CERN COMPUTER NEWSLETTER

Volume 45, issue 1 January–March 2010

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), Tim Smith (group leader, User and Document Services), Christine Sutton (*CERN Courier* editor).

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

Published by CERN IT Department

IOP Publishing

CERN participates in new OpenAIRE project

The new Open Access OpenAIRE portal is available at www.openaire.eu/.

In support of Open Access, the European Commission has launched a pilot initiative to ensure that the research publications from FP7-funded research in several fields, as well as all European Research Council (ERC)-funded research, must be made Open Access within six months from date of publication. To support this drive, they opened a call for a European-wide infrastructure to ensure that European researchers indeed have Open Access repositories available to them, and that their papers are linked to the EC funding source. The result of this call was a new European project called OpenAIRE (Open Access Infrastructure for Research in Europe) that came into existence on 1 December 2009. This pilot project will run for 36 months until the end of FP7 and is a joint effort between 38 partners from around Europe.

OpenAIRE's objectives

- Build support structures for researchers in depositing FP7 research publications through the establishment of the European helpdesk and the outreach to all European member states through the operation and collaboration of 27 national Open Access liaison offices.
- Establish and operate an electronic infrastructure for handling peer-reviewed articles as well as other important forms of publications (pre-prints or conference publications). This is achieved through a portal that is the gateway to all user-level services offered by the e-infrastructure, including access (search and browse) to scientific publications and other value-added functionality (post-authoring tools, monitoring tools through analysis of document and usage statistics).
- Work with several subject communities to explore the requirements, practices, incentives, work flows, data models and technologies to deposit, access and otherwise manipulate research data sets of various forms in combination with research publications.

The OpenAIRE portal will be the central place for EC-funded researchers and other collaborators to deposit their documents and data, or to announce their existence in any institutional or subject repository. In doing so, the authors and their project will gain visibility and ensure that their work can contribute to the benefit of Europe (and thanks to Open Access, to everybody). In addition the EC will have a powerful tool to monitor and evaluate the performance of the different projects it is funding and their compliance with the Open Access mandate.

The back-end of the OpenAIRE portal will

Editorial

be built on top of two software packages – D-NET and Invenio.

D-NET is the software infrastructure that has been built as an outcome of the DRIVER project (Digital Repository Infrastructure Vision for European Research), "a multi-phase effort whose vision and primary objective is to establish a cohesive, pan-European infrastructure of digital repositories, offering sophisticated functionality services to researchers and the general public." (From the DRIVER portal).

Invenio is CERN's Integrated Digital Library software, which is already used to serve the CERN Document Server, the soon to be released INSPIRE, the EPFL InfoScience Portal and others, and will soon power the SAO/NASA Astrophysics Data System (ADS).

Invenio is an integrated repository designed to allow the depositing, harvesting, organization and dissemination of knowledge in the form of documents and their metadata (e.g. title, authorship information, abstract, keywords, etc). It offers an enhanced search engine and collaborative tools, such as baskets (to share your findings with colleagues), alerts (so you can stay up to date with recent publications in your area of interest), comments/review (for collaborative production) and much more.

As part of OpenAIRE, CERN is hosting and maintaining the OpenAIRE Orphan Records Repository, which is the OpenAIRE special repository for all publications without an existing institutional or subject repository.

Because it is based on Invenio, this repository will provide the OpenAIRE architecture with all of the advanced features that we have available today, such as automatic reference extraction from publications, optical character recognition of scanned documents, metadata curating tools, automatic keyword/classification management tools. In addition, because this is a new use case for Invenio, the software development and advancements driven by the needs of the OpenAIRE project will in turn be available to all of the Invenio users, including all of the CERN Document Server users – a truly win-win situation.

Invenio is one of the key software packages behind OpenAIRE and recently joined D4ScienceII – another European project – proving that it has become a mature and trustworthy project, contributing to CERN's transfer of technologies and knowledge as well as to the dissemination of science worldwide.

The OpenAIRE Orphan Repository at CERN is available at http://openaire.cern.ch/.

Useful links

Invenio: http://cdsware.cern.ch/invenio/
CERN Document Server: http://cdsweb.cern.ch/
DRIVER: www.driver-community.eu/
OpenAIRE Portal: www.openaire.eu/
OpenAIRE Orphan Record Repository: http://openaire.cern.ch/

Samuele Kaplun, IT-UDS

Your views can shape the future of the CNL

We must all move with the times, and the *CNL* is no exception. We would like to hear from you, dear readers, on how you think the *CNL* should change.

- What topics interest you most? What other publications do you read?
- Should it be shorter or longer? Should it be published weekly, monthly or quarterly?
- Should it stay mainly as a paper production, or evolve into an e-mail/web newsletter?

The *CNL* survey consists of just 10 questions enabling you to help us shape the future of the newsletter – please take a few minutes of your time to help us improve the newsletter for you.

The survey is available at https://espace.cern.ch/cnl-survey/, just click on the "CNL survey 2010" link and answer the questions. The survey will be open until 30 April and your responses will only be seen by the *CNL* editor.

Thank you in advance.

CNL editor

The deadline for submissions to the next issue of CNL is

Friday 21 May

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

CNL is available on the web at http://cern.ch/cnl

The CERN School of Computing returns to the UK after 22 years

Since it started in 1970, the CERN School of Computing has visited 18 countries and it is 22 years since the school was last held in the UK, in Oxford. We are therefore happy to return there, this time to Uxbridge, 40 km west of London. The school is hosted by Brunel University, a local organizer, and will take place from 23 August to 3 September.

