

CERN COMPUTER NEWSLETTER


Volume XL, issue 4 September–October 2005

Contents

Editorial

EGEE gears up for phase two

1



Looking to the future: Fabrizio Gagliardi, EGEE project director.

Announcements & news

Computing featured in this month's CERN Courier

2

Desktop Support services need your feedback

2

Stores fulfil computing orders

3

Helpdesk takes responsibility for AIS Support

3

Desktop computing

Know the best way of accessing your mailbox from outside CERN

4

LCG news

Tier-1 networking: a US viewpoint

6

Academia Sinica drives e-science in Asia-Pacific

6

Tsai moves from the South East to the Far East

7

Conference report

International HEPiX meeting promotes greater collaboration

8

Information corner

10

Editors Nicole Crémel and Hannelore Hämmerle, CERN IT Department, 1211 Geneva 23, Switzerland. E-mail: cnl.editor@cern.ch. Fax: +41 (22) 7677155. Web: cerncourier.com (link CNL).

Advisory board Wolfgang von Rüden (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 Institute of Physics Publishing

Dirac House, Temple Back, Bristol BS1 6BE, UK. Tel: +44 (0)117 929 7481. E-mail: jo.nicholas@iop.org. Fax: +44 (0)117 920 0733. Web: iop.org.

Published by CERN IT Department



©2005 CERN

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

EGEE gears up for phase two

In September, CERN and more than 80 institutional partners in Europe, North America and Asia will submit a proposal to the European Union 6th Framework Programme (EU FP6) for a second phase of the Enabling Grids for E-science project (EGEE). The project's first two-year phase finishes in March 2006. Project director Fabrizio Gagliardi spoke to *CNL*.

What do you consider to be the main achievements of EGEE so far?

EGEE has expanded rapidly to support several different scientific and industrial application domains, from biomedical to geophysics, with more than 20 specific applications running on the Grid. This is much faster growth than we anticipated when the project was launched in April last year.

EGEE has also made a major contribution to the operations of the LHC Computing Grid (LCG) project during the past year and a half. The relationship between the two projects is symbiotic, since LCG provides an ideal flagship application for EGEE, producing groundbreaking results which inspire other scientific communities to get involved. EGEE, through its 70 partners, should be credited with helping many smaller sites to join the production Grid infrastructure used for LCG. The growth of the number of sites – now more than 160 – is well beyond the targets LCG had set prior to the EGEE project.

The impact of EGEE can also be seen in the number of Grid project proposals funded in the recent EU FP6 application round that refer directly to EGEE – 16 projects plan to use the EGEE

infrastructure, and CERN will be involved in eight of these.

What have been the most challenging aspects of running such a large project, in your view?

Well, you might expect that getting 70 partners to follow a common plan with tight deadlines for deliverables would present an insurmountable challenge. But perhaps the most surprising aspect so far, from my perspective, is that things have gone largely as planned.

According to the independent project review last February, almost all partners have performed well. And a survey of the partners themselves, to which 50 out of 70 responded, revealed that 98% rated the management of the project as either good or excellent. This result surely reflects that EGEE has benefited from lessons we learned from its predecessor, the European DataGrid project (EDG). Probably the biggest management challenge has been handling the project's progressive shift of scope – from European to truly global during the past two years. The change of the name of the project – dropping the words “in Europe” – reflects this. This sort of development requires careful negotiation with all stakeholders.

Another important challenge is balancing the interests of the high-energy physics community, where there are pressing deadlines for LCG, with the need to support other, less experienced scientific communities. Those involved in the LHC experiments are naturally reluctant to stake their entire computing model on a technology that is still at an

Announcements & news

early stage, and they need to see practical results and be constantly reassured of progress. This requires considerable effort.

What will be the main focus of the next phase of EGEE, assuming it goes ahead as planned?

I think the global scope of the project will be a main theme. Some of the recently funded EU FP6 projects are specifically targeted towards extending the reach of the EGEE infrastructure, for example to the Baltic states, the Mediterranean basin, China and Latin America. Another key activity will be promoting interoperability between different types of Grid, and the associated issue of Grid standards. The Grid is often compared to the Web, and there are indeed similarities. But the Web started out here at CERN with one well defined standard, and Tim Berners-Lee and others put a great deal of effort into ensuring that standard was maintained. The Grid has evolved differently, with multiple standards, and indeed multiple meanings attached to the same word.

I think we are at a crossroads where organizations like the Global Grid Forum are pushing for common standards. EGEE is well placed to become the prime practical implementation of an international, multi-institutional Grid, and able to set standards by example.

You have been working on EU Grid projects for a number of years, notably leading the EDG and EGEE projects. What are your plans now?
For both professional and

personal reasons, I have decided to take leave of the demanding role of EGEE project director, and move to Microsoft, where I will be EMEA technical director for the Technical Computing Initiative in Europe. Bob Jones, the current technical director of EGEE, will be replacing me next November, and I am confident this will ensure smooth continuity for the project.

