CERN COMPUTER NEWSLETTER

Volume 44, issue 2 April–June 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

IOP Publishing

CERN openlab enters phase-three period

On 2–3 April the openlab annual Board of Sponsors took place at CERN. Rolf Heuer, CERN director-general, opened this Board and officially launched the third phase of CERN openlab. During his introductory speech, he stressed the importance of collaborating with industry and building closer relationships with other key institutes as well as the European Commission. The Board was an opportunity for partner companies (HP, Intel, Oracle, Siemens), contributors (EDS – an HP company) and CERN to present the key achievements obtained during openlab-II and the expectations for openlab-III.

openlab history

Each phase of CERN openlab corresponds to a three-year period. In openlab-I (2003–2005) the focus was on the development of an advanced prototype called opencluster. CERN openlab-II (2006–2008) addressed a range of domains from platforms, databases and Grid to security and networking, with HP, Intel and Oracle as partners and EDS as a contributor. Disseminating the expertise and knowledge has also been a key focus. Regular training sessions took place throughout this period, including openlab contributions to the CERN School of Computing. The activities carried out during the last year of phase II are presented in the openlab annual report, which was published in June, and is available from the website (www.cern.ch/openlab).

openlab-III

With the start of the third phase, new challenging projects have been initiated with the partners. The projects are within four Competence Centres (CCs):

- Automation and Controls CC (ACCC): Through the ACCC, CERN, Siemens and ETM Professional Control (a subsidiary of Siemens) are collaborating on security, as well as the move of automation tools towards software engineering and handling of large environments.
- Database CC (DCC): In partnership with Oracle, the DCC focuses on items, such as data distribution and replication, monitoring and infrastructure management, highly available database services, application design, as well as automatic failover and standby databases.
- Networking CC (NCC): One focus of the NCC is a project launched by CERN and HP ProCurve to understand the behaviour of large computer networks (10 000+ nodes) in high-performance computing or large campus installations. Another activity is the Grid Monitoring and Messaging Projects in collaboration with EDS.
- Platform CC (PCC): The PCC project focuses on the PC-based computing hardware and the related software. In collaboration with Intel, it addresses important fields such as thermal optimization, application tuning and benchmarking. It also has a strong emphasis on teaching.

During the third phase, the team will

not only capitalize on and extend the successful work carried out in openlab-II, but also tackle new crucial areas. Additional team members have recently joined and the structure is now in place to collaborate and work on bringing these projects to fruition.

The openlab team

The openlab team consists of three complementary groups: the young engineers hired by CERN and funded by the partners (21 people over the past eight years), technical experts from partner companies involved in the openlab projects, and CERN management and technical experts working partly or fully on the joint activities.

The people involved are not concentrated in a single group but span multiple units. In the IT Department the CS (Communication Systems), DES (Databases and Engineering Services), DI (Departmental Infrastructure), DM (Data Management), GD (Grid Deployment), GS (Grid Support) and IS (Internet Services) groups host openlab activities, as does the ICE (Industrial Controls and Electronics) group in the EN (Engineering) Department, since the arrival of Siemens as an openlab partner (see *CNL* November 2008, p2).

The distributed team structure also enables close collaboration with computing experts in the LHC experiments, as well as with engineers and scientists from the various openlab partners who contribute greatly to these activities. In addition, significant contributions are made by the students participating in the CERN openlab Summer Student Programme; both directly to the openlab activities and more widely to WLCG, EGEE and other Grid and CERN-related activities within the IT Department.

Since the inception of openlab, more than 100 young computer scientists have participated. This summer the programme will be welcoming 14 students of 11 different nationalities.

Useful link

CERN openlab website: www.cern.ch/openlab

Mélissa Le Jeune, IT-DI (CERN openlab)

Helpdesk service takes on new support responsibilities

At the end of May, the following three support lines were handed over and fully integrated into the outsourced Helpdesk service:

- Support for Indico (Indico.Support@cern.ch).
- Network operations, known as "netops" (netops@cern.ch).
- Video-conferencing monitoring and operations (videoconference-support@cern.ch).

In practice all of the Helpdesk staff will now be able to reply to questions in these computing areas, or escalate them to the appropriate expert, with second or third level of expertise, as required.

The three e-mail addresses (Indico.Support@cern.ch, netops@cern.ch, videoconference-support@cern.ch) are now redirected to the Helpdesk (Helpdesk.Computing@cern.ch or Helpdesk@cern.ch). Phone calls to 74927 for network problems and requests will automatically be redirected to 78888 for the Helpdesk. Please note that the number 74927 may be disconnected in the future.

The Helpdesk staff have been trained and given access to the necessary tools to provide first-level assistance for the most frequent and basic problems or requests in these three areas.

The reason for these changes is to allow the specialists to focus on more advanced issues and to increase the Helpdesk efficiency and competency. In this way the Helpdesk has a global overview of all IT services, which helps, for example, when a general problem with a critical server affects several services.

The main objective is to provide a better service for you, our end-users. All of the IT computing support lines will be revised in the near future to identify other areas where a similar change would be suitable.

Remember that if you experience problems with PRMS (Remedy) tickets or the Helpdesk, you can contact the IT manager on duty (tel 163013, e-mail mod@cern.ch) or e-mail user.relations@cern.ch with any comments about the service.

Also, before submitting a problem at the Helpdesk, don't forget to consult the IT Service Status Board where all of the incidents, scheduled interventions, power cuts and service changes are recorded.

Useful link

IT Service Status Board: http://cern.ch/SSB

Nicole Crémel, IT-UDS

MS Forefront Client Security application starts to protect NICE PCs against virus attacks

By the end of this year, all NICE PCs will have MS Forefront Client Security installed. This new anti-virus and anti-malware application will replace the current anti-virus product from Symantec. The reasons for this change include the small footprint of the client application, excellent response times for pattern updates and very good integration with the existing NICE infrastructure.

The migration will be progressive (department by department) over the coming months and will be combined with the regular monthly deployment of security updates. Users will be notified about the migration by the usual CERN Alerter message announcing the monthly updates. Any users wishing to have their NICE computer migrated earlier can send their request to the Helpdesk (helpdesk@cern.ch).

Documentation for MS Forefront Client Security has recently been published and can be accessed via the following link: https://cern.ch/win/Help/?fdid=19.

CERN users who would like to install the Forefront client on their home computers are allowed to do so. The procedure and installation packages are available on the following webpage: https://cern.ch/win/Help/?kbid=051092.

Please note that there is no change for Mac OS X users. They can still use Symantec Anti-virus, which can be downloaded from DFS using the following link: https://dfs.cern.ch/dfs/Departments/it/Services/MacSupport/Utilities/Norton%20files/.

Anti-virus Services, IT-IS

Practise safe e-mail and browsing

The internet is a dangerous place with identity theft, buffer overflow, cross-site scripting, request forgery and privilege escalation. Do you need to be concerned about these? Yes. Do you need to study these technical terms to be safe? Well, you have a choice.

It is always beneficial to know more, so we encourage you to stay up to date with the risks. But if the technical jargon has already given you a headache, you may choose the easier option to be safe – stick to the best practices of web surfing. It will not guarantee that you are 100% secure (no-one is), but it will reduce the risks by an order of magnitude. This article focuses on best practices using Internet Explorer and is a follow-up to the article dedicated to Mozilla Firefox in the previous issue.

The risks

The stakes are high: if your PC is compromised, the best-case scenario is that you will have to reinstall from scratch. If you are less lucky, you may lose valuable data; or if you have access to sensitive information it may be stolen and abused by the attackers. This might threaten your professional activities or your private resources – just think of someone getting access to your bank account. If your PC is infected and compromised, you are at the mercy of the attacker.

For the attacker, your computer is a valuable resource. If the attack is successful, it will turn your computer into a bot – a robot that serves the attacker's purpose and acts on their orders. It may start sending spam and infect other computers on your network, or record your key strokes to get your passwords. There is a black market behind such activities and the attacker's motivation is money.

If the attacker is particularly interested in something you have (or can access), they may launch a targeted attack, aimed specifically at you. Although it is less common, this is a growing threat and CERN users have recently been targeted by such attacks (see http://cern.ch/it-support-servicestatus/IncidentArchive/090504-Phishing.htm).

