CERN COMPUTER NEWSLETTER

CERN - IT Division
CH-1211 Geneve 23
SWITZERLAND

April – June 2003

CERN-CNL-2003-002
Vol. XXXVIII, issue no 2

Contents

Editorial

Dear CNL readers,

It is time for the Summer edition of the Computer Newsletter of this year 2003. The main topic of this issue is the CERN openlab for DataGrid applications, with a general article describing this new model for partnership between CERN and industry, and two other articles extracted from the official CERN Press Release: first on the participation of IBM in the CERN openlab, and second on the successful completion of a major data challenge pushing the limits of data storage to tape with StorageTek.

Apart from that, you will find:

- An update on the DataGrid project with Release 2.0 in June.
- Details of the new registration procedure required to identify portable computers.
- An explanation of the different kinds of email addresses you can use to get some support, or to contact a specific service, and the consequences that depend on whether the address is connected to the general problem tracking system, or to a mailing list.
- News from the Software Development Tools Service, and their new contact email address: `swtool.support@cern.ch`.
- Our regular selection of some recent "Questions and Answers" that have been raised at the Computing Helpdesk.

Let me wish you a happy summer, and enjoy reading this new edition of the CERN Computer NewsLetter.

Best regards,

Nicole Crémel, CNL Editor, IT/User Services

1. Physics Computing

Functional description of all the Central Public Services, specifically oriented to use by the Research sector, either for interactive (central and CPU intensive) or batch processing. The main physics-computing services consist of providing computer farms, central data recording facilities and central data storage facilities. We also include in this chapter articles on Scientific Software, such as DataGrid related information, or software for the physics experiments.

1.1 The CERN openlab for DataGrid applications

Sverre Jarp and François Grey, IT / openlab

Abstract

The CERN openlab is a new model for partnership between CERN and industry is integrating and testing cutting-edge computer technologies that show promise for the LHC Computing Grid. The Board of Sponsors meeting was held on June 13 at CERN, and this article summarises key achievements that were presented to the sponsors.

The CERN openlab fits into CERN's portfolio of Grid activities by addressing a key issue, namely the impact of cutting edge IT technologies - which are currently emerging from industry - on the technological roadmap for the LCG. Peering into the technological crystal ball in this way can only be done in close collaboration with leading industrial partners. The benefits are mutual: through generous sponsorship of state-of-the-art equipment from the partners, CERN gets early access to valuable technology that is still several years from the commodity computing market that the LCG will be based on.

In return, CERN provides demanding data challenges which push these new technologies to their limits - this is the "lab" part of the openlab. CERN also provides a neutral environment for integrating solutions from different partners, to test their interoperability. This is a vital role in an age where open standards (the "open" part of openlab) are increasingly guiding the development of the IT industry.

The CERN openlab for DataGrid applications was launched in 2001 by Manuel Delfino, then IT Division Leader at CERN. After a hiatus during which the IT industry was rocked by the telecoms crash, the partnership took off in September 2002, when HP joined founding members Intel and Enterasys Networks, and integration of technologies from all three partners led to the CERN opencluster project. In April of this year, IBM joined, and the Annual Sponsors meeting held at CERN on June 13th was an opportunity to present the results obtained so far to top managers representing the industrial partners.

CERN opencluster

At present, the CERN opencluster consists of 32 Linux-based HP rack-mounted servers, each equipped with two 1 GHz Itanium-2 Intel processors. Itanium-2 uses 64-bit processor technology, which is anticipated to displace today's 32-bit technology over the next few years. As part of the agreement with CERN openlab partners, this cluster is planned to double in size during 2003, and double again in 2004, making it an extremely high-performance computing engine. Very recently, IBM joined the CERN openlab, contributing advanced storage technology that will be combined with the CERN opencluster (see next article).

For high-speed data transfer challenges, Intel has delivered 10 Gbps Ethernet Network Interface Cards (NICs) which have been installed on the HP computers, and Enterasys Networks has delivered three switches equipped to operate at 10 Gbps and with additional port capacity for 1 Gbps.

Over the next few months, the CERN opencluster will be linked to the EDG testbed, to see how these new technologies perform in a Grid environment. The results will be closely monitored by the LCG project, to determine the potential impact of the technologies involved. Already at this stage, however, much has been learned that has implications for LCG.

For example, thanks to the preinstalled management cards in each node of the cluster, automation has been developed to allow remote system restart and remote power control. This development confirmed the notion that - for a modest hardware investment - large clusters can be controlled with no operator present. This result is highly relevant to LCG, which will need to deploy such automation on a large scale.

