

CERN COMPUTER NEWSLETTER

Volume 42, issue 5 November–December 2007

Contents

Editorial

EGE group manages global complexity 1
Computing articles featured in this month's *CERN Courier* 2

Announcements & news

Next-generation EVO replaces VRVS for global communication 3
Authentication to TWiki at CERN has changed 4
VPN service will close 4
Some services require primary CERN account 4

Desktop computing

Single sign-on facilitates authentication at CERN 5
Managing your CERN account 7

Grid news

ATLAS: the data chain works 8
European Grid Initiative presented at EGEE'07 8
Visualizing the state of your grid with GridMaps 9
LHC@home server sets up new home in the UK 9

Technical brief

Web applications security: risks and countermeasures 10
Twelve steps to improving control systems cyber security at CERN 11

Conference & event reports

CHEP focuses on LHC computing in last meeting before start-up 14
Second multithreading workshop attracts CERN programmers 15
Workshop will forge links with financial sector 15

Information corner

It's all change at the bookshop 16
Computer security: think before you click! 16
Calendar 16

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

Produced for CERN by IOP Publishing Dirac House, Temple Back, Bristol BS1 6BE, UK. Tel +44 (0)117 929 7481. E-mail joseph.tennant@iop.org. Fax +44 (0)117 920 0733. Web iop.org.

Published by CERN IT Department

©2007 CERN

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

IOP Publishing



EGE group manages global complexity



Members of CERN's EGE group, which manages a consortium of 91 institutions worldwide.

In the last of our series of editorials profiling the different groups within the IT department, *CNL* speaks to Bob Jones, the director of the Enabling Grids for E-science (EGEE) project and leader of the EGE group in IT. We ask him about the challenges his group faces managing the worldwide consortium of 91 partner institutions that are participating in the EGEE project.

What is the role of the EGE group?

The EGE group at CERN hosts the management activity for the EGEE project. This includes myself; the technical director of EGEE, Erwin Laure; and the project's administration, which looks after all the relations we have with the consortium of 91 partners concerning administrative and financial matters. The EGE group also handles all relations with the European Commission on behalf of the consortium, including, for example, the regular reviews of the project.

Our group also manages dissemination and outreach for EGEE, and we are responsible for the use by the EGEE partners of CERN-based IT tools – such as Project Progress Tracking (PPT), Indico, and Engineering Data Management Service

(EDMS). Members of the group also work on related projects, such as the Diligent project for grid-based digital libraries; and we run the Related Projects Liaison Office, which manages relations with other grid projects and application projects around Europe. In total this represents 16 people. In addition, a host of fellows and associates are registered with our group but they work in the different technical activities inside the EGEE project.

How many people are involved in the EGEE project?

What people don't always see at CERN is that EGEE is a multisite, multipurpose infrastructure. While the most prolific users are the physics community for the LHC, more than 5000 people are registered to use the infrastructure, and we estimate that more than 12000 scientists benefit from its existence. About 500 people work full time for the EGEE project today; to put that in perspective it means we manage the equivalent of about 20% of CERN personnel.

Through the PPT project-planning tool we know that more than 1100 people have worked for the project in one way or another over the nearly four years of its

existence. The European Commission and external reviewers have repeatedly praised the EGEE project for its ability to manage such a complex international effort. Indeed we are now used as a role model for the administration of other large EU projects.

What have been the highlights of the EGEE project over the last year?

People often think of EGEE in terms of either the infrastructure or the middleware. These are just two out of 11 activities in the project. There's also training, applications support, outreach, work on standards and international collaboration with other grids, work with the business community, and so on. For example, two EGEE training events are held every week somewhere around the world. These have been mainly end-user and applications interface training, but as the middleware is deployed on other infrastructures we are now getting many requests for site manager training, and from private companies.

Nevertheless, the scale and level of service offered by the infrastructure remains a key measure of the project's success. We ended the European Data Grid project four years ago with 1000 CPUs at 20 sites. We now have more than 40000 CPUs in 250 sites across 48 countries, and about 12 PB of storage. We anticipate that the processing and storage capacity of the EGEE Grid will easily double before the LHC start-up.

The quality of the service of the infrastructure has improved significantly over the last year. This is thanks to improvements in the operational procedures, and a suite of monitoring tools and testing techniques that have been developed in the project, some of this in collaboration with industry. The throughput of the EGEE Grid is now more than 100000 jobs per day, and 68% of that is related to the LHC. But the remaining 32% represents a dramatic increase for other scientific disciplines. These are accelerating in their use of EGEE relative to the LHC users.

Perhaps the most impressive result over the last year, though, is that we are managing this much larger and more complex infrastructure with the same manpower as in the first two-year phase of EGEE, which ended in 2006. So the infrastructure is increasing rapidly but not the manpower to maintain it.

What challenges face the EGEE project over the coming year?

Regarding infrastructure, one challenge for the future is that many of the other sciences have been piggybacking on resources made available for the LHC. As the LHC comes online it is likely that capacity will be used, so we are working with the other disciplines to ensure that they contribute more resources to the infrastructure.

There are some issues here. Do these disciplines have the personnel to connect more resources? Does the gLite middleware support all the platforms that are used in other disciplines? At the moment gLite is constrained by the Linux distributions it works on. For example, in the life sciences we have partners who have resources but only if they can run on Windows. This is why we have been working with new EGEE business associates to better support the Windows platform.

With the gLite middleware there has been a restructuring this year to take out redundancies and unnecessary connections and constraints. The goal is to factorize gLite into more independent components that can be ported and deployed in a simpler fashion.

We have recently had confirmation that the European Commission will fund a third stage of the EGEE project, starting next spring. The key mission of the project in this third phase will be to assure the continued operation of the production infrastructure through to 2010, which is particularly crucial for the physics community during the LHC start-up. In parallel with this, the ambition is to achieve a sustainable grid infrastructure, with the vision of a permanent European Grid Initiative (EGI) that would take over from

EGEE at the end of its next phase.

Much of the work in EGEE-III will be aimed at restructuring the consortium of partners to ensure a long-term grouping based on national grid initiatives. An EGI design study project was launched in September, with the support of 37 countries. The project, of which CERN is a founding member, will design the operational and governance model for a sustainable grid infrastructure. This will be a challenging political process, and it has to converge rapidly to ensure the future of the grid infrastructure on which the LHC depends.

Another challenge is to encourage industry to take up this technology. On the one hand many companies are interested in developing products and services based on our grid standards, and this could help support the grid infrastructure in the longer term. On the other hand the big companies like Microsoft, Google and Yahoo! are not interested in grid standards since they see a competitive advantage in having a private infrastructure and selling services on it. EGEE wants to go for the open standards approach, to create a market in the same manner as has happened for the web. If Tim Berners-Lee had not insisted on open standards what value would the web have had? We've been here before, so let's try to do the right thing again.

Computing articles featured in this month's CERN Courier

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

Computing News

● ATLAS data chain passes full test

The first "end-to-end" chain is a success.

● GridPP tackles Tier-2 performance problems

Team assesses readiness of UK clusters.

● Fermilab passes new milestone for data traffic

Tier-1 site records 6 PB of data on tape.

● European supercomputer network ramps up to 10 Gbit/s

DEISA boosts speeds between its 11 sites.

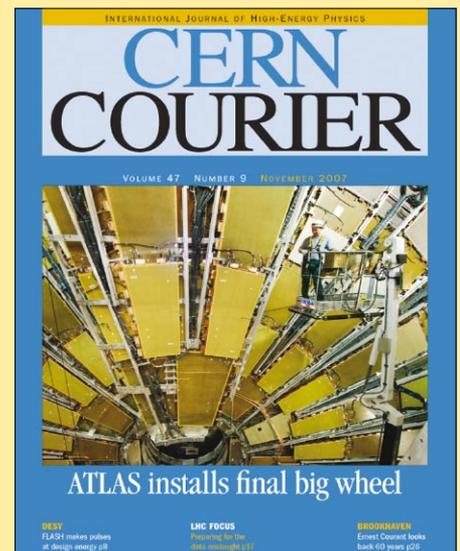
● Gigabit Ethernet links CERN with Mumbai

TIFR to benefit from undersea cables.

Feature articles

● LHC computing stability emphasized at CHEP'07

This year's conference focused on getting ready for the arrival of data from the LHC.



● CASTOR2 rises to LHC's data storage challenge

A report on the latest developments in CERN's mass-data storage system.

Calendar of events

Next-generation EVO replaces VRVS for global communication

The LHC and other high energy physics (HEP) programmes face unprecedented challenges in their need for scientists located at sites around the world to work collaboratively on data analysis and other activities during the construction, commissioning and operation of their experiments. Physicists throughout the world may be separated from the experimental site, and from their colleagues, by up to 12 time zones. For the last 10 years the web-based Virtual Room Videoconferencing System (VRVS) has been a mainstay for remote collaboration throughout the HEP community, but this is about to change as VRVS is replaced by the next-generation system, Enabling Virtual Organizations (EVO).

Virtual room videoconferencing

VRVS first went into production in 1997. It was developed and managed by Caltech and its partners; funded by the US Department of Energy and the National Science Foundation; and has been supported by CERN's IT department.