E-mail

E-mail has become one of the main ways of attacking your privacy. You are more and more likely to receive scams, which are fake messages asking you to give away your passwords and details. Attackers try to make these e-mails look legitimate, but the IT Helpdesk has no reason to ask for passwords. Never send any confidential information by e-mail.

Other malicious e-mails contain links to infected files. Simply clicking on the link will put your PC at risk, so it is safest to ignore unwanted e-mails. When in doubt, try alternative ways to confirm the contents of the e-mail, for example, call the sender to ask them about the message by phone.

The colour codes on the web address bars indicate how safe the website is.

Internet Explorer

Due to its market share of nearly 70%, Internet Explorer is an obvious target for the attackers, who are constantly trying to find new security vulnerabilities and ways to exploit them. To fix such vulnerabilities, regular security updates are published. Security problems and the updates that fix them are a fact of life and we need to make sure that our operating system, web browser and all of the applications are up to date. It is a bit like a race – will the problem be fixed before it is exploited? This is why anti-virus software is so important – it brings an additional layer of protection because it can block malicious software trying to exploit a security issue before the problem in the application is fixed.

At work

If you have a standard NICE PC managed by the IT Department, your machine has up-to-date anti-virus software. The operating system, Internet Explorer, and all of the supported applications (such as Microsoft Office) are regularly updated. You can contribute to keeping everything secure by installing the security updates as soon as they are advertised by IT.

However, if you decide to install an application that is not supported by IT, for example Apple Quicktime, then you are responsible for keeping it up to date. We will do our best to inform you about security problems in such applications and these alerts must be taken seriously. Browsing the web is dangerous if you have outdated software on your PC.

At home

Having an infected computer at home puts your work computer at risk. The virus may spread if you use USB memory or open the same file on both computers.

The security of your home system depends on you. You should have an up-to-date anti-virus and schedule regular system scans. Note that as a CERN user you can download the anti-virus software for home use (https://cern.ch/winservices/Help/?kbid=051092). To keep your Windows updated, you should make sure that automatic updates are enabled (right-click on "My Computer" "Properties" "Automatic Updates") and configured to download and install security updates as soon as they are available. Similarly, you have to keep all of the applications up to date.

Security status bar

Internet Explorer 7 has a new security status bar, which lets you quickly assess whether you should trust the website that you are currently visiting. This is the same address bar where you see the website address, but it has a colour code. Typically, the green background means that the site is safe and can be trusted.

The white background means that there is no information about this website. If the address starts with "https://" and there is the padlock icon, then all communication within this website is encrypted.

A yellow or red background means that there is something suspicious about this site. We strongly advise you to stop visiting such websites.

Your responsibility

You are responsible for the security of your computer, because the way you use it determines whether it is secure or at risk. The following guidelines will allow you to enjoy the internet in a conscious – and much safer – way.

Security best practices

- Never send any passwords via e-mail. E-mails asking for passwords are scams. The CERN Helpdesk staff will never ask you for your password.
- Think twice before you click on links sent to you by e-mail or through instant messaging, especially if you did not ask for them. They may lead you to websites that look legitimate (like your e-banking website), but are prepared by attackers to steal your passwords. A similar risk may be related to links from social networking sites, such as Facebook.
- Always read before you click. Internet Explorer will ask for your confirmation before doing anything risky. But if you never read the contents of the pop-up windows you will not benefit from this type of warning.
- Don't install software from the internet

unless you know and trust the vendor. Never run any executable (.exe) files that you did not request – in most cases it is an attempt to take control of your PC. The same is true regarding unwanted browser add-ons, like multimedia players. If you are not sure whether you need something then you probably don't, so never accept a request to run (or even download) files that you did not ask for.

- Never enter any sensitive information (passwords, credit card numbers, etc) unless the communication with the website is encrypted (small padlock icon and address starting with https://). Restart Internet Explorer before you begin; after you finish, use the site's log-off function and then close Internet Explorer.
- Never enter any sensitive information on a PC that you don't know, for example, in an internet cafe. Someone may be recording everything that you type in order to steal your passwords. The same applies to your PC if you suspect that it has been infected with a virus.

There is a new "Secure e-mail and web browsing" course designed to show you how to detect and avoid typical security pitfalls encountered when e-mailing and browsing the web. It is designed for non-technical users of Internet Explorer and Outlook. The course is free and more information is available from the training catalogue. Visit http://cta.cern.ch/cta2/f?p=110:9 or e-mail Technical.Training@cern.ch for dates of sessions.

Useful links

NICE Services – anti-virus for home use: https://cern.ch/winservices/Help/?kbid=051092
IE7 Security Status Bar: www.microsoft.com/windows/products/winfamily/ie/ev/security.mspx

Sebastien Dellabella and Michal Kwiatek, IT-IS

CERN document server improves integration

Fig. 1. A typical browser window with the CDS toolbar installed. This shows the easy search facility and the RSS subscription link.

Fig. 2. A browser window with the Cooliris add-on installed. Shows a 3D wall of CDS multimedia that can be navigated in any direction.

If you are using a web browser such as Firefox or Internet Explorer, then you have probably noticed that they can be extended in different ways through add-ons (also called extensions or plug-ins). The CERN Document Server (CDS) team is making new features available that interact with these add-ons to increase the usability and visibility of its services.

CDS Toolbar

Similar to the Google Toolbar that you may already be familiar with, you can now install a CDS toolbar for your browser. This will give you access to the powerful CDS search engine to find the document that you are looking for quickly, without having to open the CDS search interface page. The toolbar will enhance your browsing experience with the following features:

- Context-menu integration – right-click on an author name on any webpage to open a search in CDS.
- Auto-recognition of ISBN/ISSN book and periodical reference numbers – you can see if a resource is available in the library.
- Future integration with major digital content websites including Amazon, Google Scholar and Wikipedia.

The CDS toolbar is currently available for Firefox and Internet Explorer browsers and uses the LibX technology.

Zotero

Have you ever had to compile your personal bibliography while writing a resumé/CV or a scientific article? Maybe you have been using CDS personal baskets to keep the list of articles together, but now you need a way to aggregate articles from other document servers or websites in the list as well? Thanks to the Zotero Firefox extension, it is possible to do this with only a few mouse clicks. If you visit a website containing publications or scholarly information on CDS and other Zotero-powered websites, Zotero enables you to save the material, it keeps it organized and allows you to export the information to your favourite format (including BibTeX).

RSS

If you have selected a particular collection or performed a search query on CDS, in addition to the popular e-mail alert notification system you can now stay up to date with recent papers on CDS via personalized RSS feeds from that collection or query. After entering the query that you are interested in, just click on the RSS icon either in the location toolbar or in the "Interested in being notified about?" search box. You will then subscribe to a feed that will automatically alert you as soon as new papers on this topic appear on CDS. Figure 1 shows a typical browser window that is using the CDS toolbar and shows you where to click for this RSS feed.

Cooliris

The CDS now supports a new and futuristic way to browse its multimedia content: just install the Cooliris plug-in in your favourite browser. The CERN multimedia content will then be available as a three-dimensional wall that you can navigate in any direction, as shown in figure 2.

Useful links

CERN Document Server: http://cdsweb.cern.ch/
CDS Toolbar: http://libx.org/editions/download.php?edition=4F46CD81
Zotero: http://www.zotero.org/
Cooliris: http://www.cooliris.com/

The CERN Document Server Team, IT-UDS

DFS architecture changes to enable faster backups

CERN users working with the NICE/Windows Services might have heard of "DFS storage", and they probably even use it without realizing. So where does the "G:" drive come from?

What is the Distributed File System?

The G: drive that is available on all centrally managed CERN Windows machines is in fact a view of data stored on a dozen data servers, generated by a specific service called Distributed File System (DFS).

DFS is a service that provides a single point of reference and a logical tree structure for file system resources that may be physically located anywhere on the network. This single point of reference at CERN is represented by the path \\cern.ch\dfs and is automatically mounted as the G: drive on all Windows machines by the central management system.

Using DFS to share resources across the network facilitates the navigation to the data, as everything is seen as a single drive. In addition, an access control system is provided by the built-in Windows "File Security and Access Rights" methods available from any Windows Explorer and also from http://cern.ch/winservices.

What is the current DFS infrastructure at CERN?