Several major physics software packages have been successfully ported and tested on the 64-bit environment of the CERN opencluster, in collaboration with the groups responsible for maintaining the various packages. Benchmarking of the physics packages has begun and first results are promising. For example, PROOF (Parallel ROOT Facility) is a version of the popular CERN-developed ROOT software for data analysis, which is being developed for interactive analysis of very large ROOT data files on a cluster of computers. The CERN opencluster has shown that the amount of data that can be handled by PROOF scales linearly with cluster size: on one cluster node, it takes 325s to analyse a certain amount of data, and only 12s when all 32 nodes are used.

Data challenges

One of the major challenges of the CERN opencluster project is to take maximum advantage of the partners' 10 Gbps technology. In April, a first series of tests was conducted between two of the nodes in the cluster, which were directly connected ("back-to-back" connection) through 10 Gbps Ethernet Network Interface Cards. Transfer reached a data rate of 755

MB/s, that is, approximately 3/4 of the maximum attainable bit rate of the interfaces. The transfer took place over a 10 km fibre and used very big frames (16 KB) in a single stream, as well as the regular suite of Linux Kernel protocols (TCP/IP). The best results were obtained when aggregating the 1 Gbps bi-directional traffic involving 10 nodes in each group. The peak traffic between the switches was then measured to be 8.2 Gbps. The next stages of this data challenge will include evaluating the next version of the Intel processors. The CERN opencluster equipment also contributed to the data challenge carried out in May, where storage-to-tape rates of 1.1 GBps were achieved for periods of several hours by Bernd Panzer and his team. In order to simulate the LHC data acquisition procedure, an equivalent stream of artificial data was generated using 40 compute servers. This data was stored temporarily to 60 disk servers, which included the CERN opencluster servers, before being transferred to the tape servers. A key contributing factor to the success of the data challenge was a high performance switched network from Enterasys Networks with 10 Gbps Ethernet capability, which routed the data from PC to disk and from disk to tape The technology that IBM brings to the CERN openlab partnership is called Storage Tank®. Conceived in IBM Research, the new technology is designed to provide scalable, high-performance and highly available management of huge amounts of data using a single file namespace, regardless of where or on what operating system the data reside. (Recently, IBM announced that the commercial version will be named IBM TotalStorage® SAN File System.) IBM and CERN will work together to extend Storage Tank's capabilities so it can manage of the order of a petabyte of data and provide access to the data from any location worldwide. As part of their openlab sponsorship, IBM has just installed a 28TB TotalStorage system in the Computer Centre.

Storage Tank (ST) goes beyond traditional cluster file systems in two key ways. First, it can support file sharing between heterogeneous computers running different operating systems, whereas typical cluster file systems only support homogeneous (single operating system) sharing. Second, ST employs policy-based storage management to simplify and centralize storage management for all the data in the enterprise. Administrators specify policies for how backup, restore, migration and allocation are to be performed. ST enforces these policies automatically, without intervention.

An open approach

While many of the benefits of CERN openlab for the industrial partners are of a technical nature, there is also a strong emphasis in CERN openlab's mission on the opportunities that this novel partnership provides for enhanced communication and cross-fertilisation between CERN and the partners, and between the partners themselves. Top engineers from the partner companies collaborate closely with the CERN openlab Technical Team. This team includes staff from the ADC and CS groups as well as support from the FIO group. As part of the sponsorship, HP is funding two CERN fellows to work on the CERN opencluster. CERN openlab also organises thematic workshops on specific topics of interest, bringing together leading technical experts from the partner companies, as well as public First Tuesday events on general technology issues related to the CERN openlab agenda, which attract hundreds of participants from the industrial and investor communities. A new meeting room next to the User Area of the Computer Centre, nicknamed the openlab openspace, was inaugurated earlier this year. This is being actively used for VIP visits to the Centre.

A CERN openlab student programme has been created, which will bring together teams of students from different European universities in July in August - 11 students in all - to work on applications of Grid technology. And CERN openlab is actively supporting the establishment of a Grid Café for the CERN Microcosm exhibition - a web café for the general public with a focus on Grid technologies, including a dedicated website that will link to instructive Grid demos.

Efforts are ongoing in the CERN openlab to evaluate other possible areas of technological collaboration, with current or future partners. The concept is certainly proving popular, with other major IT companies expressing interest in joining. This could occur by using complementary technologies to provide added functionality and performance to the existing opencluster. Or it could involve launching new projects that deal with other aspects of Grid technology relevant to the LCG, such as Grid security and mobile access to the Grid.

