CERN COMPUTER NEWSLETTER

Volume 44, issue 1 January–March 2009

Contents

Editor Natalie Pocock, CERN IT Department, 1211 Geneva 23, Switzerland. E-mail cnl.editor@cern.ch. Fax +41 (22) 766 8500. Web cerncourier.com/articles/cnl.

Advisory board Frédéric Hemmer (head of IT Department), Alberto Pace (group leader, Data Management), 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

©2009 CERN

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

IOP Publishing

CERN is host to the UN satellite imagery group

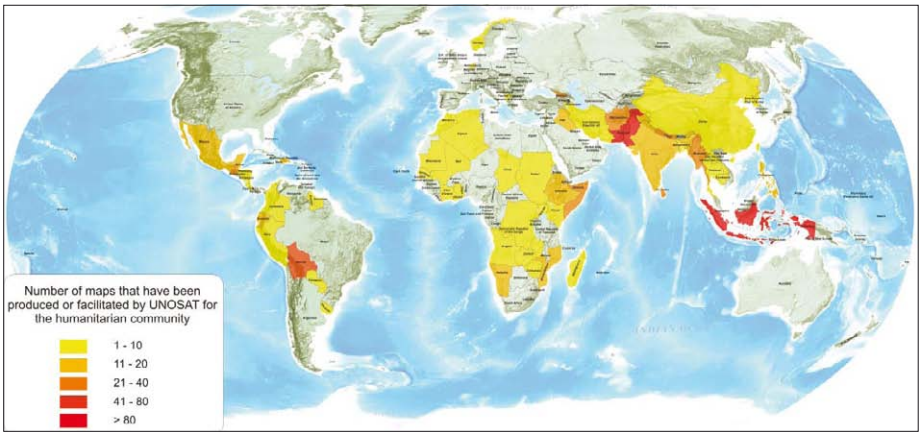

A global snapshot of UNOSAT rapid mapping activities between 2003 and 2008.

CERN has hosted UNOSAT, the Operational Satellite Applications Programme of the United Nations Institute for Training and Research (UNITAR), since 2002. With the renewal of the hosting agreement in December 2008, this relationship has developed, evolving from a hosting arrangement to the beginning of a promising partnership.

UNITAR has a mandate from the UN General Assembly to deliver innovative training and to conduct research on knowledge systems and methodologies. Through adult professional training and technical support, the institute contributes to developing the capacities of tens of thousands of professionals around the world using direct and distance learning.

UNOSAT is a technology-based programme of UNITAR, part of the institute's research department. Created initially to explore the potential of satellite Earth observation, this programme has developed specific mapping and analysis services that are used by various UN agencies and by national experts worldwide. UNOSAT's mission is to deliver integrated satellite-based solutions for human security, peace and socio-economic development, and the most important goal is to make satellite solutions and geographic information easily accessible to an increasing number of UN and national experts who work with geographic information systems (GIS).

The UNOSAT team combines the experience of satellite imagery analysts, database programmers, and geographic information experts with that of fieldworkers, geologists and development experts. This unique set of skills gives the UNOSAT team the ability to understand the needs of a variety of international and national users and to provide them with suitable solutions anywhere and anytime. Anywhere because thanks to CERN/IT support, UNOSAT is able to handle and store large amounts of data and transfer maps right to the middle of an ongoing humanitarian crisis via the Web; anytime because UNOSAT is possibly the only UN programme available for work 24 hours a day every day of the year.

In simple terms, UNOSAT acquires and processes satellite data to produce and deliver information, analysis and observations to be used by the UN or national entities for emergency response, to assess the impact of a disaster or conflict, or to plan sustainable development in the face of climate change. The main difference between this programme and other UN undertakings is that UNOSAT uses high-end technology to develop innovative solutions.

One of these innovations was the creation in 2003 of a new humanitarian

Editorial

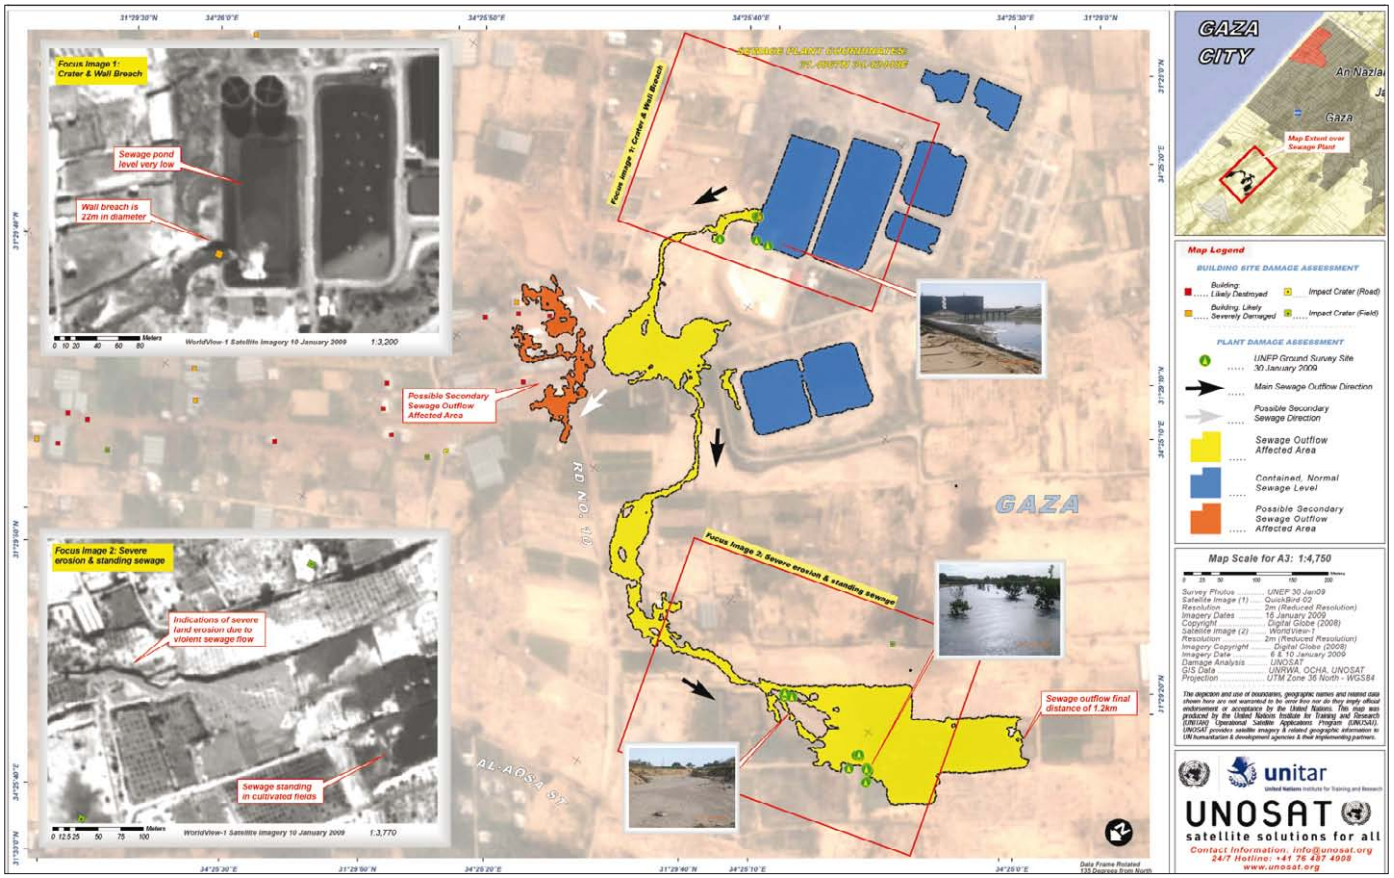

UNOSAT satellite mapping and analysis of a sewage treatment plant in Gaza. Areas of damage and sewage outflow are detected.

rapid mapping service that is today fully developed and has been used in more than 100 major disasters and conflict situations, and has produced more than 900 satellite-derived analyses and maps.

This work implies the rapid acquisition and processing of satellite imagery and data for the creation of map and GIS layers which are then used by the headquarters of UN agencies to make decisions, and in the field during an emergency response to coordinate rescue teams and assess the impact of a given emergency. This type of map was of great use in the aftermath of the Asian Tsunami for example, and in responding to the Pakistan earthquake in 2005. Similar maps were also used to monitor the impact of the conflict between Israel and the Hezbollah in Southern Lebanon, and most recently during the Middle East crisis in Gaza.

There are tens of less publicized crises every year in which the UN is involved because of humanitarian consequences on thousands of innocent civilians in developing countries. UNOSAT supports the work of relief workers and NGO volunteers with timely and accurate analysis of a situation on the ground, and responds to requests from the field for particular geographic information.

The work of UNOSAT is not only about emergencies, although the maps you can see on the Web all refer to humanitarian assistance. This publication policy is linked to the need for humanitarian workers to access maps from various field locations by connecting via Internet or satellite telecommunications to download the maps prepared by UNOSAT at CERN. Yet a large number of maps and analyses are not publicly available on the UNOSAT website because they are part of project activities done in partnership with UN agencies such as the UN Development Programme, the International Organization for Migration, and the World Health Organization.

Once an emergency is over, the work of the UN continues with assistance to governments in rehabilitation and reconstruction. UNOSAT remains engaged beyond the emergency phase by supporting early recovery activities that are undertaken to help local populations get back to normal after a disaster or a conflict. Satellites can be helpful in these circumstances: think of the work required to reconstruct an entire cadastre, for example, without appropriate geographic information, or to plan the rehabilitation of road and rail networks without accurate information on the extent of damage suffered.

UNOSAT experience in mapping and analysis and its innovative methodologies are regularly transferred to the outside world thanks to training modules and events that are organized by the UN, or directly by UNITAR. For example, at CERN UNOSAT hosts and trains national experts from Nigeria, Indonesia, and Nicaragua, to mention a few recent cases. These experts follow intensive training sessions of two weeks during which they sojourn at CERN. In other cases, it is UNOSAT that sends their trainers abroad to develop knowledge and provide technical support to developing countries. All the experts trained by UNOSAT become part of a global network of technicians who can be connected to work together in case of need.

