and a large fraction of the Tier1–Tier1 and Tier1–Tier2 targets were also achieved.

ATLAS started late as the amount of data generated in their dress rehearsal was rather less than expected. However, using simulated data starting from week three rapidly showed the rates mentioned above. They also validated the use of SRM v2 (the common mass storage interface) and the Tier1 storage system setups. They achieved most of their milestones despite the early problems and external dependencies.

ALICE and LHCb also achieved their data rate targets with sustained rates of 80MB/s and 70MB/s, respectively, over several weeks. LHCb tested bulk file deletion with SRM v2. They have tested most of their full computing model, despite the new version of Dirac only being available just before the start of the test.

In conclusion we can state that the February exercise has been a success, with relatively few issues showing up. Some problems of communication – e.g. slow reporting of problems outside of working hours – show that although processes were in place they were not well advertised or used. These points, together with a prioritized list of issues in the storage systems and other middleware services will be addressed for the May challenge. All four experiments expressed the desire to keep running at this level from now on. It is important that the full 2008 resources are in place at the Tier1 sites in time for the May phase so that the complete system can be tested at the full 2008 rates.

Ian Bird, Harry Renshall and Jamie Shiers, IT Department

ETICS software project keeps enforcing quality assurance

The infrastructure for Testing, Integrational and Configuration of Software (ETICS) project was started two years ago as a spin-off from the Enabling Grids for E-science (EGEE) project. It was initially co-funded by the European Commission for two years and has now received a new grant for a second phase, the ETICS 2 project, from March 2008 to February 2010. The project is coordinated by CERN and has eight partners: CERN, INFN, Engineering SpA, 4D Soft Ltd, University of Wisconsin-Madison, SZTAKI, Forschungszentrum Jülich (FZJ) and VEGA GmbH.

The major goal of the ETICS project is to assist the software engineering activities of small and large software projects by providing tools and services to build, test and evaluate the quality of their software. The major objective inherited from the EGEE project was to handle complex software products developed and deployed on distributed infrastructures. Building and testing sophisticated software products, such as the gLite middleware that was developed by EGEE or the gCube software produced by the DILIGENT/D4Science project, is a major challenge. It requires ways of managing complex dependencies; building on many different platforms; deploying clients and services located on the Internet in different sites; collecting metrics; and producing reports to help the certification teams spot and address existing or potential issues.


ETICS helps to manage this complexity by providing web-based tools and command-line application programming interfaces (APIs) to describe the software, build it, produce packages for different platforms, test the deployment, and perform static, dynamic and functional tests. In addition to providing tools and guidelines, we directly manage a distributed computing infrastructure located at CERN, the University of Wisconsin and INFN, where projects using the ETICS services can test their software on more than 40 different combinations of operating systems, CPU architectures and compilers.

Who's involved?

ETICS 2 is a distributed project. Our partners are located in Europe and the USA and we have started collaborations with commercial companies to explore the use of Grid middleware and applications for business activities. While the most prolific users are the Grid middleware developers from EGEE, we have more than 200 users



The ETICS 2 team during the All-Hands meeting, which was held at CERN in February.



Above: the ETICS build and test configuration management application. Right: an example of trend analysis.



from about 25 projects registered in the ETICS database. ETICS 2 is now expanding its scope to include new infrastructures like DEISA (with support for the UNICORE middleware) and new engineering communities like the Aerospace Engineering sector.

We use the PPT project-planning tool developed at CERN to manage the internal project effort. Over the past two years 37 people have collaborated in one way or another with the development and operation of the ETICS infrastructure, and more people are joining from new infrastructures and communities.

At CERN, the ETICS 2 project is run by six people, including the project director Alberto Di Meglio. Five of them work in the Grid deployment group and one in the EGEE project office, dealing with the administrative management of the project.

Major achievements

As a spin-off of the EGEE project, ETICS inherited the single-platform build system that was originally used for the gLite middleware and a simple set of quality assurance tools. However, since the beginning ETICS has been addressing one of the basic challenges that many large-scale, open-source software projects have to deal with, namely the need to set up complex software lifecycle management infrastructures and processes to guarantee acceptable levels of quality, interoperability and maintainability.

The major achievement of the ETICS project has been the deployment of a managed distributed infrastructure with more than 40 different platforms and a rich set of command-line and web clients to configure, build and test small and large software projects. This has been

made available to a reasonably large user base (25 projects and more than 200 users). A rather sophisticated build and test configuration management system helps the users (developers, integrators or testers) to record detailed information for thousands of build and test configurations. A repository service can also be used to store and publish build and test reports and packages.

Another important achievement is the possibility of running complex test suites involving the use of many different interoperating nodes in a fairly automated way. This is a major challenge in the distributed and Grid software development world. With ETICS you can define “multi-node” tests composed of services, clients and test scripts, and then automatically allocate and configure the required nodes (computers), deploying the software as needed.

The execution flow is managed by a set of synchronization commands to be used in the test scripts, and all logs and test results are consolidated and stored in the repository for later analysis. In this way, the test environment is always reproducible, the test conditions are well under control and the effort needed to manage the testbeds is reduced.

Looking ahead

During the first two years most of the ETICS users came from projects already involved in Grid activities worldwide. Even so, it has not always been easy to convince developers and testers to trust a new build and test system that claims to run everything in an automated way. Indeed, collecting all user requirements, implementing the services and managing the test infrastructure has been a very challenging task. We are very grateful to the projects that have been using the system since its infancy – especially EGEE, DILIGENT and OMII-Europe, who have provided invaluable feedback to make the system reliable. The fact that they are still using ETICS proves how worthwhile the project is.

Now that a stable environment for building and testing software is deployed and running, our next major goal and challenge is to extend this to more infrastructures. We have started collaborations with the gLite CREAM development team in INFN and the UNICORE development team in FZJ to add support for running ETICS on new types of Grids (currently ETICS uses Condor from the University of Wisconsin). Our aim is to turn ETICS into a standard service that is provided by the Grid infrastructures to

their users.

Within CERN, another important goal is to propose ETICS to software development projects outside the typical Grid development world. We have started collaborations with other groups in the IT department to see what benefits projects such as Quattor or CASTOR can receive from using ETICS. We would also like to work with teams outside the department to share our experiences and improve the system even more.

Finally, one of the major challenges is to make ETICS and its quality assurance process a standard. ETICS has a quality assurance model called Grid-QCM (Grid Quality Certification Model), which was developed by the Engineering SpA. It is based on a number of existing ISO standards, but is more practical and easier to apply in dynamic environments such as scientific communities.

We know it will be a difficult task, but we are convinced that the value of collaborations like ETICS, EGEE or LCG is also in the definition of open standards and the creation of new markets and opportunities.

● For more information about the ETICS project, visit www.eu-etics.org.

Alberto Di Meglio, IT/GD

Les Robertson looks back at six eventful years at the helm of the LHC Computing Grid

In this special feature, iSGTW chats to Les Robertson, who recently stepped down after six years at the head of the Large Hadron Collider Computing Grid (LCG).

In the beginning

Les Robertson arrived at CERN in 1974 to fix a problem. The European physics research laboratory had just purchased a new supercomputer. The problem, says Robertson, was that it didn't work.

“At that time customers fixed their own operating systems,” he explained. “I arrived as an operating systems expert and stayed on.”

Twenty-seven years later, Robertson began work on an entirely different problem. Preparations for the Large Hadron Collider (LHC) were well under way, but the computing resources required to handle LHC data had been left behind.

A hole in the funding bucket

“Computing wasn't included in the original costs of the LHC,” Robertson explained. “The story was that you wouldn't be able to estimate the costs involved, although the estimates we made at the time have proven to be more or less correct.” This decision left a big hole in funding for the IT that was crucial to the ultimate success of the LHC.



Les Robertson in 2001 (left), just before the LCG project started, and as he was in 1974 (right) when he first arrived at CERN. Images courtesy of Les Robertson.

“We clearly required computing,” said Robertson, “but the original idea was that it could be handled by other people.”