At CERN, all of the information hosted on DFS is accessible either via the G: drive or via the path \\cern.ch\dfs. Data accessible through DFS are stored on several servers that are backed up every night onto magnetic tapes.

The new DFS infrastructure is built on 10 servers with disk arrays containing 15 data disks of 1 TB each.

- Six servers are dedicated to DFS Workspaces, Experiments and Departmental data (the corresponding folders are called Workspaces, Departments, etc).
- Four servers are dedicated to Home Directories; the corresponding folder is \\cern.ch\dfs\Users and is directly accessible through the "My Documents" Windows shortcut as well as through the normal G: drive.

Some other servers are also used to store DFS data:

- 10 servers are dedicated to specific data, such as the Media Archive that hosts the official CERN video files (six servers), some specific experiments and services needing to host a large amount of data.

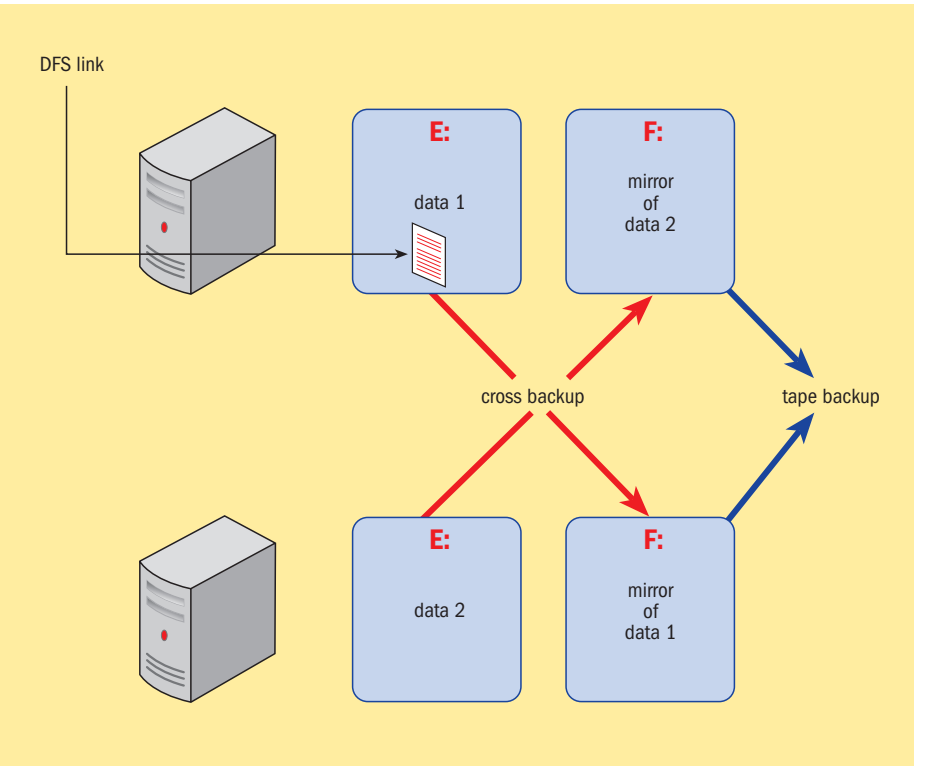

Fig. 1. The Cross-Backup architecture concept groups the file servers in pairs and mirrors all of the data on each server pair. The increased storage needed is added at low cost.

What is new in the DFS infrastructure?

Recently, while the periodic hardware renewal was being handled, a change in the infrastructure was introduced: the Cross-Backup architecture.

The idea of Cross-Backup is to group the file servers in pairs and mirror all of the data on each server pair. The required disk space is of course doubled, but the price of storage has dramatically decreased in recent years, allowing this architecture to be set up at a low cost.

For each server member of a pair, the idea is to split the disk space into two partitions: an active partition and a mirror partition. The mirror partition's goal is to be the backup of the active partition for the other server in the pair.

The backup activity takes place at different schedules depending on the type of server and the data hosted on the active partition. Large partitions containing mostly archived data are backed up every night, whereas Home Directories containing daily work and small files are often backed up in real time.

Figure 1 shows the Cross-Backup architecture concept. This new architecture means that:

- The mirror partition can be backed up onto tapes at any time without overloading the access to the active partition.
- An online backup-recovery mechanism can be provided from the mirror partition.
- There is faster service recovery in case of a disaster: if a server fails to respond, a swap of the DFS link to the mirror partition allows a full recovery in just a few minutes.

This new architecture has been in production since April and provides a better quality of service through the use of these mirror partitions. Data restoration and disaster recovery can be handled with greater efficiency. In case of data loss, a simple call to the Helpdesk (tel 78888) will launch the restoration procedure and the data will be recovered quicker than before.

Useful link

DFS services: https://cern.ch/winservices/Services/DFS/Default.aspx

Bruno Lenski, IT-IS

Europe's Grid infrastructure shifts from EGEE to EGI

Within little more than a year the organization of Europe's Grid infrastructure will be transformed. For EGEE's user community, its new face, the European Grid Initiative (EGI.eu), will mean more of a shift in "behind the scenes" focus and responsibility than a radical transformation in the software and services that they currently use. This future model will be primarily sustained by national funding, with each national grid initiative (NGI) supporting and maintaining its own resources as part of a European partnership and integrated with international partners by supplementary funding from the European Commission.

A lot has recently happened to make this plan a reality and in the next few months new developments will continue. The current picture is as follows:

- In January representatives of the emerging NGIs in Europe formally endorsed the "EGI Blueprint", which sketches the form of the EGI model, a signal that the community is moving in a united direction.
- In March Amsterdam was selected as the location of EGI's coordinating organization, EGI.eu, and the local hosts (NCF, SARA and NIKHEF) have become fully involved in the transition planning.
- EGEE has revised the project's second-year plans in alignment with the EGI Blueprint, to prepare for optimal handover to the EGI.eu and NGIs in May 2010.
- Athens hosted the first meeting of groups interested in forming Specialized Support Centres (SSCs). SSCs form part of the EGI model to support user communities and will fall within the scope of an upcoming EU funding call. A further meeting is planned on 1 July in Paris for these groups and others involved in EGI proposals, to present work plans for broader discussion. CERN, along with other HEP organizations, intends to form an SSC and is working to ensure that the broader physics community is involved.
- The ARC, gLite and UNICORE middleware teams have held four meetings since December 2008 (in Rome, Munich, Oslo and Geneva) to understand the interactions between the EGI Middleware Unit and its main software providers. A further meeting is planned before the Paris meeting in July to discuss the content and structure of a combined proposal.

In parallel, the board managing the preparations for EGI has nominated candidates for the coordination of the EGI project proposal editorial board and the interim EGI.eu project director. The following persons have been agreed:

- Laura Perini (INFN, Italy), coordinator of the EGI project proposal editorial board, focusing on Operations and the EGI.eu construct part/proposal.
- Cal Loomis (LAL/IN2P3, France) coordinator of the EGI project proposal editorial board, focusing on the Applications and Scientific Support Cluster part/proposal.
- Steven Newhouse (CERN), interim EGI.eu project director.

Next steps

Over the next six months the e-infrastructure community within Europe will be focusing on two main activities.

- Preparing and submitting the EGI, Application and Middleware proposals for the upcoming EU funding call.
- Finalizing and then starting the migration towards the EGI model.

If you are a user or a provider of the infrastructure, what will all of this mean for you? Hopefully very little.

"The key objective in all of these discussions that we are having about EGI is to maintain the level of service currently delivered to the EGEE user community," said Steven Newhouse, EGEE's technical director. "Also, by starting to make these changes now within EGEE we can phase them in gradually over the next year and help EGI in its learning process".

Stay tuned to www.eu-egee.org and www.eu-egi.org for more updates.

Danielle Venton, IT-EGE

A version of this article was originally published online in the Spring 2009 EGEE Newsletter, available at http://eu-egee.org/newsletter/spring09/spring09.html.

MATLAB offers powerful simulation capabilities

Researchers from disciplines as far apart as lasers and finance have a new computing tool at their fingertips: MATLAB can now run on Enabling Grids for E-sciencE (EGEE) computing power.

The software is a high-level language and interactive environment that enables users to perform computationally intensive tasks faster than with traditional programming languages such as C, C++ and Fortran. Widely regarded as a powerful piece of simulation software, for use in everything from optimizing rocket launch control settings to vector analysis, it is now fully compatible with any grid computing system using gLite middleware.