All the work of UNOSAT is made possible by the agreement between UNITAR and CERN. The support of CERN, visible or less visible as it may be, is of fundamental importance. The recognition and even the awards that UNOSAT enjoys in return for its relentless work go in part also to all those at CERN who help us and support us technically and institutionally. The future holds the concrete possibility for CERN and UNOSAT to work together on joint ideas and activities linked to computing and programming.

Useful links

UNITAR: www.unitar.org or play the video at www.unitar.org/media
UNOSAT: http://cern.ch/unosat/
To contact the UNOSAT team, e-mail unosat@unitar.org.

Francesco Pisano, manager of the UNOSAT Programme, United Nations

Göttingen will play host to the CERN School of Computing in August

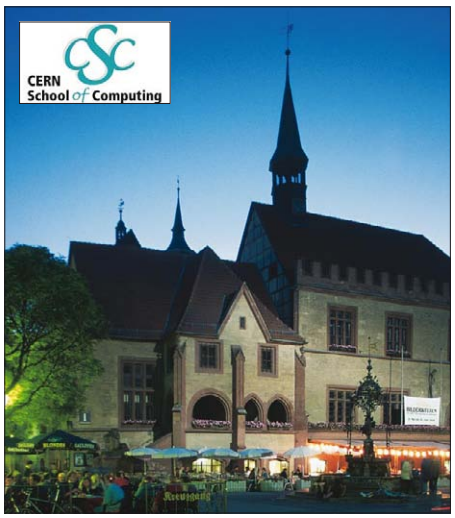

The beautiful Wilhelmplatz in Göttingen will be enjoyed by CSC attendees.

This year's CERN School of Computing will be held on 17–28 August in Göttingen, Germany, and is organized by CERN in collaboration with the Georg-August Universität Göttingen and DESY. The school is aimed at postgraduate students and research workers with a few years' experience in scientific physics, computing or related fields. The deadline for registration is 4 May.

Special themes this year are as follows:

- **Data technologies** – presenting state-of-the-art technologies and options for data storing and management in especially demanding environments, through lectures and practical exercise sessions.
- **Base technologies** – addressing the most relevant underlying technologies for software development security, networking, hardware architecture and virtualization. Lecturers come from the US and CERN to teach theory and organize practical work.
- **Physics computing** – focusing on informatics topics specific to the HEP community. After setting-the-scene lectures, it addresses ROOT and data analysis technologies. It also offers a range of practical exercises on the topics.

Further information can be found on the CERN School of Computing website (www.cern.ch/CSC) where you'll find more details about the programme, practical information about this year's school, and details of how to apply, including the application form.

Note: 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

Training on computer security keeps cyber attackers at bay

Cyber attackers are not slowing down in their efforts to attack and so we must redouble our efforts to keep them at bay. To help in this never-ending struggle, the Computer Security Team is organizing two technical training courses in the spring.

- **Secure e-mail and Web browsing** The goal of this entry-level course is to show how to detect and avoid the typical security pitfalls encountered when e-mailing and browsing the Web. The focus is on Outlook and Internet Explorer. If you wish to learn which e-mail attachments are OK to open, and to understand how worrying an "expired certificate" warning really is, then this course is for you!
- **Developing secure software** This half-day course is aimed at software developers, both for Web applications and regular software. It introduces the main security principles, such as least-privilege and defence-in-depth, and discusses security in different phases of the software development cycle. The emphasis is put on implementation: most common pitfalls and security bugs are considered, followed by advice on best practice for secure development. Not sure how (or why) to sanitize user input, or whether SQL injection affects you? Then do not hesitate to register for this course!

Both courses, although not hands-on, will be interactive and full of real-life examples. The courses will be published soon in the *CERN Training Catalogue* (http://cta.cern.ch/cta2/f?p=110:9). In the meantime, please feel free to contact Sebastian Lopienski (sebastian.lopienski@cern.ch) for more information.

Let's work together to make the CERN computing environment more secure.

CERN Computer Security Team

EGI board backs EGI Blueprint

The EGI Blueprint, which details the implementation of the EGI Organization's activities and the first phases of the development of a sustainable grid infrastructure in Europe, has been endorsed by the EGI Policy Board, which consists of representatives of national grid initiatives from 39 countries.

The European Grid Initiative (EGI) aims to establish a sustainable grid infrastructure in Europe, and to move from a structure based on short-term funded projects to a long-term service, for the benefit of the research community. EGI – a partnership between national grid initiatives (NGIs) and a coordinating body, the EGI Organization (EGI.org) – will begin its full operations in 2010. The EGI Blueprint document is a proposal designed to determine how to establish this long-term sustainable grid infrastructure. It presents a vision of the transition towards the new EGI model and includes the relevant requirements for the implementation, operation, user interaction with, and management of the corresponding infrastructure, as well as the preliminary budget outline.

At the Prague meeting on 20 January, the EGI Policy Board acknowledged the significant progress that has been made in the production of the final EGI Blueprint version. This document took into account the extensive feedback provided by the EGI Policy Board and other experts, leading to its endorsement by a large majority. It will now serve as the basis for the construction of the EGI.org institution, and for the submission of the EGI – and other related – proposals to the European Commission.

The EGI Policy Board also authorized the EGI_DS project (the EGI Design Study in charge of defining and implementing the new organization's activities) to create a task force open to experts nominated by the EGI Policy Board members, to determine the best options for EGI funding and for the preparation of the EGI proposal.

"The fact that the EGI Blueprint is now accepted by so many countries as a basis for the implementation of a European grid infrastructure is a tremendous step forward", said Ludek Matyska, project director of EGI_DS. "Eight countries have presented bids to host the EGI Organization. This also shows the commitment of governments to take this next step for a common research infrastructure," he added.

The EGI Policy Board allowed EGI_DS to prepare a proposal describing the process for the appointment of an EGI.org director, and to begin this process as soon as the location of the EGI.org headquarters had been decided. The location was announced at the EGI Policy Board meeting held in Catania, Italy, on 2 March. Amsterdam was selected as the host city ahead of seven other European cities that also expressed interest in hosting the EGI Organization. The choice of the location of the EGI.org headquarters is a further and decisive step towards the implementation of a sustainable European grid infrastructure.

Useful link

EGI: www.eu-egi.eu

Damien Lecarpentier, EGI-DS

E-groups takes over from SIMBA

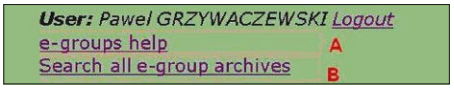

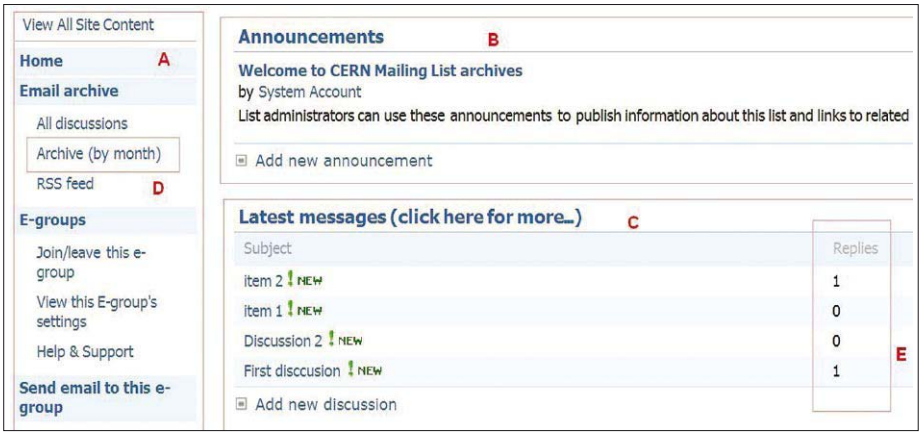

Clockwise from top left. Fig. 1. The "E-mail properties" tab in the e-group application. Fig. 2. The main view of the e-group archive with the relevant parts: A, B, C, D and E. Fig. 3. The top banner of the e-groups application with links to the help (A) and search all archives (B) facilities.

In December 2008 e-groups started to replace SIMBA as the application to manage mail distribution lists at CERN. Via the e-groups application you can manage groups for one person, mailing lists, and all properties of dynamic mailing lists populated from a set of criteria.

In addition to basic functionality, e-groups provides a new archiving system that organizes messages in threads, provides more efficient search and navigation facilities in the archives as well as the possibility of subscribing to RSS news feeds. This article summarizes these features of the new archiving system.

Enable/disable archiving

An e-groups' archive can be enabled by the e-group owner or by a member of an admin group. This action is performed via the "E-mail properties" tab in the e-groups application. When an archive is enabled, all messages sent to the group are also stored in the archive system.

The important change in comparison to the previous system is that disabling the archive or deleting the e-group does not mean the archives will be deleted. After the group has been deleted, archives remain on the server, and are automatically reattached to the group if the group is recreated and archiving is again enabled.

Archived items are accessed from the "View archives" link on the "E-mail properties" tab, see figure 1.

Main page of archive system

The view of the main archive page is divided into two main parts. On the left side (figure 2, A) there is a navigation bar giving quick access to the main actions. In the centre, users can see announcements (figure 2, B) and the last 10 messages that were sent to the archive (figure 2, C).

Browse and search the archive

To facilitate reviewing archives that contain many messages, a new display mode has been implemented: "Archive by month" in the navigation bar (figure 2, D). This mode gives quick access to older messages by month or by year.

The number of messages displayed per page can also be set and messages can be grouped into discussion threads. In addition, the number of messages that are in the thread is displayed (figure 2, E). When a list member replies to a message sent to them via the list, the answer to the original message is added as a "reply" and appended to the thread. Clicking on the main message displays the message plus all the replies.

There is a search facility available at the top of the e-groups page. Here you can search for keywords in several parts of the messages: subject, sender, text of the message and message attachments. Users can also search the entire archive of all the lists that they are a member of by selecting "All archives" in the drop-down list next to the search box, or through the link "Search all e-group archives" in the top banner of the e-groups application, shown in figure 3, B.

Accessing archives

E-groups' archives can be accessed using a Web browser, or alternatively via an RSS feed reader that checks for any new messages sent to the archive. RSS allows you to be notified about new messages in the archive without having to visit the archive's website. On the e-groups help page, you can find information on how to configure the most popular RSS feed readers for Internet Explorer, FireFox, Thunderbird, Outlook and Safari.

More information on the new archive system and on the e-groups application can be found using the help link in the e-groups application banner (figure 3A). Here you will find an extensive user guide and FAQ.