The school is aimed at postgraduate students and research workers with a few years' experience in scientific physics, computing or related fields. The programme is organized along three main themes.

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

Examples of questions that will be answered at the school include:

- How to bridge Grids and Clouds using virtualization technology?
- Is it possible to simplify LHC physics analysis using virtual machine?
- How can reliable storage services be built from unreliable hardware?
- Why are tapes still used in high-energy physics data storage?
- How can I write code for tomorrow's hardware, today?
- Do you want to see your software with attacker's eyes?
- Can you hack your own code?
- Do you know what "code injection" and "integer overflow" have in common?
- What are the key statistical methods used in physics data analysis?

As always, since 2002, the prestigious CERN School of Computing Diploma will be delivered to the participants who succeed in the final optional examination, complemented by the Brunel University Certificate of Credits.

To respect a successful tradition established during the last decade, in addition to the 50 hours of tuition, a sports programme will propose two to three hours of sport every afternoon for those interested, featuring a rich palette of physical activities.

Further information can be found on the CERN School of Computing website (www.cern.ch/CSC) where you can find details about the programme, practical information about this year's school and details of how to apply, including the application form. The deadline for all applications is 3 May.

François Flückiger, CSC director, IT Department

IT creates two new videoconferencing rooms

The current trend at CERN to use videoconferencing more and more in everyday collaborative interactions has resulted in the recent creation of two new videoconferencing venues by the IT Department. These are available at 31-S-023 and 31-S-027 (in addition to the IT auditorium that was modernized in 2008). These new rooms are equipped following the standard proposed by IT that has been widely adopted across site: a Tandberg Edge 95 MXP codec capable of HD (720p) videoconferencing. This uses the same interface for H.323-based services, EVO and telephone conferencing. The rooms can be booked via Indico, where the bookings will be validated by the IT secretariat.

IT-UDS-AVC has used the know-how and experience acquired during the last two years in the site-wide refurbishment project for these rooms, not only for the standard technical equipment but also regarding the physical infrastructure. The colours have been harmonized to give a better video effect and the room acoustics have been improved, all resulting in a more comfortable and better user experience.

IT-UDS-AVC currently manages almost 50 standard videoconferencing venues site-wide for the four LHC experiments, BE, TE, IT departments and the DG Office.

Room 31-S-027 has videoconferencing facilities with a capacity of 10 people.

All of the components (codecs, PCs and projectors) are centrally managed and configured to reduce incidents and decrease time in setting up sessions. Support is provided centrally via the Helpdesk and the large majority (80%) of support interventions are done remotely. The service includes remote assistance, troubleshooting, short tutorials and hands-on sessions. The monthly videoconferencing activity for CERN and the LHC now covers more than 2000 meetings (half of which are via equipped meeting rooms and the rest from desktops).

Useful link

Audiovisual and Collaborative Services website: http://cern.ch/it-multimedia/

Audio Visual and Conference Rooms section, IT-UDS

New mobile phone subscriptions

A recent renegotiation of the commercial conditions with our mobile telephone operator allows us to deploy new GSM services, reduce communication costs and put in place a new subscription system.

E-mail to SMS service
This service allows you to send SMS messages directly from your CERN e-mail account. It has now been extended to cover all Swiss numbers. For further details see http://cern.ch/sms.

Multimedia message service
MMS messaging will be activated by default on all CERN subscriptions by the end of March. This service allows users to attach an image, video or audio recording to a text message. All of the details about configuring this new service on CERN mobile phones will be published on the website, at http://cern.ch/mms, in due course.

New rates
New rates for all mobile services came into effect on 1 January. All tariffs have decreased and are available at https://cern.ch/gsm/tariffs.

New subscription system
To simplify mobile communication costs and follow the evolution of industry standards, a new CERN GSM subscription system will be put in place in March. The different combinations of the current subscriptions and options will be redefined according to three professional and one private subscription types: Basic, Basic+, Full and Private.

The different functionalities of these subscriptions are summarized at https://cern.ch/gsm/telecom/types_subscr.htm.

Within this new system, please note:

- Data-transfer services (GPRS/EDGE/UMTS) will be enabled on all CERN subscriptions.
- There is only one type of subscription (Full) with international roaming rights.
- There is a new private subscription that is restricted to users paid by CERN and who were granted a GSM phone for professional use. It is not possible to obtain a new Private subscription if the user does not have a professional subscription.

Reporting in CET and HRT will be updated to reflect these changes and include both communication costs and access rights.

- List of calling rights: https://hrt.cern.ch/hrt/PhonePrivs.
- Professional or private call details: https://hrt.cern.ch/hrt/PhoneDetails.

An automatic e-mail alert system will be put in place to warn budget-code holders and the users concerned if a subscription generates excessive costs when compared with CERN's average costs.