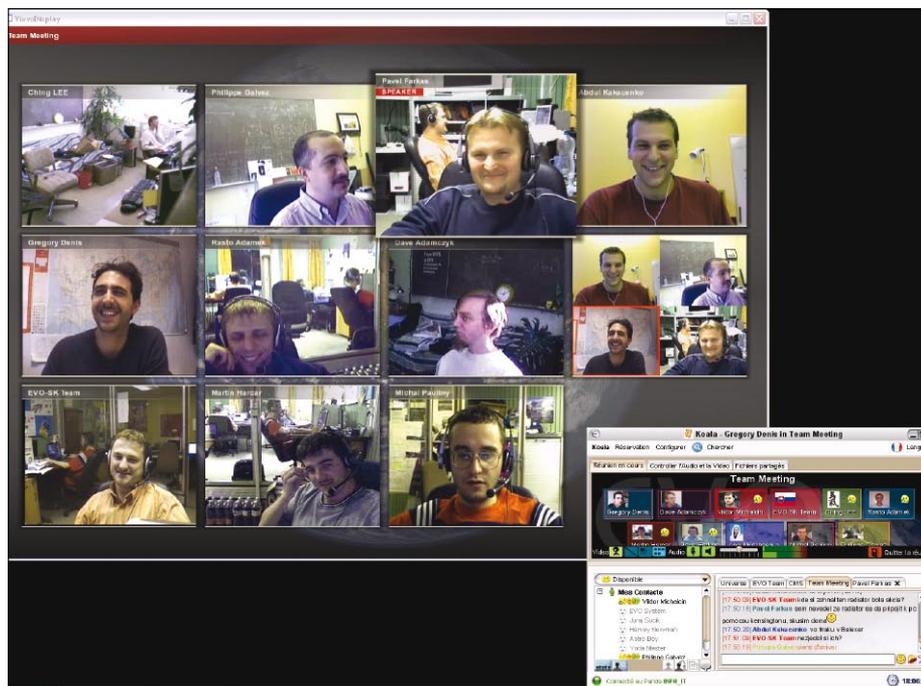
Since its introduction the system has provided a cost-effective global collaboration service for the HEP community. Until recently VRVS was unique in its overall scalability and its ability to support real-time collaborative sessions across a variety of working environments, including individuals using desktops, laptops or mobile devices; work groups in small or medium-sized conference rooms; teams in control rooms; and large groups in auditoria.

The VRVS user community has grown exponentially, by approximately a factor of two each year for the last 10 years, and now includes users in 130 countries. VRVS has thus become a standard part of the toolset used daily by a large sector of HEP, and is used increasingly by other programmes.

VRVS has met the milestones and expectations of its original development and deployment plan, and is widely recognized worldwide as a useful and efficient collaboration facility. However, as it was used on a bigger scale the VRVS system began to show its limitations, especially during recent large meetings involving many sites.

Introducing EVO

EVO was officially released in June 2007, and it represents a major step forward in providing the HEP community with a more effective collaborative system. It has greater functionality, scalability and



A snapshot of an EVO session showing the OpenGL video window (top) and Koala GUI.

robustness, including resistance to network and end-system problems. One key change is that the collaboration infrastructure has been partially automated by deploying a set of intelligent software agents in every element of the system, including the end-users' computers and the networks interconnecting them. The software is based on Caltech's MonALISA system (Monitoring Agents in a Large Integrated Services Architecture. See <http://monalisa.caltech.edu>). This change has transformed the collaboration service offered to the physics community into a resilient, adaptive global system.

Compared with VRVS, EVO includes a better integrated and more convenient user interface, it has a richer set of features with higher resolution video and instant messaging, it is more adaptable to all platforms and operating systems, and overall is more efficient and robust.

All of these aspects are particularly important as we approach and then enter the start-up period of the LHC, because the community will engage in an unprecedented level of daily collaboration. There will be intense demand for long-distance scheduled meetings, person-to-person communication, group-to-group discussions, broadcast meetings, workshops, and continuous presence

at important locations such as control rooms and experimental areas. It will be crucial that the collaboration tools are totally integrated into the physicists' working environment. EVO and its future upgrades are designed to provide a cost-effective system that will continue to scale and adapt to the current technologies, and meet the demands of the community in the LHC era with a relatively modest level of manpower dedicated to operations and support.

A unique collaboration system

EVO provides a unique, unified collaboration system for both the LHC community and the physics community at large. It integrates several widely used video and audio technologies. Through a series of developments capitalizing on the team's 12 years of accumulated experience and expertise, the system now includes advanced visual features that provide a unique level of high resolution, interactivity, and information management during working meetings that are not available elsewhere.

On EVO systems running Windows – and soon the other major operating systems as well – the live videos and portions of the users' desktops are embedded in an OpenGL graphics window where they “live”

Announcements & news

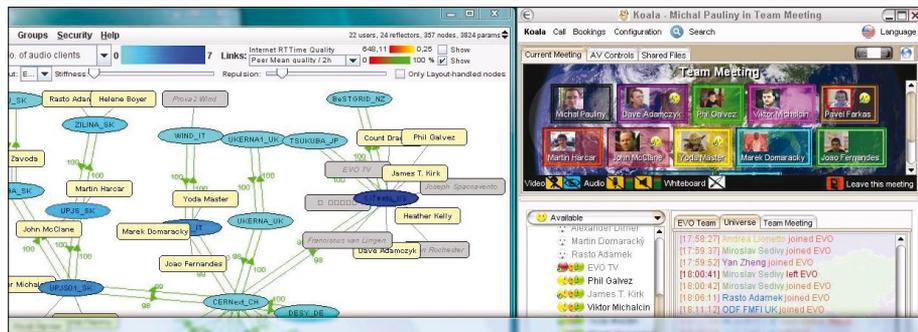
in a three-dimensional space. This enables a collaborative session to be viewed on a conventional desktop or projected, and also to be viewed in 3D, permitting a large amount of streamed and static information to be “stored” and presented as needed. A large virtual space is available on a normal desktop screen and this can be projected in medium and large rooms.

As mass-market display screens and graphics card technologies advance, the amount and richness of the information that can be handled at low cost in these sessions will continue to increase. At the moment a typical desktop set-up might include a PC with a 30 inch display with 4 million pixels for a few thousand dollars, or a large plasma display for workgroup meetings, or a high-resolution projector for large meetings.

At the other end of the scale, users can also access EVO meetings using the standard, non-digital telephone system (POTS) through gateways, the first of which have been installed at CERN and Caltech. Gateways at other HEP labs and universities are planned and will be added as the service is scaled up. This will enable a user to simply call one of the gateways closest to his/her location and save the cost of the call. Users can also access and connect to an EVO meeting via an H.323 endpoint or a multiple control unit, such as the ESnet system or Lyon’s HERMES system, for example. EVO also supports the session initiation protocol (SIP), which is the main VoIP protocol deployed today in industry.

Key features

The main basic features of EVO include: support for the three main operating systems (Windows, Linux and MacOS); IPv6



An EVO session showing the OpenGL 3D video window (bottom) and Koala GUI (top).

and Multicast integration; playback and recording functions; a shared-files function; whiteboard; and a Poll/Vote function.

Application programming interfaces will be provided soon for external development/integration, such as Indico (in progress), Shibboleth authentication (done), and Grid EE certificates (in progress). In this way users will be able to access the EVO service with a single sign on by simply using their institute

or collaboration login.

The team is now encouraging VRVS users to move to EVO as soon as possible, as the support for EVO exclusively will begin in December. More information on EVO, its many features and functions, and how to get started, can be found at <http://evo.caltech.edu>. Questions and comments are welcomed by the EVO team, and can be sent to evosupport@vrvs.org.

Philippe Galvez and Harvey Newman, Caltech

Authentication to TWiki at CERN has changed

On 18 September the method of accessing the CERN TWiki service changed. Users must now authenticate with the CERN Single Sign On procedure (see <https://cern.ch/authentication>), using their CERN login instead of their AFS username and password.

To log on to TWiki with the CERN login, you can use your CERN account, e.g. the existing MAIL, NICE or AIS/EDH login, or your “external account” if you are a non-CERN external user, as for Simba mailing

lists or Indico. External users without a CERN account may obtain an external account from <http://cern.ch/cernaccount/Externals>. One advantage of the new procedure is that a user whose AFS account has expired will still be able to access TWiki. Information about how to manage your CERN account can be found at <http://cern.ch/cernaccount>.

Please note that only registered TWiki users can now edit pages with TWiki. If you are not yet registered you will be prompted to do so before editing a page. The user details will be hidden from the internet and only be accessible to authenticated users.

CERN TWiki Support team (twiki.support@cern.ch)

VPN service will close

Access to CERN using the Virtual Private Network (VPN) service will be discontinued on 29 January 2008. This is because of continued incidents and the growing security risks posed by the service. New registrations will no longer be accepted.

Users are requested to stop using VPN immediately and instead use the methods recommended for connecting to CERN from the internet. These are outlined at <http://cern.ch/security/Internet>.

Background information to the closure of the service can be found at <http://cern.ch/security/vpn/vpn-closure.asp>.

CERN Computer Security Officer

Some services require primary CERN account

We would like to remind users who have several CERN accounts, and have two or more different login names to access the central applications, that they can only use their primary CERN account to access certain applications.