At Microsoft, I expect to learn a lot about how industry perceives the future of Grid technology and distributed computing in general. I am impressed by the interest that companies like Microsoft are showing in scientific and technical computing and hope to make an important contribution towards a wider adoption of Grid technology in this field. Industrial interest in Grids is high, and industry must play a key role in promoting Grid standards and providing Grid services – this is not just a job for the academic community. Also, together with the IT Department management, I initiated efforts to look beyond EGEE, and plan a longer-term and more stable funding for an international Grid infrastructure. This would be along the lines of the GÉANT network for high-speed communications.

It might seem premature to discuss what comes after a second phase of EGEE before it has even started, but experience shows that the funding system operates on these timescales. I am looking forward to seeing this effort come to fruition in the coming years, and hope that CERN can continue to play a leading role in this field.

Computing featured in this month's CERN Courier

The articles listed below appear in the September 2005 issue of *CERN Courier*. Full-text articles and the rest of the issue's contents are available at www.cerncourier.com.

Computing News

● CrossGrid project concludes

The project that developed Grid-enabled solutions closes.

● PHENIX experiment uses Grid to transfer 270 TB of data to Japan RHIC uses Grid for new data.

● W3C backs XML Key Management System

Security framework gets new interface through XKMS.

● GÉANT2 set to boost research in Europe

Pan-European communications infrastructure steps forward.

● US launches weekly online newsletter

How to get the free newsletter, *Science Grid This Week*.


● New service 'deep-mines' databases

EVL Xtra cross-searches online collections for free.

● QCDOC computers study quarks New 10 teraflop computers are commissioned for QCD.

● Scientific Linux 4.0 offers new features

Operating system is upgraded.



● High-performance APE computers are introduced to US

European supercomputing technology is shown to the US.

● Latest release of Grid middleware runs on Mac OS X NMI-R7 is released.

● WAH compression speeds up searches

New technique seeks out rare events in particle collisions.

● US inaugurates the Open Science Grid

Experiments get under way as OSG opens for business.

IT products and calendar

Feature article


● Running the Grid on lite

Middleware helps to make the Grid more user-friendly.

Desktop Support services need your feedback

A very simple User Satisfaction Survey is now included in the Web interface for end-users when a ticket (generated by the Problem Tracking system used in IT) is set to "fixed" by the support staff.

At the end of the mail that is sent to the user when a ticket is fixed there is a link to a Web page starting with: <http://cern.ch/helpdesk/problem/CT...> (the exact message is shown in the picture). By selecting this link you will open a Web page with



all the details concerning the problem and its treatment, and at the top of this page you are now offered the opportunity to rate the way in which your problem has been treated. Five

levels (very poor, poor, average, satisfactory, very satisfactory) are proposed and all you need to do is to select a level.

We thank all the users who take a few minutes to open this

link and provide some quick feedback. It is also possible for users to reply to the mail and give more specific comments – especially useful when their level of satisfaction is poor or very poor. All this will help us to improve the Helpdesk and Desktop Support services.

If you have any other comments please feel free to e-mail the manager on duty at mod@cern.ch.

Nicole Crémel and Roger Woolnough, IT/UDS

Stores fulfil computing orders

The activity of the IT Department in charge of computing-hardware procurement was transferred to the CERN Stores in June 2005. The well known PC-Shop in building 513 has closed and all orders of standard computing material should now be made directly to the Stores, using the EDH Material Request.

This article describes some of the new procedures in place and suggests how users can profit from the improved service.

Ordering new material

To order new material, fill in the EDH Material Request available at <https://edh.cern.ch/Document/MAG/>.

Users can choose from any of the items available in the Stores catalogue at <https://edh.cern.ch/edhcat/Browser>. The Computing Hardware section begins at SCEM with the prefix 80. Note that some equipment is still published with a SCEM number starting with 12 but this should change soon and all computing items will be in section 80 of the catalogue.

The CERN Stores are only responsible for the delivery of computing hardware to the user's desk, and within a reasonable amount of time. While the IT Department does not expect to get involved in the logistics of hardware handling (provision, ordering, transport, billing, etc), it remains responsible for installing and supporting these items.

Installing and using new material

Having bought computing equipment from the store, users should e-mail Helpdesk@cern.ch

for a support technician to install the hardware.

As an alternative each computing item on sale from the Stores will have a data sheet with instructions on how to install and use the device. The user can follow the instructions to install the device and call the Helpdesk in case of difficulties.

Once an order is placed, the material is immediately shipped to the user's desk using the standard Store logistics (available twice a day). The user will receive with the material all the information about opening a ticket (work order) to arrange for the device to be installed at their desk by Desktop Support, if necessary.

Ordering non-standard material

The procedure to order non-standard material is unchanged. The user should fill in the EDH Internal Purchase Requisition form at <https://edh.cern.ch/Document/DAL/>.