Table 1. Subscriptions

Service (VS) subscription		Professional subscription			Private subscription	
		Basic	Basic+	Full	Private	Cost
Abonnement	Voice subscription	X	X	X	X	CERN
	Data subscription	X	X	X	X	Included in subscription
Calls	Between CERN numbers in Switzerland	X	X	X	X	Included
	External calls from Switzerland	No	Local only*	X	X	CERN or private**
	Making or receiving calls on roaming	No	No	X	X	Private
SMS	From Switzerland to Switzerland	X	X	X	X	Included
	From Switzerland to international	X	X	X	X	Private
	SMS on roaming	No	No	X	X	Private
	E-mail to SMS (to Swiss numbers)	X	X	X	X	Included
MMS	From Switzerland to Switzerland	X	X	X	X	Included
	From Switzerland to international	X	X	X	X	Private
	On roaming	No	No	X	X	Private
Data transfer	In Switzerland	X	X	X	X	Included
	On roaming	No	No	X	X	Private

*Local-only means access to phone numbers 022, 021, +33450 from Switzerland. **An interactive voice response system (IVR) may ask the user to confirm the type of call, «1» for professional or «2» for private.

SIM-card change for all CERN mobile phones
The ageing technology of our current SIM-card generation has already caused problems with the activation of some new services on our operator's network. We have also noticed an increasing rate of mechanical failures due to the deterioration of SIM-card electrical contacts.

Therefore, it is necessary for CERN to undergo a general SIM-card replacement campaign for all mobile phones. This operation will take place throughout March. At the same time, the new subscription system will be deployed according to the following schedule.

Users should have collected their SIM cards between 1 and 15 March, as indicated in the e-mail that they received on 15 February. They will have been supplied with instructions on how to carry out the SIM-card replacement.

Between 23 and 26 March
SIM cards will be migrated every night in batches of 1500 cards. Users will receive an e-mail on 17 March to warn them about the exact day of this migration. A last alert will be sent by SMS on the day of the migration. GSM numbers for "piquet" services will be migrated on 29 and 30 March during working hours. In case of problems, please contact the Telecom Lab immediately.

1 April
The Telecom Lab will resume its normal activities from this date. The opening hours of the Telecom Lab have already been extended and it is open from 8.00 a.m. to 6.00 p.m. every weekday (e-mail labo.telecom@cern.ch, tel 72480, Bat 58 R-017). In addition, the Telecom Lab team will be reinforced throughout this transition phase to answer any questions that you may have.

Useful link
GSM website: http://cern.ch/GSM

Anna Raczynska-Tartivel and Rodrigo Sierra, IT-CS

Updated SPAM filters reduce unwanted e-mails

The flow of SPAM (unsolicited e-mail) targeting CERN mailboxes is constantly increasing: about 96% of the 1.5 million mails received daily at CERN is successfully filtered SPAM, but spammers are getting smarter and detecting it is becoming ever more difficult.

To address this evolution, a new SPAM-detection software engine will be deployed. The aim is to decrease the number of SPAM messages that are delivered to your inboxes and SPAM folders.

Users can notify the Mail Service about SPAM messages by attaching the SPAM message as described here: http://cern.ch/mail/Help/?fdid=31.

To configure the spam filtering level for your CERN mailbox, go to: http://cern.ch/mail and click on "SPAM fight".

Thank you for your collaboration.

Pawel Grzywaczewski, IT-OIS (Mail Services)

The daily statistics for e-mail messages, including the amount of SPAM received.

Printer drivers evolve to become universal

A printer driver is a piece of software that allows computer applications to print without having to know the technical details of each printer model. Over the years, printer manufacturers have multiplied the number of driver versions available to cope with the evolution of their printer models. Each new model required a new specific driver, leading to a complex range of drivers to maintain. Recently, to reduce this complexity, the concept of universal print drivers has been introduced by the majority of printer manufacturers. The idea is that each vendor provides only one universal driver that will handle their complete range of printer models.

Therefore, to install new printer models on the CERN network, the CERN printing infrastructure has to evolve to this new scenario. There are added benefits, such as increasing availability and standardizing printers as well as the printing process itself. In addition, as more and more users are moving to 64-bit operating systems, the printing service has to evolve to support such universal 64-bit drivers.

New print servers will be introduced and printers will gradually be moved to this printing infrastructure. This migration is currently planned for May 2010 and, for economical reasons, by default all printers will be set to print in black and white on both sides of A4 paper. As before, users will be able to change the default print settings on their computers. A flyer will be distributed shortly to explain how to redefine these print settings.

Useful link

Print Service website: http://cern.ch/service-print/

Bruno Lenski, IT-OIS, and Natalie Pocock, IT-UDS

The EGEE community enters the era of EGI

From April 2010 the grid infrastructure, which has been nurtured and managed by Enabling Grids for E-sciencE (EGEE), composed of sites in every corner of Europe, is being coordinated by EGI.eu, the central organizing kernel of the European Grid Infrastructure (EGI).

Many details that we have been wondering about over the past months are becoming defined, and some aspects that we thought we understood are being revisited. A meeting at the beginning of February between all of the activities in EGEE covered a lot of ground. A detailed document explaining the outcomes from this meeting and the current status of the transition is available at https://edms.cern.ch/file/1060619/1/EGEE-III-MNA1.5-1060619-v0_1.pdf.

The EGI statutes were approved by the founding member states. The central organizing body, EGI.eu, has been established in Amsterdam and is now able to start advertising positions (www.egi.eu/cms/about/jobs) and hiring personnel. Because this can take several months, at the start of May EGI.eu will likely have only a skeleton staff in place. Priority will be given to filling positions in the senior management team so that National Grid Initiatives (NGIs) and users will have someone as a contact point.