"Our motivation for incorporating this tool was rather straightforward. MATLAB is one of the most popular and important general-purpose scientific software tools. A large number of grid users from a multitude of scientific disciplines have been requesting for a long time to be able to run MATLAB applications within EGEE's gLite infrastructure," said Vangelis Floros, who helps to coordinate and manage the EGEE user community and application support activities.

"Through our collaboration with The Mathworks (creators of MATLAB) we managed to achieve a very efficient integration between the two platforms – MDCS and gLite – on the technical level. The software is still in the trial phase but we are working closely with them to define a grid-friendly licensing model in order to be able to offer it very soon for production-level exploitation."

Sample cases

The software can be used in a number of places. A few examples include:

- Finance: Spyros Skouras of the Athens University of Business Economics and Business, is examining financial market data in extremely fine detail.

"To examine a market's microstructure, we use data that are timestamped at the millisecond level and measure tens of variables for thousands of financial instruments," said Skouras. "Such a fine-grain view can shed light on issues of relevance to academics as well as professional traders who can translate milliseconds into profits. But it also means that we work with unusually large data sets that are not easily manageable."

Skouras recently began using MATLAB on EGEE in his research, which he says is well suited to grid technology because the same analysis is performed on data for each instrument. He believes that many researchers working in finance, economics, econometrics, statistics and other fields may become users of EGEE if they become aware of how accessible the grid is, now that the interfacing can be done with

MATLAB tools they are familiar with.

• Making better lasers: research fellow Pavlo Ivanov from the University of Bristol, is seeking to improve fibre-optic telecommunication networks by making better semiconductor lasers with the use of MATLAB on EGEE.

"Making a new laser is very expensive," said Ivanov. "Improving design by trial and error would be a disaster. Laser researchers need to optimize their designs prior to fabrication and MATLAB can help me understand how the lasers are behaving."

Laser-created light is used to carry information on fibre-optic communication lines. If Ivanov can, as he hopes, build a high-power, single-mode laser, it could translate into better quality and longer-distance telecommunications. Ivanov hopes to present his most recent findings at the upcoming EGEE'09 conference in Barcelona this September.

Useful links

EGEE'09: http://egee09.eu-egee.org/
MATLAB: http://www.mathworks.com/products/matlab/

Danielle Venton, IT-EGE

• This article was published online in iSGTW on 10 June

Kerberos5 signals end of an era

In the past, the way to authenticate in AFS was a command called klog (or klog.krb). When we switched our authentication mechanism to kerberos5, we announced that klog would become deprecated.

Over time we have convinced most users to employ the kerberos5 command "kinit" and the number of klogs has decreased a lot. However, we still see a tail of klog usage, which we would like to bring to zero.

We have identified five main ways that klog is being used today.

1. Users typing klog because they always obtained/refreshed their CERN AFS tokens in this way. On centrally maintained machines we already provide a wrap-up around klog, reminding users that klog is deprecated and kinit should be used instead.

2. Users who have to use klog to get AFS tokens for other, external AFS cells. This will continue to work. For example, the wrapper described above does not interfere with such requests.

3. Users requesting a CERN AFS token from outside CERN. There are ways to achieve the same result with kerberos5's kinit.

4. Users who save their password on some local disk and use "klog –pipe" to automatically obtain an AFS token within some script. For such users a replacement is already available and we encourage them to move to kinit as soon as possible.

5. People using an old unmaintained script "reauth" to get their token automatically refreshed. We are providing a replacement (k5reauth) with improved capabilities.

How do you replace klog?

For case 1, habit: it is clear that kinit can be used instead even if things are not as well integrated as they were in the klog times. Two commands are needed: kinit then aklog. So, to help people, CERN distributions of Linux provide /usr/sue/bin/kinit, a script available on the top of the search path, which packages the two commands together.

For case 2, external cells: until the target external site moves to kerberos5, the following command:

```
klog –cell xxxx user
```

continues to work. This is the only case where a user is invited to type klog.

For case 3, remote klog into CERN: the author can successfully create a replacement based on kerberos5 at some remote site (e.g. in2p3.fr). The setup is as follows:

• Have a kerberos5 client installed (generally comes preinstalled in all modern Linuxes, but with a default configuration describing a non-existing EXAMPLE.COM realm).

• Have "aklog" installed (this comes out of openAFS distribution). At worst, copy it locally.

• Have a "good" krb5.conf file, describing the CERN.CH realm. If this is not distributed system-wide in /etc/krb5.conf, then get a copy from any CERN lxplus node and save it locally.

• Create a two-line script, which runs:

```
kinit user@CERN.CH
aklog -c cern.ch -k CERN.CH
```

Watch out for the order (-c before -k).

• If you run with a local copy of krb5.conf, add a first line to the script above:

```
KRB5_CONFIG=local path to krb5.conf;
export KRB5_CONFIG
```

For case 4, saved password: you can simply replace:

```
cat file | klog -pipe
```

with:

```
cat file | /usr/sue/bin/kinit
```

or alternatively:

```
klog –pipe < file
```

with:

```
/usr/sue/bin/kinit < file
```

However, keeping a plain text password is not safe and is strongly discouraged. Kerberos5 offers a similar but slightly safer solution called a keytab. A keytab contains a one-way encrypted copy of the user's password, which kinit can understand, but that cannot be used anywhere else. Getting hold of someone's keytab can do lots of damage, as in the saved password case, but at least it cannot be used to visit EDH or NICE.

When user *nopc3* wants to create a keytab for their account on their local machine, they have to run (assuming AFS and the arc command are installed):

```
arc –h afskrb5m kas ext nopc3 >
safe_local_file
chmod 600 safe_local_file
```

Then user *nopc3* can replace on their local machine:

```
klog –pipe nopc3 <
local_password_store
```

with:

```
/usr/sue/bin/kinit –k –t safe_local_file
nopc3
```

But remember that a keytab must never be exposed to thieves. So the *safe_local_file* must be readable only by the user.

For case 5, reauth: we are providing a replacement called k5reauth. This script works differently from the old reauth. It used to be started from a shell as a disconnected child, which was then "forgotten". It is not uncommon to see several copies of reauth refreshing a token on an lxplus node, which is now useless because the "token owner" has logged off days before.

We designed k5reauth to start a sub-process and watch it. When the child process terminates, k5reauth stops. The started process could be a shell, so it could be left running until the user logs off. k5reauth also accepts keytabs, so it could be used in cron jobs.

At the time of reading, this utility should already be available on SLC machines.

Useful link

http://consult.cern.ch/service/afs/

Bernard Antoine, IT-FIO

CERN's TIM Video Viewer manages PS tunnel access

Since March the operators in the CERN Control Centre (CCC) have used the Technical Infrastructure Monitoring (TIM) Video Viewer (TVV) to provide access to the PS tunnel. It is now planned to use this application from this summer to control access to the SPS tunnel.

The TIM system was developed by the ASE group in the GS Department as a monitoring system for CERN's technical services. The system is able to collect data from different industry-standard hardware and software components, such as Programmable Logic Controllers (PLCs), Object Linking and Embedding (OLE) for Process Control (OPC) servers or CERN proprietary protocols. The core part of TIM is implemented with Java 2 Enterprise Edition (J2EE) technology to create a highly available, reliable, scalable and flexible control solution. It has been in production since 2006 and currently processes more than 1.7 million monitoring values per day. Several client applications were built on top of TIM for processing, visualizing and analysing these data.

A new video system was required in the PS tunnel because all of the old analogue cameras were replaced by digital models. The decision to use the TIM system as a framework for the video system was made after a careful evaluation of proposals from external companies, which were all deemed too complex in design as well as too expensive.

Using the existing TIM infrastructure, the GS-ASE-SSE section was able to develop the TVV within a very short time period. It is completely written in Java and designed so that the underlying video viewer can be exchanged easily if needed in the future. The integration into TIM allowed the creation of a remote interface that provides visual feedback on linking process data with video streams. The GUI design was kept simple and intuitive but it is still able to display up to five video streams of calls from people waiting for access. Furthermore it shows the operators how many people have registered to enter or leave the tunnel. For security reasons the TVV does not have any control functionality. It is completely event-driven by signals coming from the access doors or from a button control panel in the CCC. All of the control signals are provided through PLCs to TIM, and are then forwarded via the Java Messaging Service (JMS) to the TVV without any perceptible time delay. The video streams are MPEG4-encoded for superior image quality as well as low resource consumption.