In this case, the Stores, IT and the Helpdesk will not be involved in the purchasing, installation, support or warranty-handling processes.

Troubleshooting hardware and software problems

Hardware or software problems should be diagnosed by a support technician. Users should e-mail Helpdesk@cern.ch.

The technician will try to solve any software-related problem, or will completely re-install the software. If the technician identifies a hardware problem, he will advise the user on how to return the material to the Stores for repair (see below).

Material return

Material ordered in error (in "as new" condition), or faulty material needing repair under warranty, should be returned to the CERN Stores using the form at www.cern.ch/MaterialReturn.

To ensure quick handling of the return, and for repair under warranty, the EDH Material Request number and a budget code must be used.

Returning hardware that has been modified, tampered with or damaged for repair under warranty may incur repair costs. Costs may be also generated if the problem stems from incorrect software configuration, or if hardware is out of warranty.

To avoid misunderstandings concerning repair costs, users should consult a support technician before returning material to the Stores whenever the diagnostic is not evident.

Modifications to the Computing Hardware catalogue

The IT Department, in collaboration with user representatives, is responsible for defining the functional specification needed to satisfy the majority of users and for deciding what items should be in the Stores catalogue.

However, services wishing to add their own equipment to the catalogue may do so, provided they ensure technical support and are able to define after-sales return processes.

Improved services

The motivation behind this change has been to reduce the waiting time experienced by users when ordering computing

material. In fact, there will be no boxes in transit from the IT buildings and the standard internal transport will use the normal distribution used for CERN Stores material.

Users needing an item that is in stock at CERN can even get it immediately if they collect it personally from the Stores. For example, the standard CERN PCs are planned to be kept in stock.

We foresee that at any one time the CERN Stores will offer more than one desktop model and a few portable models. Multiple desktop models are necessary to address conflicting specifications between simple and silent desktop computers and high-performance engineering workstations. Similarly, different models may also be necessary to address different needs of the Linux or Windows communities. The IT Department will still be responsible for coordinating the certification process, software compliance analysis and tests from Windows and Linux support teams.

The Computing catalogue in the Stores will also propose a set of accessories and peripherals. These include monitors, memory upgrades, additional hard disks, various USB or PCMCIA devices (memory sticks, wireless cards, Bluetooth adapters, etc), CD or DVD writers, additional power supplies for notebooks, docking stations, replacement or second batteries, and carrying cases. As a general strategy, USB devices will be favoured in the future.

Christian Isnard and Alberto Pace, IT/IS

Helpdesk takes responsibility for AIS Support

As announced in *CNL* January–March 2005, since the start of this year the Administrative Information Services (AIS) Helpdesk has been staffed by Serco, the same firm already responsible for desktop support (see "Desktop Support moves to IT/UDS" at www.cern.ch/cnlart/2005/001/6).

As of 1 July 2005, following a

six-month training period for all the Helpdesk staff, AIS Support is now fully integrated into the general Helpdesk configuration. It is part of the general Desktop Support contract (currently outsourced to Serco, which has held the contract since 1 July 2004).

There is no longer a specific "AIS Support" contract, and all

Helpdesk staff will now reply to AIS questions, or forward them to the appropriate expert when needed. Users can still use the old entry AIS.Support@cern.ch, or phone 79933. Although this AIS-specific number will still work, and calls will be automatically redirected to 78888, 79933 may be disconnected in the future.

Please remember that if you experience problems with Remedy tickets when they are passing through the Helpdesk, you may at any time contact the IT manager on duty (mod@cern.ch), or e-mail user.relations@cern.ch with any comments you have on the service.

Nicole Crémel and Roger Woolnough, IT/UDS

Know the best way of accessing your mailbox from outside CERN

People often ask what is the recommended method to access their CERN mailbox from outside. Until recently the CERN Mail team recommended using Outlook Web Access (OWA). This was the simplest solution, since it needed no special configuration, and it allowed access to their mailbox, contacts and calendar. However, OWA should not be considered a day-

to-day mail reader, as it is more of a casual tool.

Another solution has been to log on to the CERN network through authentication (e.g. virtual private network – VPN) and use the same mail clients as when working at CERN. However, given the security implications, users are now formally discouraged from using VPN as a general solution to

access their mailbox.

It is also possible to configure your local mail client to use IMAP/SMTP protocol to access the CERN mailbox, with the outgoing (SMTP) server name set to mmm.cern.ch. However, this is not recommended for Outlook XP users as the settings are complex and do not always work. For Outlook users not satisfied with OWA, and who

need to access a more powerful mail client from outside CERN, the current recommended solution is to use Outlook 2003 and configure the client to use “RPC over HTTP”, as will be explained in this article.

Technical background and requirements


Outlook 2003 uses the MAPI protocol to connect to the mail

Update the “Microsoft Exchange Server” field

Input your NICE username in the “User Name” field

Check “Use Cached Exchange Mode”

Do not click “Check Name”!