Useful links

https://e-groups.cern.ch/
https://espace.cern.ch/e-groups-help/

Pawel Grzywaczewski, IT-IS

The deadline for submissions to the next issue of CNL is

15 May

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

Add-ons make surfing with Mozilla Firefox more secure

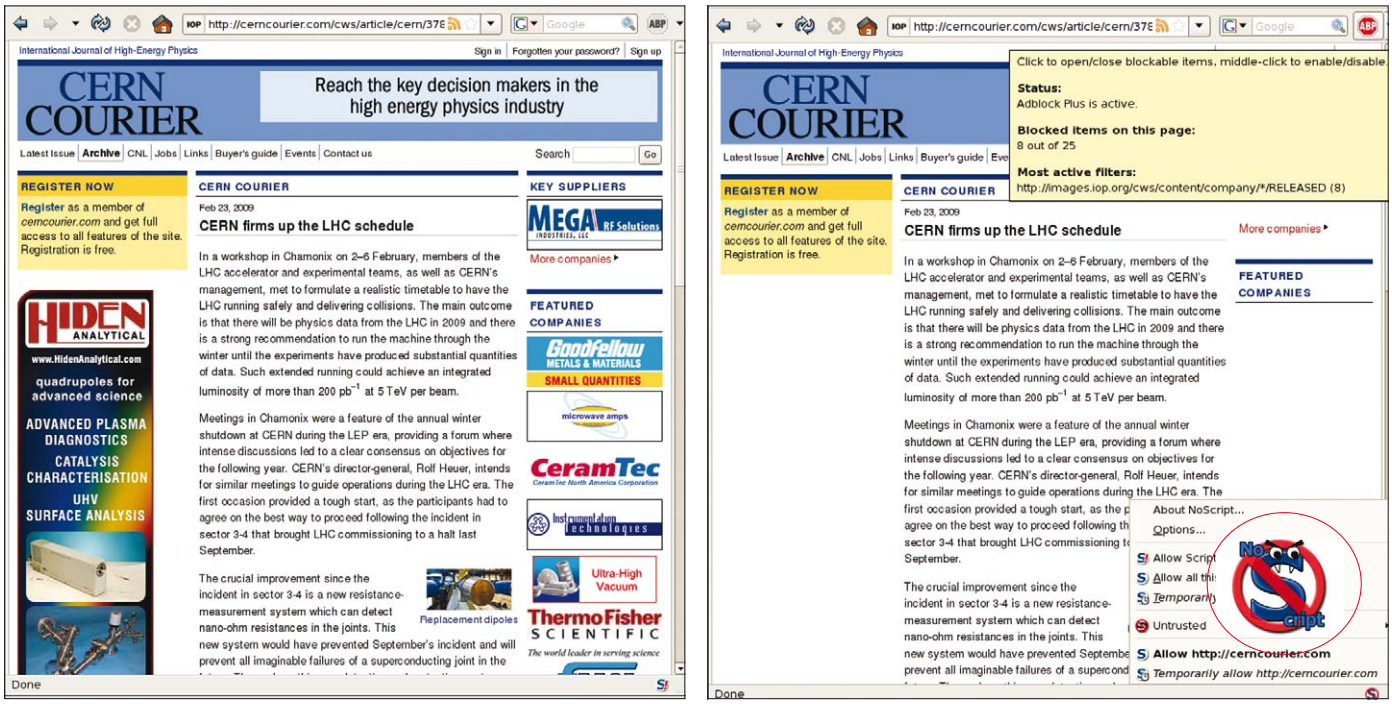

Left: original Web page with adverts. Right: scripts and graphics are blocked. More fine-grained options are available to the user.

Imagine visiting a website and being told first to download and run a program from that site to "enhance your viewing experience". This should set your teeth on edge – why would you hand over complete control of your PC to an untrusted party? If your teeth are not affected, please read the computer security recommendations again at http://cern.ch/security.

Unfortunately, for most of us this is a daily experience – even if your browser typically hides this. Modern websites make extensive use of "scripting" – small programs that are provided by the originating site. Common examples are JavaScript (technically this is ECMAscript which has nothing to do with Java the programming-language, it's simply a clever marketing ploy), Active-X (Microsoft-only), Java applets, Silverlight, etc. Even supposedly harmless media content such as "Flash" or "PDF" are in fact full-blown programming languages.

The only thing preventing full-blown infection of our machines by random websites is that these languages tend to be executed in a restricted environment, the so-called "sandbox". However, implementing a truly secure sandbox is very hard – witness the constant flurry of browser updates – and frequently "usability" conflicts directly with "security". Scripts are also notorious for causing accessibility problems – a screen reader (for the blind) will usually not be able to cope with them, and user preferences (e.g. for large fonts with high-contrast colours) may be ignored.

In some cases the additional functionality is beneficial to the user. For example, large Web applications have become far more responsive since they don't need to reload a Web page on every minor change from the user, to the point where complete "Office" suites are now feasible as a Web service.

On the other hand, a lot of the scripts simply provide visual fluff, serve the more annoying type of advertising or are being employed to track user behaviour. In the worst case, scripts will actively attack, for example by trying to exploit known vulnerabilities in the browser or media viewers ("drive-by download" – no user interaction required), or by stealing or re-using the user's Web credentials. Some attacks are known as "cross-site-scripting" (or XSS), where scripts from one website exploit a programming error on a different site to attack the users there. More specialized attacks go by "cross-site request forgery" (known as XSRF), where an ongoing authenticated Web session is being (ab)used by a third party (think Google getting you to buy books via Amazons' one-click system, without you noticing).

Lastly, the JavaScript in question might not come from the original page owner: user-supplied-content sites (blogs, forums, comments) often do not properly sanitize the provided material, and a swarm of website programming errors mean that often the owners of a website have no idea about the offending code they are propagating.

Different browsers have different defence mechanisms against such attacks and this article concentrates on the Mozilla Firefox browser, which is the default browser on Scientific Linux CERN. Internet Explorer will be featured in the next CNL, in June.

Within the browser

Obviously, the sandbox mechanism ought to contain malicious software. However, Firefox has a particular problem since the browser's user interface itself is written largely in JavaScript. This means it needs to make a clear separation between browser code (with the privileges of a normal application, i.e. read and write files, connect to remote machines or launch other executables) and the scripts coming in from remote websites. Historically, a

lot of Firefox "exploits" were aimed at the boundary between the two.

Firefox has the ability to turn off JavaScript and Java applets (Edit→Preferences→Content), and will not understand other languages such as Flash or PDF, at least not without the help of a plugin (which can be turned off). However, this "all-or-nothing" solution is not very satisfactory as some sites might be "trusted", some sites are simply unusable without their scripts, and Firefox itself does not provide a more fine-grained mechanism.

Extensions to the rescue!

Firefox "extensions" are small add-on programs that add new functionality to Firefox. They are usually stored in the user's Firefox directory and are available for free from the Web. While "downloading and running programs from the Web" to solve the problem of "downloading and running programs from the Web" sounds weird, please remember that ultimately the whole of Firefox is distributed this way and the issue here is trust. The extensions discussed here are widely used (and reviewed), are distributed from the official Mozilla website and have well-known authors. Please do not take this article as a general invitation to install whatever extension you find on some shady website!

NoScript – selective whitelisting

This rather popular extension implements the missing functionality for allowing scripts from just some "trusted" sites and not from others. It also tries to prevent various attacks even between these trusted sites.

The extension is unobtrusive: a corner icon informs the user whether a website contains scripts, or whether all or some of them are allowed or all are blocked. The extension deals with JavaScript, as well as dangerous (since programmatic) media types. Allowing a website either temporarily or permanently takes a single mouse-click. At the same time, the extension has various powerful whitelisting options for more advanced cases.

NoScript is under active development to address new attacks, and hence is frequently being updated (something that Firefox will propose to do automatically). Some settings are required for CERN websites, these are documented under http://cern.ch/twiki/bin/view/LinuxSupport/NoScriptProblems.

Adblock Plus – no more "ads" (less malware too)

However if general advertisements are considered harmful, wouldn't we all be blind by now from traditional text media ads? In fact, the security problem associated with Web advertising is not the visual offences committed to attract our attention, or the occasional inappropriate content on otherwise harmless sites, or the privacy-violating tracking of users across the Web, but rather that the Web advertising industry breaks the "trust" relationship between a user and the website.

The key problem is "syndicated advertising": websites sell off a part of their Web page to other companies, who then sell it off again etc. and all of this via automated placement tools. This means that the website owner may have no idea who provides the publicity currently displayed on their site and the user will not be able to tell either. This mechanism has been abused in the past to quickly infect thousands of visitors to otherwise reputable sites, including high-profile sites such as *MySpace*, NHL.com, Canada.com and *The Economist*. In conclusion, browsing only to "safe" websites to avoid getting exposed to malware is an illusion.

The Firefox "Adblock Plus" extension implements blacklist-based filters even for non-scripted content such as images, and as such it is the natural complement to NoScript. Known advertising content is filtered (unless the user explicitly allows it, e.g. to support a particular site), this cuts down on the exposure to malicious content. As a positive side effect, pages typically load quicker and become much more readable.

A downside is that a lot of free Web content is being financed by advertising, and such filtering eventually might harm revenues to a point that some sites will have to shut down. At least at work, CERN users should not be a major contributor to such revenue, but the decision whether to be subjected to these adverts is with the user.

Default at CERN? Why not?

One major drawback is that while these extensions do make Web surfing safer, they tend to break things. A lot. And in sometimes not-so-obvious ways, and at inconvenient times. For example imagine filling out a longish Web form, only to discover that the final Submit button requires JavaScript, and that a reload will wipe your input.

The user needs to be much more aware of what they are doing on the Web, and constantly take decisions on whom to trust. As such, these extensions are a potential support nightmare – the CERN helpdesk would not be able to cope with questions for every non-working website.

Another point is that these extensions tend to get updated quickly, and in bursts – this would not fit well in the standard IT update cycle (and take a lot of manpower). A centrally maintained extension installed on a machine prevents users from installing a more recent version themselves.

Lastly, the decision on whether to trust a site depends on the impact of an infection for a particular user. Taking these decisions centrally will lead to a bad compromise: too open for the truly security-conscious, too restricted for the people who'd just like "to get their job done". As an example, for technical reasons often a complete domain has to be trusted, even if only a single website merits that trust.