A view of the control screens. The TIM Video Viewer is in the bottom row, third from left.

A close-up image of the TIM Video Viewer showing the five possible streams.

Due to the success of the TVV for the PS access, the SPS tunnel access will also start using this system during the summer, before the LHC begins.

Useful links

TIM homepage: http://timweb

TIM Video Viewer user manual: http://timweb/wiki/doku.php?id=documentation:tim-video-viewer

Matthias Bräger, GS-ASE

TWiki increases in popularity

The TWiki collaborative web service at CERN (see *CNL* issue 1, 2007) has gained a lot of popularity over the last few years, notably in the LHC experiments. There are now almost 6000 active users and the service has become part of the day-to-day toolbox for web documentation.

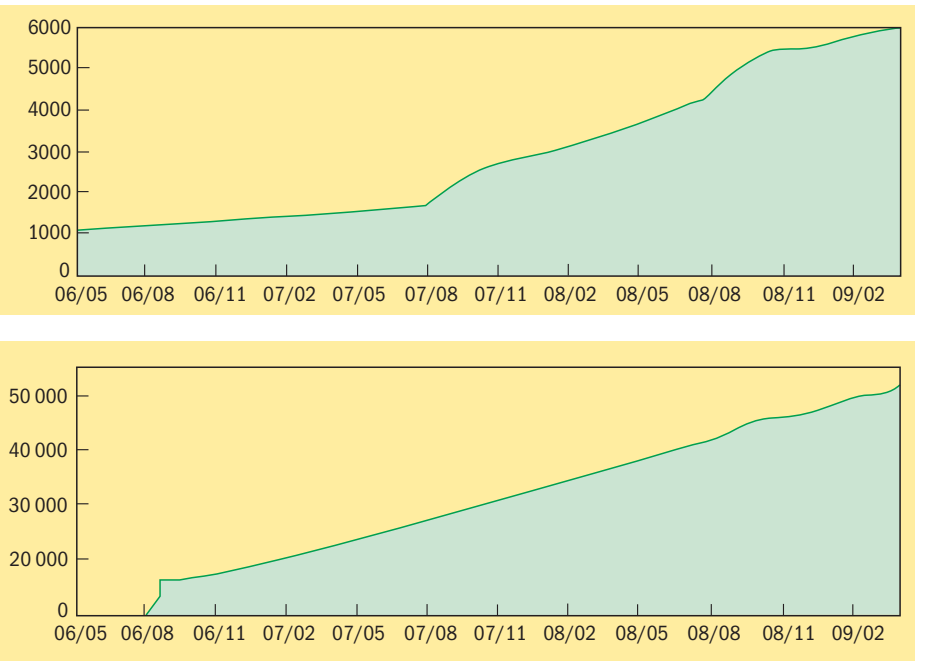

Charts showing the evolution of user (top) and topic (bottom) numbers since May 2006.

Who uses TWiki and why is it so popular?
The TWiki service hosts more than 150 "webs" for dedicated projects and these are used by collaborations such as the LHC experiments, software projects, departmental groups and even social groups. The Wiki concept of shared editable webpages is ideal for documentation in collaborative projects. Wikis are like the original web browser/editor invented by Tim Berners-Lee in 1990. His idea was to create a democratic medium where everybody could contribute knowledge.

TWiki is a Wiki implementation that is easy to use and works on any browser. It has attracted user groups as it allows them to be more efficient by improving the exchange of information. Examples include:

- Replacing a static website.
- Sharing ideas and designs for software projects: virtual blackboards have proven ideal tools for sharing ideas.
- Building physics work books: TWiki allows users to collaborate and create large documents that can be converted to PDFs.
- As a knowledge base like the support web or an FAQ.
- Tracking actions and problems: TWiki has many plug-ins that help adapt to a page to clickable lists and spreadsheets.
- As a web document management tool and archive: thanks to the revision possibilities, you always have the history of changes and can go back to earlier versions of pages.

What are the recent changes?
The popularity of TWiki has made it a victim of its own success and the pages are accessed every second. With the increased load, new servers were installed in 2008, which improved the performance. The different ways that users access their data has also affected performance. To address this we have optimized parts of the functionality, such as the RSS feeds.

Several plug-ins were also updated and installed, such as the image plug-in, which controls the display and alignment of images using an easy syntax and support for server-side resizing and thumbnailing.

A major change came when TWiki adopted the CERN Single-Sign-On service with a common login for all CERN web applications. This makes it easy to authenticate to TWiki for anyone who is a CERN user. To edit documents in TWiki, or be a member of restricted groups, you need to be registered as a TWiki user. Registration is easy by filling in an online form. TWiki uses the CERN login e-mail address for authentication.

What happens if a user changes e-mail address?
When a user changes the e-mail address associated with their CERN login, the TWiki registration needs to be updated and the Helpdesk or twiki.support@cern.ch should be contacted. Further integration between TWiki registration, CERN accounts and E-groups is planned for the end of the year.

Can lightweight registered users use TWiki?
Yes. However, for security, this type of user is vetted before being allowed full access.

How do you use TWiki?
Go to https://twiki.cern.ch/. Here you will find a 20 minute tutorial, a link to register, user documentation and an online support web. A user should have a CERN account and be registered in TWiki to edit pages.

Can pages be protected?
Open, free-form editing is the essence of Wiki culture, which has made TWiki a success. However, if a user wants to protect a page or a group of pages it is possible to restrict access. This is a built-in feature.

What editors are available?
Currently, TWiki topics can be edited within the internet browser by using a simple text editor or a WYSIWYG editor for those who do not like typing text with the Wiki notation. The next version of TWiki will have a new and improved WYSIWYG editor.

When will the new version be available?
TWiki evolved greatly during 2008, resulting in several software releases in two Wiki branches. The latest version was released in April this year and offers many enhancements and new features. Our focus has been to provide stability for the service but if tests prove positive, we can expect an upgrade to this release during the summer.

Can you summarize the new features?

- Usability enhancements including a new and improved WYSIWYG editor.
- Fast SQL style query language for programmable search.
- New restore feature allowing easier rollback and policing of page changes.
- Improved login name/wiki name mapping to enable easier integration with external authentication mechanisms.
- Improved Application Programming Interface (API) for plug-ins.
- Enhancements to many plug-ins.
- A new default skin and easier templates.
- The Table of Content feature is enhanced.
- Security enhancements, reducing the risk of cross-site scripting.
- 300 bug fixes since the current version.

Where can I get more information?
A list of available documentation can be found at https://twiki.cern.ch/. For further enquiries e-mail twiki.support@cern.ch.

Peter Jones and Nils Høimyr, IT-DES

Nightly build and test system supports LHC experiments

The Software Development for Experiments Group (SFT) in the CERN Physics Department, together with the Data Management Group of IT are hosting the LCG Applications Area (AA), which provides several projects that form the most fundamental building blocks for LHC experiment software. Examples for such projects are the ROOT analysis framework, the Geant4 simulation toolkit and projects for data handling such as POOL, COOL and CORAL.

These projects that are developed within AA are augmented by a set of "external" software packages to cover a wide area of functionality, such as mathematical libraries, databases, graphics libraries, scripting languages, Monte Carlo Generators, etc (around 80 at the moment). Most of the external packages are retrieved from sources outside CERN, e.g. Python, Boost, Qt, but also from within e.g. the Grid client software provided by IT-GD. All of these constituents are built and tested with the same version on several platforms, i.e. a combination of operating systems (SLC4, SLC5, Windows, Mac OSX), architectures (i686, x86_64) and compilers (gcc 3.4, gcc 4.0, gcc 4.3, VC7.1, VC9). Being in the transition phase between two SLC releases and two windows compilers, a total of 20 different platforms are provided (figure 1).