In 2001 these “other people” had not stepped forward. “There was no funding at CERN or elsewhere,” Robertson explained. “A single organization could never find the money to do it. We realized the system would have to be distributed.”

CERN began asking countries to help. The charge was led by the UK, who contributed a big chunk of e-science funding, closely followed by Italy, who continues to supply substantial funding to CERN. Germany also donated a chunk of funding, and then other

countries followed suit. “This money gave us a big boost,” Robertson said. “It allowed us to create something much bigger.”

The Grid

In 1999 Harvey Newman from Caltech had initiated the Monarc project to look at distributed architectures that could integrate computing resources for the LHC, no matter where they were located. At around the same time, Carl Kesselman and Ian Foster carved a spot on the world stage for the Grid.

“Their book motivated the idea of doing distributed computing in a very general way,” Robertson explained. “It stimulated everyone's interest. We decided to ride the wave. But the Grid has not become the panacea. It has become 250 different things, which has led to benefits and also to problems. Standards haven't emerged in the way we expected, nor have off-the-shelf products.”

Democratic and global

One of the successes of the LCG has been the involvement of multiple centres from around the world. “Different countries, universities, laboratories...We have over 110 Tier2 centres up and running, some big and some very small, but all delivering

resources to the experiments,” Robertson explained. “Many of these are computing centres that haven’t been a fundamental part of the experiments environment before, and we’ve all put a lot of effort into working as a collaboration, sensitizing people to what will be required when the first data starts to arrive. The advantage is that all these centres are now involved in the experiments and so there are many options for injecting new resources when they are required.”

Challenges so far: the big three

When asked about the challenges he faced as head of the LCG project, Robertson laughs wryly. “There were several big problems,” he said, “and they were all a bit the same.”

● **Money.** “Funding was certainly a problem, and the UK and Italy were especially important in providing people to get us started. We also benefited enormously from EGEE and OSG and their predecessors. As far as equipment is concerned, with the exception of ALICE, we have what we need for the first couple of years. After that there’s still a lot of work to be done to build up resources as the data grows.”

● **Collaboration.** “I was surprised by the intensity of competition within the high-energy physics (HEP) community. This is collaboration, not a project with funding for people, so we all have to agree on what we do. People have had lots of good ideas, but in the end you have to do the practical thing. Achieving resolution has been harder than I expected.”

● **Distant deadlines.** “When the end is far away, there’s a temptation to think of sophisticated, clever ways of doing things. But this is difficult when there is little experience and you don’t actually know what you need. Over the past year, the LHC has drawn closer to startup and this situation has changed. People have started to realize that we have to use what is available, because we want to do physics and we need a solution.”

Our challenges for the future

● **Immediate: stabilizing operations.** “The future of high-energy physics and



Some of the centres involved in the WLCG (clockwise from top left): the French Tier1 centre in Lyon, the Asian Tier1 in Taipei, the University of Wisconsin Tier2 in the US, and the Cern Tier0 in Switzerland. Images courtesy of IN2P3, ASGC, UW-Madison and CERN.

the Grid depends on what comes out of the LHC. It is very important that the LHC produces something quickly and that the Grid operations stabilize rapidly.”

● **Mid-term: managing the data.** “We are able to physically move the data around quite efficiently, but the challenges of data placement and management are still being proven. How do we distribute the data, and how do you find out where the data is? There are some enormous challenges yet to come.”

● **Long-term: managing energy requirements.** “Computing has been getting cheaper and cheaper, but now the costs are actually going up because of the power requirements. The cost of supplying energy will affect all large-scale computing. We will have to invest heavily in ways to improve efficiency.”

Countdown to startup

So is Robertson confident all will go according to the LCG plan when the first proton beams race through the LHC? He’s hoping! “There’s a lot of work still to be done,” he said. “This is new, this idea that you start a machine and the computing required is not all in the same place as the machine. It hasn’t actually been done before. When the beams come we don’t know what will happen. Things will be chaotic, people will want things we didn’t expect. But HEP is showing that this highly distributed environment is usable. Physicists are no longer dependent on CERN for all the funding or CERN deciding on priorities. We’ve created a democratic environment where you can plug into computer resources wherever you find them. In principle, that was the real goal of the Grid.”

Cristy Burne, iSGTW

● This article was published online in iSGTW on 27 February.

If you want to be informed by e-mail when a new CNL is available, subscribe to the mailing list cern-cn1-info.

You can do this from the CERN CNL website at <http://cern.ch/cnl>.

Announcements and news

Mail servers must register to accept e-mail from the Internet

The following announcement is intended for administrators of a mail server (e.g. sendmail, postfix, etc). It concerns e-mails that are sent from the Internet to addresses of the following type: somemailbox@somehost.cern.ch.

Mail server managers are requested to register their servers so that they can accept e-mail from outside CERN. In future, the CERN mail infrastructure will only relay messages from outside CERN to officially registered mail servers. This rule only applies to messages sent from the Internet. There will be no change with respect to

messages sent from inside CERN.

If you are responsible for a mail server that accepts e-mails from outside CERN, please check <http://cern.ch/mail/Help/?kbid=191090>, where you can find information about the new rule and check to see if your host is already registered. If you wish to register a mail server then please send a message to mail.support@cern.ch.

This rule will have been gradually enforced from 20 February onwards. Thank you for your co-operation.

CERN Mail Services

New spam detection for all CERN mailboxes

The flow of spam (unsolicited e-mail) targeting CERN mailboxes is constantly increasing – 96% of the two-and-a-half million mails received each day at CERN are defined as spam, but spammers are getting smarter and detecting it is becoming ever more difficult.

To address this evolution, a new spam-detection software engine will be deployed at the beginning of April. It will not affect how CERN users currently deal with spam – the spam configuration page, the “CERN spam” folder and the Outlook “report spam” button will continue to work exactly as before.

Because the detection of spam is based on probabilities, some legitimate messages can be incorrectly detected and labelled as spam. We therefore invite all users of CERN mailboxes to regularly check their “CERN spam” folder for such messages. To prevent incorrect filtering of legitimate messages, it is possible to set up a whitelist (list of trusted e-mail addresses that bypass the spam filter) and adjust the spam filtering level from the CERN mailbox configuration pages (see links below).

Useful links

For more information about spam, please visit <http://cern.ch/mail/antispam>. To configure the spam filtering level of your CERN mailbox, go to <http://cern.ch/mail> and click “spam fight”. Thank you for your collaboration.

CERN Mail Services

School of Computing 2008 to be held in Norway

This year’s CERN School of Computing (CSC), organized in collaboration with the Gjøvik University College and the University of Oslo, will be held from 25 August to 5 September in Gjøvik, Norway.

The school is aimed at postgraduate students and research workers with a few years of experience in scientific physics, computing or related fields. The deadline for registration is 1 May.

Special themes this year are:

● **Grid technologies.** This track will deliver some unique theoretical and hands-on education on some of the most advanced Grid topics.

● **Base technologies.** This theme will address some of the most relevant underlying technologies for software

development, security, networking, hardware and architecture.

● **Physics computing.** This track will focus on the particular challenges that the high-energy physics community is facing for LHC computing.

For further information about the school see the CSC website at www.cern.ch/CSC. This includes details about the programme, practical information about this year’s school, and details of how to apply (including the web-based application form).

The financial support from the European Union (EU) ended in 2007. The CSC organizers can therefore no longer offer EU grants to any of the participants.

**François Flückiger, CSC director,
IT Department, CERN**

Changes to support for engineering tools

This article documents a number of changes to the support structure, which will be necessary to cope with the predicted reductions in engineering support staff. These changes have already been announced to the steering committees for the services concerned (CAEC, DTF and ELEC). The support changes will have taken effect by the end of March, and include the following:

Software engineering support will stop

The following tools will be available “as is” as long as the product is fully functional: Adobe Framemaker, Altova XMLSpy, Borland Together, Grammatech Codesurfer, IBM Rational Products, NAG software libraries, Parasoft C++test, Parasoft Insure++, Syncro Oxygen, Windriver SNIFF, Wingware WingIDE etc.