Altogether, this means that we will not be able to make this the default browsing environment. However, we would like to strongly encourage knowledgeable users to take advantage of the additional protection outlined above.

Useful links

Mozilla Firefox Add-ons: http://addons.mozilla.org
Computer Security: http://cern.ch/security

Jan Iven, IT-FIO

If you want to be informed by e-mail when a new 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.

Follow best practice for Windows file/folder security management

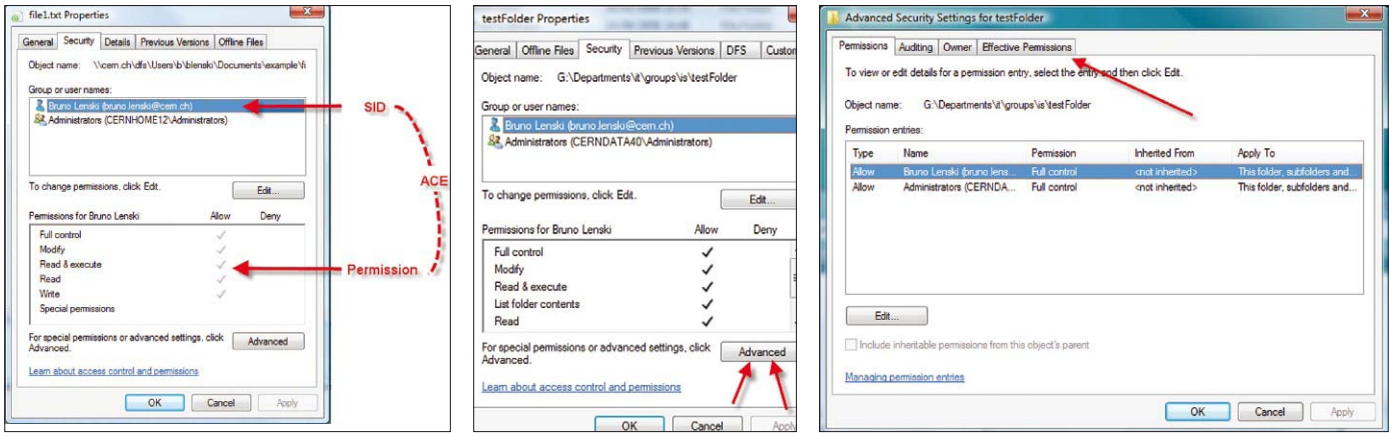

Clockwise from above. Fig. 1. Access a file's security information by selecting the file, right clicking on the mouse and selecting "Properties". To verify if a given user/group has access to a folder, open the properties of the folder and select the security tab. Fig. 2. Click on the "Advanced" button. Fig. 3. Open the "Effective Permission" tab. Fig. 4. Click on "Select ..." and enter a user/group name in the pop- up screen. Fig. 5. Click on "OK" in the pop up screen to display what permissions are granted.

Nowadays, security concerns are increasing everywhere and we need to control who is accessing our information and when it is accessible. Windows provides a built-in way to access information and check user privileges to decide who can and who cannot access a document/folder.

This article describes best practice in using Windows built-in access control and how to avoid problems accessing files on a local PC or on DFS folders.

Each file/folder in Windows has security information. This information is accessible by selecting the file, right clicking on the mouse and selecting Properties (figure 1).

The Security tab contains an SID (a group/user Name) and a set of permissions (ACE). Changing these permissions is done via the edit button.

Here are the best practice rules for permission management:

- **Never remove the administrator entry**: The administrator entry is used by the system to access file information. If the administrator rights to access the file are removed, a side effect could be that back-ups on this file are not carried out anymore. Moreover, if you need to ask the Helpdesk to recover a file, then the process will take longer and be more difficult.

Not allowing administrator access to your data does not make it more secure, instead it puts your data at risk.

- **Do not use the "deny" permission**: If you do not want someone to access a folder/file, it is better to not give them the "allow" permission, rather than "denying" them access. This is due to the resolution method of access. Deny permissions are taken into consideration first. For example, if you are a member of IT-IS group and you deny folder access to IT-IS group but grant your login full access on the folder, then you will be denied access.
- **Use groups to give access to multiple persons**: If you want to grant access to a given file/folder to members of your group, you should use the group permission. For example the default groups "Users <dep>-<group>" works for all CERN groups (in the previous example, the group used. was "Users IT-IS").

You should not add permissions to each user one by one. If someone leaves or joins the group then you would have to modify the permissions for every single file: this operation is time-consuming and error-prone.

Adding someone to a group is a straightforward operation done at a structural level (not the file level) and built-in, dynamic groups are automatically updated, i.e. a newcomer/departure will automatically be reflected in the corresponding "Users <dep>-<group>" according to the HR database.

All existing groups are available via the Win Services page: https://cern.ch/winservices/Services/GroupManager/GroupManager.aspx. And in e-groups: http://e-groups.cern.ch. These pages also allow you to create and manage your own groups.

- **Check user permissions in case of doubt**: To verify if a given user/group has access to a folder, open the properties of the folder and select the "Security" tab. Then:

1. Click on the "Advanced" button (figure 2).
2. Open the "Effective Permission" tab (figure 3).
3. Click on "Select ..." and enter a user/group name in the pop up screen (figure 4).
4. Click on "OK" in the pop-up screen to display the permissions (figure 5).

For more information on file security and managing Access Control Lists (ACLs), please consult the "Managing ACL" page in the help pages of win-services website or in the following PDF: https://cern.ch/winservices/Help/Contents/Images/Security%20How-to/ACL_helpPage_v1.0.pdf.

Bruno Lenski, IT-IS

Grids and clouds go head to head

So, we've all heard about clouds and grids. But what are the pluses and minuses of each approach? For that matter, just what exactly is a cloud? A just-published paper goes into this and more, with comparisons of grid computing and cloud computing.

Written by former iSGTW editor Cristy Burne (now of GridTalk) and available as a GridBriefing at the GridTalk website, the four-page document dissects the benefits and drawbacks of each approach, and draws on opinions and sources as diverse as Hewlett Packard and the European Software Association. There is also a quick-reference chart that goes into such meat-and-potato issues as:

- Who provides the service?
- Who uses the service?
- Who pays for it?
- Where are the computing resources?
- Why use them?
- What are they useful for?
- How do they work?

	Grids (e.g. EGEE)	Clouds (e.g. Amazon)
What?	Grids enable access to shared computing power and storage capacity from your desktop.	Clouds enable access to leased computing power and storage capacity from your desktop.
Who provides the service?	Research institutes and universities federate their services around the world.	Large individual companies.
Who uses the service?	• Research collaborations. • "Virtual organizations," comprising researchers located around the world.	• Small to medium commercial businesses. • Researchers with generic IT needs.
Who pays for the service?	Governments: providers and users are usually publically funded research organizations.	The cloud provider pays for the computing resources; the user pays to lease them.
Where are the computing resources?	In computing centers distributed across different sites, countries and continents.	In the cloud provider's private data centers, which are often centralized.
Why use them?	• You don't need to buy or maintain your own personal computer center. • You can complete more work and tackle more difficult problems. • You can share data with your distributed team.	• You don't need to buy or maintain your own personal computer center. • You can quickly access extra resources during peak work periods.
What are they useful for?	Grids were designed to handle large sets of limited duration jobs that produce or use huge quantities of data.	Clouds best support long-term services and longer-running jobs.
How do they work?	Grids are an open source technology. Resource users and providers alike can understand and contribute to the management of their grid.	Clouds are a proprietary technology. Only the resource provider knows exactly how their cloud manages data, job queues, security requirements and so on.

Grids and clouds, side by side. Details in GridBriefing. Image courtesy of GridTalk.

Pointers

The GridBriefing has pointers to other essays and studies of grids and clouds, including a new, invaluable, much longer e-Infrastructure ReflectionGroup (e-IRG) White Paper written by Fotis Karagiannis of the Athens University of Economy and Business.

Seven topics were selected and examined in-depth in the document: grid and cloud computing; security; education and training; global collaboration; sustainability of the computing-related e-infrastructure; remote instrumentation; and virtualization.

These hot topics were chosen after several rounds of consultation with experts belonging to the e-infrastructure community, and were presented to the e-IRG delegates at the e-IRG workshop in April 2008, in Zürich.

The initial reactions of cloud bloggers such as Markus Klems (*Cloudy Times: Random Thoughts of Markus Klems*) are positive, saying: "Instead of describing theoretical features of grids and clouds, the authors take a look at two concrete implementations: the EGEE grid and the Amazon cloud (EC2+S3). This approach seems reasonable as it avoids long-lasting, fruitless discussions about how to define clouds and grids...Simplicity is an architectural design choice that I believe to be the major success factor of Amazon-style cloud computing. Easy-to-use interfaces allow third-party developers to hook into the cloud and build their own frameworks and tools on top of it."

Or, as the e-IRG authors put it: "In the medium term, the greatest potential benefit of cloud, as proposed by Amazon, is probably not the service itself, but its interfaces and usage patterns."

For more details, see the e-IRG *Facebook* groups site and the press release.

Useful links

GridTalk: www.gridtalk.org/
e-IRG *Facebook* groups site: www.facebook.com/pages/Espoo-Finland/e-Infrastructure-Reflection-Group/49818607909
e-IRG press release: www.isgtw.org/pdfs/Press%20release%20WP-prefinal-1.pdf

Dan Drollette, iSGTW

- This article was published online in iSGTW on 11 February.

Gridipedia is one-stop shop supporting European business

Gridipedia is a knowledge and toolset repository that consists of grid service components and best practices to support European businesses with the uptake of grid technologies.

The Gridipedia website is currently seeing large growth and is populated with real grid applications from the BEinGRID project, which is focused on specific business processes and addressing current customer needs and requirements. It is now seen as public source of information on application of grid technologies in a business context.

The Gridipedia repository contains information on applying grid technology to business. Image courtesy of gridipedia.eu.

According to Santi Ristol, services area manager at Atos Origin, "Gridipedia is the European grid meeting point of the future – a repository for business case studies and best-practice guides where grid players can find solutions and partners."

Gridipedia differs from other grid repositories

Gridipedia currently holds a wealth of business-oriented information based on the insights of two specialist teams-one technical, one business-that have been analysing successful grid implementations in real business scenarios.