We need to revisit how users will be cared for in the EGI. At the time of the all-activity meeting, feedback from the European Commission indicated that several proposals that focused on continuing support for our existing users' needs (including ROSCOE, CUE, TAPAS and SAFE) will not be funded. This means that discipline-specific and generic community application porting, training, dissemination and user support is particularly affected. It is likely that the existing, well established user communities will have sufficient critical mass to continue, but the situation is not so clear for those under formation.

The proposed EGI-InSPIRE project (the new funding mechanism for the infrastructure as a whole) intends to continue some user-focused services. These include the application registry (available at http://appdb.eu-egee.org/), which will be enhanced to distinguish between applications and tools, and linked to the RESPECT program (Recommended External Software for EGEE CommuniTies, http://technical.eu-egee.org/index.php?id=290). A new set of user-focused metrics will be developed to help EGI understand the user experience and the scientific gateways will be ported to the EGI domain.

Training will now be primarily an NGI responsibility, with EGI.eu coordinating NGI activity and maintaining a registry of trainers and a repository to which the EGEE training material will be migrated. The EGEE's training activity has been expanding the network of accredited trainers to ensure that there is at least one trainer in each NGI.

The infrastructure operations also require some attention. Interaction with the NGIs for technical and coordination matters needs to be improved because currently only the EGI official contact is known for each NGI. The variation between NGIs means that there is no single recipe for operating an NGI and so the EGEE operations activity will offer a set of guidelines that will need to be tried and adapted for each NGI. The transition will be simpler for single- or dual-country Regional Operations Centres (ROCs), such as Italy, France, Germany/Switzerland, UK/Ireland, Portugal/Spain, compared with many-country ROCs, notably central Europe and south-east Europe.

A series of meetings is being organized to follow-up on these points. EGEE will share more details as they become available.

The upcoming EGEE User Forum in Uppsala on 12–15 April will be an important milestone in gauging the transition and collecting the community together one last time under the umbrella of EGEE. The forum will host several EGI sessions to provide updates on the transition status. For details visit www.eu-egee.org/index.php?id=681.

Danielle Venton, IT-DI

TIM Web Viewer allows simple browsers to display views

Since February 2010 the Technical Infrastructure Monitoring (TIM) tool suite, which is used to monitor and control CERN's technical infrastructure, has been complemented by another application: the TIM Web Viewer. A simple web browser window is now enough to display the TIM views, which were previously only accessible via dedicated interfaces. TIM views, a set of animated synoptic diagrams, help the CERN Control Centre (CCC) operators and the equipment specialists to visualize the state of the monitored technical services.

The TIM monitoring system was developed by the ASE group in the GS Department. The system collects real-time 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 reliable, scalable and flexible control solution. It has been in production since 2006 and currently processes more than 2 million non-redundant monitoring values per day. Several client applications were built on top of TIM for processing, visualizing and analysing these data.

The TIM Web Viewer is the latest addition to the TIM tools suite, which also includes the TIM Viewer and Tim Video Viewer (*CNL* issue 2, 2009). This new application is not the only one in charge of displaying TIM views: the primary TIM Viewer application, a desktop java application heavily used by CCC operators, also relies on these views to monitor and control various technical equipment and installations on the CERN site. As well as other functionalities, the TIM Viewer enables interaction with the monitored installations via control commands, which has forced the implementation of solid security constraints and its access is restricted to the Technical Network (TN). The aim of having a tool exclusively for monitoring purposes that does not require special authorization and that is accessible from the General Network (GN) led us to develop the TIM Web Viewer.

The TIM Web Viewer is a web application accessible from the GN, which means that users can monitor the different technical services from anywhere at CERN simply by connecting their devices (laptop, PDA, iPhone, etc) to the CERN network. In case of on-site interventions, the operators and equipment specialists can easily access the application through their PDAs. Because the TIM Web Viewer is a thin client application, it is a good choice for these and similar devices with limited performance capabilities and resources. Therefore, in the next months, personal safety systems will also start to make use of the TIM Web Viewer for on-site data display at the LHC access points.

Close-up image of the TIM Web Viewer application in a PC web browser.

Due to the variety of devices that might access the application and their different characteristics, the Tim Web Viewer offers the possibility to adapt and configure its user interface (UI). This configuration can be made via two methods: direct UI interaction or URL parameterization. The direct UI interaction can be performed by using the different buttons provided in the application toolbar and the tab at the bottom right corner of the view (figures 1 and 2). A more powerful customization is supported by additional parameters in the application's URL, for example changing the refresh rate, hiding/showing the toolbar or direct access to a specific view. By building up the application URL, users can decide which configuration best suits their needs to embed some TIM views in their own applications.

The TIM views are organized into categories and can be browsed via the TIM Web Viewer. The list of available views might vary depending on where the request comes from; two different access scopes have been defined: all requests from the TN and from some authorized clients have access to the full set of TIM views, whereas requests coming from the GN only have access to a subset. The reason for maintaining these two scopes is to control access to TIM views that might contain sensitive data.

Using the existing TIM infrastructure, the GS-ASE-SSE section was able to develop the TIM Web Viewer in a short time period. This was possible because the diagrams within TIM are based on IBM's iLog Visualization technology and consequently follow the Model-View-Controller architectural pattern. In this way, a web interface could be added to the application model by making use of some existing custom iLog modules. This means that we can display all developed TIM views in the browser without making any changes to them, and easily integrate the web interface with the back-end application.