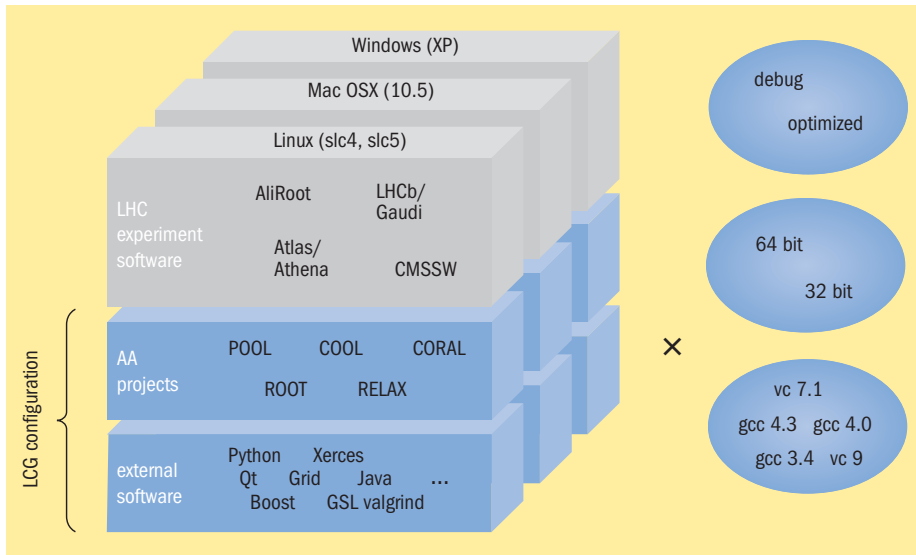

Fig. 1. The LHC software stack showing the 20 different platforms currently provided.

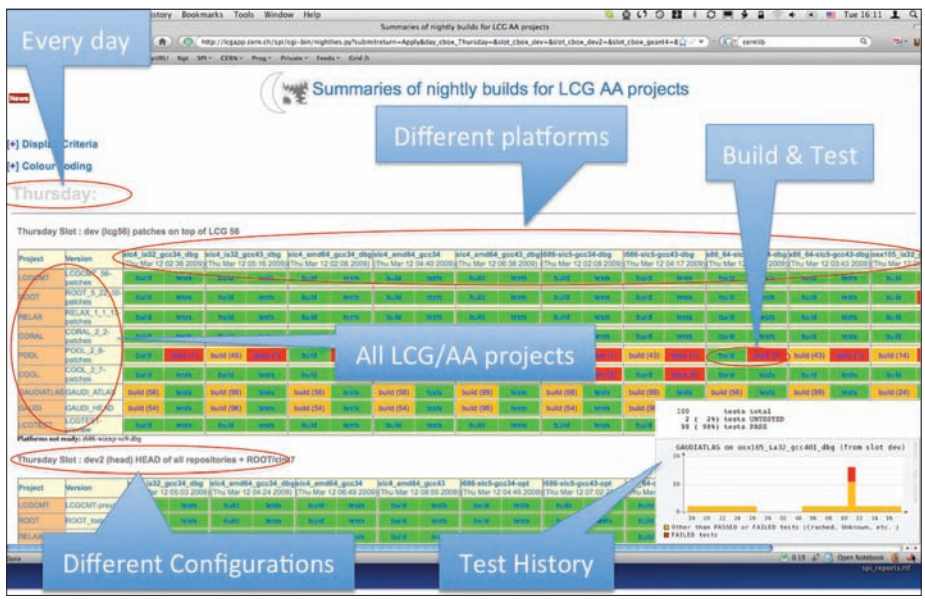

Fig. 2. The SPI nightly builds overview webpage showing results for AA projects.

Releasing AA software

On the requests of LHC experiments this software stack is released all together in so called "LCG Configurations", which are sets of versions of all packages (projects and externals) working together across all of the provided platforms. With this model several challenges for AA developers and LHC experiments emerge:

- LHC experiments would like to test early the latest code changes of the AA projects against their own applications.
- AA software developers need to check and validate their code changes on the different provided platforms.
- The change/upgrade of an external package needs to be validated against the AA stack and experiment code.

The LCG nightly builds

To facilitate all of these tasks a "nightly build and test system" has been implemented by the Software Process and Infrastructure project (SPI), which works in three "dimensions".

- All projects are compiled every night. Usually the building and testing starts around midnight and results are provided later in the morning. The produced binaries and results are kept for one week.
- Projects are being compiled in different configurations, i.e. a set of versions of AA packages that are supposed to work together. A configuration can be set up e.g. to test patches for a production release, the repository "HEAD" of all projects, the upcoming LCG configuration, etc.
- Different platforms. As mentioned above, the compilation of AA software is provided on different platforms. A build machine for each operating system is provided (e.g. for Linux by IT-FIO) and the nightly builds are executed on it.

The nightly build system provides the binary products to experiments for their integration testing and feedback to AA developers about their changes. The results of all building and testing steps are summarized in a webpage that gives a quick overview of the status of the different

platforms for each configuration and day (figure 2). The binary products of the build step are provided in afs (see link below) where they can be further picked up by the users. For example, the Atlas and LHCb experiments are using these products to build their own nightly builds on top of these installations. With this "chain of nightlies" we are able to spot any potential problems in the whole LHC software stack early and take appropriate actions.

The nightly build scripts are implemented in Python allowing easy deployment of the tool across different operating systems. The steering of the nightly builds configuration is done via an XML file and the actual build instructions are encapsulated in Python "builder classes", making it easy to extend the system for other users. An open source product, the CMT tool is used for configuration management and software building. This facilitates the handling of dependencies between AA packages and also provides the build instructions for each package. The use of CMT is optional; the nightly builds would also work without it.

Use of the nightlies

The nightly builds are a tool that greatly helped with the process of releasing AA software. Before this, every project had built its own release one after the other discovering possible problems and inconsistencies very late in the process. Nowadays, as we have the overview of all projects on all platforms, when it comes to releasing the AA software we "rubberstamp" the state of the whole AA stack once satisfied with its quality and release immediately. The actual release of the projects is done with exactly the same build instructions as the nightly builds, thereby guaranteeing that the software is built in exactly the same way.

The use of the nightly builds (also for releasing the AA stack itself) allowed for the production of an "LCG Configuration" within hours instead of in the order of days (sometimes weeks). This fast turnaround will be crucial once the LHC data taking has started and any bug fix to the software has to be provided as quickly as possible.

The nightly build system itself is not only used for AA project testing but also by other parties at CERN. It has been adopted by the Geant4 and LHCb Offline Computing team for its software testing showing that the tool can be easily extended for users outside AA. Moreover each of the parties is not only using the software but also contributing with new developments and ideas to the overall project leading to mutual benefits.

Useful links

LCG-AA Software Elements: http://lcgsoft.cern.ch
Nightlies webpage: http://lcgapp.cern.ch/spi/cgi-bin/nightlies.py
AFS path to nightly build binary products: /afs/cern.ch/sw/lcg/app/nightlies
CMT tool: http://www.cmtsite.org
Geant4 nightly summaries: http://tinyurl.com/klvfoj
LHCb nightly summaries: http://cern.ch/lhcb-nightlies
PH-SFT homepage: http://cern.ch/ph-sft

Stefan Roiser, PH-SFT for the "LCG nightly build developers", **Ana Gaspar** (LCG/SPI), **Yves Perrin** (LCG/SPI), **Victor Diez** (Geant4) and **Karol Kruzelecki** (LHCb offline)

Training on computer security keeps cyber attackers at bay

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

Secure e-mail and web browsing

The aim 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 and discusses security in different phases of the software development cycle. Topics include:

- threat modelling and risk assessment;
- protection, detection, response;
- security through obscurity;
- main security principles (defence-in-depth, least privilege principle etc);
- security in different phases of the software development lifecycle – architecture, design, implementation, testing, deployment, maintenance;
- implementation (coding) – common pitfalls and security bugs, advice on best practice for security development.

There is also a forthcoming course dedicated to Java programming.

These courses, although not hands-on, are interactive and full of real examples. They are available from the CERN Training Catalogue, but you can contact the Computer Security Team (Computer.Security@cern.ch) for more information.