● No new tools will now be purchased for software development and there will be

no new versions installed for the listed applications.

● Application support will stop. E-mail accounts and related Remedy categories will be closed.

● Support: some tools will be handed over to main users and documentation has been prepared for these cases.

Users are encouraged to refer to the Software Development Tools website, where information about installation and licensing can be found. See <http://engineering-software.web.cern.ch/engineering-software> and follow the link to Software Development. Users are also encouraged to refer to the manufacturer’s support pages for specific application support.

Electronic engineering support will be reduced

● Current installations and tools will continue to be supported.

● Application support will stop. Users

should get help through other users, forums, mailing lists, suppliers, local experts, etc.

Mathematica support will be reduced

● Current installations and tools will continue to be supported.

● Application support will stop. Users will have to contact other users, user forums, mailing lists, suppliers, local experts, etc.

CAEC/ELEC client workstation support will be reduced

● Testing, evaluation and providing installations of all tools will continue but only for the standard NICE environment.

● Only standard NICE installations and PCs will be supported. No on-site visits to solve problems on individual configurations.

In case of problems, the user will have to reinstall the standard NICE environment and the relevant engineering tools.

DES Group

Collaboration just gets better as EVO takes over from VRVS

After more than 10 years of operation, the virtual room videoconferencing system (VRVS) will be shut down on 15 April, giving way to the next-generation enabling virtual organizations (EVO) system. From this date all planned VRVS meetings should be migrated to EVO.

VRVS, created at Caltech and operated in collaboration with CERN and other partners at laboratories and universities throughout the world, was the first web-based collaboration tool available for the physics community when it started in 1995.

It later led the way for other web-based collaboration services that were based on similar ideas (albeit with more limited scope), which are still available today. VRVS's worldwide spread can be summarized with the following statistics: 30,500 registered users from 151 countries working on 126,000 host computers and other IP-based conferencing devices; 123 servers (the so-called "reflectors") forming a core infrastructure spread across 41 countries to direct and manage the network traffic; and more than 57,000 meetings representing a total of 23 years of meeting time (or, more than 100 years of meetings during "prime time").

The EVO (evo.caltech.edu) service will now take over the responsibility to provide a robust and efficient collaboration service for the LHC as the experiments are about to start. The migration from VRVS to EVO started a few months ago and is shown in

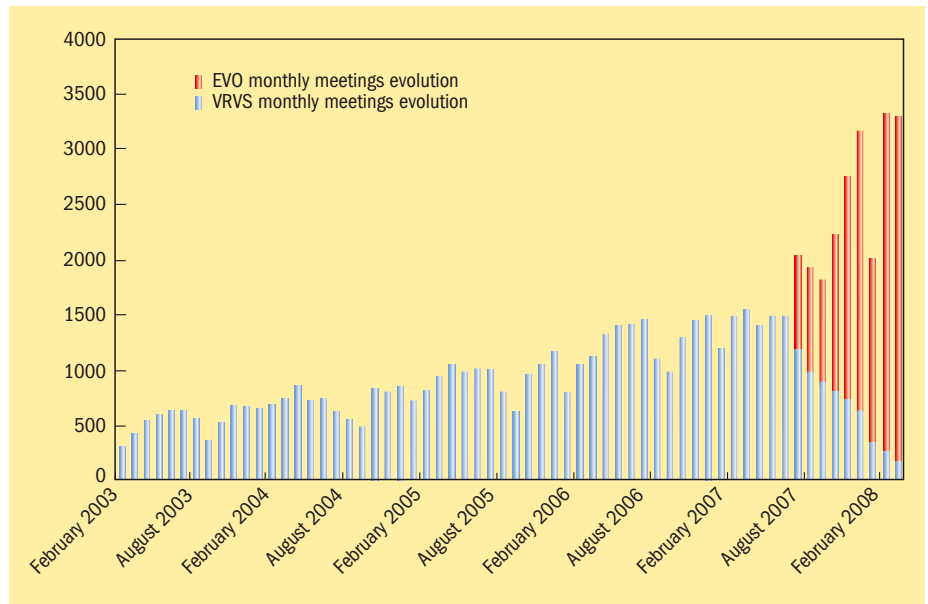


Fig. 1. A graph showing the recent migration from the virtual room videoconferencing (VRVS) system to the next-generation enabling virtual organizations (EVO) service.

figure 1, which shows an impressive rise in the number of EVO meetings each month as we approach the first LHC runs.

At CERN in particular, in addition to the widespread use of EVO on individuals' desktops and laptops, there is an ongoing effort by CERN-IT in collaboration with the ATLAS and CMS experiments to make all the conference rooms EVO-compatible by

completely refurbishing them.

The results are already evident, as all of the main rooms used by ATLAS and CMS for conferencing are now EVO-certified. For more information about which CERN rooms are ready to use with EVO, please check <http://it-multimedia.web.cern.ch/it%2Dmultimedia/collaborative/rooms/>.

EVO Caltech Team

User Support services welcomes a new face

Roger Woolnough, a well known figure from the Helpdesk and User Support services, is retiring and leaving CERN in April after 37 years of fruitful collaboration. This issue of the *CNL* gives us the perfect opportunity to thank Roger warmly for his excellent work and pleasant personality, to say goodbye and to wish him all the best for his new life after CERN.

At the same time I would like to inform users that Nick Ziogas (formerly of the IT/AIS team) will take over the "Helpdesk and User Support" section from Roger. Nick will also be an active member of the IT Manager on Duty team, who act behind the Helpdesk to coordinate this first-line service (outsourced to an external company) with the various service managers within the IT department.

One important task for the IT Manager

on Duty is to maintain the IT Service Status Board (<http://cern.ch/SSB>) and keep it up-to-date with information for users on current problems and forthcoming changes. Other members of this team include Catherine Delamare and myself.

I note with pleasure that I have spent more than 10 years with Roger in the User Support section, as proven by an old *CNL* article published in September 1998, when the Helpdesk was still referred to as the UCO (User Consultancy Office). The title of this article was "Service Improvements at the User Consultancy Office" and those of you who are nostalgic can browse the article on the web at http://cern.ch/cnlart/232/art_uco.html.

Some of the other articles from past issues of *CNL* that Roger and I wrote together show how this service has evolved

over time. For example:

- "Central Computing Helpdesk – a Front-line Service" (<http://cern.ch/cnlart/2000/003/helpdesk/>).
- "The UCO is dead, Long Live the Computing Helpdesk" (cern.ch/cnlart/2001/001/uco-mail/).
- And more recently, "Computing Helpdesk – a New Strategy" (<http://cern.ch/cnlart/2002/003/helpdesk/>), where the role of Manager on Duty (known as the MoD) was explained in detail to the users.

An important challenge for the User Support team this year is the renewal of the contract for desktop computing support that the IT department provides for the whole of CERN. A call for tender for this new contract is to be issued this summer so that the new contract can begin in July 2009.

Nicole Crémel, IT/UDS

GridTalk: spreading the word

Over the last 18 months, the weekly newsletter *International Science Grid This Week* has become one of the best ways of finding out what's happening in the Grid world. And for anyone who has ever tried explaining Grids to the general public, CERN's *GridCafé* website is invaluable. From 1 May these two online success stories will become part of a new EU project called GridTalk that will spread the word about the achievements of Grids in science.

Sarah Pearce, the GridTalk project manager, explains: "The impact of Grids has rapidly expanded beyond that which can be disseminated by individual projects. During its two years, GridTalk will tell the success stories of Europe's e-Infrastructure to three main audiences: policy makers working in government, science and industry; scientists with an interest in Grid computing; and the general public. We will coordinate the dissemination outputs of Enabling Grids for E-science (EGEE) and other European Grid efforts, so that their results and influence are reported widely in print and online."