Useful links

TIM homepage: http://timweb
TIM Web Viewer: http://timweb/tim-web-viewer
TIM Web Viewer documentation: http://timweb/wiki/doku.php?id=documentation:tim-web-viewer

Marta Ruiz Garcia and Matthias Braeger, GS-ASE

Top tips for TWiki editors

Editing a TWiki page is easy: just open an edit session, write some text and save. One thing to remember is that it can be more than just a page of text and there are features that can enhance your documentation. Here are some tips to help you to make the most of your TWiki experience.

Turn your page into a slide show
The SlideShowPlugin converts your topic into a slide show presentation. Simply add %SLIDESHOWSTART% to the topic, a few headers to split the pages and some bullets as typically used for presentations. Add the tag %SLIDESHOWEND% to finish the presentation area and your page is ready for an online slide show.

Table of contents
The variable %TOC% will automatically create a table of contents for a topic based on the headers of the topic. To place the table of contents on the right of the page use the following at the top of your page:

```
<table align="right"><tr>
<td>%TOC%</td></table>
```

Roll-back to a previous version
A "Restore topic" feature has been added to the "More topic actions" menu to easily restore an older version of a topic.

Editing with WYSIWYG
The latest version of TWiki has integrated the TinyMCE WYSIWYG editor. This is a fast, highly functional editor and supports a wide range of browsers. The CERN TWiki default editor is "raw edit" but both are available. Each user can set the preference for the editor button by setting the EDITMETHOD in the user homepage.

```
Set EDITMETHOD = wysiwyg
```

Print a workbook in PDF
The PDF button found at the top right of each page will generate a PDF version of the page. It is also possible to generate a PDF document containing all of the descendants of the base topic as separate chapters. For example, if you create a ParentTopic, then create topics FirstChild and SecondChild with ParentTopic as their parent topic, then you can create GrandChildOne with FirstChild as its parent. You get a tree:

```
ParentTopic
  - FirstChild
    - GrandChildOne
    - GrandChildTwo
  - SecondChild
    - GrandChildThree
```

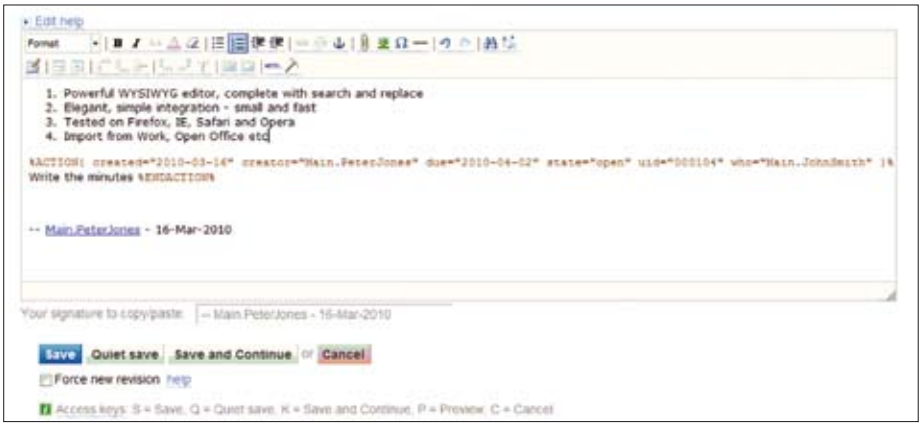

The TinyMCE WYSIWYG editor is a fast editor that supports a wide range of browsers.

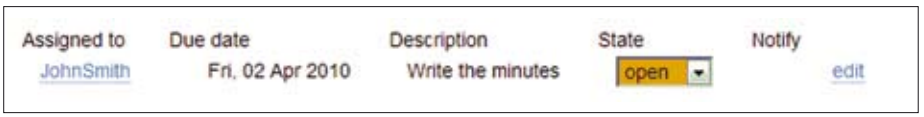

The ActionTracker plugin allows you to change the status directly on your TWiki site.

If you add "?pdfrecursive=on" to the URL parameters, all of the topics will be rolled into the PDF.

Tree view of your documents
You can also dynamically create a site map for a given web or for a selection of topics. The following example returns a site map for all topics that were created under the topic ParentOfMyFaqs in the Sandbox web.

```
%TREEVIEW{web=" Sandbox "  topic="
ParentOfMyFaqs " formatting="ullist"}%
```

Shared blackboard
Collaborate with your project team by using TWiki as a virtual blackboard. Set up a few user topics: SubBlackboard1, SubBlackboard2, SubBlackboard3 and SubBlackboard4. Now with the %INCLUDE% variable make a main blackboard topic that contains each sub-blackboard.

This allows several people to write on the same page simultaneously.

```
<table>
<tr><td>%INCLUDE%{"SubBlackboard1"}
<td>%INCLUDE%{"SubBlackboard2"}
<tr><td>%INCLUDE%{"SubBlackboard3"}
<td>%INCLUDE%{"SubBlackboard4"}
</table>
```

Search functionality
There are several search methods. The box at the top right of each TWiki page uses the CERN search, which indexes TWiki pages every day and returns fast results.

Each web has a topic called WebSearch providing a search that returns an alphabetical list of pages found for the given keywords or regular expressions.