This is the case for accessing EDMS or AIS applications such as EDH, and also for accessing the Listbox site (<http://cern.ch/listboxservices>) for mailing lists. Users who try to access this site with a secondary CERN account will see an “access denied” screen and a message inviting them to use their primary account and login name.

As a reminder, the primary CERN account is the only one that can be e-mail enabled, or that can access AIS applications.

Please also see the article “Managing your CERN account” on p57, which details CERN’s policy on account management.

The User Support team

Single sign-on facilitates authentication at CERN

Signing on to a service such as e-mail involves entering a username and password: your credentials. What happens behind the scenes when you enter your credentials is called authentication: the computing service in question establishes that you are the owner of a valid account and gives you access to it. Managing who gets access to which computing services, and with what privileges, is called authorization. Typically, the service manager authorizes the access and privileges of individuals and groups.

For authorization to work, information about individuals has to be associated with the groups they belong to and roles they play in the organization – a process known as identity management. This ultimately involves information managed by the organization's HR service. This situation is summarized in figure 1.

Several years ago in the course of an average working day a typical CERN user would have to authenticate many times using many different credentials. First a login and password were required to unlock the Windows desktop, then the user had to type in another login and password to unlock the Linux desktop or session. Credentials were also needed to read e-mails, use administrative applications, and submit a CHEP presentation on Indico, for example. Today the situation is better but still not optimal. Many applications use the same credentials – an improvement.

CERN authentication

The aim of the CERN authentication service is to provide a single sign-on for CERN applications, using a unique central account database.

The first step to achieving this was to reduce the number of user databases. Cutting the number of login and password pairs to remember has a direct impact on user satisfaction and on security: it is easy to remember one password.

Centralizing the user database in one location avoids leaks and reduces potential security holes. It also helps to provide extended services like external accounts management (using a lightweight registration process for non-CERN users). The centralized database of CERN users also provides groups and roles membership, to enhance access rights to resources.

To further improve security, the

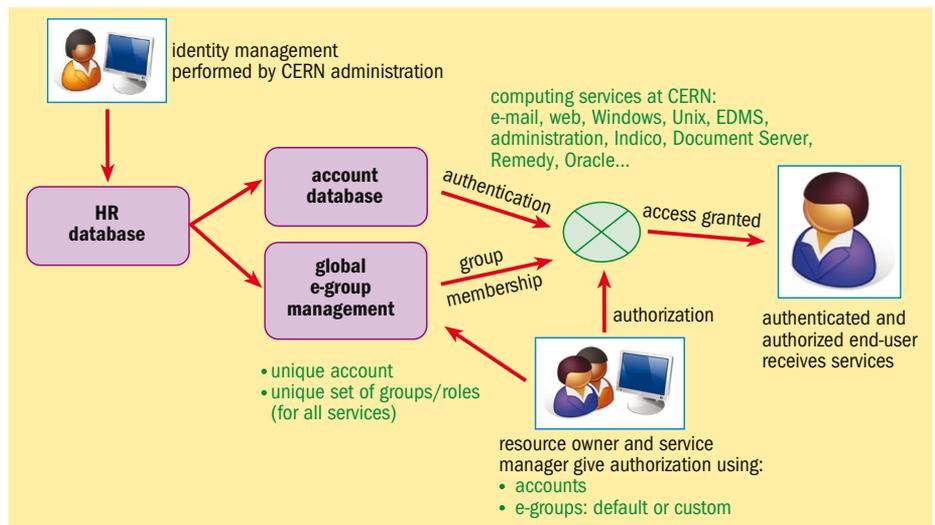


Fig. 1. An overview of the identity management system at CERN, and the players involved.

centralized database means that access by a user to all applications can be blocked with just one click. Users can give others access to resources by means of permissions and delegation instead of sharing credentials: shouting a password in the corridor is no longer an option.

Authentication methods and user data

The CERN Single Sign On (SSO) infrastructure provides different authentication methods, with different security levels:

- Classic web forms, where users type in the login and password. Depending on the browser features, the login and password can be saved locally for later use.
- Windows integrated authentication. The current Windows session token is reused to authenticate the user, without having to retype credentials. The security of Windows desktop sessions has also been increased by deploying a strong screen-lock policy: screen savers with password locking have been forced on all Windows desktops with a short time out, to avoid users leaving their screens unlocked for long periods of time.
- Certificates. Authentication can be made using a certificate provided by the CERN Certification Authority (CA) or any grid-trusted CA member. Depending on the security level required, SmartCard tokens with pin codes can also be used.

The calling application can select all or some of the authentication methods. For example, to ensure maximum verification,

experiment controls can request operators to authenticate only with certificates on SmartCards.

The single-sign-on infrastructure can also provide information about the user needed for the calling applications. All account information associated with the user is returned to the calling application, such as name, e-mail address and building. Membership of groups and mailing lists is also returned, so that the calling application can rely on central group membership to handle access control. Instead of a “per application dedicated role system” we now have a centrally managed group management on which all applications can rely.

Authentication overview

Figures 2–4 show what users can expect to see with the CERN SSO service:

- The user opens a website that requires authentication and is redirected to the SSO form (figure 2).
- His/her details and groups membership are available to the application, and can be used, for example, to restrict access to some pages (figure 3).
- Finally, clicking on the Logout button will disconnect the user from all opened applications (figure 4).

Identity and service providers

The SSO system has two components: the Identity Provider and the Service Provider. The Identity Provider checks the

Desktop computing

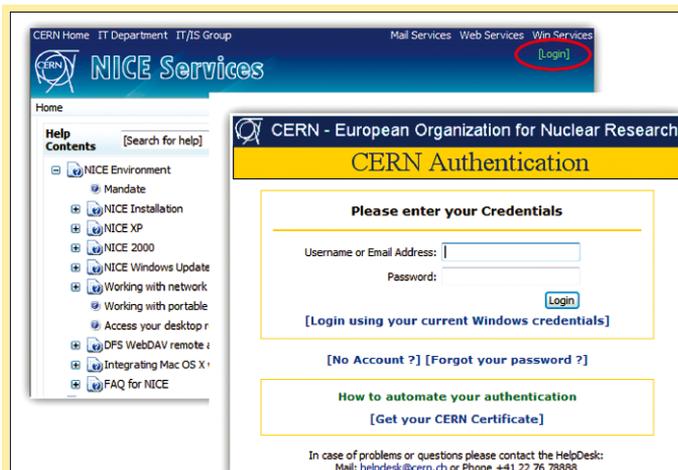


Fig. 2. The Single Sign On form makes it easier to authenticate.

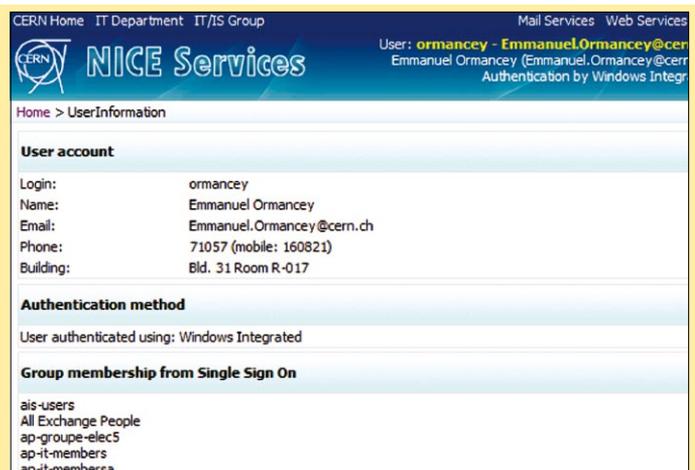


Fig. 3. User information is available to the calling application.

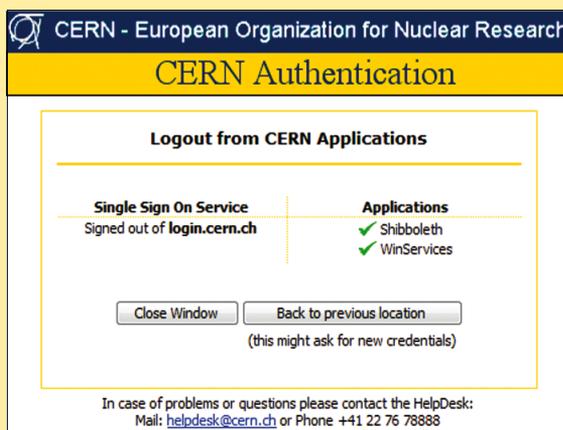


Fig. 4. Logout from Single Sign On and applications.

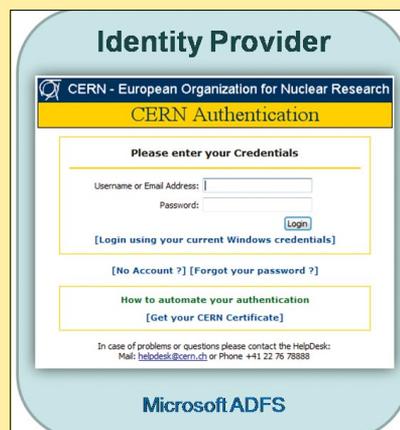


Fig. 5. The Identity Provider form.