For GridTalk, everything comes in threes. The project has three partners – Queen Mary University of London, CERN and APO (a French web design company). It aims to reach three audiences. And it also has three main products:

Informing policy makers

Many e-infrastructure policy activities produce reports that would be of interest to a wider community, for example on standards, security or progress towards a sustainable infrastructure. Working with Europe's e-infrastructure projects, GridTalk will produce articles and briefings written in jargon-free language that provide timely summaries of policy-oriented reports or discussion of key issues. The outputs will be aimed at non-technical policy makers in government and industry, scientists and the public.



GridCafé

The GridCafé website (www.gridcafe.org) was produced by CERN in 2003 to tell the public about Grids, and has been nominated for Pirelli International and Webby awards. GridTalk will maintain and extend the GridCafé website, keeping it at the cutting edge of Grid dissemination. It will also continue the Gridcast project, where scientists at Grid events blog about their experiences. A new goal will be to produce a map of the Grid, giving a human

face to the Grid, allowing users to listen to podcasts from Grid sites worldwide and watch interviews with researchers.

By integrating this with Imperial College's more technically oriented Real Time Monitor, audiences will get a visual demonstration of what Grids are doing and who is using them.

International Science Grid This Week (iSGTW)

iSGTW (www.isgtw.org) is a successful electronic newsletter, published weekly, which informs more than 3000 readers in 100 countries about scientific Grid computing. *iSGTW* was formed as a joint project between Open Science Grid in the US and EGEE. GridTalk will allow *iSGTW* to cover more European Grid projects and to expand its resources section, providing information and support for scientists working with Grids or considering becoming involved with them.

GridTalk will work closely with projects funded through European and national means to raise the profile of their results. e-Infrastructure projects will make their outreach products available to GridTalk, and in turn will use GridTalk's material to engage new audiences with their objectives and achievements.

As well as outreach across Europe, GridTalk will coordinate activities with international partners, both in areas with well established Grid infrastructures such as the US, and in developing areas such as Latin America and Asia. *iSGTW*, through its co-funding with Open Science Grid, has already started this process.

So, if you want to let the world know about your work with Grids, GridTalk is here to help. Whether you have a suggestion for an article for *iSGTW*, a project that you think should be covered on the *GridCafé* website, or a site that needs to be on the map of the Grid, contact Sarah Pearce at s.pearce@qmul.ac.uk and we will help spread the word.

DILIGENT: from digital libraries to virtual research

In August 2007 Greece saw more wildfire activity than any other European country in the last decade. The fires, which involved scores of separate blazes, killed 64 people and destroyed half a million acres of vegetation. Now, thanks to data collected by organizations such as the European Space Agency (ESA), we have a detailed record of this event.

Such Earth observation data is collected all the time, all around us. For instance, to continue the wildfire example, two

ESA satellites – ERS-2 and Envisat – are continually monitoring the Earth's surface, sensing hot areas: spots exceeding 308 Kelvin at night are classified as "burning" fires. The current and archived maps are available online as part of ESA's *World Fire Atlas*.

Such data is invaluable for agencies like the European Environment Agency, European Union and the United Nations; however, these data quantities are just a drop in the proverbial data bucket when it

comes to ESA's Earth observation.

Over the next 15 years ESA alone expects to collect and process around 12 petabytes of data. To manage these mammoth amounts of data, ESA is turning to Grids, and to the DILIGENT project in particular.

Virtual research environments

DILIGENT, which stands for "A testbed Digital Library Infrastructure on Grid Enabled Technology", is now testing new software that can help to create dynamic,

on-demand virtual research environments.

Called gCube, DILIGENT's software allows users to search for and collect data from distributed and dissimilar resources, and to then process or analyse that data using the requested tools and services. The digital library or research environment that has then been created can be saved for later use or shared with others.

ESA are participants in DILIGENT's testbed phase and the new software will help them to access huge amounts of space-based data and information more efficiently, complete distributed calculations, collaborate across distances and prepare technical reports, which will be especially welcome.

Living documents

"ESA's data and Earth observation information contributes to many regular national and international reports on the status of the environment," said Veronica Guidetti, ESA's work package coordinator for DILIGENT. "These are huge, complex documents; they can be referred to as "living documents" because of their need for periodic updating."

Guidetti says that DILIGENT's technology could save many hours' work. "For instance, if I had to contribute to a yearly technical report on Europe's environment, with updated statistics and maps using the same calculations as previous years, but with updated input, I would like to submit fast reprocessing of entire archives of data with just one click of a button," she explained.

DILIGENT is currently rolling out the first full release of its gCube software, which aims to help organizations rapidly adapt to new requirements, implement new

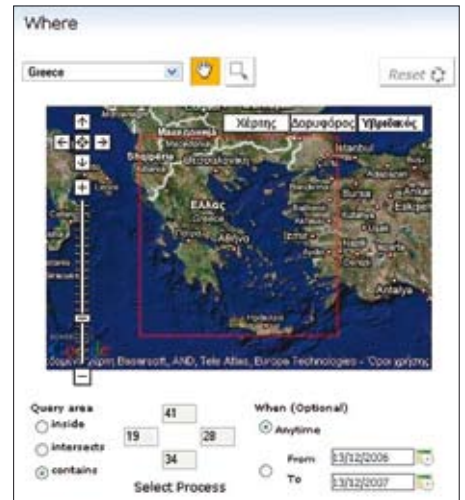


Fires in Greece during August 2007 killed 64 people and destroyed around half a million acres of vegetation, leaving a scarred battlefield. Thanks to the continual collection of Earth observation data, scientists now have a detailed record of this event. Image courtesy of the European Space Agency.

workflows and reprocess data archives.

Since gCube and DILIGENT are Grid-enabled, making use of the shared resources of Enabling Grids for E-sciencE, organizations can also access a large amount of computing power for a fraction of the cost.

"This reduces the cost of ownership, and at the same time increases the availability of exploitable data information, while still preserving user integrity and confidentiality," explained George Kakaletis, DILIGENT validation manager and research associate at the University



Software produced by Grid-powered archiving project DILIGENT can make it much easier to search through the many databases of information being collected by agencies all over the world. DILIGENT can help compile, analyse, process and connect this data to other distributed data libraries. Image courtesy of DILIGENT.

of Athens in Greece. "We're hoping our findings will become fundamental to the future of the Grid."

The sequel

DILIGENT celebrated its project sequel as DILIGENT for Science or D4Science, which started on 1 January. DILIGENT services will be available to the Grid-interested public once the system is in the production stage.

Danielle Venton, EGEE

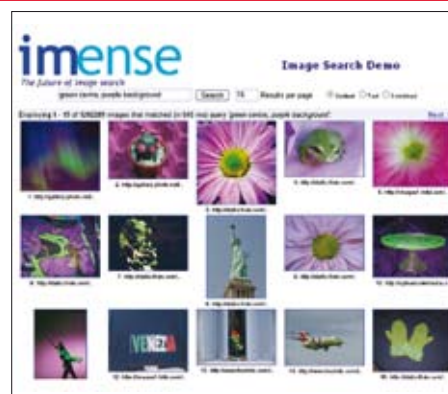
● This article was published online in *iSGTW* on 6 February.

The business of Grid computing: what can Grids do for you?

Enabling Grids for E-sciencE (EGEE), Europe's largest Grid for a wide range of applications, is working with industry to help businesses new and old to capitalize on the open-source Grid technologies that have already been developed.


Stephanie Parker, who chairs EGEE's Industry Forum, says that Grids have a lot to offer. "The biggest selling point is that Grids can save you money: companies can access more computing power without shelling out for new hardware. Grid delivers major value by improving business performance, by helping to get products to market faster, and by enabling companies to do new things."

Many companies are working to develop software that enables new technologies for computing Grids. An excellent example is Imense, of Cambridge, England, a start-up



Imense, of Cambridge, use gLite middleware to run their content-based image-retrieval technology. This particular result came from a search for "green centre, purple background." Image courtesy of Imense Ltd.

company that uses EGEE's open-source gLite middleware to run their content-based image-retrieval technology. "Our work with the Grid let us demonstrate that our software can handle millions of images, at a time when we were a small company



Grid computing provides the power required to run complex applications, like this financial simulation, an application of Avande Italy. This screenshot shows a rendering of a volatility matrix. Image courtesy of Avande Italy.

and could not supply the computing power needed ourselves," explained David Sinclair, one of the founders of Imense. "This in turn impressed the investors, and led to us winning funding for our company."