The %SEARCH% variable can be embedded into a TWiki page to create a dynamic search. Searches can be of different types; among others, TWiki offers keyword search, regular-expression search and a flexible SQL-like query language. Inserting the following line into a TWiki topic will return the seven most recently changed pages in the current web.

```
%SEARCH{ "\.*" scope="topic"
type="regex" nosearch="on"
nototal="on" order="modified"
reverse="on"  format="| $topic |
$wikiusername  | $date |" limit="7" }%
```

Access control with egroups
Since November 2009, egroups can be used for access control on the web level or in individual topics. For example, to allow read access only to members of the egroup myegroup:

```
* Set ALLOWTOPICVIEW = myegroup
```

Action tracking
It is possible to track actions from the browser with the ActionTracker plugin, e.g. by changing their status inline. This can also be used for managing personal to-do lists and for highlighting things.

```
%ACTION{ who="JohnSmith" due="2 Apr
2010" state="open"  }% Write the minutes
%ENDACTION%
```

Report a problem
E-mail twtool.support@cernSPAMNOT.ch.

Peter Jones, Alex Bernegger, Nils Høimyr, IT-PES

Attackers are on the prowl for your security weaknesses

The academic freedom at CERN demands a rather open computing environment, which is reflected in the liberal nature of the CERN Computing Rules (http://cern.ch/security/ComputingRules). Depending on your project(s), you have the liberty to use whatever programming language you wish to use, be it to program your physics analyses, your controls software or your web application. For the CERN Computer Security Team, however, such a liberal environment poses an interesting challenge, because this freedom of choice is hard to control and even more difficult to secure. Even more interesting, this challenge has to maintain a justifiable balance between CERN's academic freedom and its security.

You are responsible for securing *your* systems

To keep this balance, you are obliged to assume your share of computer security at CERN: you are responsible for the security of your systems, your services, your programs. In order to help you, the Security Team has prepared three new means of finding security weaknesses in your AFS folders, your code and your web applications.

1. Find *your* credentials before attackers do

Protecting sensitive information is an important cornerstone of computer security. Just as misconfigurations or vulnerabilities in applications can be exploited to attack computer systems, disclosed sensitive information presents a potential security risk too. For example, the public disclosure of passwords or private keys might allow an attentive attacker to directly enter CERN computer systems like LXPLUS or Windows Terminal Servers. In particular, AFS is a worldwide-accessible file system, so sensitive information stored there must be well protected.

Control *your* AFS access rights

During the past months, many credentials have been exposed on public and unprotected AFS folders. In many cases this is due to users having accidently misconfigured certain AFS access rights by mistakenly using the LINUX-based "chmod" commands, where "fs setacl" should have been used. Previous *CNL* articles explained how data can be protected by AFS' built-in access control mechanisms (see "AFS revisited: Controlling access" in *CNL* issue 3, 2008).

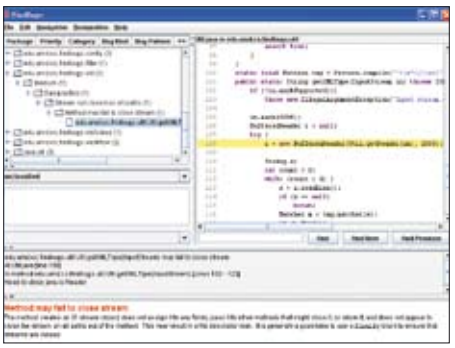

A screenshot of Findbugs, a Static Source Code Analyser available for Java.

In addition, the Security Team provides a tool to check (and if desired to correct) the access control lists on a user's home directory. The tool is aware of the structure of home directories, such as the default locations of private keys, and checks the access restrictions accordingly. This ACL checker is available from http://cern.ch/security/AFS.

Enforcement of default AFS access rights

To identify information disclosure early, the Security Team, in collaboration with the AFS Team, has started to scan proactively AFS folders at CERN. The scans are supposed to identify world-readable files containing sensitive information like private SSH keys. Going one step further, a campaign will be started during the next months to enforce restrictive settings for a well defined set of AFS folders. Once applied, public read access to, for example, your AFS home folder (~/*), to your private folder (~/private) and to your hidden folders (~/.*) will be disabled, and kept disabled. Furthermore, public write access to any of your AFS folders will be withdrawn, if the same folder also has public read access.

2. Find *your* bugs before attackers do

Another important part of improving computer security at CERN is to make the attackers' job harder by minimizing the number of bugs and exploitable flaws in newly developed software. While it is very hard to write perfectly secure software, there are some simple means to improve the reliability and security of your piece of software.

One of the easiest ways is using Static Source Code Analysers. These tools look at the code without executing it, but point out what they consider to be potential weaknesses. A typical example of what such tools can find is calls to the "gets" C function. This function is inherently insecure and can lead to buffer overflows. Specially crafted user input values can for instance allow an attacker to access or modify confidential data or even take control of any computer executing that piece of vulnerable software.

Of course, such automated source-code analysis tools are only able to find a fraction of bugs present in a piece of code. Still, you are strongly recommended to use them to find at least some of the existing weaknesses (of any type, not only security related) – especially given how easy it is to install and run these tools.Because these tools need to understand the code, they are language specific.

Static Source Code Analysers at *your* disposal

The Security Team has evaluated a number of such tools, for various programming languages (C/C++, Java, PHP, Perl and Python) and has compiled a shortlist of recommended analysers, selected for their ease of use, their simple configuration and their established benefit in pointing out potential bugs.

This list of tools is available at http://cern.ch/security/CodeTools, along with advice on installation, configuration and recommendations on how to fix the most common errors, as well as pointers to websites and books containing further information. Most of these tools are available in SLC yum repositories, making it easy to install and start using them.