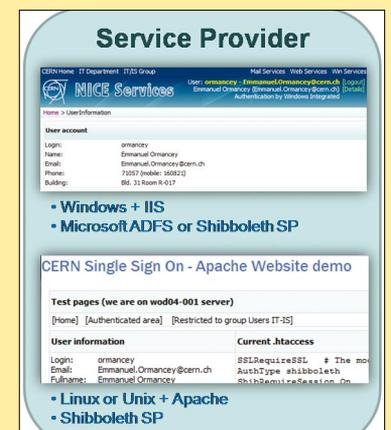


Fig. 6. The SSO Service Provider.

identity of users and supports various authentication methods (figure 5). It can be seen as the server side of the SSO system. It verifies credentials and loads user information for the calling application, using so-called “claims” or “attributes”. Note that each application can retrieve a different set of attributes, depending on the confidentiality level required.

The CERN implementation of the Identity Provider is based on Microsoft Active Directory Federation Services (ADFS). ADFS is an extension of the Microsoft Active Directory central authentication database, using the WS-Federation Passive Requestor Profile standard, based on SOAP (Simple Object Access Protocol) web services (see <http://technet2.microsoft.com/windowsserver/en/technologies/featured/adfs/default.mspx>).

Authentication is a highly critical service so the production service is hosted in the Computer Centre's critical area, on load-balanced servers to maximize availability.

The Service Provider is the client side of the SSO system (figure 6). It enables users to access web applications after they have been successfully verified by the Identity Provider. The Service Provider can be a Microsoft Internet Information Service (IIS)

module, an Apache module or an application module – some Java classes, for example.

On Windows- and IIS-hosted websites an IIS ADFS module comes with Microsoft Windows 2003 R2 as an additional component. It can easily replace an existing authentication system using basic or NTLM authentication.

On Linux-, Unix- and Apache-hosted websites, the Shibboleth Apache module can be used. Shibboleth is an open-source middleware software (see <http://shibboleth.internet2.edu>). The upgrade is quite simple because the classic Apache configuration files for access control are still used and can be kept; only the authentication module needs to be replaced.

In both cases moving to an SSO solution is a straightforward procedure. Only small configuration changes are required, and it is transparent for all applications that rely on the default authentication systems provided by web servers.

Non-web applications

The CERN SSO classic service can only be used for web-based applications. For those that are not web based, a dedicated SOAP web service is provided for verifying credentials, and retrieving user

information and group memberships.

This service requires some coding on the client side: a SOAP client, sending credentials to the web service and decoding return codes accordingly, is needed. Note that this implementation is not standard but a CERN-specific implementation based on SOAP standards.

The next steps

Today many of CERN's central services already use the SSO system on Windows and Linux platforms. Several pilots are also running on various Linux distributions as well as on Solaris for different services, and interest in the system is growing quickly. The fact that interaction between Windows and Unix is working opens up many possibilities, and it is the only way to achieve a truly viable single-sign-on solution.

At the same time, CERN also needs to improve and clarify the management of accounts. It is necessary to clean up, and ensure that each user has only one account for all services. Finally, it is important that control over access to resources is centrally managed. Figure 1 shows an overview of the way in which accounts and groups will soon be managed at CERN.

Emmanuel Ormancey, IT/IS

Managing your CERN account

Over recent years the IT department has been streamlining CERN users' access to all central computing services. The goal is to converge on a unique CERN account, which will increase computer security and simplify the maintenance of accounts. Administrative Information Service (AIS) applications, MAIL, Simba, the Engineering Data Management Service (EDMS), Windows and others already use this common CERN account, and the Andrew File System (AFS) accounts will soon be synchronized as well. The following paragraphs explain how CERN and external users can manage their unique account via a standard web browser.

CERN users

In order to manage their unique account,

CERN users should first go to <http://cern.ch/CERNAccount> (figure 1). Here they can set their password, sign the computing security rules, or view information about their account and services.

It is possible to add, delete or modify an account by clicking the Edit/Add/Delete Account link, which will redirect the user to the appropriate AIS site.

As the CERN account is the unique central account for CERN applications, changing a password will cause a global change for all CERN applications, including access to AFS.

Clicking the Check Account Status link will show the user's status and other information, such as group membership and mailing list ownership (figure 2). The green box at the top of the CERN Account

Status page shows available identifiers that can be used to authenticate on CERN resources: these are usually the user's login or his/her registered e-mail address. The CERN Account Status page also shows various service-related information, from web sites owned by the user to available space in his/her home directory (figure 3).

External users

External (non-CERN) users who need access to a CERN resource have to follow the lightweight registration procedure based on their e-mail address to create an external account. They can do this at the CERN Portal for External Accounts, at <http://cern.ch/externals> (figure 4).

Emmanuel Ormancey, IT/IS

The screenshot shows the 'CERN Authentication' page. It has a navigation bar with links: 'Check Account Status', 'Change Password', 'Forgot your Password?', and 'CERN External Accounts'. The main content is divided into two columns: 'CERN User Accounts' and 'CERN External Accounts'. Under 'CERN User Accounts', there are sections for 'Operations' (Check Account Status, Change or Reset Password, etc.), 'Advanced features' (Manage Groups, etc.), and 'Technical Side' (Behind CERN Authentication, etc.). Under 'CERN External Accounts', there are instructions on how to register and links for 'Register a new account', 'Forgot your password', and 'Manage your account'. Below this is a 'Service Specific Settings' section with details for 'Mail Services' (mail addresses, server, etc.), 'Spam' (filtering activated), 'NICE Services' (LDAP Path, Home directory path, etc.), and 'Web sites owned by ormancey' (a table listing sites like http://cern.ch/alerter, http://cern.ch/alerts, etc.).

Fig. 1 (top, left). The home page, for managing a CERN account.
 Fig. 2 (top, right). Account details are shown on the status page.
 Fig. 3 (bottom, left). Information about user services is available.
 Fig. 4 (bottom, right). The CERN Portal for External Accounts.

The screenshot shows the 'CERN Account Status' page. It has a navigation bar with links: 'Check Account Status', 'Change Password', 'Forgot your Password?', and 'CERN External Accounts'. The main content is divided into several sections: 'Credentials' (Login: ormancey, User Principal Name: Emmanuel.Ormancey@cern.ch), 'User account' (User name: Emmanuel Ormancey, Status: ACTIVE, etc.), 'User details' (Department/Group: IT/IS, Owner CCID: 596928, etc.), 'Groups and lists' (Group membership, Mailing List Ownership, etc.), and 'Mailing List Membership'. A red box highlights the 'Available identifiers' section, which contains a green box with the text: 'User can authenticate on CERN resources using CERN\ormancey or Emmanuel.Ormancey@cern.ch'. Below this is a 'Welcome to the CERN Portal for External Accounts' section with 'Usage / Intended Users' and 'Lightweight Registration Process' instructions.

ATLAS: the data chain works

In September the particle physics experiment ATLAS went “end-to-end” for the first time.

When the LHC is turned on, physicists worldwide will be waiting by their computers. They’ll be expecting express and non-stop delivery of massive amounts of data, streamed in a virtually seamless sequence direct to their doorstep.

And in September ATLAS used the LHC Computing Grid (LCG) to prove that this data distribution to physicists across the globe will be possible.

From cosmic ray to you

“We did the whole thing,” smiles ATLAS’s Kors Bos. “We mastered the entire data chain for the first time, from measurement of a real cosmic ray muon in the detector to arrival of reconstructed data at Tier-2 computers all over the world, with all the steps in between.”

“We measured about 2 million muons over two weeks,” says Bos. “We spent the first week setting up, but in the last week we started getting something useful. The detector measured real particles, real analysis happened at sites across Europe and the US, and it happened in quasi real time. The whole machine worked without human intervention.”

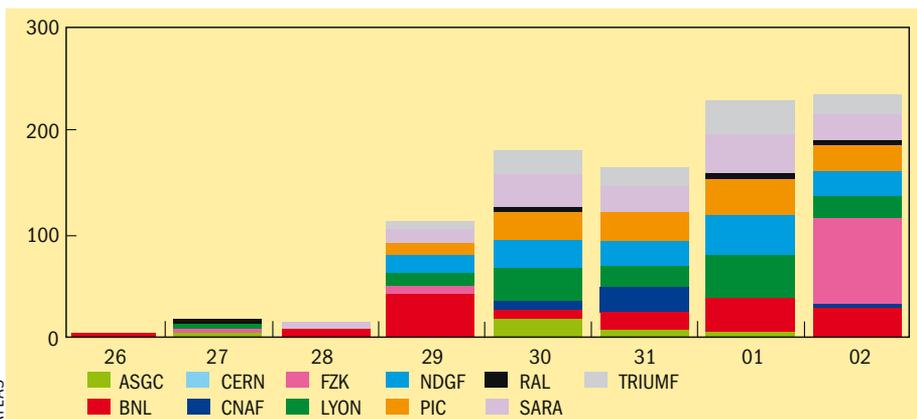
Quasi real time, says Bos, will be more than sufficient to keep the chain moving.

“It’s good that we could achieve this transfer in hours, but with the real data there will be a greater delay. We will need to do regular calibrations, and the processing and subsequent data transport will have to wait for that. The data chain is designed to cope with this delay and there are sufficient disk buffers at the Tier-0 stage to keep the data for as long as several days.”