The EGEE Business Associate Programme

is an important component in EGEE's strategy for the take-up of EGEE Grid technologies by business and industry. "EGEE is currently working with six companies, all of which are experts in technical development and the Grid computing market," said Parker. "We're also bringing on board the world's oldest bank, the Italian-based Banca Monte dei Paschi di Siena."

Collaboration

Among these companies is Avande Italy, who are working with EGEE on their Grid technology and applications. "To fulfill the computing power needs of complex projects, such as finance simulations, we created a proprietary Grid architecture based on .NET," explained Raffaele Sgherri from Avande Italy. "We are working with EGEE to implement their standards in our

solutions, targeting full interoperability with gLite to provide more flexibility and features to both systems."

EGEE's Industry Forum reaches out to private and public organizations that have not yet looked at Grid for innovation. The forum interacts with current and new users, highlighting the benefits of Grid through user case studies and evaluating adoption potential across diverse verticals.

EGEE's Business Days also provide a primary setting for encouraging collaborations with organizations. "Businesses learn about how gLite might meet their specific needs," explained Sy Holsinger from Trust-IT. "EGEE learns more about what businesses are asking for. It's a win-win situation." Businesses can also see how they can benefit from inexpensive training services to accelerate adoption.

These meetings are valuable in

understanding how to clear any technical and non-technical hurdles to adoption through the support of EGEE's Industry Task Force and the Focus Group for commercial adoption. "If companies have a specific problem, my job is to find a person in EGEE who is working in that area and can help," Holsinger said. "In this way we can work together to optimize solutions."

The next EGEE Business Day, entitled "Showing companies why they need Grid and how they can do it", will be held on 21 May in London, with the support of the UK Science and Technology Facilities Council. This event explores the opportunities for open source, with particular reference to the energy, finance and pharmaceuticals industries.

Danielle Venton, EGEE

● This article was published online in *iSGTW* on 12 March.

New Rome Tier2 centre prepares for CMS data

The Rome CMS Tier2 regional centre, hosted by INFN, moved to its new location in December 2007. A new hall was built for the centre, as well as for the ATLAS Tier2 centre.

CMS computing is a distributed activity, organized in a multi-Tier hierarchical structure. Data are processed at the Tier0 centre at CERN, then shipped to Tier1 centres via high-speed networks. These provide CPUs for data reprocessing and simulations, and archive data on tapes. Part of the data is then copied to Tier2 centres, where Grid services automatically dispatch jobs submitted by physicists.

The INFN Rome Tier2 centre is equipped with four water-cooled racks made by German company Knuerr. Three more racks will have been installed in March 2008, but the hall can host up to 14 racks. Racks are fed by water stored in a three cubic metre reservoir and chilled to 12 degrees. Water enters heat exchangers inside the racks, cooling hot air produced by the CPUs, and exits the rack at 18 degrees. Such a modular technology allows for a reduction of electrical power and temperature stability, keeping the hall at room temperature.

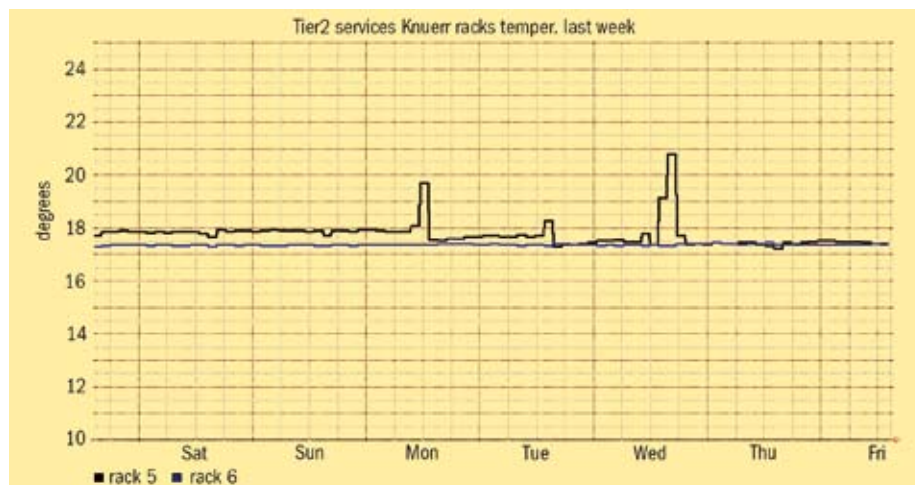
The centre is equipped with tools for remote operations and monitoring, as well as security devices such as fire and flood sensors. Potentially catastrophic events generate automatic emergency procedures, and in case of power failure we are able to automatically shut down the machines cleanly, thanks to the Uninterrupted Power Supply (UPS). Any event is logged and notified to administrators via e-mail, and SMS notification is coming soon.

Luciano Barone and Giovanni Organtini, PH Department

● This article was published online in the *CMS Times* on 18 February (<http://cmsinfo.cern.ch/outreach/CMSTimes.html>).



Massimo Nuccetelli (left) and Fabio Pellegrino, two of the technicians who helped in the realization of the new Rome centre, standing by the two currently installed CMS racks.



A plot of the temperature inside the CMS racks during one week is impressively stable. The peaks correspond to the opening of the rack doors to mount new servers.

A precious stream of data: from Point1 to analysis sites

Lots of things go into the ATLAS detector: electricity and many types of gases and fluids flow into it, proton beams circulate through it, and a lot of intellectual power is injected into it! What we get out from Point1 is a precious stream of information from which we hope to distil our future understanding of particles and forces.

What does this stream of data look like? It is composed not only of event data, but also of other information (metadata) like the status from all sub-detectors, the LHC and the DAQ system, all of which are required to interpret the event data.

Figure 1 gives an overview of data and metadata flowing first from Point1 to Tier0, and then further to Tier1/2 and the analysis sites. The (blue) event data stream is a composite stream as it carries not only physics events meant for analysis, but also partial event data for detector calibration, and a selection of events for express monitoring of data quality and other more technical data streams that can capture rare error situations.

The physics events are subdivided further into streams according to classes of trigger signatures, e.g. electron/photon, muons/B-physics, jets/taus/missing energy, minimum bias. An event can have more than one trigger signature, and then it is flowing in several streams. This we call “inclusive” streaming (as opposed to “exclusive”, where multi-signature events go to a separate overlap stream). Quantitative studies of the expected multi-signature overlaps have established the feasibility of having more than one raw physics stream.

Data streams don't flow continuously, but are quantized into so-called luminosity blocks – typically 1 minute long, during which conditions like instantaneous luminosity and detector status are stable. The event filter (third level of the ATLAS trigger) subfarm outputs form data files per luminosity block and stream, as shown in figure 2. On the hardware level, data flow again serially in the end via optical fibres, such as from Point1 to Tier0 to Tier1/2/3 sites.

Streaming continues into the results from the subsequent data-processing steps. Event summary data (ESD) form the same streams as the raw data, and analysis object data (AOD) can subdivide these streams further as required. It is important to note that the global overview of all events can be maintained with event tags, a tiny (1 Kbyte) subset of the event data carrying the most important event




Fig. 1. An overview of data and metadata flowing from Point1 to Tier0, then Tier1/2 sites.