RATS, shorthand for Rough Auditing Tool for Security, covers C, C++, PHP, Perl and Python. It mainly targets calls to commonly misused or exploited library functions. It is a very fast tool, available on Linux and Windows.

For C/C++, David Wheeler, a renowned IT security expert, provides Flawfinder, which looks for risks of buffer overflows and race conditions. It is unfortunately only available on Linux. **Coverity** is a security company with extensive experience in C/C++ static analysis, responsible for finding many bugs in major open-source projects, such as the Linux kernel or implementations of samba and is contracted by the US Department of Homeland Security and Yahoo among others. Currently, the PH/SFT group is arranging an agreement with Coverity,

which will enable CERN to use this tool.

For Java, **FindBugs** will find various security-related and non-security-related errors, for example vulnerabilities to SQL injection. It is available both as an Eclipse plug-in and as a stand-alone Java application.

Pixy will review PHP code and warn against risks of SQL injection and cross-site scripting. It follows execution paths in PHP code, looking for insecure use of values coming from user input.

The **Perl::Critic CPAN** module will raise warnings for many risky Perl idioms and, used in conjunction with Perl's tainted mode, it should help to produce secure Perl code.

As for Python, the most important thing is simply to keep your version of Python up to date with security patches (for SLC machines, it is done for you – the security patches are back-ported to older versions of Python). One step further is to use **pychecker**, which is very good at finding various types of potential bugs (not only security related).

3. Find *your* web application vulnerabilities before attackers do

Websites and web applications are the part of the CERN computing infrastructure that is directly visible and accessible from outside the organization. This means that attackers can see them too and often actively browse the websites maintained by you, and search for vulnerabilities such as cross-site scripting, code or SQL injection, local or remote file execution, cross-site request forgery, etc. Successfully exploiting such weaknesses in websites or web applications could result in web defacement, unauthorized access to information or functionalities, stolen users credentials, etc.

Last year, the CERN Computer Security Team started various initiatives to address this potential issue. Several technical training courses are now regularly offered for developers, including web and PHP developers. In co-operation with the Web Services Team, defaults for newly created centrally hosted websites were changed to be more secure (e.g. made visible only to CERN-authenticated users, using only static pages). These settings can of course be opened if needed for a particular website.

Scanning for *your* vulnerabilities

Soon, the Security Team will start vulnerability scanning of all CERN websites. The goal of this scanning is to detect security vulnerabilities and to improve the quality of CERN sites. All websites and web applications at CERN, visible on the internet, or on the General Purpose Network (the CERN office network) will be scanned regularly. Issues found will be reported by e-mail to the relevant website owners and must be fixed in a timely manner. Website owners may also request a one-off scan of their site or web application, by sending an e-mail to Computer.Security@cern.ch.

These web scans will be performed with automated tools. Several web application scanners were evaluated, and two of them (W3AF and Wapiti – both open-source tools) were finally chosen. Such tools first crawl a given website, looking for all scripts and their arguments. Then, they provide specially crafted values to each argument of each script, to detect those that are vulnerable to one of the common web application attacks. As with the static source-code analysis tools discussed above, such an automated approach cannot guarantee to detect all problems, but it is good enough to pinpoint the most common.

These web scans at CERN are designed to limit the impact on the website being scanned. Nevertheless, in very rare cases these scans may cause undesired side effects, e.g. generating a large number of log entries, or crashing particularly badly designed or badly implemented web applications. If a website is affected by these security scans, it will also be susceptible to more-aggressive scans that can be performed at any time by a malicious attacker. Such web applications should be fixed and also additionally protected (e.g. by restricting their visibility).

More security measures to come

Our list of projects mitigating security weaknesses is much longer. For example, in the near future we will start to review the configuration of CERN's outer perimeter firewall. Therefore, current firewall openings of your devices will be reviewed and you will be asked to consider closing those that are no longer needed.

Let's work together to find the right balance to make CERN a more secure place while maintaining its academic freedom.

If you have any questions or comments, contact Computer.Security@cern.ch.

CERN Computer Security Team

Can you solve our computer security quiz?

How many bugs can you find in this code? E-mail your findings to us at Computer.Security@cern.ch and the person who finds the most valid bugs will win the book *24 Deadly Sins of Software Security* by Michael Howard, David LeBlanc and John Viega.

The deadline for replies is 30 April. The winner will be announced in the next issue of *CNL*.

```
/* Safely Exec program: drop privileges to user uid and group
 * gid, and use chroot to restrict file system access to jail
 * directory. Also, don't allow program to run as a
 * privileged user or group */
void ExecUid(int uid, int gid, char *jailDir, char *prog, char *const argv[])
{
  if (uid == 0 || gid == 0) {
    FailExit("ExecUid: root uid or gid not allowed");
  }

  chroot(jailDir); /* restrict access to this dir */

  setuid(uid); /* drop privs */
  setgid(gid);

  fprintf(LOGFILE, "Execvp of %s as uid=%d gid=%d\n", prog, uid, gid);
  fflush(LOGFILE);

  execvp(prog, argv);
}
```

• Courtesy of Barton Miller, University of Wisconsin, Madison, US

The CNL is evolving – have your say in our reader survey
https://espace.cern.ch/cnl-survey/

Siemens openlab team lays emphasis on cyber security analysis for industrial control systems

The growing use of Ethernet and TCP/IP in industrial devices (replacing dedicated networks) has led to the necessity to reach a higher level of security against common threats on Ethernet cable. These threats can be deliberate (attackers), collateral (viruses and worms) or accidental (misconfigurations). Moreover the introduction of more IT functionalities into process-control devices gives us more reasons to perform security analysis to find any possible weak points.