Click “OK”

Fig. 1, top: the correct format for configuring Outlook 2003 for external access. Fig. 2, bottom: users will see this message during the process.

server, through a port 135 RPC call. However, port 135 is quite often closed by well known ISPs for security reasons. RPC over HTTP embeds the MAPI RPC protocol into the HTTP layer, using standard HTTP's port 443, usually open everywhere. This allows Outlook 2003 to connect to CERN Mail servers from anywhere without VPN. To use Outlook 2003 with RPC over HTTP the requirements are:

- Your operating system must be Windows XP SP2.
- You need to install Outlook 2003 (part of Office 2003).

Outlook 2003 configuration

Next you will need to configure Outlook 2003 to use RPC over

HTTP, by modifying your Outlook profile.

Step 1

- Click the Mail icon in the Control Panel, then:
- Show profiles – select your profile and click “Properties”.
- Click “E-mail Accounts”.
- Select “View or change existing e-mail accounts”.
- Select “Microsoft Exchange Server” and click “Change”.

Step 2

- Type “cernxchg01.cern.ch” in the “Microsoft Exchange Server” field (see figure 1).
- Check “Use Cached Exchange Mode”.
- Type your NICE username in

the “User Name” field.

- Do not click “Check Name”!

Step 3

- Click “More Settings”.

The next four steps will occur only if you are connecting from outside CERN.

Step 4

- In a few seconds you will get an error message (see figure 2).
- Click “OK”. A window will appear (see figure 3).

Step 5

- Click “Cancel”.
- The advanced settings will eventually appear. Select the “Connection” tab (see figure 4).

Step 6

- Check “Connect to my Exchange mailbox using HTTP”.
- Click “Exchange Proxy Settings...” (see figure 5).

Step 7

- Type “mmm.cern.ch” in the URL field.
- Check “On fast networks...” and “On slow networks...”.
- Select “NTLM Authentication”.
- Click “OK”, then “OK”.
- Click “Next”, then “Finish”, then “OK”.

Run Outlook 2003

- When asked, type in your Mail username (in the form CERN\login) and your password.

Emmanuel Ormancey, IT/IS





Fig. 3, top left: click “Cancel”. Fig. 4, top right: select the “Connection” tab to produce this box. Fig. 5, bottom: the final stage of configuration.

Tier-1 networking: a US viewpoint

US-based physicists are preparing to exploit LHC data. The LHC programme requires substantial computing, which is supplied in a tiered model. Apart from the Tier-0 Centre at CERN, LHC computing is organized into Tier-1, Tier-2 and Tier-3 Centres. Therefore most computational resources are based away from CERN. These remote resources are essential for exploiting the LHC data. In the US, LHC computing has generated an unprecedented reliance on transatlantic and US national networking. There have been considerable early successes as preparations continue.

LHC computing models allow data to move over global distances. For example, the experiment's computing models require the distributed Tier-1 Centres to cache data sets generated at CERN and other Tier-1 Centres. Tier-2 Centres can download analysis data sets from any appropriate Tier-1 Centre. Additionally, groups of physicists on different continents can replicate data sets of many terabytes interactively to work on a specific physics topic.



CERN data received at the CMS Tier-1 production dCache storage system.

What's new is the level of reliance and the distances involved. There is at least one intrinsic problem using networks at long distances. To obtain full throughput in any network, at least (bandwidth $\times$ delay) bytes must be kept in "in flight" within the network itself. To see what this means for global LHC computing, consider that transatlantic networks are currently 10 Gbit, and the speed of light constrains latencies to be more than 100 ms. Therefore, a full 10 Gbit channel implies more than 1 Gbit of data (or approximately 120 MB) in flight. Today, 120 MB is acceptable

buffering for a single channel for network equipment, or even host computers. Aside from the intrinsic problem of exploiting any network at these rates and distances, common TCP protocol stacks are very sensitive to any kind of data-packet loss.

The US physics community anticipated this and has worked on these constraints since early on, and has been demonstrating the feasibility of the global LHC model. The HEP community, via the DATATAG project and the Caltech-led US LHCnet, has been involved in provisioning high-quality transatlantic networking from the Starlight open optical

exchange in Chicago directly to CERN. Fermilab procured and lit a dark fibre to Starlight, which in part allowed the large-scale US CMS Tier-1 Centre to use the high-rate network. We paid particular attention to avoiding packet loss and packet re-ordering in every component of the path, including the last-mile networks within the facilities. This early provisioning allowed full integration into the production storage systems, which is not a trivial activity.

The figure demonstrates the quality and care of this integration. It shows the CMS Tier-1 production dCache storage system at Fermilab ingesting tens of terabytes from CERN each day for many days. These transfers were interrupted only by use of the link for bandwidth challenges for the SC04 computing exhibition.

The plot also shows that more than 1 TB/h has been attainable for practical systems for some time. It also illustrates the technical potential for interactive transfer of terabytes of data for ad hoc physics analysis over global distances.