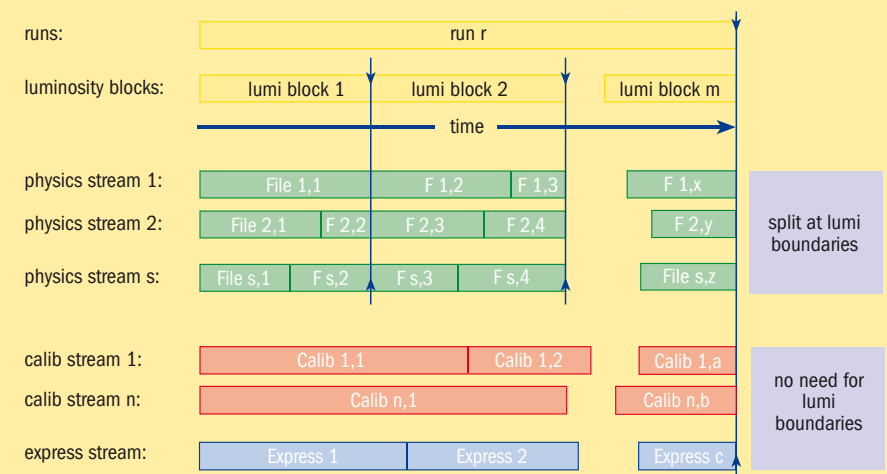


Fig. 2. The data files per luminosity block and stream at Tier1/2 sites.

classification information, e.g. the trigger decision, together with the storage location of the event.

One motivation for having more than one raw-physics data stream is flexibility: it allows processing of just part of the huge data volume with priority if required. But flexibility is also required in the definition of the streams that may have to change if triggers are redefined, e.g. as a consequence of initial beam conditions. This is why the streams are defined in terms of trigger signatures together with the

trigger menus themselves, in the trigger configuration database. The contents of this database is part of the metadata flow – which brings us back to the gross picture.

In a materialistic view, each Watt into the LHC, ATLAS and supporting infrastructure at CERN yields about two or three Bytes per second of raw data flowing out of Point1.

Hans von der Schmitt, Max-Planck-Institut für Physik

● This article was published online in the ATLAS e-News on 25 February (<http://cern.ch/atlas-service-ews>).

New web-application tool boosts usability of JavaScript and CSS

Following the renewed interest in the browser side of web applications, as shown by the success of AJAX (Asynchronous JavaScript and XML) as the current buzzword, there has been consolidation surrounding the development of the technologies that made this new wave of web possible in the first place: JavaScript and dynamic html.

In this article we will describe how we were able to quickly improve the usability of some existing applications, with the help of a modern JavaScript library (jQuery) and careful design.

The CERN training applications (the back-office for human-resources personnel and the public catalogue are available at <http://cta.cern.ch/cta2/f?p=110:9>) have been developed using Oracle Application Express (APEX), a rapid web-application development tool for the Oracle database.

APEX operates as a thin layer between the database and the web (hence its former name: HTMLDB), and provides lots of nice features to build your web application, but still allows you to retain full control of the database and browser, and harness their full potential.

A powerful engine

The APEX engine generates the HTML to be sent to the browser via a simple yet powerful template engine that effectively separates the application logic from the presentation layer.

Inserting an external JavaScript library into an APEX application is as simple as adding it to the page template. JQuery is a remarkably elegant, fast and compact (15Kb, cached on subsequent requests) library that really simplifies the most tedious – and error prone – uses of



Fig. 1. The “smart” search box.

JavaScript, and works across different browsers.

JQuery lets you traverse and select elements in the DOM (document object model) tree of your page by using some “selectors” very similar to the cascading style sheet (CSS) ones and allowing you to loop over the desired nodes (html elements) without having to write a single JavaScript loop. The retrieved nodes can be arbitrarily manipulated, and new behaviours (e.g. actions on double-click or right-click) injected.

Thanks to the APEX templating mechanism, we easily created a semantic layer of style sheets, assigning different CSS classes to components that play different roles, creating a distinction not based on appearance (which is what CSS is normally used for), but on meaning.

This way JQuery can quickly “understand” the page and add behaviours that are relevant to the current page, as context-sensitive menus.

One of the main challenges of the web interface for the CERN training application back-office is the management of

screen real estate: users wish to have a comprehensive view of the data, but the amount of information contained in a page can easily fill several screens.

Simple selection

To allow the user to hide and show any content region with a click, we have to select the regions we need, and for each of them assign a new behaviour corresponding to a single mouse click. This can literally be done in a single line of code thanks to the powerful JQuery selection mechanisms and the flexibility of JavaScript’s enclosures (in-line anonymous functions).

Similarly, the various buttons and their corresponding CSS classes will tell us what actions are available on the various regions that make up the page. This information can easily be extracted and presented in the form via right-click context menus conforming to the standard behaviour of desktop applications, where right-clicking on most items will bring up a menu showing what operations can be performed on that item.

The training catalogue interface was enriched in a similar fashion with hierarchical dynamic menus that let the user access all the courses in just one click, and with a “smart” search box offering suggestions depending on what the user has typed (see figure 1).

References

Oracle Application Express: <http://apex.oracle.com/>.
JQuery: <http://jquery.com/>.
CERN training catalogue: <http://cta.cern.ch/cta2/f?p=110:9>.

Giovanni Chierico, IT-AIS

The deadline for submissions to the next issue of the CNL is

23 May

Please e-mail your contributions to cnl.editor@cern.ch

If you would like to be informed by e-mail when a new issue of CNL is available, subscribe to the mailing list cern-cnl-info. You can do this from the CERN CNL website at <http://cern.ch/cnl>.

Genome centre hosts HEPiX event

The HEPiX Fall 2007 meeting was held at the Genome Sequencing Center at Washington University in St. Louis, Missouri, from 5–9 November.

It was probably the first HEPiX meeting not to be hosted by a high-energy physics site. Gary Stiehr, the principal organiser, had previously worked at Fermilab and was familiar with HEPiX from his time there. In his new role at the Genome Sequencing Center he recognized that his institute and those represented in HEPiX share many common challenges, so he invited HEPiX to hold their meeting at the site. Further proof of this common ground was the presence of several representatives of the UK-based Sanger Institute for Genome Research.

Common themes

The attendance in St. Louis was around 70, which is a healthy number for a US meeting, and this was further boosted during some sessions by staff members from the local institute. Following the successful format first introduced at the HEPiX meeting in CASPUR, Rome, in 2006, this meeting had been planned along the lines of specific themes, each with one or two conveners to promote the talks. As a result, there were a total of 61 talks submitted.

In general, the meeting was well organized. The meeting room was well equipped with power plugs and wireless capacity, the hotel was a few minutes walk away, and the agenda, although pretty full, still allowed lots of time for informal discussions. What follows is only a brief outline of the week; readers are recommended to consult the overviews of topics of particular interest at <https://indico.fnal.gov/conferenceDisplay.py?confId=805>, and look under the agenda link.

The welcome address was given by the Genome centre's director Rick Wilson, who gave a talk about the work at the centre, which was created 17 years ago and is considered one of the premier genome research centres in the world. He indicated some of the techniques used and their associated IT aspects. Increasing data samples and the use of parallel processing are greatly expanding the centre's IT needs (the image top right shows a modern genome sequencing device), and across the road from the centre is a large construction project (pictured right), which will house their new computer centre.

As usual, the meeting started with reports from the different sites represented. DAPNIA and LAL reported on a new project to investigate 40Gb networking; SLAC reported on ongoing preparations in their transformation from a high-energy physics site to one supporting



Processing increasing data samples at the Genome Sequencing Center requires equipment like this modern genome sequencing device. Courtesy of Alex Lossent.



The construction site of the Genome Sequencing Center's new computer centre in St. Louis. Courtesy of Andras Horvath.

multiple disciplines.

Computer centre capacity was a common theme. Many sites are coping by replacing older systems with newer, much more power-density efficient ones. DESY is another site in transition as HERA comes to a close and photon and laser physics are introduced.

There was a session where CERN's IT/IS group reported on their plans for gradually introducing Vista. The most notable quote from the speaker was that, like the replacement of audio cassettes by CDs (and now MP3), you will have to migrate sooner or later so why not sooner. Apart from CERN, almost all sites reported that they were "ready", but only CERN seems to be actively promoting Vista's use, even if only on new PCs.