In conclusion, CERN openlab puts a new twist on an activity - collaboration with leading IT companies - which has been going on in the IT Division for decades. Whereas traditionally, such collaboration was bilateral and focused on here-and-now solutions, CERN openlab brings a multilateral long-term perspective into play. Judging by the enthusiastic response of the industrial partners at the Board of Sponsors meeting, we can expect the CERN openlab to develop into a powerful and highly visible component of the Division's strategy in the coming years.

See also the two articles (extracts from the official CERN Press Release):

1. IBM Joins CERN openlab for DataGrid applications.
2. CERN Breaks Gigabyte/s Storage-To-Tape Barrier With StorageTek.

1.2 IBM Joins CERN openlab for DataGrid applications

François Grey and François Fluckiger, IT/DI

Abstract

This article is an extract from the official CERN Press Release which was published in April 2003, and available from the CERN Public Web site (`http://public.web.cern.ch/public/`) under the link "Press & Media".

On 2nd April 2003, in an official CERN Press Release (`http://cern.ch/info/Press/PressReleases/Releases2003/PR02.03EOpenlabIBM.html`), CERN and IBM announced that IBM is joining the CERN openlab for DataGrid applications (`http://www.cern.ch/openlab`) to collaborate in creating a massive data-management system built on Grid computing.

IBM's innovative storage virtualization and file management technology, will play a pivotal role in this collaboration, which aims to create a data file system far larger than exists today to help scientists at CERN understand some of the most fundamental questions about the nature of matter and the Universe.

Conceived in IBM Research as Storage Tank®(*), the new technology is designed to provide scalable, high-performance and highly available management of huge amounts of data using a single file namespace regardless of where or on what operating system the data reside. IBM and CERN will work together to extend Storage Tank's capabilities so it can manage and provide access from any location worldwide to the unprecedented torrent of data - billions of gigabytes a year - that CERN's Large Hadron Collider (LHC) is expected to produce when it goes online in 2007. The LHC is the next-generation particle accelerator. It will recreate - on a tiny scale - conditions that existed shortly after the Big Bang, enabling researchers to answer outstanding questions about what the Universe is made of and the laws that govern its behaviour.

The very same CERN spirit that invented the World Wide Web in 1990 is now developing a new application for the Internet - Grid computing - that will push technology limits with its data processing requirements for the LHC. CERN openlab is a collaboration between CERN and leading industrial partners, which aims to create and implement data-intensive Grid-computing technologies that will aid the LHC scientists. Because the same issues facing CERN are becoming increasingly important to the IT industry, the CERN openlab and its innovative partners - which include Enterasys Networks, HP and Intel - are eager to explore new computing and data management solutions far beyond today's Internet-based computing.

By 2005, the CERN openlab collaboration with IBM is expected to be able to handle up to a petabyte (a million gigabytes) of data, which is equivalent to the information stored in 20 million four-draw filing cabinets full of paper, or 500 million floppy disks, or 1.5 million CD-ROMs.

As part of the agreement, several leading storage management experts from IBM's Almaden Research Center and Haifa Research Lab (Haifa, Israel) will work with the CERN openlab team. In addition, through its Shared University Research (SUR) programme, IBM will supply CERN with the system's initial 20 terabytes of high-performance disk storage, a cluster of six eServer xSeries systems running Linux and on-site engineering support and services by IBM Switzerland. The SUR award is valued at 1.5 million dollars for the first year.

Storage Tank employs policy-based storage management to help lower its total cost of ownership. Clustering and specialized protocols that detect network failures enable very high reliability and availability.

In this initiative, IBM is following a collaboration strategy initiated in 2001 with the European Union-sponsored European DataGrid project, which is also led by CERN.

(*) Storage Tank is a registered trademark of the IBM Corp.

1.3 CERN Breaks Gigabyte/s Storage-To-Tape Barrier With StorageTek

Bernd Panzer-Steindel, IT / ADC

Abstract

This article is an extract from the official CERN Press Release which was published in May 2003, and available from the CERN Public Web site (`http://public.web.cern.ch/public/`) under the link "Press & Media".