Don Petravick, Fermilab

Academia Sinica drives e-science in Asia-Pacific

The Academia Sinica Grid Computing Centre (ASGC) in Taipei is currently the only LCG Tier-1 Centre in the Asia-Pacific area, with 400 KSI2K computing capacity, 50 TB disk space and a 35 TB tape library dedicated to the LCG. Since 2004, Academia Sinica has provided the services of a regional operation centre (ROC), site monitoring, virtual-organization (VO) support, middleware deployment, certificate authority (CA) and global Grid-user support (GGUS – mainly first-line support and FAQs). The centre supports not only Tier-2 sites locally, but also Grid operations in South Korea, Singapore and other Asia-Pacific countries that are not supported by other Tier-1 sites.

To support service and data challenges, a maximum



Grid tutorial at the Academia Sinica Grid Computing Centre.

1.6 Gbit/s transmission rate was achieved in the 2 Gbit network bandwidth between CERN and ASGC in June 2005. During the

CMS service challenge, ASGC received 20 TB of data from CERN at an average rate of 56 Mbit/s from 14 July to 14 August. The

ASGC Tier-1 Centre provided 12% of the LCG-2 computing jobs, second only to the 14% of CERN in the ATLAS data challenge in 2004. Academia Sinica will work closely with Tokyo University and other Tier-2 sites in this region for the ATLAS and CMS service challenges in the near future.

ASGC is engaging in collaboration and sharing of information by taking advantage of e-science applications in the Asia-Pacific area. ASGC is also working with different partners to help form and support application-driven e-science communities in the Asia-Pacific region, to improve the next-generation research infrastructure and build up the e-science applications. Hosting the International Symposium on Grid Computing (ISGC) since



The ASGC team with director Simon Lin (front row, second from right).



The data centre at the ASGC with 400 KSI2K computing capacity.

2003, Academia Sinica provides a platform for the co-operation of the Asia-Pacific region, Europe and America.

In addition to applications from high-energy physics, such as the ATLAS and CMS applications running locally, the

centre also supports the biomedical VO in EGEE. Local VO services for users of the TWGrid include bio-informatics, atmospheric-sciences applications, earth sciences and digital libraries. Another VO for general applications in the

Asia-Pacific region is currently being constructed.

In total, 30 staff take care of Grid technology, application development, operations, deployment and support for the users. For phase two of the EGEE project, ASGC will extend its

partnership to the areas of middleware re-engineering and integration (JRA1), training (NA3) and outreach (NA2), while keeping a strong presence in operations (SA1) and application support (NA4).

Eric Yen, ASGC

Tsai moves from the South East to the Far East

Min Tsai is an excellent example of how Grid computing helps international collaboration. In September he will move back to Taipei, leaving CERN where he has been acting as liaison for the South-Eastern European Grid-enabled e-Infrastructure Development (SEE-GRID) project. He will take up the role of deputy manager of the first Asian Regional Operations Centre (ROC) for the Enabling Grids for E-science (EGEE) project at Academia Sinica in Taipei.

Tsai's path to Grid computing is unusual. Neither a computer scientist nor a physicist, he studied mechanical engineering in the US then worked in industry in telecommunications and networking. Hired because of his networking background, he arrived at Academia Sinica just over a year and a half ago to help set up Grid computing there. His first assignment was to go to CERN and learn hands-on about Grid operations by working in the Grid Deployment group, and after six months he was offered the liaison position for SEE-GRID.

Providing deployment support for the SEE-GRID partners has proved an excellent way for him to learn about all aspects of the Grid. When he didn't know the solution to a problem himself,



Heading east: Min Tsai is moving from CERN to EGEE, Taipei, where he will also support Grid partners in Singapore and South Korea.

someone in his vicinity at CERN could usually help him, providing a very fast turnaround. Partly as a result of this, the SEE-GRID project has been able to rapidly deploy production services at sites in the 10 participating countries. Tsai emphasizes the importance of the SEE-GRID

partners themselves in supporting each other. In particular countries already belonging to an EGEE federation, such as Greece, Romania and Bulgaria, have provided support to new partners such as Croatia, Macedonia and Turkey.

The Asian ROC was officially

launched this spring and has a staff of seven with support from other teams. Tsai's responsibilities as deputy to Chiang Gen-Tao include supporting partners in Singapore, South Korea and locally. There are discussions with Japan about joining this ROC, and other countries in the region are rapidly getting involved in Grids, so his SEE-GRID liaison experience should prove highly valuable.

One of his goals this autumn will be to establish a Core Infrastructure Centre (CIC) in Taipei. This means taking responsibility for monitoring, problem tracking and the operation of core Grid services. Other goals include improving the availability of Grid services in the Asia-Pacific region – which is critical for user acceptance – and increasing communication and coordination with regional partners.