Most sites, especially those that are members of the Worldwide LHC Computing

Grid Project, reported on plans for new acquisitions in the ramp-up towards LHC production, but a large proportion of these also reported increasing pressure on computer centre cooling and power requirements, and there was a session devoted to this subject. To cope with the required expansion in computing power, SLAC has been experimenting installing one of the new "computer centre in a shipping container" from SUN Microsystems, and this was presented as a series of time-elased photographs showing the ground being prepared and the container being delivered and installed. The jury was still out on whether they will acquire more such self-contained systems. BNL, Rutherford, NERSC and the Genome centre hosting this meeting are among the sites following a more traditional route, with new centres either planned or in construction. NERSC are also experimenting with using DC power at the delivery point.

Wish list

Two of the gurus behind the movement to support and further develop AFS (the Andrews File System) after the end of support from IBM had been invited to give a presentation at the meeting. They related their wish list of developments for which they seek backing, particularly financial backing. In an offline discussion with them it became clear that they could be funded to implement high-energy physics-desired features, so it was agreed that AFS experts from the various labs should look at the OpenAFS wishlist (<http://www.openafs.org/roadmap.html>) and try to agree a

Conference and event reports

package of features of interest. This could then be priced and the labs could discuss cost-sharing, although whether a package of desired features could be agreed and appropriate cost-sharing arranged remains a doubt.

Other invited speakers included Jim Williams from Oracle Corporation, who discussed the recently discovered problem related to silent corruption of data on disks and how this could be prevented or at least signalled, and a most interesting presentation transmitted from Fermilab by Don Holmgren on lattice quantum chromodynamics (QCD) IT challenges.

There was a talk from GSI on performance testing of Lustre; in a presentation described by the speaker as a “graveyard of numbers”, he showed how different configurations affect performance and how kernel tuning can improve throughput. Lustre is being further evaluated at other sites and Andrei Maslennikov described some of the plans of the Storage Working Group that he leads.

A new topic for HEPiX was virtualization. The first talk was a warning from a CERN speaker on the potential pitfalls of virtualization, including problems of scalability and relying on invalid assumptions. Then a series of speakers showed some examples of virtualization in practice. These included high-available clusters at INFN Rome, along with load balancing and dynamic provisioning among worker nodes; high-available clusters at Fermilab; dynamic clustering and again

high availability at Karlsruhe; resource optimization by running independent virtual services at CERN; and an exercise at TRIUMF, which is using Globus Virtual Workspaces to make use of free CPU cycles on remote systems across Canada, where they would otherwise not have access (for example because the host operating system does not support LCG software naturally). Xen seemed to be the virtual operating system of choice in most instances.

Data monitoring

Introducing the monitoring track, James Casey of CERN quoted Rule 1 of monitoring – you can’t manage what you don’t measure. James listed the various tools typically used for monitoring services at different levels – applications, middleware, local resources – but noted that getting these systems to talk to each other is difficult. He is leading an LCG working group to build an easily extensible scheme with special plug-ins for Grid services, and to make the resulting data readable by the standard Grid monitors. The prototype is built round Nagios, with the intention to later integrate it with Lemon. Graphic Grid maps have been developed to allow different views of the monitoring data – geographic, virtual organization, trends, site availability, etc – and within these Grid maps there are links that allow you to drill down to investigate particular problems or issues.

Another working group is starting to

look at CPU benchmarking for high-energy physics. Like most sites, INFN currently uses SPECint2000, usually represented as SI2K, but this is now declared obsolete by SPEC and was replaced by CPlint2006. INFN performed some tests and compared the results, and found that the CERN-modified SI2K tests most closely resembled the SPECint2006 numbers published by the vendors. The result was to propose to INFN to stop using SPECint2000 tests; for the next year use the CERN modified version, but the year after move to SPECint2006, and the best solution is the SPEC INT 2006 RATE test measured by the gcc compiler – the RATE test is claimed to take account of scalability of multi-core architectures. Other sites, notably FZK and CERN, reported similar trends towards using SPEC2006 as the base of their future benchmarking.

Lastly, Troy Dawson, one of the two principal authors at Fermilab of the Scientific Linux distribution used widely throughout high-energy physics worldwide, gave an update on its status. He had just published a new Version 3 release, planned to be the last in SL3. With the just-published release of 5.1 from Redhat, Fermilab had started work on this and expected to have the first distribution available within a week or two.

The next meeting will be held at CERN from 5–9 May, once again planned along the lines of specific themes, and the meeting after that will be in Taipei from 20–24 October.

Alan Silverman, IT-DI

EGI workshop in Rome: defining the European Grid infrastructure of the future

In mid-March, the European Grid Initiative (EGI) held a workshop at the Italian Ministry of University and Research in Rome, which gathered more than 80 representatives of the National Grid Initiatives (NGIs).

The European Grid infrastructure, which is now operational over many countries, was established through nationally and European-funded projects such as EGEE (Enabling Grids for E-science). As this project-based funding cannot be extended indefinitely, the EGI organization aims to construct a sustainable Grid infrastructure in Europe based on national Grids.

The EGI Design Study (EGI_DS) project was launched in September 2007 with the support of the European Commission’s 7th Framework Programme. The project will continue until the end of November 2009. According to the current plans, the EGI organization will begin its functions in early 2010. The EGI_DS has nine principal partners and is already endorsed by 38 National Grid Initiatives.

The discussions at the workshop were



Participants at the European Grid Initiative workshop that was held in Rome in March.

based on a draft definition of functions (collected by the INFN-led work package 3) and draft ideas of a legal structure (provided by the CNRS-led work package 4) summarized in preliminary documents provided by the CERN-led work package 5. The draft propositions of the EGI_DS consortium gave rise to a lively discussion

between the participants. The feedback from the representatives of the NGIs supporting the EGI_DS will guide further work, leading to a “final draft” or “blueprint” describing the EGI in more detail. This blueprint is expected to be ready for further discussion at the next EGI_DS workshop at CERN on 30 June.

At the end of the Rome workshop, the representatives of the NGIs gathered in a closed “policy board” session overseeing and steering the project. The chairman of the policy board, Gaspar Barreira of Portugal, stated after the workshop that: “The body represents and puts together all the National Grid Initiatives in Europe. It follows and endorses the entire process of designing the future European Grid infrastructure. The interoperable permanent Grid computing framework will be for the benefit of the increasing community of users.”

● For more information on the EGI_DS, see <http://web.eu-egi.eu/>

Jürgen Knobloch, IT/EGE

EGEE User Forum looks to build bridges in the science community

This year's EGEE User Forum took place on 11–14 February and attracted more than 300 people to see 19 demos, 36 posters and attend 109 presentations. More than 20% of the accepted abstracts came from the physics community, but Europe's largest scientific Grid is stretching its wings and beginning to fly.

While the size of the infrastructure continues to expand, the event highlighted the need for the increasing number of user communities to connect their own resources to ensure that the application needs do not exceed available capacity. The Grid communities' thoughts have moved from examining the infrastructure itself to looking at the most appropriate higher level services, and this shows that the gLite middleware has established itself as a key platform that is capable of supporting a broad range of users and services.

The tools and techniques developed by the high-energy physics (HEP) community are also finding their feet amongst other disciplines. GANGA (the easy-to-use front-end for job definition and management, originally developed to meet the needs of ATLAS and LHCb), AMGA (the gLite Grid Metadata Catalogue) and the Dashboard have all been adopted by the life-sciences community, and have also been used in other disciplines including telecoms and climatology.

Challenges ahead

One of the founding cases for EGEE and the Grid was presented by Fairouz Malek of CNRS, in a plenary on the "Challenges and success of the HEP GRID". The computing demands of the LHC are presenting an unprecedented challenge, with more than 15 petabytes of data to be generated and processed each year. Such a large amount of information coming from the four experiments will require more than 60,000 of today's fastest CPUs. To meet this demand EGEE is working with other Grid initiatives, such as the Open Science Grid in America, to form the basis of the Worldwide LHC Computing Grid Project, WLCG. The baseline service is now in production, and continuously increasing in capacity and service quality.