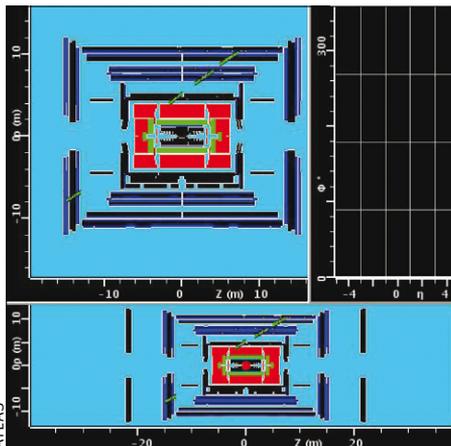
Domino data chain of events

The data chain is triggered only when a particle enters the ATLAS detector and produces a signal in its sensitive layers.

“We know these particles must be muons,



Terabytes of data were moved from the Tier-0 site at CERN to 11 Tier-1 sites across Europe, North America and Asia. Transfer rates reached the expected maximum in early September.



Tracks recorded in the muon chambers of the ATLAS detector can now be analysed simultaneously across the globe.

because everything else is stopped by the metres of clay above the detector. Only muons can get through, and only muons can trigger the detector,” Bos explains.

Once the particle has been detected, the domino data chain is set in motion. The raw data is sent to the Tier-0 centre, the Computer Centre at CERN, where

it is recorded onto tape before being sent to a different part of the centre for reconstruction.

“The reconstructed data also goes to tape,” says Bos. “It is then exported to all the Tier-1s, and when it arrives it is exported by the Tier-1s to their Tier-2s. When it arrives at the Tier-2s the physicists can pick it up and say, ‘wow, look at that’.”

Bos says that while researchers from different Tier-2 centres will be analysing the data in different ways, they will all be analysing the same data.

“The data analyses are bound to be different, but if you say ‘give me the data from run 123, event 345 on Sunday morning’, they should all be able to provide the same thing.”

This data chain will be used to share ATLAS data with more than 1900 physicists from more than 160 participating institutions in 35 countries. The massive data transfer requirements of the LCG are supported in Europe and the US by the Enabling Grids for E-science and Open Science Grid infrastructures.

Cristy Burne, iSGTW

This article was published online in *iSGTW* on 19 September.

European Grid Initiative presented at EGEE’07

The EGEE’07 conference on 1–5 October attracted more than 600 delegates, and it felt like most of them attended Tuesday’s workshop on the status of the European Grid Initiative (EGI).

The room was bulging before the session had even started. When Dieter Kranzlmüller of Johannes Kepler University Linz, Austria, began with an overview of the EGI Design Study (EGI_DS) project, he was speaking to an audience straining to see more.



Many delegates at the EGEE’07 conference attended the workshop on the EGI project.

This first EGI_DS workshop marked a very concrete step towards the realization of the EGI project: an endeavour that has already garnered the support of more than 37 national grids and that aims to produce a pan-European grid infrastructure.

The EGI Design Study kicked off on 1 September 2007 and, although yet to be officially approved, it counts on being supported by the European Commission’s (EC’s) 7th Framework Programme.

“We’re aiming to establish the conceptual set-up for a sustainable grid infrastructure in Europe,” said Kranzlmüller. “This will be driven by the EGI_DS project, which we hope will be supported by the EC. We finished our negotiations with the EC on 27 September and are now just waiting for a draft contract.”

The 27-month project has a strong emphasis on coordination and sustainability, and will rely heavily on input from the collaborating parties, Kranzlmüller said.

“Many countries have launched or are in the process of launching national grid initiatives (NGIs),” he said. “EGI aims to coordinate the integration and inter-operation of these NGIs, moving towards

a long-term sustainable initiative less dependent upon EU-funded project cycles.”

“We will be under considerable time pressure and will need to work together to produce a complete blueprint for the EGI by June 2008,” Kranzlmüller told attendees. “Expect a lot of e-mails and a lot of communication. EGI_DS will depend on your feedback to refine the design of the EGI organization, so it can emphasize support to each NGI and each use case.”

EGI is scheduled to begin operations in early 2010, providing overlap with EGEE-III.

EGI Design Study in a nutshell

What form will the EGI Design Study take?

- Work package 1: project management

(lead: GUP Linz, Austria).

- Work package 2: EGI requirements consolidation (lead: GRNET, Greece).
- Work package 3: EGI functions definition and roadmap (lead: INFN, Italy).
- Work package 4: study of EGI legal and organizational options (lead: CNRS, France).
- Work package 5: establishment of EGI (lead: CERN, Switzerland).
- Work package 6: EGI promotion and links with other initiatives (lead: CSC, Finland).

For more information about the EGI_DS project, please contact the EGI Design Study team at contact@eu-egi.org.

Cristy Burne, iSGTW

This article was published online in *iSGTW* on 3 October.

Visualizing the state of your grid with GridMaps

Which of your grid sites are operational at any instant? How much does each site contribute to which virtual organization? How can you have all this information and more at your fingertips?

With GridMaps – a top-level grid services monitoring visualization produced as a collaboration, initiated by EDS fellow Rolf Kubli, between CERN openlab and EDS, and undertaken within the Grid Deployment group at CERN.

“This style of data monitoring requires much less space than conventional tables or bar charts,” says EDS senior engineer Max Böhm, architect and developer of the GridMap prototype. “It also allows you to page quickly through different points or focus on different metrics to find correlations, patterns and failure mechanisms.”

The prototype GridMap interface [see <http://gridmap.cern.ch/gm>] is simple, interactive and super user-friendly:

- Grid sites or services are represented by rectangles; the size of the box indicates



The GridMap prototype shows the status of sites in the Worldwide LHC Computing Grid.

the number of CPUs at that site.

- The status of each site is represented by a different colour: red means the site is down; orange indicates a degraded level of service; green means all systems are go, go, go.

The visualized GridMap data comes from the underlying Service Availability Monitoring (SAM) framework, a tool for monitoring services at grid sites.

With a click of the mouse GridMap can shift views to show data from different virtual organizations and geographical perspectives. Users can also plumb deeper into each site.

“This same type of visualization can be used for top-level, regional, and virtual-organization-specific views,” says Böhm. “It has very broad application and can be adapted to suit specific needs. The size of each box can be adapted to indicate the number of jobs running, or the colour could indicate site availability... The possibilities are exciting.”

CERN openlab is a framework for evaluating and integrating cutting-edge IT technologies or services in partnership with industry. EDS, a leading global technology services company, is one of several corporate partners and contributors to the 2007 programme.

Cristy Burne, iSGTW

This article was published online in *iSGTW* on 17 October.

LHC@home server sets up new home in the UK

Researchers in the UK are gearing up for an influx of help, with the recent arrival of the successful LHC@home volunteer computing project at Queen Mary, University of London.

Forty thousand people from more than 100 countries have already contributed the equivalent of about 3000 years on a single computer to LHC@home, which is migrating from its first home at CERN.

Running off the BOINC platform, LHC@home uses volunteer computing power to model the progress of subatomic particles travelling at nearly the speed of light around Europe’s newest particle accelerator, the Large Hadron Collider (LHC).

Lyn Evans, head of the LHC project, says that “the results from this initiative are really making a difference, providing us with



The LHC@home project enables volunteers to run simulations of physics inside the LHC.

new insights into how the LHC will perform”.

LHC@home is a collaboration between CERN; the Helsinki Institute of Physics; the Niels Bohr Institute in Copenhagen; Queen Mary, University of London; and TRIUMF in Vancouver, and will now be managed by physicists from the GridPP project in the UK.

GridPP’s Neasan O’Neill explains: “We started trial running LHC@home from a computer server in the UK in June, and have spent the last few months working with the physicists who use the data it produces. Now, with the official launch of the UK base for the project, we’re ready to fully exploit this fantastic resource.”

Sarah Pearce, GridPP

This article was published online in *iSGTW* on 17 October.

Web applications security: risks and countermeasures

Reports of vulnerabilities in web applications have risen sharply. Exploits are easy to develop and targets are easy to find. It is therefore important that we adapt to these threats, when programming and also when we use these applications.

Web applications are commonly used from a web browser and they cover a range of activities, such as e-banking, webmail, online shopping, community websites, blogs, vlogs, network monitoring and bulletin boards.

In recent years the development of such applications has been considerable, and today rich internet applications offer complex, real-time interactions with users. For instance, web operating systems such as eyeOS offer many functionalities that were previously available only with traditional operating systems.

While web applications have become ubiquitous, they also present new security risks. It is important to identify and understand these risks when developing, hosting or simply using these applications.

Security risks

There are two main reasons that web applications are vulnerable to attack.

First, it is generally difficult for the service manager to keep up to date with security patches. This is a common issue for services in general, but it may be particularly challenging for web applications. This could be improved by better design and packaging but it is often impossible to upgrade web applications automatically. This requires the service managers to actively monitor announcements by the application vendors. It may also be necessary to customize the application to meet the community's needs, and this may result in additional delays to an upgrade.