Tsai underlines the importance of face-to-face meetings, such as the quarterly meetings in SEE-GRID. He said, "Telephones and e-mail are very useful, but really getting to know the people involved at different sites makes a big difference when it comes to helping them with deployment or production issues."

François Grey, IT/DI, CERN

International HEPiX meeting promotes greater collaboration



Eyes down: the audience at the latest HEPiX meeting, hosted by the Grid Computing Centre Karlsruhe and the Institute for Scientific Computing.

HEPiX is a global organization of service managers and support staff that provides computing facilities for the high-energy and nuclear physics (HENP) community. All operating systems used by HENP are covered, including Unix, Linux and Windows. HEPiX holds regular meetings for participants to present recent work and future plans, and to share experiences.

The latest HEPiX meeting was hosted by the Grid Computing Centre Karlsruhe (GridKa) and the Institute for Scientific Computing (IWR) at Forschungszentrum Karlsruhe (FZK) on 9–13 May 2005. The meeting was well organized and the room was functional, sited just outside the FZK main gate and well equipped for portables. Sessions were videoed and video-conferences were arranged for three of the sessions.

About 100 people attended the meeting, and the usual US sites were well represented; we think that more HEP sites are coming to see the value of attending HEPiX meetings, and that sites are sending more people. Also the workshop on batch systems at the end of the week attracted a number of new attendees who clearly added

value to the discussions.

Prof. Dr Reinhard Maschuw, the director of the laboratory, opened the meeting. He said that five years ago he would have been surprised to see us because FZK is not known as an HEP site; rather it is an interdisciplinary centre running research programmes in various fields, but not specifically HEP. FZK's range of interests shares the need for good computational facilities, hence the recent interest in supporting Grid computing for German groups participating in the LHC. They are a Tier-1 Centre for the LHC Computing Grid (LCG), and participate in the various LCG Service and Data Challenges. Those at FZK were pleased to host the meeting, and that HEPiX recognizes their participation in the world of HEP computing.

Meeting highlights

The following are a selection of highlights from the meeting. The theme of the week seems to have been "collaboration".

- Guy Wormser, chairman of the International High-Energy Physics Computing-Coordination Committee (IHEPCCC), asked if HEPiX would be willing to act as technical advisers to IHEPCCC on specific questions where it has

expertise. Two immediate examples were given: the status of Linux in HEP, and the idea of a Virtual Organization for HEP physicists. The meeting agreed with an earlier board decision to accept this role with certain caveats, and actions were agreed to respond to IHEPCCC's first two questions.

- There were reports and discussions on the most recent successful HEP collaboration – the production, distribution and widespread acceptance of Scientific Linux. Discussions focused on which versions would need support in the years up to the start-up of LHC.

- There was a brief discussion of a limited collaboration being set up, on the suggestion of Jos van Wezel, to discuss the data-storage needs of Tier-2 sites. This should happen over the summer, chaired by Roger Jones of Lancaster, and be reported to the next HEPiX in the autumn.

- An open discussion on the possibilities of a HEP collaboration on monitoring tools ensued. This followed a CERN talk on Lemon, showing what it could do, and a Fermilab talk on the successful experiences of NGOP, each of which monitor in-production node clusters totalling more

than 2000 nodes. SLAC invested some effort in Nagios, but they found serious drawbacks (i.e. scalability limits).

- In most sites we seem to be focusing on a finite number of tools for system-administration tasks. These include ROCKS (public domain) and Quattor (CERN) for system installation, Ganglia and Nagios (both public domain) for monitoring and management tasks, dCache (DESY/Fermilab) for data caching, and SpamAssassin for fighting spam mail.

- The workshop at the end of the week, organized by Tony Cass, covered all the batch systems in the major, and some not so major, HEP sites. It raised some issues concerning the evident gap between local job scheduling and Grid scheduling. It was unfortunate that there was no Grid-scheduler author present to listen to the issues of Grid scheduling as expressed by the site administrators who have to make grid scheduling co-exist with local environments. Nevertheless there was lively discussion and plenty of ideas to reflect upon.

- For a full list of the talks, presentations and workshops at HEPiX, see p9.

Alan Silverman, IT/DI

Conference report



Presentation on Scientific Linux, a common Linux solution for HEP sites.

Workshop on Batch Systems (chair: Tony Cass, CERN)

A key aim of this workshop was to enhance communications between users and developers of local-resource scheduling systems and Grid-level resource scheduling. At present, HEP sites use a variety of local-resource schedulers to maintain fine-grained control over heterogeneous local resources according to site policies. At the Grid level, however, the scheduling often treats all resources at a site as homogeneous and equally available to all Virtual Organizations (VOs). Can we improve Grid-level scheduling to reflect the fine-grained level of control locally? If so, how? If not, what will be the impact on local operations in the future?

How are Local Batch Schedulers used at HEP Sites?