All content must pass stringent tests for submission, ensuring that it really delivers on its promises. By ensuring that all hosted material has business relevance and can be traced to a real use scenario, Gridipedia avoids the "technology-push" problem of many repositories. As a business-pull repository, it also seeks out key contributions from third parties.

Useful links

Gridipedia: www.gridipedia.eu/
BEinGRID project: www.beingrid.eu/

iSGTW

- This article was published online in iSGTW on 11 February.

Detector safety system delivers a bespoke protection solution

IT-CO was given a real challenge when it was mandated by the Joint Controls Project Steering Committee to produce a detector safety system (DSS) for the LHC experiments.

Major requirements

Many meetings and fruitful discussions were required to specify and agree the complete needs of such an important safety system. Here are some of the major requirements for the DSS:

- It should have autonomous detector protection (i.e. without human intervention), executing "actions" in response to the detection of "alarms". An alarm, in turn, arises when a logical combination based on the values of one or more "sensors" becomes true. In DSS jargon, this chain (sensors→alarms→actions) is called the "alarm-action matrix". A DSS action is similar to an interlock, and typically results in switching off part of the detector equipment.
- Human intervention by experiment operators would then be needed to rearm the system and bring the detector back into operational mode.
- The DSS should follow the "passive safety" approach; this means that if for any reason the DSS cannot have active control over what is happening, the detector should automatically be put into a safe state.
- The DSS should provide uninterrupted safety, 24 hours a day for the next 20 years.
- The experiment operators wish to be able to modify the existing alarm-action matrix, or to extend it by adding new sensors, alarms or actions, without stopping the DSS.
- The alarm-action matrix should be completely re-evaluated every second.

Finally, many requirements were specified that defined the way a user would interact with the DSS, to configure and monitor the alarm-action matrix.

It was decided to split the DSS into two parts: a programmable logic controller (PLC)-based front-end (FE), and a PC-based back-end (BE).

The main duty of the FE would be to cyclically process the alarm-action matrix, looking for alarms to be triggered and actions to be taken; but another important task would be to monitor its own status.

The BE would run a PVSS-based application, which is a supervisory control and data acquisition (SCADA) system from ETM (a company now owned by Siemens), to let the user configure and monitor what

Fig. 1. The two detector safety units (DSUs) of the DSS prototype. About 30 DSUs have been installed at the experiment sites.

is going on in the FE. The FE should be able to run independently from the BE, preserving its full functionality even if the latter were dead.

Keywords for the FE are: availability, robustness, simplicity, reliability and speed; and for the BE: flexibility, user-friendliness, homogeneous look and feel, intuitiveness and reliability.

The front-end hardware infrastructure

To fulfil the required high availability, the FE is based on a redundant pair of Siemens PLCs, which constantly compare their status using an optical link. The two PLCs are hosted in two different racks (detector safety units, DSUs in DSS jargon).

These DSUs, together with other PLC-less DSUs, also contain I/O modules cabled to patch panels, to allow the connection of sensors and actuators. The DSUs are interconnected through a redundant Profibus network. Each DSU contains a redundant pair of power supplies, and is powered through an uninterrupted power supply (UPS) system which guarantees more than one hour of up-time in case of a power failure, see figure 1.

In principle, all the elements are duplicated, and the DSS FE system should tolerate "single-point" failures without any adverse consequences. Sensors and actuators connected to the system are the responsibility of the experiments ; and as we will see below, the DSS offers the possibility of implementing some redundancy here as well, but cannot impose it.

The data-driven approach: identical software everywhere and modularity

One or more of these systems will be required per LHC experiment. To build a system that could be partially reconfigured by the experiments at any time without any service interruption plus be maintained for 20 years with the minimal possible effort, two things are essential:

- every DSS specimen should run the same, identical software;
- and the FE software should be unaware of the specificities of the different alarm-action matrices. In other words, the FE code should not know any details about the alarm-action matrix.

Rather, these details have to be transformed (by the BE) into data "tables", stored in the PLC memory, on which the FE code would cyclically perform a few basic operations (comparing, counting, etc.). Every DSS item (sensor, alarm, actuator etc.) has a unique place in one of these data tables. In this way, when adding a sensor, an alarm, an actuator, or modifying their definitions, only the data tables are modified, not the FE code.

On the BE side, every DSS item is represented as a "PVSS datapoint". This can be described as a C structure with several different fields. Some of these fields contain "parameters" and have to be "mirrored" into the data tables defined in the FE memory; other fields contain "status", to be mirrored from those FE tables. The PVSS "S7 driver" protocol implements the mirroring of the data between the FE Siemens PLC and the BE PVSS project.

As an example of the basic operations performed by the FE, one of the parameters for an analogue sensor is the "alarm-high

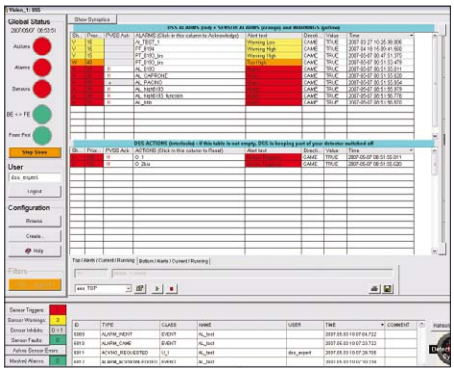

Fig. 2. The DSS user interface shows the DSS front panel and alert-event screen.

threshold" value, formatted as a 16-bit integer. The corresponding status is a single bit which shows if the analogue value is above this threshold. The FE code that deals with analogue sensors, loops on all the slots, and for each occupied slot, checks if the value of the sensor is above or below the alarm thresholds and sets the status bits accordingly.

For an alarm, the parameter structure is composed of up to eight pointers containing the sensor's slot identifiers, and a counter to specify how many of these corresponding sensor-status bits should be true for the alarm to be triggered. It is this flexibility that enables the user to implement redundancy by duplicating the most critical sensors. The status will tell if the alarm is triggered or not, if an acknowledgement request came, etc.

Another table contains the links between alarms and actions (known as actuators). An entry in this table contains the identifiers of an alarm and an actuator, as parameters, plus a delay; the FE code loops on all the slots, and for each slot if the alarm has triggered, it will set the actuator to the safe "0" state (possibly after a delay, when defined).

For an actuator slot, the parameters contain a bit to inform the FE code if the user has asked for an actuator reset (to be able to restart the equipment protected by it); however, the FE will only reset the actuator if all the alarms linked to the actuators are in the non-triggered state. This feature automatically protects the detector against human error.

The DSS user interface

The two major tasks of the DSS user interface are configuration and monitoring. Configuration (i.e. adding/deleting/modifying DSS items) can be performed by the user either through a set of sophisticated Wizard-like panels, or by running a line-oriented text configuration script through a dedicated panel. Furthermore, at the end of a configuration session, the user can run a tool to check the consistency of the modified DSS system, namely to check that the parameters contained in the FE memory tables and in the BE datapoint elements are identical.

The monitoring facility is built around the PVSS alert-event screen (AES). This enables the user to see any abnormal conditions detected by the sensors, any triggered alarms, and the protective actions that have been taken. Through the AES the authorized user can also acknowledge alarms and ask for actions to be reset (this is completely safe, as explained earlier).

An alternative view is provided by the synoptic facility, where the user can structure the experiment as a tree of "locations", with each DSS item assigned to one location. The synoptic panel can display the location layout, with graphical symbols representing the status of the items attached to that location. Every location "child" of the current parent is also displayed as a symbol, with colours summarizing the status of everything defined in the location sub-tree. By clicking on a location symbol, the synoptic panel will display the newly selected location.

The main DSS panel also displays a "logbook", showing the latest events that have taken place, together with the relevant user actions, see figure 2. This logbook is stored in an Oracle database, and is very helpful in understanding the reasons behind a sequence of events.

The users can also be notified of alarms and sensor warnings via SMS or e-mail, which is useful in case of an unattended control room. The BE also performs periodic checks on the DSS status, and informs the user in case of problems. For example, the UPS systems that power the DSUs are monitored and the users are notified five minutes after they go into battery mode. In this way, the users have time to react and bring the power back before the UPS discharges completely and the relevant safety actions are taken.

Operational experience

The first detector safety system was delivered to the CMS experiment in 2003 and was first used operationally in 2004. Since then, this CMS system and the one installed later have been working smoothly, protecting the entrusted equipment 24 hours a day, to the satisfaction of our users.

Table 1

	ALICE	ATLAS	CMS	CMSX	LHCb
detector safety units (DSUs)	7	7	6	10	3
analogue sensors	322	4	6	20	82
digital sensors	82	582	431	1540	246
alarm conditions	195	581	417	520	297
alarm→action links	~220	~2000	~770	~1850	~1150
actions	223	309	232	513	191

A DSS is installed in each LHC experiment. The figures relate to the number of components.

While the FE software is rarely modified, the BE software has constantly been evolving to add new features and to follow the evolution of PVSS (2.12.1 to 3.1 to 3.6SP2). We currently have five fully operational and complete DSSs installed (two for CMS). Table 1 shows the number of different parts in these systems.

Acknowledgements

We would like to thank the many people who, through their experience, dedication and helpfulness, contributed to the success of the DSS project, during the design, implementation and operation phases. In particular, we wish to highlight the contributions of all the members of the DSS steering committee who met for a year to define the requirements for the DSS, after Wolfgang Tejessy first identified the need for such a system.

The members of the DSS advisory committee, chaired by Wayne Salter, additionally aided the project by constantly ensuring that the DSS implementation went in the direction required by the LHC experiments. We would also like to note the contribution of Bruce Flockhart (originally the DSS project leader) and Sascha Schmeling, who were part of the initial IT-CO DSS development team.

Furthermore we would like to thank our contacts in the experiments (Andre Augustinus and Pascal Blanc from ALICE, Helfried Burckhart, Heidi Sandaker and Fernando Baltasar Pedrosa from ATLAS, Christoph Schaefer and Alain Meynet-Cordonnier from CMS, Rolf Lindner and Laurent Roy from LHCb) for their patience at each modification or improvement. Their many suggestions and feedback significantly contributed to the constant enhancement of the DSS. Finally, we would like to thank the "DSS Piquet", who are the first line of intervention in case of problems, and who guarantee us a good night's sleep.