Second, web applications are often easy targets for attackers. As a relatively recent development, they use non-mature code compared to traditional network services. Unfortunately exploits – malicious code that exploits software vulnerabilities – are generally easy to prepare, remotely executable, cross-platform, and require no compilation. This helps attackers to design effective and scalable automated attacks. Vulnerable installations can be found quickly, easily and silently by using search engines to detect known vulnerable patterns, generally filenames, of specific web applications.

Not only critical services are at risk

It is important to understand the threats against which web applications need to be protected. Several years ago attackers were mostly attracted by challenge or fame, but since then attacks have become more professional and many attackers are now motivated simply by money (e.g. phishing, spam, extortion, distributed denial of service attacks (DDoS), and click fraud). Some malware kits have also become professional: user friendly, modular (it is possible to buy extra components), and some even include one year of support. For more information, an interesting article about the Storm Worm is available at www.schneier.com/blog/archives/2007/10/the_storm_worm.html.

A result of this professionalization is that attackers need bandwidth and platforms to operate from. To obtain and maintain a sufficient amount of compromised resources, attackers usually choose the easier target. No matter how insignificant a particular service may be, it is worth something for an attacker: an unmaintained photo album may be worth as much as a critical production service.

For instance, once a web application has been successfully compromised the attacker may gain the ability to:

- cause damage to the reputation of the organization running the service;
- execute code remotely with the privileges of the user running the web server – this is sufficient to run botnets and spam engines, and enables the attacker to try to spread the attack further on the system or against third-party services;
- access all the information hosted on the web server;
- change the content of the website (defacement);
- delete files or damage web services provided by the host (denial of service).

Shift in the vulnerability trends

There are different types of web application flaws, but the most common are caused by a lack of user input validation. Data coming from the client must be filtered to ensure that no malicious content is passed to the server. If this is not performed correctly, the resulting vulnerabilities include:

- Cross-site scripting (XSS). If a parameter is not sanitized correctly an attacker may be able to control the content of the vulnerable page. To perform the attack

the victim is often tricked into clicking on a malicious link, but this is not always necessary for the attack to be successful.

- SQL injection. Some parameters retrieved from the user's web browser may be used to perform database queries. If a parameter passed from the user to the database is not correctly filtered, an attacker may attempt to execute arbitrary SQL commands and/or to gain privileges on the web application.

- Remote file inclusion (RFI). By exploiting insecure calls to local files, such as templates, an attacker may attempt to upload arbitrary code on the server. The resulting payload (for example, a shell written in PHP) may be executed with the privileges of the web server.

CVE (Common Vulnerabilities and Exposures) is an online dictionary of registered vulnerabilities that provides yearly statistics about known problems (see <http://cwe.mitre.org/documents/vuln-trends/index.html>). In 2001 the number one vulnerability affecting software was the well known buffer overflow; XSS, SQL injection and RFI were almost non-existent. In 2006 a clear shift towards web applications is visible: XSS is the number one vulnerability, SQL injection number two, and RFI number three. Buffer overflow has moved to fourth place.

Recommendations

Common ways to mitigate the risks related to web applications are detailed below.

If you are developing web applications:

- do not trust any content, parameter or variable coming from the web browser, and be sure to properly sanitize all input before using it;
- check all input by design, even if it is not directly visible to users;
- use the validation functions provided by your environment and avoid relying only on your own filters;
- if you use a development framework that provides some input filtering, do not solely rely on it;
- keep your framework up to date with security patches – it can be a target as well;
- beware of the information revealed or echoed by error messages/pages;
- requiring (re)authentication for privileged operations is always a good idea;
- keep your support lists private – this may prevent leaks when a vulnerability

is reported to your team;

- aim to reduce the exposure of your web application and avoid directly exposing it in the site firewall; for off-site access, whenever possible, require that users connect via a gateway system, such as Windows Terminal Services or the LXPLUS Linux cluster, as documented at <http://cern.ch/security/internet>.

If you are using web applications:

- take careful note of security warnings from the web browser;
- whenever possible, disable Javascript/Flash/ActiveX; for example, for Firefox using the NoScript extension at <http://noscript.net> may help;
- avoid following links to sensitive portals, such as for e-banking, and type the URL

by hand if possible;

- whenever feasible, log out as soon as possible and/or close the browser when the session is completed;
- use the more secure HTTPS protocol instead of HTTP if available.

More information is available from <http://cern.ch/security>.

Romain Wartel, Computer Security team

Twelve steps to improving control systems cyber security at CERN

Modern accelerator and experiment control systems have, over the last few years, been based increasingly on commercial off-the-shelf products such as VME crates, programmable logic controllers (PLCs), supervisory control and data acquisition (SCADA) systems, on Windows or Linux PCs, and on communication infrastructures using Ethernet and TCP/IP.

Despite the benefits that come with this (r)evolution, new vulnerabilities are inherited too: worms and viruses can spread within seconds via the Ethernet cable, and attackers are becoming interested in control systems. Unfortunately control PCs cannot be patched as quickly as office PCs. Even worse, vulnerability scans at CERN using standard IT tools have shown that commercial automation systems often lack even fundamental security precautions: some systems crashed during the scan, while others could easily be stopped or have their process data altered [1].

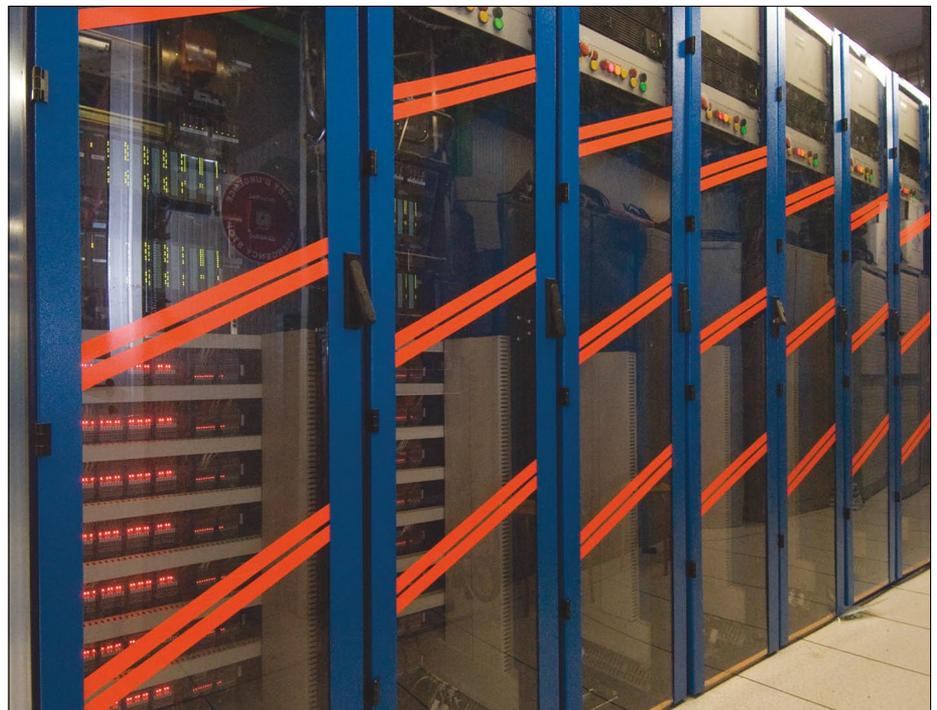
Several recent incidents in industry have unfortunately proved that the risks coming from security breaches are no longer fiction, and the resulting consequences can be severe:

- In 2005 “a round of internet worm infections knocked 13 of DaimlerChrysler’s US auto manufacturing plants offline for almost an hour, stranding some 50 000 auto workers as infected Microsoft Windows systems were patched” [2].

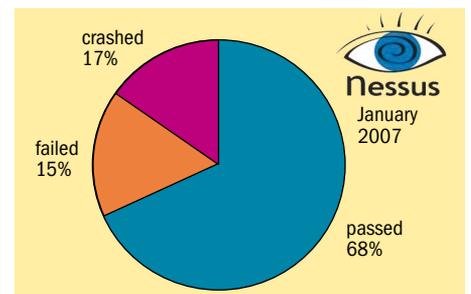
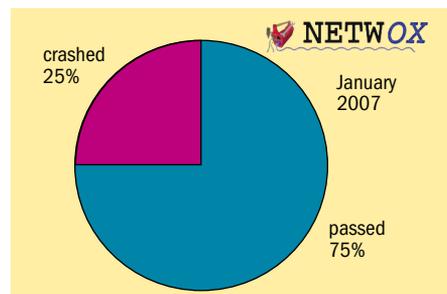
- In 2006 a nuclear power plant in the US had to be shut down after a “data storm” on its internal control system network stopped the control of two cooling water pumps. The storm was apparently caused by a malfunctioning PLC [3].

- A report in 2007 stated that “US critical infrastructure [is] in serious jeopardy”, and that the US “electrical service, transportation, refineries and drinking water are at serious risk from very simple hacker attacks” [4].

- “Researchers [at the US Sandia National Laboratories] who launched an experimental cyber attack caused a [power] generator to self-destruct...” [5].



The Detector Safety System used in the LHC experiments is one of many control systems.



Results of vulnerability tests performed by CERN on PLCs using standard penetration tools.