On 28 May 2003, in an official CERN Press Release (`http://cern.ch/info/Press/PressReleases/Releases2003/PR06.03EStoragetek.html`), CERN announced the successful completion of a major data challenge aimed at pushing the limits of data storage to tape. Using 45 newly installed StorageTek** 9940B tape drives, capable of writing to tape at 30 megabyte/s, the Data Challenges team in IT Division was able to achieve storage-to-tape rates of 1.1 gigabyte/s for periods of several hours, with peaks of 1.2 gigabyte/s - roughly equivalent to storing a whole movie on DVD every four seconds. The average sustained over a three day period was of 920 megabytes/s. Previous best results by other research labs were typically less than 850 megabytes/s.

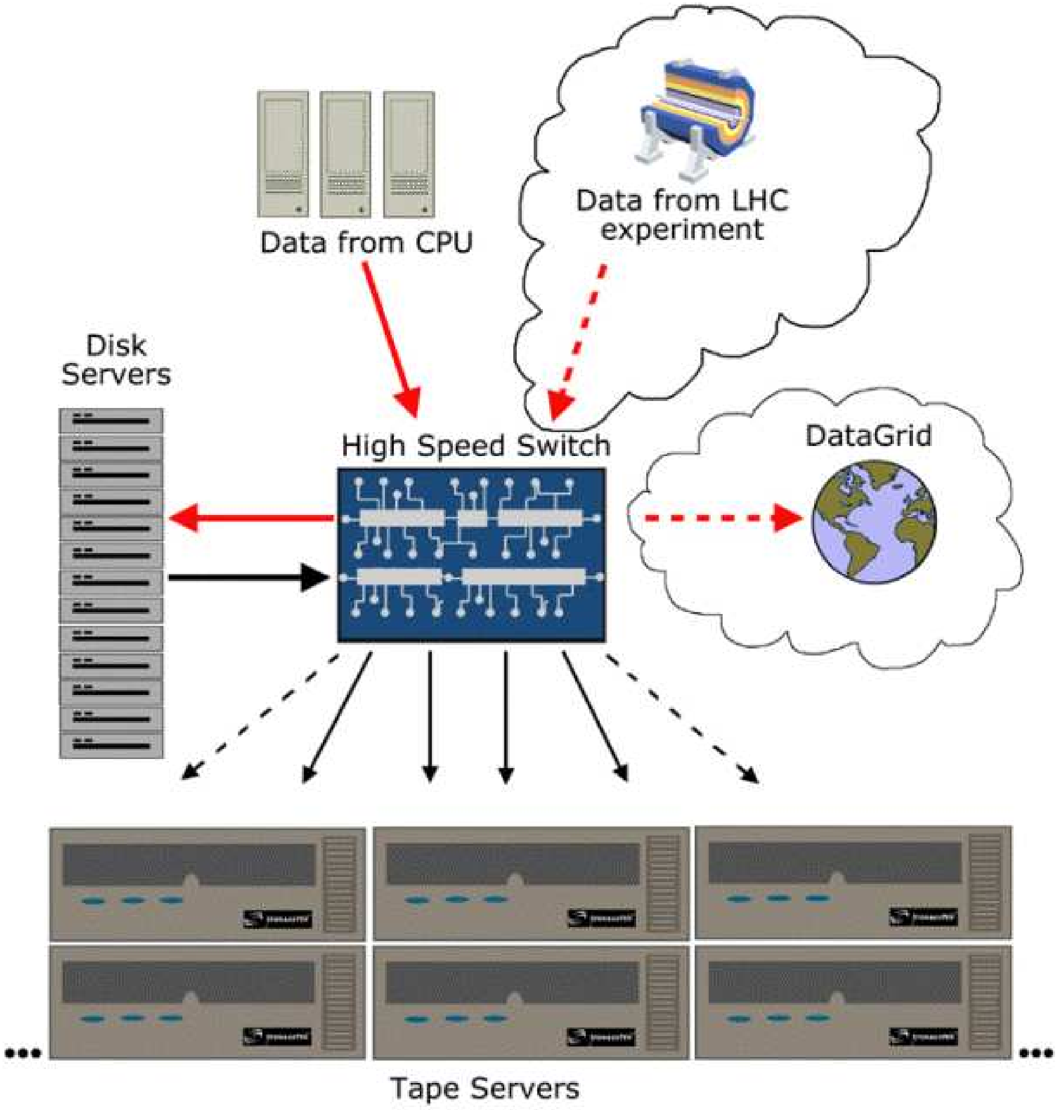

Figure 1.1: Schematic of the data setup for the storage-to-tape challenge. Data, in this case generated by 40 compute servers, is temporarily stored to disk, then copied to the 45 StorageTek tape servers. Once the Large Hadron Collider is running at CERN, the data will come directly from the experiments and a copy will be distributed onto the DataGrid that CERN and partners are currently developing.