Useful link

More information on the DSS is at https://edms.cern.ch/nav/CERN-0000011111

Giulio Morpurgo, IT-CO (now EN-ICE) and Stefan Lueders, IT-DI, with Oliver Holme and Jeronimo Ortola Vidal IT-CO (now EN-ICE)

AFS revisited: understand groups

These "AFS revisited" articles are intended to act as short reminders. Here we look at standard AFS commands, common pitfalls and some tips and tricks to get more out of AFS at CERN.

In AFS, the number of entries on an access control list (ACL) for a directory is limited to about 20 (in fact the total number of available slots on the ACL depends on the size of the individual items). However, it is quite a common case scenario that more users need to be assigned certain rights for a directory. In addition, it is quite tedious to maintain the same list of users on several ACLs, so AFS addresses this problem by the concept of AFS groups, which can then be added to ACLs.

Every user can create a group using the *creategroup* subcommand of the pts tool suite. Taking our sample user Pam, the command to create a group called pam:friends would look like:

```
$ pts creategroup –name pam:friends
```

which will return a line such as

```
group pam:friends has id –103690
```

to confirm that the group has been created. Groups created by users have to follow a naming convention: they must be prefixed by *‹owner›:*, otherwise AFS will refuse to create the group. While normal users have to follow this convention, administrative rights in AFS allow the creation of (almost) arbitrary group names. That's why you may find groups on ACLs that do not follow this naming convention.

The number of groups a user can create is limited to 20. You can list the groups you own by the *listowned* subcommand:

```
$ pts listowned –nameorid pam
Groups owned by pam (id: 7302) are:
pam:friends
```

Since this group has just been created, it is empty. One can check the contents of a group with the *membership* subcommand:

```
$ pts membership –nameorid pam:friends
Members of pam:friends (id: –103690) are:
```

The very same command can also be used to perform the inverse look-up, i.e. answering the question "which groups am I a member of?":

```
$ pts membership –nameorid pam
Groups pam (id: 7302) is a member of:
party–committee
michael:friends
```

Note that you can check your own memberships, but not usually other users' memberships. On the other hand, the default settings on group creation allow you to see the members of an arbitrary group. Fine-grained control on who can do what with a group is available, please see the *setfields* subcommand to find out more on how to set the individual attributes of a group.

Since an empty group is not very useful, the *adduser* subcommand populates the group:

```
$ pts adduser –user jim –user dwight
–group pam:friends
$ pts membership –nameorid pam:friends
Members of pam:friends (id: –103690) are:
jim
dwight
```

and *removeuser* removes group entries:

```
$ pts removeuser –user dwight –group
pam:friends
$ pts membership –nameorid pam:friends
Members of pam:friends (id: –103690) are:
jim
```

Once you no longer need a group, you can delete it by means of the *delete* subcommand:

```
$ pts delete –nameorid pam:friends
```

Note that this will delete the group, but not clean up your ACLs! The (numerical) entry for the deleted group will stay on the ACL. So, make sure you clean up your ACLs whenever you delete a group. Managing ACLs has been covered in the first article of this series. If you need to do that recursively on a directory tree, have a look at the *afind* command to which you can pass the corresponding *fs setacl* command.

In addition to user defined groups, AFS comes with three predefined system groups:

- *system:anyuser* (literally all users, even unauthenticated ones);
- *system:authuser* (all users who have successfully authenticated in the local cell);
- and *system:administrators* (the AFS administrators).

You may find one or more of these groups on your AFS space ACLs. For instance, *system:anyuser* is usually listed on your public directory ACL to grant everyone read access to this folder. In addition to the anonymous groups *system:anyuser* and *system:authuser* that come with AFS, there are some additional groups where users are not explicitly listed. One example is *cern:nodes* which contains all machines with an IP address on the CERN network. These anonymous groups in conjunction with ACLs that were too open have been misused by outsiders in the past. Please have a look at http://cern.ch/security/afs for more details on how to set your ACLs so that they can protect your AFS space from such attacks.

Three common pitfalls

If access is denied because a user is not a member of any of the groups listed on the ACL, adding him to one of the groups may not solve the problem immediately. For the current connection, the file server already has a list of the groups that the user belongs to and there is no way of notifying the process that the membership list has changed. The user must re-acquire AFS credentials, for instance logout and login again to establish a new connection to the file server and force it to retrieve an up-to-date version of the membership list.

It is also important to keep in mind that the *removeuser* and *delete* subcommands serve completely different purposes. *Removeuser* is meant to manipulate the member list of an AFS group, the *delete* subcommand will try to erase the user or group entry from the underlying AFS database. Since most AFS commands allow short forms that omit the actual parameter names, some caution is required when using these subcommands.

AFS also allows IP addresses to be used on ACLs. Before they can be added to a group however, a user-type entry for the machine has to be created in the AFS databases. To help machine administrators create such entries, *cern-config-keytab* has been extended by the *–I* option. Furthermore, IP addresses cannot be added to the ACLs directly, but they have to be added to groups, which are then put on the ACLs. Since hosts cannot log in again to force an update of the membership mentioned above, a server-side timeout of a maximum of 2 hours may have to expire before the change takes effect. Remember that IP addresses on ACLs allow anyone with access to that machine to access the AFS space according to the ACLs! There is therefore quite a serious potential security risk here.

Useful links

The *openAFS* website provides information on all AFS commands: www.openafs.org/manpages

The *CERN AFS User Guide* is a much more comprehensive manual, available at: http://cern.ch/service/afs

Arne Wiebalck, IT-FIO

openlab tests Oracle VM for database service virtualization

Why database virtualization

There is a combination of factors that leads to increased interest for virtualization of database servers: number of servers, power usage, cost, new platforms with multi/many cores, high availability and ease of management. Database servers often do not use the full CPU/server capacity. This can be for reasons such as separate test-and-development servers or because spare capacity has to be kept available in case of increased load or for high availability.

In addition, several database services can often not be merged onto a single database or physical "machine" as they need to be on different versions of the database software or operating system. We also need to "protect" each of the database instances from the others so that memory can be guaranteed for a given instance.

Having a smaller number of servers will obviously reduce the cost of hardware, maintenance and electrical power. The power usage is especially important due to the constraints on "critical power" in computer centres as well as the impact on a cooling system.

We believe that consolidation with virtualization will become more and more interesting, in particular given that processors will continue to have an increasing number of cores in the medium term future. Another attractive opportunity with virtualization is the possibility to relocate and distribute virtual machines on different nodes depending on the load and intervention needs. This leads to higher availability while having fewer constraints on the operations to be performed.

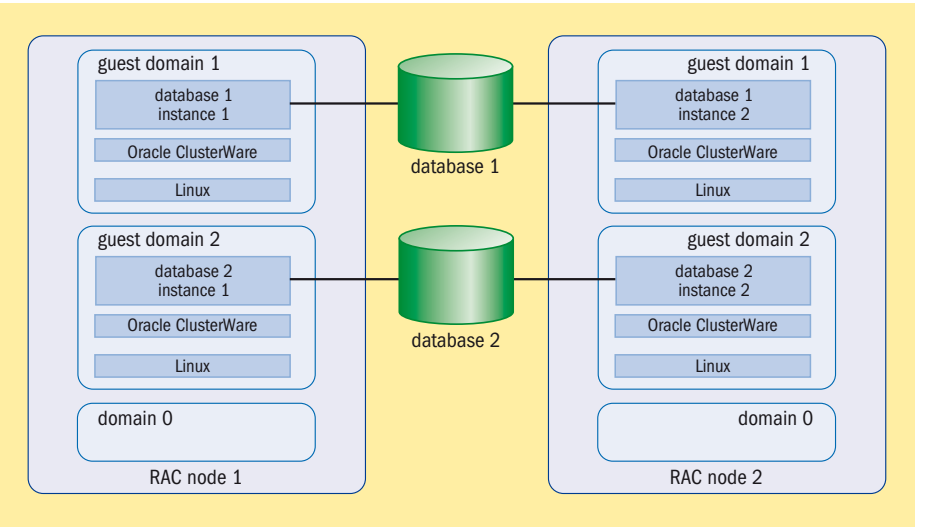

Fig. 1. Fighting underutilization by combining several clustered databases on two servers.

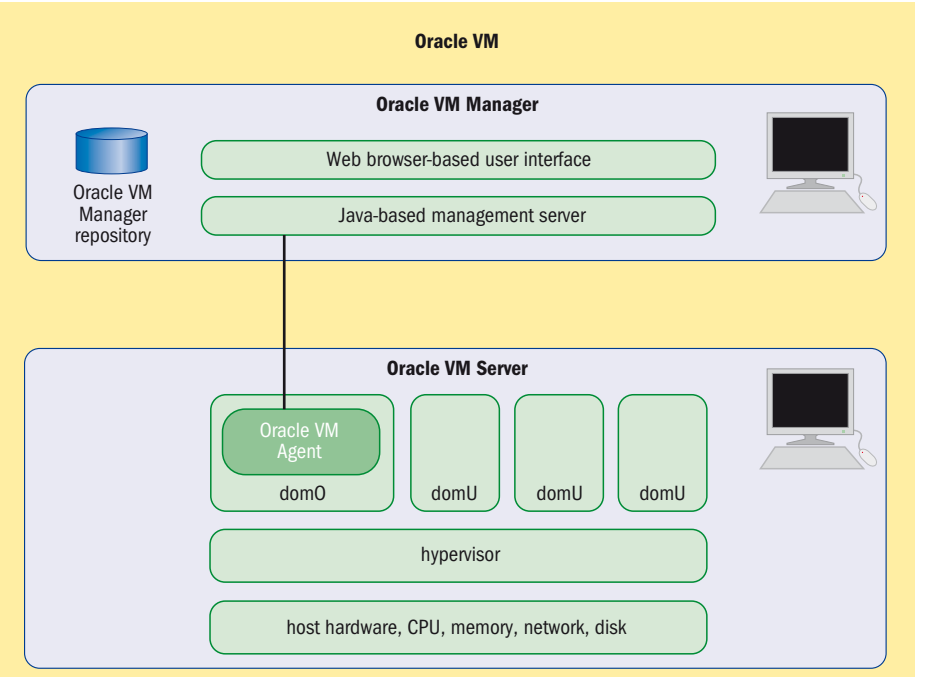

Fig. 2. Architecture: Oracle VM Manager and Oracle VM Server are installed on separate servers with communication provided by the management server and Oracle VM Agent.