Subject	Speaker
Introduction	Tony Cass
Batch Scheduling at SLAC (LSF)	Neal Adams
Batch Scheduling at RAL (Torque/Maui)	Martin Bly
Batch Scheduling at LESC (SGE)	David McBride
Batch Scheduling at JLAB	Sandy Philpott
Batch Scheduling at IN2P3 (BQS)	Bernard Chambon
Batch Scheduling at FNAL	Ken Schumacher
Batch Scheduling at DESY (SGE)	Andreas Haupt
Batch Scheduling at CERN (LSF)	Tim Bell

Local and Grid Schedulers: Status and Plans

Subject	Speaker	Organization
Some Thoughts on How (Not) to Organise Local and Grid Scheduling	Jeff Templon	NIKHEF
BQS for the Grid	Yves Fouilhé	IN2P3
BQS and the GRID: Problems and Solutions	Julien Devémy	IN2P3
LoadLeveler	Roland Kunz	IBM
LSF	Christopher Smith	Platform Computing
PBS	Ralf Eichmann	Altair
Sun Grid Engine	Andreas Haas	Sun
Condor	Todd Tannenbaum	University of Wisconsin
EGEE Batch Local ASCII Helper (BLAHP)	David Rebatto	INFN

Developing a Common Batch Interface

GLUE Schema: Status & Plans	Laurence Field	CERN
A Proposal for Standardising the Working Environment for a LCG/EGEE Job	David Bouvet	IN2P3

List of talks and speakers

We include here a full list of all the talks and speakers at HEPIX. For a summary of all talks and site reports, please see <http://iwrwww1.fzk.de/hepix/hepix.html>, where the overheads and videos of almost all talks can be consulted online. On the same website you can also read the detailed trip report written by Alan Silverman at: http://iwrwww1.fzk.de/hepix/HEPIX_Spring_2005_FZK.pdf.

Site reports	Speaker	Organization
GridKa	Holger Marten	FZK
CASPUR	Andrei Maslennikov	CASPUR
TRIUMF	Corrie Kost	TRIUMF
SLAC	Chuck Boeheim	SLAC
FNAL	Lisa Giacchetti	FNAL
CERN	Helge Meinhard	CERN
RAL	Martin Bly	RAL
GSI	Walter Schoen	GSI
LAL	Michel Jouvin	LAL/IN2P3
BNL	Chris Hollowell	BNL
INFN Tier 1	Andrea Chierici	INFN

Talks and presentations	Speaker	Organization
Lemon Monitoring	Harry Renshall	CERN
Host and Service Monitoring at SLAC	Alf Wachsmann	SLAC
Experiences Building a Distributed Monitoring System	Marc Mengel	FNAL
Security Update	Bob Cowles	Stanford
Enstore	Michael Zalokar	FNAL
Fresh Results on Lustre, AFS...	Andrei Maslennikov	CASPUR
USCMS Update on Use of Ibrix Fusion FS	Lisa Giacchetti	FNAL
Collaboration IHEP-HEPIX	Guy Wormser	IN2P3
Experience Deploying an Xrootd Data Serving Infrastructure at RAL	Chris Brew	RAL
FNAL SATA Experiences	Lisa Giacchetti	FNAL
Scientific Linux, Status Report	Jack Schmidt	FNAL
CERN Linux, Status Report	Jan Iven	CERN
Fermi Farms Opteron Evaluation	Ken Schumacher	FNAL
Anti Spam Activities at GARR	Michele Michelotto	INFN
SPAM – Is There Still a Problem?	Wolfgang Friebel	DESY
FNAL Central Email Systems	Marc Mengel	FNAL
WIKI – a Collaborative Documentation Tool	Christopher Huhn	GSI
Managing CERN Desktops with SMS 2003	Michel Christaller	CERN
Plone: FNAL Experiences with Content Management	Marc Mengel	FNAL
Cross Platform Single-Sign-On using Client Certificates	Alberto Pace, Emmanuel Ormancey	CERN
DESY Registry – Cross-Platform User Administration	Dirk Jahnke-Zumbusch	DESY
A New Graphical Management Tool for PBS	Michael Gehle	FZK
USATLAS Tier 1 Facility and Service Challenge	Dantong Yu	BNL
LCG Tier 2 at LAL and DAPNIA	Michel Jouvin	IN2P3
SC-Overview	Sophie Lemaitre	CERN
LHC Service Challenges: Tier 2 Planning	Sophie Lemaitre	CERN
LCG File Catalogs	Jean Philippe Baud	CERN
File Transfer Service and Software	Sophie Lemaitre	CERN
Disk Pool Manager	Jean Philippe Baud	CERN

German students learn from CERN IT Department

The University of Applied Sciences (FH) in Wiesbaden, Germany, has started to work with the CERN IT Department to give computer-science students the opportunity to do their final-year thesis in a big laboratory.

Starting with two students this year, Prof. Dr Detlef Richeter hopes to expand this collaboration in the future. "In a Europe that is becoming more open, it will be important for the students to gain international experience and to prepare them for mobility in their later career," he said.