The significance of this result, and the purpose of the data challenge, was to show that the CERN's IT Division is on track to cope with the enormous data rates expected from LHC experiments. These experiments will produce data at rates in excess of 100 megabytes/s, and Alice alone is expected to produce data at rates of 1.25 gigabytes/s.

In all, the LHC experiments are anticipated to spew out over 10 petabytes of data a year, which will be stored on tape as well as being distributed around the world onto disk, for subsequent analysis using advanced "Grid" technologies for distributed computing and data storage. The data will contain information about the result of protons colliding in the accelerator at unprecedented energies, and recreating for a brief instant the extreme conditions that existed just after the Big Bang. Scientists will spend years sifting painstakingly through this data, in an effort to better understand the fundamental laws that govern matter in the Universe.

While waiting for the LHC to be completed, IT/ADC generated an equivalent stream of artificial data, using 40 compute servers. This data was stored temporarily to 60 disk servers before being transferred to the StorageTek tape servers (see Figure 1.1). A data compression factor of 1.3 was deliberately chosen during this data challenge, as this is characteristic of the compression that can be achieved with real experimental data.

Besides the StorageTek equipment, a key contributing factor to the success of the data challenge was a high performance switched network from Enterasys Networks with 10gigabit/s ethernet capability, which routed the data from PC to disk and from disk to tape. This switched network is part of the

CERN opencluster, an advanced computer cluster also involving technology from HP and Intel.

About the CERN opencluster CERN opencluster is the first common project in the CERN openlab for DataGrid applications (http://www.cern.ch/openlab) , a partnership with industry. The CERN openlab is a response to the new level of intensive industrial collaboration needed to solve the unprecedented computing challenge of the Large Hadron Collider project, currently under construction at CERN. The current partners in the CERN openlab are Enterasys Networks, HP, IBM and Intel. The CERN opencluster currently involves 64-bit processor technology from Intel, advanced servers from HP, and a 10 gigabit switching environment from Enterasys Networks.

About StorageTek (http://www.storagetek.com/) StorageTek, a worldwide company with headquarters in Louisville, Colo., delivers a broad range of storage solutions for digitized data. StorageTek is leader in virtual storage solutions for tape automation, disk storage systems and storage networking and is a voting member of the SNIA.

1.4 European DataGrid Getting Ready for Release 2.0

DataGrid Project Office, CERN

After the 2nd EU review, held at CERN on February, (see "European DataGrid: a Successful 2nd EU Review" (http://ref.cern.ch/CERN/CNL/2003/001/DataGrid/) in the first 2003 CNL issue) more experiments, in particular ALICE and LHCb, joined the effort to evaluate the 1.4 release of the EDG software. The results, along with the ones from CMS and ATLAS, are available in a public project document, where the main lessons learned from these evaluations are summarized, together with recommendations for future developments: https://edms.cern.ch/file/375586/1.3/wp8-D8.3-0119-3-1.pdf.

The team is now getting ready for the EDG 2.0 release, to be deployed on the application testbed during June 2003. EDG 2.0 is a major milestone for the HEP community, since it represents the synchronization point with LCG and it is an important element for the LCG-1 release foreseen for July 2003. Moreover, it will provide the opportunity to synchronize with US projects (via VDT) and NorduGrid in terms of a common version of Globus and Condor and on the use of the GLUE schema and compatible information providers.

The major highlights of the release include:

- migration to Globus 2.2 delivered via VDT (bugs correction and compatibility with other grid projects)
- port to linux RH 7.3 (with gcc.2.95)
- new data management tools (EDG implementation of RLS and EDG replica manager) to overcome scalability issues with basic replica catalog
- inter-operability with US grids via use of GLUE schema as a basis for information sharing
- new information service (R-GMA) to overcome scalability and performance issues with MDS
- improved fabric management and network monitoring tools
- more scalable and reliable resource brokering and job submission chain (includes new version of Condor-G)
- new storage element service providing basic, consistent interface to disk and mass storage systems

A number of the software components developed by EDG for release 2.0 have been adopted for the initial LCG-1 grid facility for the LHC experiments. They include the Replica Location Service (RLS), the Resource Broker (RB), the data management tools and the gateway between the grid and the local fabric. Further modules such as R-GMA and VOMS (backward compatible fine-grained authentication mechanism) and extensions to RLS will be also used by LCG later in the year.