Database virtualization at CERN openlab

In the CERN openlab Database Competence Centre, we started looking at virtualization in early 2006. The first project was aimed at testing the Oracle database in the Xen virtualized environment. The main focus was on the possibility to run a single-instance database in a virtual machine (VM), as well as testing different Xen schedulers to allocate CPU slices to different virtual machines. On successful completion of this project, we decided to expand our tests, and include Oracle Real Application Clusters (RAC) in such an environment.

RAC – an Oracle technology widely used at CERN and in the HEP community – allows several physical machines to work as a single clustered database, thus providing scalability and high availability. High availability is the main driving force behind implementing RAC in less workload-intensive infrastructure databases, which also leads to underutilization of the database nodes. Bearing this in mind, we were attracted to the idea of running RAC in a virtualized environment, thus utilizing more of the hardware without compromising the high availability features (see figure 1).

The concept of RAC on VM was proved to be successful by openlab summer student Maria Leitner in 2007. She created a cluster database inside a single physical host with Red Hat Enterprise Linux 4 as a base operating system. The set-up was quite stable and the performance loss

due to the virtualization layer was at an acceptable level.

In November 2007, Oracle announced its virtualization solution: Oracle VM, which consists of three parts: Oracle VM Server, Oracle VM Manager and Oracle VM Agent (see figure 2).

Based on Red Hat Enterprise Linux (RHEL) and open-source Xen, Oracle VM Server is the base virtualization software to be installed on a bare metal server. The Oracle VM Agent is the application, installed together with the Oracle VM Server and is responsible for communication with Oracle VM Manager Web-based GUI.

Oracle then certified this single instance (November 2007) and RAC (September 2008) databases in the Oracle VM-based virtualized environment, which was the step we were waiting for.

It was decided to continue with our tests and compare RAC performance on Oracle VM versus the native Xen set-up we tested before. Together with openlab summer student Andrei Dumitru in 2008, we created two set-ups on identical hardware, using Oracle VM and pure Xen, included in RHEL 5. The stability of both set-ups proved to be good and the performance of the Oracle VM set-up was 5 to 20% better than on pure Xen, depending on the workload. Tests of live VM migration with an active database showed only a few seconds downtime with no session-state loss. This is an outstanding addition to the high availability features initially provided by RAC.

Following these successful tests, we are now working closely with IT-FIO to integrate this Oracle VM solution into the CERN fabric management system – ELFms.

In summary, these Oracle VM features, their ease of use and the official certification have directed our choice towards Oracle VM as a virtualization solution for our virtualized cluster databases for the future.

Useful links

CERN openlab: http://cern.ch/openlab (with openlab summer student reports in Technical Documents section)
Oracle VM: www.oracle.com/virtualization

Anton Topurov and Eric Grancher, IT-DES

A great success: 4th EGEE User Forum attracts 600 attendees

By the sea, among lemons and sunshine, more than 600 members of the European and international grid community – representing 45 countries – gathered in Catania, Sicily, Italy for the 4th Enabling Grids for E-sciencE (EGEE) User Forum co-located with Open Grid Forum 25 & OGF – Europe's 2nd International Event.

Home base for EGI announced

The most eagerly awaited news to come from the event was the announcement that Amsterdam was selected as the host city at the European Grid Initiative (EGI) policy board meeting on Monday 2 March, ahead of seven other European cities that also bid to host the central organizing body of EGI.

"When I heard the decision I thought, 'Yes! We did it.' This will be a great opportunity for Amsterdam, hosting such an important international organization and we especially hope that the staff of EGI.org will benefit from our beautiful city," said David Groep, NIKHEF, member of the successful bidding team from Amsterdam.

"I'm very grateful to the people who voted for Amsterdam. I know how hard the selection process was. There were other strong contenders that were also excellent options. Of course," he said "I think Amsterdam was the best choice. I hope the location in Amsterdam will help EGI be a great success. We'll do our best to ensure the success of EGI, to keep the users happy so that we get great science."

Some 600 members of the European and international grid community (representing 45 countries) attended the 4th EGEE User Forum. Image courtesy of EGEE.

New memorandum of understanding with BalticGrid-II

On Wednesday 4 March, the project directors of EGEE-III and BalticGrid-II, Bob Jones and Ake Edlund signed a Memorandum of Understanding detailing BalticGrid-II's involvement with EGEE on an infrastructure and application level as the European grid community moves towards the European Grid Initiative.

"More than just an MOU," said Edlund, "this document involves a work plan and concrete targets with names and dates."

"This is an important step for the European grid community," said Jones. "as we work towards a united e-infrastructure in all European states in preparation for EGI."

Conference and event reports

Achieving grid interoperation through standardization: OGF endorses new proposed standard

In another landmark for grids, OGF announced on Tuesday 3 March that it endorses the GLUE 2.0 specification as a proposed standard (www.ogf.org/documents/GFD.147.pdf). The specification delivers the long-awaited common information model of grid entities. This document is a product of the international grid community, with contributions from the largest grid infrastructure projects and their middleware providers, such as EGEE, Open Science Grid, TeraGrid, NorduGrid, NAREGI and practical experiences from the science collaborations around the Large Hadron Collider (WLCG).

Balazs Konya, NorduGrid technical coordinator and co-chair of the GLUE working group noted: "During recent years the grid community has been working very hard to reach convergence on how grid entities are modelled and described. The nonexistence of a common information model has always been a major obstacle for interoperability. The release of the GLUE 2.0 specification as an OGF proposed standard is a major achievement of the grid community. As one of the founders of the GLUE working group the Nordugrid collaboration is naturally committed to adopting GLUE 2.0 through its ARC middleware. "This will allow us to provide standards-based interoperability for our users with several grid infrastructures, including EGEE, the world's largest multidisciplinary grid."

Putting a human face on the grid: GridGuide unveiled

On Monday 2 March, the GridTalk project launched a new website, *GridGuide* (www.gridguide.org). A work in progress, *GridGuide* allows visitors to explore an interactive map of the world, visiting a sample of the thousands of scientific institutes involved in grid computing projects. Sites from 23 countries already appear on *GridGuide*, offering insider snippets on everything from research goals and grid projects to the best place to eat lunch and the pros and cons of specific jobs.

"We're thrilled to see *GridGuide* bringing grid sites to life," said Bob Jones. "This website shows the human aspects of grids by highlighting how people from all over the world are contributing to the success of grid computing."

Worldwide grid helps to fight heart disease

On Tuesday 2 March an EGEE press release announced the latest work from the Cardiogenics Consortium. The project investigates the genetic causes of one of the world's biggest killers – coronary artery disease – using the EGEE infrastructure. Their work will be published in the March issue of *Nature Genetics*. EGEE's infrastructure enabled the researchers to do two years' work in fewer than 45 days. This allowed them to identify possible genetic candidates for the causes of a disease that kills more than 2 million people a year in Europe alone.

- Note from the side-lines: The consensus was that the local food, while delightful, was not helping anyone's efforts at dieting. In the words of one delegate: "At least we have a few months to recover before EGEE '09 in Barcelona."

Useful links

BalticGrid: www.balticgrid.org/
European Grid Initiative (EGI): www.eu-egi.org/
EGEE: www.eu-egee.org/
GridGuide: www.gridguide.org

Danielle Venton, IT-EGE

A word from Vangelis Floros, conference organizing committee chair

This was a very successful event: we attracted more than 600 people. Catania was a warm host – literally and metaphorically. The food was great, the wine uplifting and the social events gave great opportunities for networking and lots of formal and informal discussions.

The mixture of people, from many scientific and application communities, standardization groups and collaborating projects proved very successful. Although co-hosting the event made it much more complicated to organize, the final results compensated for all the effort of the past months. Speaking about organization, you know we started the whole process last June at OGF23 in Barcelona, and since then so many people have been involved and have done their best to make this event possible. It was a challenge for us all, but the final results made us happy and proud.

No two events are the same. Each user forum, each EGEE conference has had that certain something, a particular feeling that grows inside you and remains there for a long time. The 4th User Forum is no exception. The general feeling I will remember is this anticipation of grid communities about the future of grid infrastructures in Europe and the future of this technology in general.

The anticipation around EGI will continue to grow, as the end of EGEE-III approaches. The announcement of the host city of EGI.org was an important milestone for EGI. It was an historic moment for the European grid community and Amsterdam. The choice of city is just the start of an avalanche of developments anticipated for the coming months. There is still much work ahead for the transition, and here in Catania this feeling of anticipation and alertness dominated the atmosphere.

As expected, we had lots of discussions about cloud computing which, although it remains a very talked-of term, I think it is slowly being better understood. We are starting to realize the potential of clouds, and experiment with the technology delivering the first (still modest) results. Together with green computing, I believe these will be the two topics that will attract most of the attention of people coming from high-performance and distributed computing communities.

Our next rendezvous is in Barcelona, for the final conference of the EGEE project. My mind is already on the 5th User Forum, wondering how the European grid community will look a year from now. We will be opening the call for bids to host the 5th User Forum in the next few weeks.

Look out for the March issue of CERN Courier

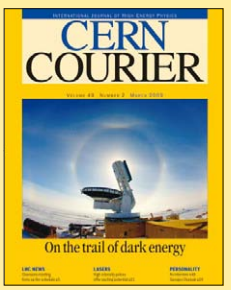

Taiwan hosts HEPiX Fall Meeting

The HEPiX Fall Meeting was held at ASGC in Taipei, Taiwan on 20–24 October 2008.

With an attendance of around 85, from both "traditional" HEPiX sites and from many Austral-Asian sites, this meeting marked the first incursion of HEPiX into Asia since its creation at CHEP in Tsukuba, Japan in 1991. The meeting was organized expertly by the Academia Sinica Grid Computing (ASGC) staff to the extent that the agenda was already full around one month before the event, and last-minute submissions could only be accepted by extending sessions and indeed the overall meeting length. All slides, some formal papers and many photographs are available at http://indico.twgrid.org/conferenceTimeTable.py?confId=471.

HEPiX participants pose for a photograph at the Academia Sinica site in Taipei.