Protecting the critical infrastructure

Partially driven by these incidents, but more by the fear of terrorism after 9/11, industry and government authorities have begun to review the consequences of a security incident on the so-called “critical infrastructure”. The critical infrastructure is those industrial sectors

upon which everyday living depends, such as electricity providers, oil and gas companies, water and waste plants, chemical and pharmaceutical factories, and the transport sector. The demand for critical infrastructure protection has led to a consolidation of worldwide efforts with respect to “control systems cyber security”



The Computer Management Framework scheme is used in the CERN Control Centre.

implemented by the IT/FIO and IT/IS groups: Linux For Controls (L4C) for CERN Scientific Linux 3 and 4, and the Computer Management Framework (CMF, see <http://cern.ch/cmfc>) for Windows XP and Windows Server 2003 [12].

Both schemes give control system experts the flexibility to manage their control PCs in a centralized manner, i.e. using a centralized installation scheme. While the standard operating systems, applications and their patches are provided centrally by the corresponding IT service providers, it is up to the control system expert to configure a PC, schedule external interventions (e.g. upgrades or patches), and validate changes before they are applied to PCs that control sensitive equipment. Such a scheme can also make it easier to re-establish and recover a system after a security incident. By passing the flexibility to the experts they take responsibility for securing their own PCs too.

Today CMF is used widely for control PCs, such as for the Windows consoles in the CERN Control Centre, and has been accepted by the controls community at CERN. It has become the standard way to install any Windows PC at CERN, and can additionally be used to configure Windows Terminal Servers. L4C, on the other hand, is used extensively to set up local Linux computing farms for the LHC experiments. For example, at the moment CMS is configuring about 1000 servers through L4C. LHCb uses L4C to install their diskless systems.

● **Step 7:** Make sure that all your PCs are fully patched, the local firewall is enabled, and that the anti-virus program is kept up to date.

● **Step 8:** Use CMF and/or L4C to manage your control PCs. Both provide flexibility and offer you the mandatory security.

Authentication and authorization

Although authentication and authorization is an important pillar in the “defence-in-depth” approach, CERN has previously lacked a coherent solution for control systems, office PCs and web applications. The recent single sign-on project addresses these issues (see “Single sign-on facilitates authentication at CERN” on p5).

The Accelerators and Beams department’s Controls group has implemented role-based access control [13] for the accelerator controls and operations; user credentials are needed for authentication (i.e. to identify the user), while role assignments define the authorization level (e.g. which actions that user can take). ATLAS and CMS are following a similar route. Furthermore, ALICE has implemented a solution based on Windows credentials using SmartCard technology [14], a direction that has been tested by the IT/IS group. In a broader approach, IT/IS is deploying a CERN-wide Single Sign On portal with a CERN certification authority (<http://login.cern.ch>).

● **Step 9:** Use passwords with sufficient complexity; avoid using obvious words.

● **Step 10:** Never tell anybody your password, not even the Computer Security team, and do not write it down.

User training

A series of awareness campaigns and training sessions have been held for users, operators and system experts at CERN. Monthly CNIC meetings provide a forum for questions and discussions (<http://indico.cern.ch/categoryDisplay.py?categId=691>). Additional dedicated discussions have been held with many system experts (see the CNIC TWiki for details: <https://twiki.cern.ch/twiki/bin/>

[view/ath/CNIC/WebHome](https://twiki.cern.ch/twiki/bin/view/ath/CNIC/WebHome)). The TWiki also provides minutes of discussions with users outside CERN, such as other high energy physics laboratories, and with major players in industry. There are links to major standards and guidelines.

Furthermore, CERN has raised aspects of control system cyber security at several conferences and workshops, such as the CS²/HEP workshop [15]; interacted with major vendors of control systems; and is now leading the European Information Exchange on SCADA Security (EuroSCSIE) with members from European governments, industry and research institutions that depend on and/or whose responsibility it is to improve the security of SCADA and control systems.

● **Step 11:** If you have any questions, join the CNIC users exchange meetings held every last Thursday of the month, or contact CNIC-Coordination@cern.ch.

Incident response and system recovery

Even with a stringent security policy, incidents can never be prevented completely. Therefore incidents on a domain have been and will be handled jointly by CERN’s Computer Security team and the appropriate domain administrator. If the domain administrator cannot be reached, the acting computer security officer, after receiving approval from the acting IT department head, has the right to take appropriate action in a justified case of emergency.

After an incident has been analysed, the CMF and L4C central installation schemes enable the appropriate system expert to recover the system rapidly.

● **Step 12:** If you follow steps 1–11 the domain administrator or Computer Security team may never need to contact you.

Personal responsibility

The continuing integration of common IT technology into control systems means that IT security vulnerabilities and cyber attackers end up threatening control systems and, therefore, CERN’s operation and assets. However, control systems require a different approach to security to that which is appropriate for office systems.

The CNIC Security Policy for Controls document presents a thorough set of rules to secure CERN’s control systems. The implementation uses a “defence-in-depth” approach that is based on network segregation, central installation schemes, authentication and authorization, user training, incident response and system recovery, and security auditing. However, each user plays a major role in protecting CERN’s assets and operations, so please make sure that you have reviewed all the steps mentioned above.

Stefan Lueders, IT/CO

You will find a list of references in the web version of this article.

CHEP focuses on LHC computing in last meeting before start-up

Computing in High Energy and Nuclear Physics (CHEP) is a major series of international conferences for physicists and computing professionals in the fields of high energy and nuclear physics, computer science and information technology.

CHEP'07 took place in Victoria, Canada, on 2–7 September and was attended by some 470 people from all over the HEP world. It was preceded by a two-day workshop on the Worldwide LHC Computing Grid (WLCG), attended by nearly 200 people. This was the last CHEP conference before the LHC start-up, and most talks covered the preparations for the different components of LHC experiment computing, from online data taking and event processing to event reconstruction and analysis.

To continue with statistics, 429 abstracts were submitted from 1208 authors spread across the seven programme tracks: online computing; software components, tools and databases; computer facilities, production grids and networking; collaborative tools; distributed data analysis and information management; event processing; and grid middleware and tools (figure 1). Owing to this large number of contributions, around 50% of the offers were given as posters in two sessions. Each session was on display for two days, with fixed times when the poster authors were available by their posters for discussions.

Looking ahead

The evolution of computing technology over the next few years was covered in several plenary talks. One such talk was “Towards petascale and exascale computing” by Jim Sexton from IBM Research – IBM was one of the conference’s gold sponsors. Although Sexton believes Moore’s law still holds, we are approaching fundamental limits; every time we want to take a step forward we need to make work-arounds, and much more performance gain is coming from parallelism in the cores. IBM’s Blue Gene solution is an example of this, with configurations ranging from one to 72 racks and up to 294 912 cores. However, although adding cores to increase the power may still be possible, writing programs that make optimal use of such computers is difficult.

Memory will be the dominant cost for upcoming computers, and this theme was taken up by another sponsor when Steve Pawlowski from Intel gave his talk on multicore processors. He gave examples of how memory can be brought logically and physically closer to the processing

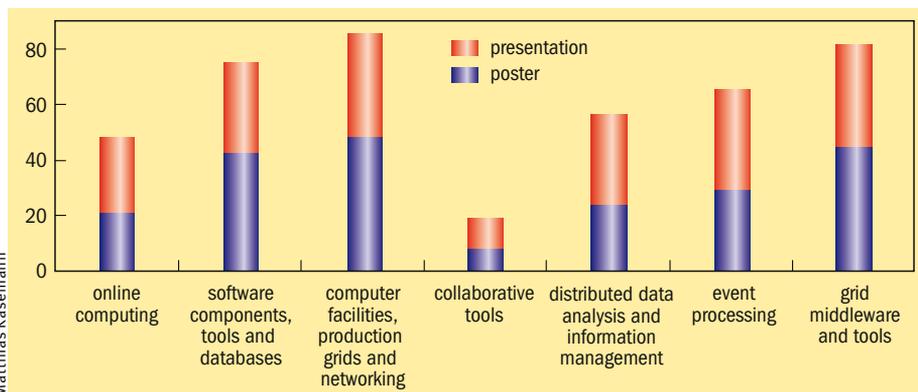


Fig. 1. Some 429 abstracts were submitted from 1208 authors across seven subject areas.



Sverre Jarp gives a plenary talk at CHEP'07.

units. Like the IBM speaker he talked of increased performance through more parallelism, but improving raw chip performance is offering diminishing returns. In a plenary session on “LHC software and CPU architecture” Sverre Jarp (CERN/IT) recommended increasing the level of instruction parallelism in LHC software to help compilers produce more effective machine code. We need to improve the multithreading capabilities of the applications and, more generally, to simplify and restructure the code. As we approach billion-transistor processors “we should increase the agility of our software structures”, said Sverre.

Another gold sponsor talk, by Lim Goh of SGI, covered some aspects of the physical expansion of computing: how computer facilities are gradually running out of capacity in terms of space, power and even weight. SGI has produced a rack system packed so tightly that it exceeds the weight rating of many computer floors!

A squeeze on funding?

Returning to today’s realities, several speakers noted that LHC experiments have benefited from an unprecedented level of support from grid projects backed by

national funding agencies, although Miron Livny from the University of Wisconsin, US, believes the days of such “easy funding” may be drawing to a close.