Looking beyond the needs of the LHC, and into the future, a plenary talk presented by Dieter Kranzlmüller, project director of the EGI Design Study, set out the steps towards a European Grid Initiative. The ultimate goal is a Europe-wide sustainable, production Grid infrastructure that has no need to depend on short EU



David Manset (right) and Jérôme Revillard (left) from maat Gknowledge, creators of the winning "Health-e-Child" demo. Photo courtesy of Y Legré (www.healthgrid.org).

project funding cycles. The production Grid is already here, it is being used by applications not just from Europe but further afield. But projects such as WLCG depend on the infrastructure that EGEE provides and need to know it can survive. A significant investment has been made in the Grid, both in monetary terms and personal effort, from funding bodies and from users, and this investment needs to be protected. The current funding model cannot offer the protection the Grid needs for long-term success, and the EGI Design Study is looking for solutions.

"The experience of EGEE is an invaluable asset in the preparation of the EGI organization. The results of EGEE and the input from the EGEE experts are of major importance in solving the many challenges of setting up a sustainable organization for the operation of the European Grid," said Kranzlmüller.

There are already 38 National Grid Initiatives in Europe who want to support this effort, and a consortium of nine partners, including CERN, have formed the EGI preparation team, to conceive and set up a new organizational model of a sustainable pan-European Grid infrastructure. For more information on the European Grid Initiative and the Design Study, please visit <http://web.eu-egi.org/>.

At the user forum, the demonstration and poster session plays an important role in helping user groups reach out and find others who perhaps face the same problems, or who have achieved something new. The 19 demonstrations of Grid applications ranged from mathematical models of living cells to bioinformatics

portals, and tools to make using the Grid simpler and easier to control. The prize for the best demonstration, sponsored by HP, was won by "The Health-e-Child (HeC) Gateway and Case Reasoner, a Concrete Implementation and Use-Case of a gLite-based Healthgrid for European Paediatrics". Health-e-Child brings together three major paediatric medical centres with several European institutions to integrate and exploit biomedical information for improved clinical practice, medical research and personalized healthcare. Although working on a separate Grid infrastructure, Health-e-Child deploys EGEE's gLite middleware, one of an increasing number of projects to do so.

"EGEE plays an active and invaluable role by supporting the HeC community in better understanding how best to use gLite, how to gridify biomedical applications, and stay up to speed with ongoing developments," explained demo-winner David Manset, director of biomedical applications at maat Gknowledge.

He went on to explain: "In turn, HeC plays a pioneering role in applying this technology to a new application area, and feeds back new requirements to EGEE and other active groups in the community."

Of the 36 posters on display, which covered a wide range of topics including much of the science currently being done using the Grid, as well as tools, dashboards and methods for making the Grid more accessible, the one chosen by delegates to win the HP Poster Prize was "Fair Grid Scheduling" by Emmanuel Medernach of LPC, France, which looked at ways of efficiently sharing Grid resources.

"The presentations, demonstrations and posters shown testify to the wealth and diversity of the applications currently running on the EGEE infrastructure. Looking back over the three annual user forums, the basic question has changed from "if" Grids can be used to produce scientific results to "how best" to use the Grid for these ends as well as new business models," said Bob Jones, EGEE project director.

EGEE will hold its next conference, EGEE'08, in Istanbul, Turkey, from 22–26 September (see www.eu-egee.org/egee08). Continuing the theme "building bridges" this conference will provide the perfect opportunity for both the business and academic sectors to network with the EGEE communities, collaborating projects, developers and decision makers alike, to realize the vision of a sustainable, interoperable European Grid.

A new editor for the CNL

CERN Computer Newsletter 213 July - September 1993 1

I. General

I.1 A new Editor for the CNL

Tony Osborne CN/AS

After four years, the second longest period in recent memory, the editorship of the CNL will now pass from Bernd Pollermann CN/AS to Nicole Crémel CN/AS.

Some of the achievements of this period have been: composition of the CNL in SGML/BOOKMASTER subsequently followed by a move to L^AT_EX, automated procedures for article submission, the creation of X_FIND versions and its availability in the World Wide Web information scheme. A readership poll confirmed the need for the CNL in general and, specifically, the desirability to keep a printed version.

Clearly the success of the CNL belongs first and foremost to the contributors and you the readers. The continuing contributions of Michel Goossens, for technical help, and Michael Metcalf, for copy editing, are gratefully acknowledged.

I wish Nicole Crémel the best in her new role and feel sure that she can count on the community of computer users for continued support of the CNL.

The last time an article titled "A new Editor for the CNL" was published was in September 1993 in *CNL* 213 (see above).

Now, after 15 years (the longest period of the *CNL* in recent memory!), we are again announcing a change in the *CNL* editorship, which will pass from myself (Nicole Crémel, IT/UDS) to Natalie Pocock (IT/DI).

I will not detail all the numerous changes in the *CNL* during these last 15 years, just the most recent and important one. In September 2004 we started collaborating with IOP Publishing, the same publisher that produces *CERN Courier*, to perform all the production work. This resulted in a completely new *CNL* layout for our readers.

A recent survey of *CNL* readers shows that there is still a lot of interest in the *CNL*,

and in particular the printed version, which gives CERN users an opportunity to learn about the trends and changes in IT, without having to use a computer screen to access the information.

I have enjoyed being *CNL* editor immensely – which is why I continued for so long! And it is only because I will be taking on more responsibilities in the User Support area that I decided it was time to find someone new.

I can still say today that the success of the *CNL* belongs mainly to the contributors and the readers. I am sure that Natalie will be excellent as the new editor and I wish her all the best, with the continued support of the *CNL* from all the users at CERN.

Nicole Crémel, IT/UDS

Computer security: think before you click!

As a follow-up to the article in the November/December 2007 issue, here is the leaflet on computer security in French. Readers are encouraged to follow

the advice on how to keep your computer safe, because attacks are becoming more frequent and more sophisticated.

The CERN Computer Security Team

Les ordinateurs sont attaqués, même au CERN !



Pour préserver la sécurité de votre ordinateur :

- Attention aux astuces visant à voler votre mot de passe
- Ne pas cliquer sur les liens douteux dans les spams
- Ne pas ouvrir d'emails ou de fichiers attachés inattendus
- Ne pas échanger de fichiers Peer-to-Peer (P2P) (ex. avec BitTorrent)
- Ne pas utiliser les discussions en ligne (IRC - Internet Relay Chat)
- Ne pas télécharger ou installer de logiciels depuis Internet

Information sur la sécurité <http://cern.ch/security/> Règles informatiques <http://cern.ch/ComputingRules/>

Calendar

April

21–25 **WLCG Collaboration Workshop**
Geneva, Switzerland
<http://indico.cern.ch/conferenceDisplay.py?confId=6552>

May

5–9 **HEPiX Spring 2008 Meeting**
Geneva, Switzerland
<http://indico.cern.ch/conferenceDisplay.py?confId=27391>

6–7 EGEE-III Transition Meeting

Geneva, Switzerland
<http://indico.cern.ch/conferenceDisplay.py?confId=30221>

19–22 CCGrid 2008: 8th IEEE International Symposium on Cluster Computing and the Grid

Lyon, France
<http://www.ens-lyon.fr/LIP/RESO/cgrid2008>

21–25 WWW 2008: 17th International World Wide Web Conference

Beijing, China
<http://www2008.org/>

June

2–6 **OGF23: 23rd Open Grid Forum**
Barcelona, Spain
<http://www.ogf.org/OGF23/>

12–13 WLCG CCRC'08 Workshop: WLCG Common Computing Readiness Challenges Workshop

Geneva, Switzerland
<http://indico.cern.ch/conferenceDisplay.py?confId=23563>

17–20 ISC'08: International Supercomputing Conference

Dresden, Germany
<http://www.isc08.org>

30 June – 4 July High Performance Computing and Grids Workshop

Cetraro (Cosenza), Italy
<http://www.hpcc.unical.it/hpc2008/>

August

18–29 **Introduction to High Performance Computing**
Stockholm, Sweden
http://www.pdc.kth.se/systems_support/training/2008/summerschool/

25 August – 5 September CERN School of Computing

Gjøvik, Norway
<http://www.cern.ch/CSC>