The collaboration between Siemens and CERN openlab focuses on the robustness of automation devices (e.g. Programmable Logic Controllers) through a deep investigation of these devices' resistance against attacks. More specifically, the major aim of the project is the definition of a test bench and specific procedures, which allow us to perform a security mapping of devices' architecture and to simulate common attacks originating from either the internal or the external network.

Once the security mapping is complete, it is necessary to generate a detailed vulnerability report. It specifies the security breaches that need to be analysed to develop several practical and easy-to-apply solutions to fix those vulnerabilities.

Standards and guidelines can be used to help identify problems and reduce the vulnerabilities in a cyber security system. By knowing the problems and vulnerabilities, standards can be applied to cyber security systems to minimize the risk of intrusion. This is why at the beginning of our activities we compared three cyber security standards: ISA-99 (and part of the ISA-95), NERC-CIP and IEC-62351.

During the analysis of these standards we have noticed lots of congruencies and some discrepancies in the specific approaches that they suggest. At the end of this analysis, ISA-99 seems to be the most relevant standard, the only one able to face up to the wide heterogeneity of control systems (also relevant for CERN experiments).

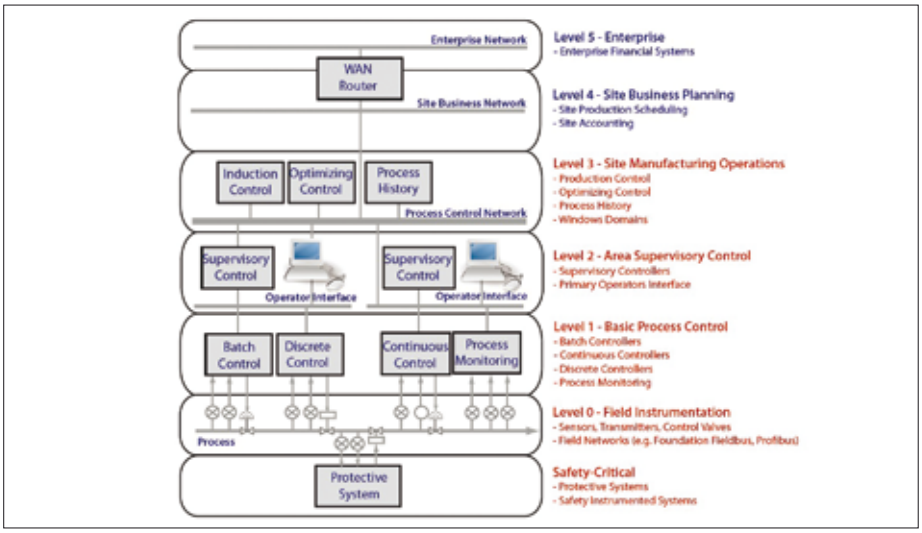

The general reference model for the distributed control systems.

This also implies that the ISA-99 approach is quite general and can only provide a theoretical (instead of practical) guidance and direction on how to establish and implement procedures (overall in the assessing phase, designing the security plan and defining the security policies). The defense-in-depth model is sustained as customer's security scheme by the ISA-99 standard too, which recognizes that some attacks will inevitably penetrate the boundaries and thus requires further protections within the boundaries.

Programmable Logic Controllers represent the lowest level in the layers architecture of control systems. As such, they are an essential link in any defense-in-depth strategy and must be considered as first-class citizens in the chain of control.

Component testing is finalized to assure that the specific component meets the required security specifications. To do this we have defined some procedures and an entire test bench, which allow us to validate the confidentiality, integrity and availability of every process-control device. In this context, one of the major problems is represented by the definition of the features and key cyber-security aspects (relevant to CERN) that must be tested, and of the minimum level of compliance that would allow us to identify whether a component is safe or not.

Unfortunately, at the moment there are no standards able to provide any criteria or specific procedures that must be followed. For this reason, we have deployed and developed specific techniques and methodologies of attacks to evaluate the robustness of process-control devices.

In the following phase, we are reporting all of the discovered vulnerabilities that need to be fixed to improve the quality and security level of these control devices, which are widely deployed at CERN.

Useful link

CERN openlab: http://cern.ch/openlab

Filippo Tilaro, EN-ICE (CERN openlab)

• This article was published in the *CERN openlab newsletter* in January 2010.

Calendar

April

7–10, **International Conference on Computer Supported Education (CSEDU 2010)**
Valencia, Spain
www.csedu.org/

7–10, **International Conference on Web Information Systems and Technologies WEBIST**
Valencia, Spain
www.webist.org/

12–16, **5th EGEE User Forum**
Uppsala, Sweden
http://egee-uf5.eu-egee.org/

13–14, **Fourth Workshop in Information Security Theory and Practice (WISTP 2010)**
Passau, Germany
www.wistp.org/wistp-2010

26–30, **The 19th International World Wide Web Conference (WWW '10)**
Raleigh, NC, US
http://www2010.org/www/

May

17–21, **Fifth International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2010)**
Angers, France
www.visigrapp.org/