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

Useful link

CERN Training Catalogue: http://cta.cern.ch/cta2/f?p=110:9

CERN Computer Security Team

The deadline for submissions to the next issue of CNL is

14 August

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

CERN celebrates another key contribution to the internet

CERN recently celebrated the 20th anniversary of the first web proposal. This was a fundamental achievement that everyone knows at CERN. Amazingly, another key contribution from CERN seems to be better known outside: the role that it played during the period 1982–1994 in the creation of the internet, i.e. the transnational transport network that supports all applications, including the web. This network, made up of switches and links, uses addresses to identify connected computers. They are called IP addresses. RIPE (Réseaux IP Européens) – the organization that allocates IP addresses in Europe today – stems from an initiative in which CERN was instrumental and is the outcome of a small meeting held at CERN.

On 7 May this year in Amsterdam, RIPE celebrated the 20th anniversary of its creation, where veterans of this epic tale were invited, including Olivier Martin, now retired from CERN, and me (see www.ripe.net/ripe/meetings/ripe-58/).

The ceremony was part of the 58th RIPE meeting, which assembled more than 500 attendees. The very first meeting also took place in Amsterdam on 22 May 1989 and was attended by just 14 people. But how was CERN involved with this?

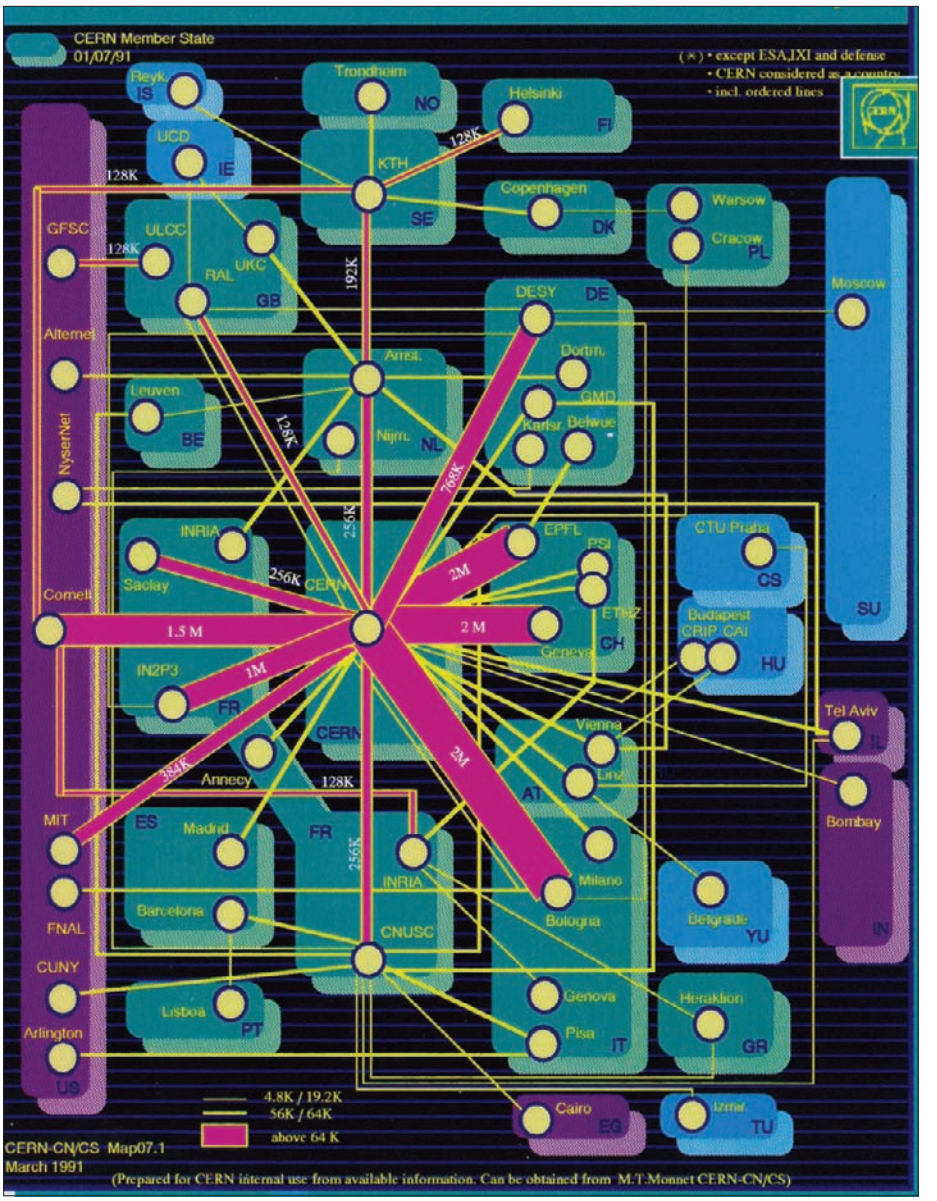

A map of the European internet in the early 1990s: a big star with CERN at the heart.

Crystals

Capitalizing on its experience gained in the 1970s with the on-site network CERNet (a genuine suite of networking technologies developed entirely at CERN, and following the principles of today's internet), CERN started a programme of external networking in the early 1980s. The major HEP centres in Europe set up direct links to CERN and soon physics centres in other regions requested direct connections as well. At that time, the centres and CERN were connected directly by means of leased circuits specially constructed by telecommunications operators. Often, these connections were ground-breaking where major technical or administrative obstacles had to be overcome. For example, the first data connection between China and the scientific world was carried out in July 1988 by means of a slow link set up between IHEP and CERN by me and my colleague Giorgio Heiman during a memorable trip in Beijing. The leased circuit involved multiple operators and it traversed China, the former USSR and eastern Europe.

Networks grow like crystals: when one hub starts crystallizing, everyone wants to connect to it and not elsewhere, as everyone prefers to be just one hop away from the heart. Therefore, non-HEP scientific centres started to set up links to CERN, often serving as relays to other domestic HEP sites.

The heart is CERN

At the end of the 1980s, CERN had become the centre of an enormous star-shaped network, one of the largest in the world. In 1989 when Olivier Martin organized the conversion of these links to IP technology, CERN immediately became the largest internet hub in Europe. Just imagine: in 1991 80% of the internet capacity installed in Europe for international traffic was terminated at CERN in Building 513. It is no surprise that the performance of Tim Berners-Lee's first web server impressed the world: it was at the heart of the European internet, benefiting from the fastest links and was just a few hops from most destinations.

In fact a year before, it had become clear

to Olivier Martin, Brian Carpenter (leader of the Communications Systems group) and me (in charge of External Networks) that the technical conversion to IP technology of the entire CERN External Network was necessary and inevitable. However, it was also obvious that the management of our giant hub would eventually become problematic with the operational arrangements of that time. In particular, back in 1988, IP addresses were allocated for the entire world by an American organization, the Network Information Centre (NIC). It was evident that eventually the European internet would not remain as a star around CERN; this is not CERN's mission. There was an obvious need for a structure to coordinate the management of the internet in Europe.

Five of the RIPE pioneers at the 20th anniversary celebrations in May. From left to right: François Fluckiger, Daniel Karrenberg, Enzo Valente, Olivier Martin and Rob Blokzijl.

Abandon of power

I had the honour of convening the very first meeting of the Coordinating Committee for Intercontinental Research Network (CCIRN) in May 1988, co-hosted by CERN and the University of Geneva. This committee was the first attempt to harmonize the inter-regional operation of the emerging worldwide research network.

The second meeting took place in October 1988 at a summer resort in Western Virginia, US. The main topic was "the future of the internet" and the Americans turned up in force, with representatives from the Department of Energy and the Department of Defense, together with Vint Cerf, co-inventor of the TCP/IP protocols. The European representatives were thin on the ground: a German representative, a British representative and me representing CERN. At that time, the European policy was opposed to the internet technology, hence the lack of enthusiasm to send participants to an internet-focused meeting from the National Research and Education Networks.

Towards the end of the meeting the Americans, led by Vint Cerf, became amazingly insistent: "It is essential that you, the Europeans, set up a structure to allocate internet addresses in Europe. It is not good that we keep doing it for you." My two colleagues did not seem too concerned, replying that they did not foresee the use of the internet in Europe. I comforted our American colleagues, replying that there were indeed several people in Europe that believed in IP and were willing to set up such a structure.

The Americans, sometimes accused of wanting to dominate the internet world, showed great respect towards the Europeans on this occasion.

The gang of six

On my return from Western Virginia, Olivier and I called a meeting at CERN with the interested parties from Europe to promote the idea of a structure to coordinate IP networks in Europe, and in particular the attribution of addresses. Time was pressing.