More aspects of the developments of EDG dedicated to HEP are illustrated in the presentations made by several members of the collaboration at CHEP (Computing in High Energy Physics Conference) in San Diego (USA). A collection of the DataGrid papers presented is offered on the web at:

http://agenda.cern.ch/fullAgenda.php?ida=a022246

But release 2.0 is not only for physics, and representatives of the HEP experiments have been working together with representatives from bio-informatics and earth observation, the other scientific fields supported by DataGrid, on the definition of common high-level grid services.

A more general picture of the overall status of EDG and its latest developments has been given during the "Sixth European DataGrid Project Conference", held in Barcelona the 12-15 May 2003. More then 160 people attended, representing all collaborating institutions, industrial partners and sister projects. The plenary presentations are available at:

http://edms.cern.ch/document/384120

For more information on the DataGRID project see: http://cern.ch/eu-datagrid/

---

2. Internet Services and Network

Everything related to Internet in a large sense: i.e., network issues and performances, all technical Web issues, electronic mail, Internet news, home access to CERN facilities, video conferencing and multi-media, network security, etc.

2.1 Hardware Address Registration Requirement for Portables

Jean-Michel Jouanigot, IT/CS and Denise Heagerty, IT/DI

Abstract

From Autumn 2003, only portable computers whose hardware addresses have been registered at CERN, will be able to connect to outlets for portables and wireless using DHCP.

Recent incidents have highlighted the need to increase security for connections of portable computers. It is currently planned that, from Autumn 2003 onward, only devices with a registered hardware address will be assigned a dynamically allocated (DHCP) IP address to connect to the network. This only applies to outlets for portable computers or wireless connections. This registration is required to identify devices and provide contact details for security and network related follow up. If hardware addresses (wireless and cable) for your portables are not already registered, please register them as soon as possible.

To register the hardware addresses of your portable computers, display the web page `http://cern.ch/register`, and select “Register Portable”. For security reasons, you can only access this page from *inside* CERN or using ACB (`http://it-div-cs.web.cern.ch/it-div-cs/public/services/acb/acb_contents.html`) or VPN (`http://cern.ch/vpn`) access from off-site, and you must sign your registration using your NICE account and password. If you do not have a NICE account you can request one, as described at `http://cern.ch/it-div/comp-usage`.

If you are not currently at CERN, and you have a CERN computer account, you can register the hardware addresses of your portable computers when you arrive (e.g., using one of the publicly available computers near restaurant 1 or in front of the computer centre in building 513). If you have VPN (`http://cern.ch/vpn`) access to CERN then you may register as described above.

For temporary external visitors who wish to connect their portable computer to the CERN network infrastructure (short visit from 1 to 4 weeks), who do not have a CERN computer account, a web form accessible from outside CERN will be offered from June; the CERN person who is responsible for the visit, will need to agree to the registration request.

2.2 Support Email Addresses: Mailing Lists or Problem Tracking?

Nicole Crémel, IT / User Services

There are many generic email addresses at CERN that can be used to contact an IT service when you need some help or support. It is not obvious for users to know all these addresses, or to know the exact email address that they should use in a particular case. Another difficulty is that these addresses can correspond to a service account, or to an entry in PRMS (Problem Report Management System - the Problem Tracking System that is used by IT Division to keep trace of all user problems and ensure a proper follow-up), or even to a mailing list, maybe involving a large number of people.

On Lxplus you can get a list of all the existing "support" email addresses (with syntax `Xxx.SUPPORT@cern.ch`) by executing the command:

```
phone support -limit 500 | grep "SUPPORT" | nl
```

Today, you get 166 results! Some of these email addresses are obsolete and not used anymore. We are currently reviewing the list and cleaning it up.

This article explains how all these "Support email addresses" behave and what users should or should not do.

Generic email addresses

When you do not know the email address of the IT service you need to contact, the recommendation is to send a mail to the "Computing Help Desk", at `Helpdesk@cern.ch` (or phone `78888`). Either the helpdesk team will be able to provide an immediate answer, or they will address the query to the appropriate service for you, sending back to you a ticket number (that you will need to include in all your correspondence related to the same problem).

A frequent mistake that is made by users is to send problems to several email addresses in parallel. Such users think that, by sending the mail to several places, they have more chance to get a quick solution. This is not necessarily true (people who see that the mail is sent elsewhere may think that they do not need to reply and will let someone else do it). Instead it can lead to further issues, especially if one of the email address is a mailing list that include many members.