The first student, Matthias Bräger, has already been at CERN for several months, and the second, Felix Ehm, is soon to follow. Both have been to India on work placements, and even if CERN is not quite as exotic as India, it is still a very special place. The international atmosphere is one of the reasons why Bräger decided to come to CERN. "It is great to have a small part in such a big research institution and to make a contribution to the larger picture," he said. "Basic research does not create profit, but it tells us more about who we are."

For his thesis, Bräger is working on the CASTOR Hierarchical Storage Management (HSM) system, where he redesigns the daemon for the tape libraries. To access data from a particular tape, the daemon checks what kind of tape drive is needed and its status. The request to mount the tape is put in a queue, which is kept in the online storage.

To avoid loss of this information in the case of server failure, Bräger is backing up the system with a database, which will store every request. He also optimizes the code that assigns tapes to tape drives. So far tapes are only mounted in drives that can both read and write, even if read-only drives are available. The improved daemon is making use of these free capacities, thus contributing to the preparation for the huge amount of data expected from the LHC experiments from 2007.

Hard-drive data could be at risk

We would like to remind all users that data stored on the local disk of their PCs (drive C:\) are not guaranteed. These data are never backed up by a central service and it is the user's responsibility to make a regular back-up of all important data.

If a virus infects a PC, often the only solution is to re-install the system, as requested by the CERN security team. It is then generally too late to back up the hard drive, as files might be infected by the virus. When a request to re-install a system comes to the Helpdesk and Desktop Support, the request is automatically registered with a disclaimer (see box below).

Let us remind you that storing data on local disks is not recommended. Local disk space at CERN is regarded as scratch that will be lost if a system needs to be re-installed. It is assumed that a Linux system runs the AFS client for central storage of permanent user files, or that a Windows system uses the central NICE servers for permanent files.

For storing large files on NICE (Windows) users must use departmental or experimental shares. Shares are backed up on a daily basis. If you need to create such an area just fill in the form at <https://dfsservices.web.cern.ch/DFSservices/>. For data up to 100 GB a valid description from the user is required as to why this extra space is needed. For more than 100 GB a detailed explanation will have to be provided to the IT NICE system administrators who will then decide about the space allocation. The explanation should include a justification from the department head or deputy, or the head of an experiment.

For large amounts of data on Linux, if AFS space is not sufficient, the CASTOR system must be used. For more details on the CASTOR architecture you should refer to the website at www.cern.ch/castor/.

We thank the users for their understanding in this matter.

Nicole Crémel and Roger Woolnough, IT/UDS

● Please note that, by CERN computing rules, users are responsible for the data stored on their hard drive. Back-ups and restores of data from your machine are therefore done on a best-effort basis.

● Veuillez noter que, de part les règles d'utilisation des équipements informatiques au CERN, les utilisateurs sont seuls responsables des données stockées sur leur disque dur. Les sauvegardes et restaurations de données sont donc effectuées sans aucune garantie.

User-support bookshop moves and diversifies

The IT user-support bookshop closed in building 513 on 11 July and moved to the Central Library area where it reopened on 25 July. In future, the choice of books will include physics, engineering and other titles to be decided according to demand, but will of course continue with a large selection of IT titles. We will also

continue to sell IT CDs. The initial selection of physics books for the launch was based on Library input.

Thank you for your continued support in this move, which was stimulated by the heavy demand shown at last year's highly successful Bookfair.

Jutta Megies and Roger Woolnough, IT/UDS; Jens Vigen, DSU/SI

Calendar

September

12–18 **XX International Symposium on Nuclear Electronics & Computing (NEC'2005)** Varna, Bulgaria, http://sunct2.jinr.ru/NEC-2005/first_an.html

18–21 **12th European Parallel Virtual Machine and Message Passing Interface Conference** Sorrento (Naples), Italy, www.pvmpi05.jeanmonnet.unina2.it

26–30 **iGrid 2005** San Diego, CA, US, www.igrid2005.org/

October

10–15 **ICALEPCS 2005** Geneva, Switzerland, <http://icalepcs2005.web.cern.ch/icalepcs2005/>

16–19 **GGF15** Boston, MA, US, www.ggf.org

November

12–18 **Supercomputing** Seattle, WA, US, www.sc05.supercomputing.org

30–2 Dec **IFIP International Conference on Network and Parallel Computing (NPC 2005)** Beijing, China, <http://grid.hust.edu.cn/npc05>

December

5–8 **International Conference on e-Science and Grid Technologies** Melbourne, Australia, www.gridbus.org/escience

7–9 **International Symposium on Parallel Architectures, Algorithms, and Networks (I-SPAN)** Las Vegas, NV, US, www.sigact.acm.org/ispan05/

The deadline for submissions to the next issue of CNL is

26 September

E-mail contributions to cnl.editor@cern.ch

To be e-mailed 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>