Experiments have begun demonstrating production-level processing on grids but not yet at the scale needed for full LHC running conditions. There is still a lot of work expected in the final year of preparation. “We are getting there slowly,” said Les Robertson (CERN/IT), speaking about the WLCG, “but from now until first beams we need continuous testing, driven by experiments, with realistic scenarios, good monitoring and measurements, and the proactive participation of all sites, developers and especially storage experts”. According to Jamie Shiers (CERN/IT), despite the considerable progress made since CHEP'04, the main areas of concern remain unchanged:

- data storage, and management and integration with experiment frameworks;
- reliability and usability, which will become critical for widespread use in data analysis, for example.

In his summary talk, Matthias Kasemann of CERN and DESY reminded the audience that we expect LHC to be operational next summer, and therefore by spring 2008 the experiments must be ready to take data. He closed the meeting by thanking the principle organizers. The next CHEP will be held in Prague in spring 2009.

As usual a detailed trip report has been written by Alan Silverman, with contributions from many people who are listed in the report. It is saved in CDS as CERN-IT-Note-2007-013 at <http://cdsweb.cern.ch/record/1057583>. You can also read a summary in the *CERN Courier* Computing News section, November issue.

Alan Silverman, IT/DI

Second multithreading workshop attracts CERN programmers

For the second time this year, CERN's programmers had the opportunity to participate in a multithreading and parallelism workshop organized jointly by CERN's openlab and Intel. The event took place on 4–5 October and attracted nearly 50 participants.

Two experienced instructors from Intel, Herbert Cornelius and Hans-Joachim Plum, presented classical parallelism theory and practice; described modern mainstream parallel programming techniques such as pthreads, OpenMP and Threading Building Blocks; and warned about common coding pitfalls. The tutors also provided details about upcoming processors, and answered questions during a Q&A session.

On the first day, which consisted entirely of lectures, three members of CERN gave presentations. Sverre Jarp spoke about the challenges that high energy physics software will face in the era of multicore systems; Fons Rademakers described the multithreading efforts in ROOT; and Eric McIntosh recalled his experience with OpenMP in Fortran in the field of computer tomography.

In the hands-on labs on day two, the participants were given assignments relating to the material presented the previous day. Some of Intel's threading products, such as VTune and Thread Checker, were demonstrated in both Linux and Microsoft Windows. The attendees were given the opportunity to use these and other tools to locate bottlenecks and fine-tune the multithreaded exercise applications.

A survey conducted after the workshop revealed that all of the respondents had had their expectations fulfilled, and more than 90% said that they would recommend



Herbert Cornelius gave a presentation about parallelism and the multicore revolution.



Hans-Joachim Plum gave advice in the lab.

the course to other people. Participants also made several useful suggestions, such as the idea of conducting a survey of the proficiency level and background of participants before the workshop. The

comments gathered will help improve future classes and fine-tune the course programme to the audience.

Since the number of subscribers was again much higher than the number of places available, another event will be held next year. A multithreading workshop is planned for spring 2008, and it will focus more on CERN-specific topics with CERN presenters leading the event.

In addition, openlab is considering organizing a workshop on performance monitoring. It would revolve around such topics as the current state of the x86 architecture, performance monitoring in general, and specific tools such as pfmmon and VTune. The class is initially planned to take place in January 2008; more details will follow in the coming months.

Andrzej Nowak, IT/DI, openlab

Workshop will forge links with financial sector

CERN will host an event on Computing for Finance on Wednesday 21 November. The event, co-organized by the Enabling Grids for E-sciencE project and CERN openlab, gathers four distinguished speakers from the banking industry and from IT companies that serve financial institutions. The talk will take place at 5–7 p.m. in the main auditorium.

The financial sector is one of the driving forces behind the use of distributed or grid computing for business purposes. The speakers will provide an insight into how different types of grid computing – from

local clusters to global networks – are being applied in the sector. In particular, the event will encourage discussion on how new developments in grid computing emerging from the world of science can benefit financial institutions.

The speakers are Michael Yoo, the managing director and head of the Technical Council at UBS (Switzerland); Fred Gedling, chief technology officer at DataSynapse (UK); Adam Vile, head of grid, HPC and technical computing at Excelian Ltd (UK); and Daniel Egloff, head of the Financial Engineering Computing Unit at

Zurich Cantonal Bank (Switzerland). As well as reviewing the state of high-performance computing in the financial sector, the speakers will describe the use of software and techniques from physics, such as Monte Carlo simulations, in the financial world.

The talks will be followed by a Q&A session and a networking cocktail for the audience and speakers. The event is organized in collaboration with the regional business network Rezonance. Participation is free and open to all. Registration is mandatory via www.rezonance.ch.

It's all change at the bookshop

In July 2005 the IT bookshop moved from building 513 to CERN's main library area. Since then the shop has branched out and now has a large selection of books on physics, engineering and statistics as well as the original computing stock.

This issue we say goodbye to Jutta Megies, who has looked after the day-to-day running of the shop for the past 10 years and is now retiring. Thank you Jutta! At the same time, Joanne Yeomans from the library (Directorate Services Unit) is taking over general responsibility from Roger Woolnough (IT department),

who started the bookshop in 1994.

With the expansion of the shop we are looking at ways to provide book orders more quickly, and in the coming issues of *CNL* we will tell you what new services have been put in place.

Some of the latest additions to the bookshop catalogue are listed below. The new edition and CD of *Numerical Recipes* from CUP has proved very popular. Thank you for your continued support and please pass on any suggestions for new titles to bookshop@cern.ch.

Roger Woolnough, IT/UDS

Computing

L Carlson, L Richardson 2006 *Ruby Cookbook* (O'Reilly).
M Fitzgerald 2007 *Learning Ruby* (O'Reilly).
M Fitzgerald 2007 *Ruby Pocket Reference* (O'Reilly).
W H Press *et al.* 2007 *Numerical Recipes* 3rd edn, with optional CD (CUP).
G Reese 2007 *MySQL Pocket Reference* 2nd edn (O'Reilly).
A S Tanenbaum, M van Steen 2006 *Distributed Systems* 2nd edn (Prentice Hall).

Physics

K Becker, M Becker, J H Schwarz 2006

String Theory and M-Theory (CUP).
F Close 2007 *The Void* (OUP).
M Dine 2007 *Supersymmetry and String Theory: Beyond the Standard Model* (CUP).
L Rossi *et al.* 2006 *Pixel Detectors: From Fundamentals to Applications* (Springer).
A Sessler, E Wilson 2007 *Engines of Discovery* (World Scientific).

Other fields

H Kleinert 2006 *Path Integrals in Quantum Mechanics, Statistics, Polymer Physics, and Financial Markets* 4th edn (World Scientific).
D Sivia, J Skilling 2006 *Data Analysis: A Bayesian Tutorial* 2nd edn (OUP).

Computer security: think before you click!

A leaflet on computer security is being produced and will be distributed shortly. It will be printed in English on one side and French on the other. The preliminary English version is reproduced below.

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

The CERN Computer Security team

Computers are under attack, even at CERN!



To keep your computer safe:

- Do stay alert to tricks to steal your password
- Don't click on suspicious Web links in Spam
- Don't open unexpected Emails or attachments
- Don't use Peer-to-Peer (P2P) file sharing (e.g. BitTorrent)
- Don't use Chat Rooms (IRC - Internet Relay Chat)
- Don't download or install software from the Internet

Security Information: <http://cern.ch/security/> Computing Rules: <http://cern.ch/ComputingRules/>

Calendar

November

26–30, **IEEE GLOBECOM 2007: IEEE Global Communications Conference**
Washington, DC, USA
www.ieee-globecom.org

26–30, **MGC 2007: International Workshop on Middleware for Grid Computing**
Newport Beach, California, USA
<http://mgc2007.lncc.br>

26–30, **Middleware 2007: ACM/IFIP/USENIX International Middleware Conference**
Newport Beach, California, USA
<http://middleware2007.ics.uci.edu>

December

8–9, **WITS'07: Workshop on Information Technologies and Systems**
Montreal, Canada
<http://zen.smeal.psu.edu/wits07>

10–13, **e-Science 2007: IEEE International Conference on e-Science and Grid Computing**
Bangalore, India
www.garudaindia.in/e-science_2007.asp

16–19, **IDEAL'07: International Conference on Intelligent Data Engineering and Automated Learning**
Birmingham, UK
<http://events.cs.bham.ac.uk/ideal07>

18–21, **HiPC 2007: International Conference on High Performance Computing**
Goa, India
www.hipc.org

January 2008

13–16, **IUI 2008: International Conference on Intelligent User Interfaces**
Canary Islands, Spain
www.iuiconf.org

22–25, **GRAPP 2008: International Conference on Computer Graphics Theory and Applications**
Madeira, Portugal
www.grapp.org

22–28, **VISIGRAPP 2008: International Joint Conference on Computer Vision and Computer Graphics Theory and Applications**
Madeira, Portugal
www.visigrapp.org

February

4–7, **WSCG 2008: International Conference on Computer Graphics, Visualization and Computer Vision**
Plzen, Czech Republic
<http://wscg.zcu.cz/wscg2008/wscg2008.htm>