I will give a precise example to explain the problem: imagine a member of the CMS collaboration who wants to report a mail related problem; this user knows the three email addresses: `cms.support@cern.ch`, `mail.support@cern.ch` and `helpdesk@cern.ch` (all valid addresses); he is not sure which email addresses should be used in his particular case, and then he sends the mail to the three email addresses in parallel (for instance "`To: cms.support@cern.ch`" and "`Cc: mail.support@cern.ch, helpdesk@cern.ch`). What can happen?

- `cms.support` (or `cms-support`) is a mailing list (`cms-support@cms-lb.cern.ch`). It should be used for problems specific to the CMS computing environment. It contains many members and there is no posting restriction. A web archive of all mails sent to this mailing list is public and available at: `https://wwwlistbox.cern.ch/earchive/cms-support/`
- `mail.support` is the generic email address to contact the Mail Service team, via the Computing Helpdesk who filter all mails sent to them (to avoid disturbing the Mail team if not necessary). Consequently `mail.support` is a direct entry into PRMS. As soon as the helpdesk will assign the query to the Mail service in IT a ticket number will be returned to you (the sender of the mail).
- `helpdesk` is the address of the Computing Helpdesk and the mail is again a direct entry into PRMS. Most probably another person at the Helpdesk will get the mail, and will also assign the query to the Mail service in IT: another ticket number will be returned to you.

You can already understand that 2 tickets have been created in PRMS for the same problem. Even worse, if members of the mailing list `cms-support` reply to the original mail and select the option "`Reply to all`", new tickets will, again, be created in PRMS for each mail reply.

To avoid such problems, and an overflow of ticket creation at the Helpdesk and in our Problem Tracking System, a basic rule is: *Please do not send mail to several generic email addresses in parallel!*

Mailing list or direct mail-feed into PRMS?

Another question that some users are asking is: *Is it possible to know if a generic email address is a mailing list, a service account, or a direct entry in the Problem Tracking System?*

An old rule was to try and reserve the character "`-`" (minus) for mailing lists, and the character "`.`" (dot) for service account: e.g. `atlas-support` is a mailing list, and `asis.support` is a service account. Unfortunately nothing was implemented to "force" the rule and, for user convenience, some addresses do not "respect" it: for instance, both `cms.support` and `cms-support` point to the same mailing list.

However, in most cases, the answer to that question (is it a mailing list or a direct entry to the service?) is provided by the "Phone Book Form" (or “xwho service”) at the URL `http://consult.cern.ch/xwho`.

For instance, using this form, do a "Start_search" on:

1. "CMS Support": You will get one entry registered in "Teams"; if you click on this entry "Cms SUPPORT" you can see that there is no account attached to it, and that the "valid mail addresses" are: `cms-support@cms-lb.cern.ch` for the "physical mail address" (or PEM) and `Cms.Support@cern.ch` for the "generic mail address" (or GEM). The syntax `cms-support@cms-lb.cern.ch` (for the PEM) corresponds to a mailing list. (N.B. Most mailing lists at CERN are on the `listbox` server, and then the syntax is: `name@listbox.cern.ch`).
2. "Mail Support": Again you will get one entry registered in "Teams"; if you click on this entry "mail SUPPORT" you can see the "valid mail addresses": `hfilter@mail.cern.ch` for the PEM and `Mail.Support@cern.ch` for the GEM.
3. "Computing Helpdesk": Still one entry registered in "Teams", and, for the "valid mail addresses": `hfilter@mail.cern.ch` for the PEM and `Helpdesk@cern.ch` or `Computing.Helpdesk@cern.ch` for the GEM.

Important:

All entries that do have `hfilter@mail.cern.ch` (e.g. "Mail SUPPORT", "Helpdesk", etc.) or `prms-rmf@remedy01.cern.ch` (e.g. "Design-Space SUPPORT", etc.) for the PEM, are "a direct entry in PRMS" (the Problem Tracking System), and then a ticket number ("`CTnnnnnn`" or "`CT0000000nnnnnn`") will be generated and returned to the user.

Otherwise, the generic address is a mailing list, or a service account. For a "service account" with a login id `srvlogin` the syntax of the PEM is `srvlogin@mail.cern.ch`, whereas for a mailing list, with a name `listname` the syntax of the PEM is `listname@xxx.cern.ch`, where `xxx` is the mailing list server. For most mailing lists at CERN today the server is called `listbox` and the syntax of the PEM is `listname@listbox.cern.ch`

On Lxplus, when you execute the commands:

```
phone -mail hfilter@mail.cern.ch | nl
```
and
```
phone -mail prms-rmf@remedy01.cern.ch -limit 100 | nl
```

You can find 79 (8 + 71) entries that correspond to a direct entry in PRMS.