This founding meeting took place at CERN in December 1988, in a small office of Building 600.

Six people participated: Rob Blokzijl from Nikhef (NL), Mats Brunel from SUNet (SE), Daniel Karrenberg from EUNet (NL), Enzo Valente from INFN (IT), and Olivier Martin and me from CERN. Rob said: "The weather was horrible. Cold, heavy rain. The room was dark and smoke-filled. We had the impression we were writing a page of history". Enzo recalled: "The tension was palpable. The ideas were popping up so fast. I needed to smoke. Every 10 minutes I had to get to the door, my arm and the cigarette outside in the corridor, my body inside." Indeed, the undertaking was technically and politically complex (IP was not at that time the technology favoured by politicians in Europe).

At this meeting, Daniel Karrenberg proposed the name RIPE for the organization that we were creating. A French name and an acronym that was also relevant in English: time was indeed RIPE for the Europeans to organize the internet in Europe.

Six months later, the first meeting of the RIPE forum took place in Amsterdam chaired by Rob Blokzijl. A bit later, an operational structure was created, the RIPE-NCC (Network Coordination Centre) in the form of an independent, not-for-profit membership organization.

It was essential that this organization was neutral and fully independent, not controlled by particular governments or commercial companies as it was to manage what we expected would become crucial resources for the internet and therefore for the society at large.

Daniel Karrenberg became the first managing director of the RIPE-NCC.

Twenty years later

Twenty years later, the RIPE forum continues at a pace of two to three meetings per year. The RIPE-NCC is still the authority that allocates internet resources and related services to internet service providers, telecommunications organizations and large corporations located not only in Europe, but also in the Middle East and parts of Central Asia. It is one of five Regional Internet Registries providing internet resource allocations, registration services and coordination activities that support the global operation of the internet. It employs 100 people and has 5000 contributing members (see www.ripe.net/membership/indices/). Even though RIPE represents the interests of its members, most of which are commercial firms, its neutrality remains a guarantee against the pressure that is inevitable when managing scarce resources.

What are the gang of six, who laid down the principles, doing now? Rob Blokzijl is still the chairman of the RIPE forum. Daniel Karrenberg, after spending more than a decade as managing director, is now the RIPE-NCC chief scientist. Enzo Valente is the chairman of GARR, the Italian Academic and Research Networks. Olivier Martin retired from CERN a few years ago. We have not been able to locate Mats Brunel.

After the launch of RIPE, CERN continued its contributions to the development of the internet in multiple fields. For example, CERN was a founding member of the Internet Society (ISOC) and of the European Backbone (Ebone). It also set up the first high-throughput internet transatlantic link with IBM – but that is another story!

Useful links

RIPE: www.ripe.net/
Le Réseau des Chercheurs Européens: www.larecherche.fr/content/recherche/article?id=14841

François Fluckiger, IT-DI

Bookshop prepares for IT Book Fair and integration with Stores

The 2008 CERN Book Fair took place on 8–10 October and was a great success. During the three-day event, we had six authors presenting their new books and eight publishers exhibiting their wares. Given this large interest, the Bookshop plans to organize an annual Book Fair, alternating a Physics and Engineering Book Fair with an IT-themed one. Therefore, the next Book Fair will focus on computing books and will take place on 3–4 September this year. Key publishers and book distributors have been invited and special side events, such as book presentations, are already being prepared. For example, Bjarne Stroustrup, who developed the C++ programming language, is scheduled to give a talk on Thursday 3 September.

Integration with CERN Stores

The Bookshop's stock will soon be integrated into the CERN Stores catalogue. Therefore, it will also be possible to purchase books via the EDH interface using a CERN budget code. Private purchases are still possible by paying directly into our UBS account. CERN users can buy books and CDs at discounted prices.

The Bookshop is located in the Central Library, Building 52/1-052, and is open on weekdays from 8.30 a.m. to 5.30 p.m. The Bookshop is staffed by Eva Papp; and since November 2008 Tullio Basaglia has taken over the management from Joanne Yeomans. We purchase books from about 15 publishing houses, however, the offer is not limited to these publishers – the staff are ready to order any title that exists in print. In addition, we welcome suggestions from the user community, especially concerning new titles that you wish to be added to the stock. You can contact us by e-mail at bookshop@cern.ch.

There are many new titles available in the Bookshop. The following list is a sample of some of these.

Computing

D M Smith 2009 *Engineering Computation with MATLAB*.
M Summerfield 2008 *Programming in Python 3: A Complete Introduction to the Python Language*.

Physics

D J Griffiths 2008 *Introduction to Elementary Particles* 2nd edition.
B Martin 2009 *Nuclear and Particle Physics: An Introduction* 2nd edition.
B Martin and G P Shaw 2008 *Particle Physics* 3rd edition.
D McMahon 2009 *String Theory Demystified*.

Accelerators and detectors

W Blum, W Riegler and L Rolandi 2008 *Particle Detection with Drift Chambers* 2nd edition.
C Grupen and B Schwartz 2008 *Particle Detectors* 2nd edition.
H Padamsee 2009 *RF Superconductivity: Volume II: Science, Technology and Applications* (available in June 2009).
J T Tanabe 2005 *Iron Dominated Electromagnets: Design, Fabrication, Assembly and Measurements*.
T P Wangler 2008 *RF Linear Accelerators* 2nd edition.

Other topics

F E Close 2009 *Antimatter*.
C Leroy and P-G Rancoita 2009 *Principles of Radiation Interaction in Matter and Detection* 2nd edition (available in June 2009).
H Schopper 2009 *LEP: The Lord of the Collider Rings at CERN 1980–2000* (available in June 2009).

The full list of current titles is available at http://cdsweb.cern.ch/collection/CERN%20Bookshop and http://cdsweb/tools/itbook.py.

Tullio Basaglia, GS-SI

Calendar

July

1–4, **European Workshop on Challenges in Modern Massive Data Sets (EMMDS 2009)**
Copenhagen, Denmark
http://mmds.imm.dtu.dk/

5–17, **7th International Summer School on Grid Computing (ISSGC 09)**
Sophia Antipolis, France
www.issgc.org

August

17–28, **CERN School of Computing (CSC 2009)**
Göttingen, Germany
http://cern.ch/CSC

25–28, **European Conference on Parallel Computing 2009 (Euro-Par 2009)**
Delft, the Netherlands
http://europar2009.ewi.tudelft.nl/

September

7–9, **Third International ICST Conference on Networks for Grid Applications (GridNets 2009)**
Athens, Greece
http://gridnets.org/

15–17, **5th International Conference on IT Security Incident Management and IT Forensics (IMF 2009)**
Stuttgart, Germany
www.imf-conference.org/

16–18, **NORDUnet Conference 2009**
Copenhagen, Denmark
www.nordu.net/conference/ndn2009web/welcome.html

21–23, **European Symposium on Research in Computer Security in conjunction with the 12th International Symposium on Recent Advances in Intrusion Detection (RAID 2009)**
Saint Malo, France
www.esorics.org/

21–25, **EGEE'09 Conference**
Barcelona, Spain
http://egee09.eu-egee.org/

Special EGEE issue of Grid computing journal

EGEE is teaming up with SpringerLink, the publishers of the *Journal of Grid Computing*, to produce a special issue entitled "EGEE applications and supporting technologies".

This will give the 96 oral presenters at the User Forum an opportunity to prepare a full paper describing their work, which will be reviewed by their peers.

The publishers are looking for 15–20-page papers describing original and significant work involving applications on EGEE, and the use and development of key Grid technologies. The areas of interest are:

- Scientific results obtained using Grid technology.
- Planned or ongoing scientific work using the Grid.
- Experiences from application porting and deployment.
- Grid services exploiting and extending the gLite middleware (job management, data management, monitoring, workflows, etc).
- Programming environments.
- End-user environments and portal technologies.
- Emerging technologies within the EGEE infrastructure (cloud, virtualization, etc).

Author instructions are available at www.springer.com/computer/communications/journal/10723.

The deadline for paper submission is 31 July 2009, final acceptance confirmation will be in March 2010 and the final edition is due to be published in September 2010.

Danielle Venton, IT-EGE

This article was published online in the Spring 2009 EGEE Newsletter, available at http://eu-egee.org/newsletter/spring09/spring09.html.