Highlights

- An obvious highlight was the good response from Asian sites. The meeting chairman, Simon Lin, had invested a lot of effort in attracting sites from across the region (not only HEP sites) and he got a very good response, the Asian representatives making up for any drop in numbers from the more traditional HEPiX sites. Preceding the HEPiX meeting with a Grid Workshop, including a day devoted to CASTOR, surely helped boost attendance. As a consequence, many new sites were introduced to the HEPiX community and vice versa. We hope to see them again, and not only in Asia. Meanwhile, many "new" sites say they will maintain their contact with us, and we have already one firm and one tentative offer to return to Asia.
- Another hot topic, as in the past few meetings, is the challenge faced by many sites in hosting their steadily growing compute farms. Solutions range from adding external boxes (e.g. SLAC), to getting more from existing space (e.g. GSI) to adding new or re-using existing buildings (e.g. FNAL). And some sites are simply stuck while funding is found or red-tape is overcome (e.g. IN2P3 and PDSF).
- Based on the number of talks offered and the time spent on it, storage remains the number one subject at HEPiX meetings. The various results presented were interesting and the Lustre community is getting more and more excited.
- Working groups: the HEPiX Benchmark WG feels it has achieved the task that it was set (establishing an agreed HEP benchmark) and has plans to document and publish its results and then wind up. The Storage WG has established a dedicated test bed at FZK and restarted various tests.
- On the non-scientific side, the local organization team was excellent, not only in the pre-meeting tasks of establishing and filling a full agenda but in the execution of the event during the week. There was a dedicated team to change overheads and switch microphones between talks; more-than-adequate and reliable networking; on-time shuttle buses; full catering on site; and excellent receptions. Taipei HEPiX will be a hard act to follow.
- Next meetings: Umea, Sweden, 25–29 May; and probably Berkeley in October or November.

The meeting started with a keynote talk on cloud computing from Fred Baker of Cisco. He was the chair of IETF from 1996 to 2001 and was involved with network standards for many years. After a brief review of computing from the 1950s to the 2010s, he explained how cloud computing is the natural next step for providing computing resources. He considers grid computing to be a way of sharing mainframe resources, but his view of cloud computing is the complete outsourcing of needed IT power to an outside supplier. He compared the various models on the market today, explaining the Google model in some detail, as far as is known from the outside. He compared the advantages and disadvantages of cloud computing for small and larger companies, concluding that the smaller firms see more on the positive side. Cloud computing is being driven by the technology and the providers but the existence of different models from different providers poses a serious risk of vendor lock-in. Nevertheless, it is a natural evolution in next-generation computing centres and there will be an increasing number of alternatives appearing in the market-place.

The traditional Site Reports were split over three mornings and gave sites new to HEPiX an opportunity to introduce themselves. Some significant tidbits:

- CERN reported on the LCG computing challenge of May 2008.
- All Fermilab production systems are now on Fermigrid.
- Fermilab is moving to ITIL practices. ASGC, CERN and others are also looking to do this.
- The various INFN sites are moving towards individual AFS cells.
- In line with its move away from high-energy physics, SLAC is changing its name although the "word" SLAC will still appear in the new name. Their GLAST satellite was successfully launched and is taking very good gamma-ray measurements.
- RAL's new computing building is almost complete but there have been hitches.
- A couple of sites, including DESY, propose not to adopt Windows Vista but to wait for Windows 7 to be released.

Sessions on storage also spanned several half-days. They included an invited talk from a local disc manufacturer on data integrity, describing mechanisms used to improve the chances of avoiding data corruption; a progress report on the latest benchmarks performed under the aegis of the HEPiX Storage Working Group; a talk from SUN Microsystems on Lustre plans for the next few releases as well as some Lustre user experiences; and a full session on the CASTOR storage suite as used in the LHC Computing Grid project.

Several centres reported on work starting or ongoing to add facilities to increase their computing capacity. The Pisa, INFN group reported in particular on heating and cooling fluid dynamic studies they had undertaken and how simulation was now being compared to reality in advance of starting a second phase of their new installation.

Troy Dawson (Fermilab) presented the latest updates for Scientific Linux and his plans for the next releases. In one discussion, an end-date was agreed for version 4 updates and support (October 2010). There was a status report from the Benchmarking group. It proposes SPECall-cpp2006 as the preferred benchmark set to use, and the slides explained how results in these units can be calculated from the SPECint and SPECfp published numbers. It was pointed out that running the tests locally is quite easy and a script is available if the SPECcpu2006 licence – which is very cheap – is available. Having achieved their primary objective – an agreed HEP benchmark – the working group will write up the results, try to present these at CHEP09 in Prague this March and wind up.

Other sessions covered virtualization, Windows, networking, and security. The social programme showed off the best of local Taiwanese culture, featuring a wonderful Chinese banquet accompanied by an orchestra playing Chinese and Western music on traditional Chinese instruments.

Useful link

www.hepix.org/

Alan Silverman, IT-DI

CERN commemorates 20 years of the Web

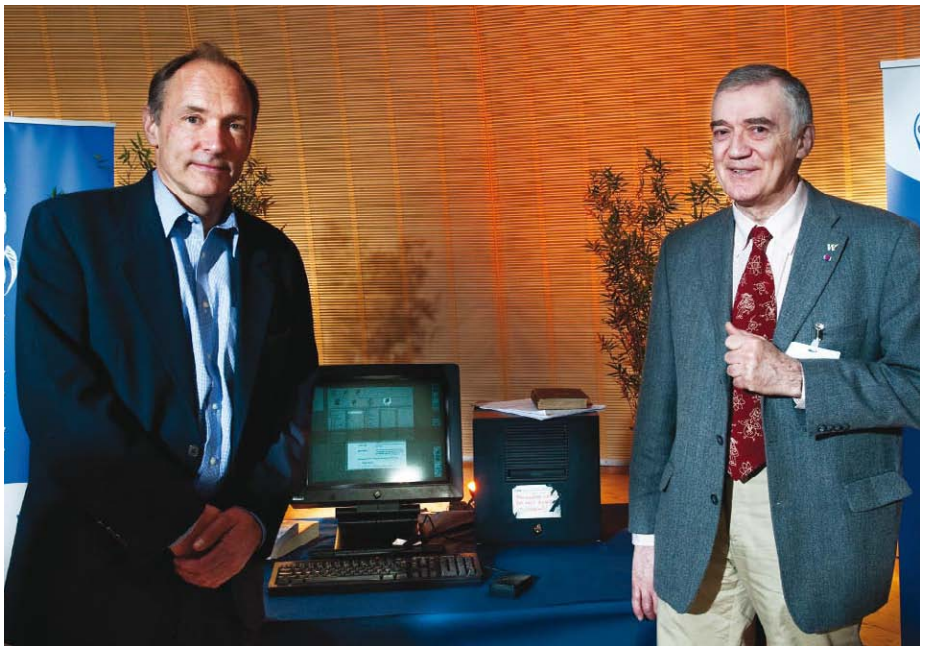

Web inventor Tim Berners-Lee (left) and early Web pioneer Robert Cailliau (right) stand on either side of the NeXT computer that acted as the world's first Web server.

On 13 March, CERN commemorated the past 20 years of the World Wide Web in a unique event that gathered together some of the early Web pioneers.

Web inventor Tim Berners-Lee returned to CERN, the birth place of his invention, 20 years after submitting his paper entitled Information Management: A Proposal, in March 1989 to his boss Mike Sendall.

This celebration was broadcast live from the Globe and, following an introduction by CERN director general, Rolf Heuer, there were short presentations looking back at some of the early history of the Web. We heard stories and anecdotes from Web pioneers such as Robert Cailliau, Jean-Francois Groff and Ben Segal.

We were then treated to a demo of the world's first Web server and original browser on the NeXT computer used by Berners-Lee back in 1990. The enthusiasm of the presenters was obvious and we were given a small insight into what it must have been like hearing these gentlemen excitedly explaining this new Web concept to others.

In his keynote speech, Berners-Lee reminisced about his days at CERN, which he considers as a "microcosm of the world", and praised the unique and diverse environment here. "CERN has come a long way since 1989, and so has the Web, but its roots will always be here." It was inspiring to hear how all this started from such humble beginnings and a desire to have agreement on just a few simple things such as the URL that have led to the Web as we know it today.

There was also a feeling of looking forward to what the Web will bring in the future, for example on mobile phones where there is a real push from the W3C Mobile Web Initiative to develop best practices for phone browsers. This theme took us through a series of talks from some of today's Web pioneers, covering diverse areas such as linking data on the Web and a social semantic Web application.

Berners-Lee also took the opportunity to talk about the World Wide Web Foundation. Launched last year, this foundation will look to fund and coordinate efforts for a single Web that is free and open.

And like Berners-Lee said: "The Web is not all done, it's just the tip of the iceberg...I'm convinced that the new changes are going to rock the world even more."

Useful links:
Event details: http://info.cern.ch/www20/
Video: http://cdsWeb.cern.ch/record/1167328
World Wide Web Consortium (W3C): www.w3.org/
World Wide Web Foundation: www.Webfoundation.org/

Natalie Pocock, IT-UDS

Calendar

April

13–17 **GridAsia 2009**
Singapore
http://gridasia.ngp.org.sg/2009/

16–23 **Int'l Symposium on Grid Computing 2009**
Taipei, Taiwan
http://event.twgrid.org/isgc2009/

20–24 **18th Int'l World Wide Web Conference**
Madrid, Spain
http://www2009.org/

21–25 **5th Int'l Conference on Networking and Services 2009**
Valencia, Spain; www.iaria.org/conferences2009/CfPICNS09.html

May

4–8 **4th Int'l Conference on Grid and Pervasive Computing 2009**
Geneva, Switzerland
http://gpc09.eig.ch/

11–15 **Networking 2009**
Aachen, Germany
www.networking-2009.org

14–15 **Open e-IRG Workshop**
Prague, Czech Republic
www.e-irg.eu/index.php?option=com_content&task=view&id=192&Itemid=1

18–20 **24th IFIP Int'l Information Security Conference 2009**
Pafos, Cyprus
www.sec2009.org

18–21 **9th IEEE Int'l Symposium on Cluster Computing and the Grid (CCGrid 09)**
Shanghai, China
http://grid.sjtu.edu.cn/ccgrid2009/

20–21 **Cloud Expo Europe 2009**
London, UK
www.cloudexpoeurope.com

25–29 **6th High-Performance Grid Computing Workshop**
Rome, Italy
www.cs.unb.ca/profs/aubanel/hpgc/

June

22–25 **TeraGrid 2009**
Arlington, USA
www.teragrid.org/tg09/

23–26 **International SuperComputing 2009**
Hamburg, Germany
www.supercomp.de/isc09/

28 June – 1 July **HealthGrid**
Berlin, Germany
http://berlin2009.healthgrid.org/