Unfortunately the mailing lists are not visible in "`xwho`" unless (for some cases only) a service definition directly linked to this mailing list was made in CCDB.

If no service definition has been made in CCDB (the Computer Centre DataBase) then the mailing list is only visible in SIMBA, which is the web interface for CERN's mailing list service, at URL: `http://wwwlistbox.cern.ch/mmm_index.html`.

For instance "Atlas SUPPORT" has no service definition. With `xwho Atlas SUPPORT` or a search on "Atlas SUPPORT" in the Phone Book (`http://consult.cern.ch/xwho`), you will not find anything, but, in SIMBA, if you make a search for "`atlas`", you will find the mailing list `atlas-support@cern.ch`.

If a service definition is registered for the mailing list, e.g. the service "Cms SUPPORT" linked to the mailing list "`cms-support`", then two generic email addresses are both valid and correspond to the same address: e.g. "`Cms.SUPPORT@cern.ch`" and "`cms-support@cern.ch`".

Conclusion

We hope that we have clarified he final destination of all the "`service-name.SUPPORT`" email addresses, and that you understand why, in some cases, an automatic reply with a ticket number is returned to you, and why not, in other cases. However, for non specific issues, and during the CERN working hours, please always use

`Helpdesk@cern.ch`

to contact the Computing Helpdesk, who will ensure your problem is correctly registered and followed-up.

About the author(s): *Nicole Crémel is responsible, with Roger Woolnough, for the supervision of operations of the Computing Helpdesk.*

3. Technical Computing

This includes: Products support for Engineering applications (Mechanical CAE, Electronics Design Automation, Mathematics Tools, etc.), software development (tools, CVS servers) and documentation (FrameMaker); Data Management and databases (Oracle); Accelerator and Controls Computing.

3.1 News from the Software Development Tools Service

Pete Jones, IT/PS

Abstract

Latest news concerning XMLSpy, Sniff and SDT support.

SDT Problem tracking

From May 2003 the SDT service has started to make full use of the Remedy Problem Management Service (PRMS). All user problems concerning SDT tools will now be handled by Remedy.

For registering problems users are asked to use the email address `swtool.support@cern.ch`. However, for general queries please continue to use the email address `Sdt.Support@cern.ch`.

Sniff (`http://sdt.web.cern.ch/sdt/Sniff/`)

An Integrated development environment for C, C++, Java, Fortran and CORBA IDL.

There is a problem related to `glibc` with Sniff 4.1 on Redhat 7.3. The problem is fixed in Sniff 4.1.1 and is now available for use. Please refer to the URL `http://cern.ch/sdt/Sniff` for more information.

XMLSpy (`http://www.xmlspy.com/`)

XMLSpy 5 is the industry standard XML Development Environment for building software applications based on eXtensible Markup Language (XML) technologies.

Version 5 Release 4 (`http://www.altova.com/press/2003-04-07_rel4.pdf`) of XMLSpy is now available and includes enhanced support for developing customized XML content editing applications. In addition, there are over 50 bug-fixes and improvements. Release 4 can be installed from the Windows Add/Remove programs feature.

Access to this software is still restricted but anyone interested in this product should contact `Sdt.Support`.

For a full list of all SDT products please refer to the URL `http://cern.ch/sdt/`.

About the author(s):

Peter Jones is the manager of the SDT service

4. Education and Documentation

All kinds of "educational" computing issues and documents, such as computing news from the Press, tutorials, a selection of questions and answers from the Computing Helpdesk, and offerings of the Computing Bookshop.

4.1 Questions and Answers from the Computing Helpdesk

Collected by the User Assistance Team in IT/US

Nicole Crémel (Editor), IT/User Services

Abstract

This is a collection of Questions & Answers that have been treated by the Computing Helpdesk managed by the User Assistance Team in the User Services group.

N.B. The number in parentheses refers to its relative numbering in the "Question-Answer" database , which is available at the URL: `http://cern.ch/consult/qa/`NNN, where NNN is the problem identifier (number).

Windows Support

Question [3310] – Delete traces of all documents and web pages accessed on a PC

Windows keeps histories of the Web pages and documents I have recently accessed. How can I easily get rid of these traces in W2000 or W-XP?

Answer

Windows XP:

```
- Right-click Start and select Properties
- Start
- Customize
- Click Clear List (this clears the start menu from showing recently
  used programs)
- Click the Advanced tab
- Click Clear List (this clears pointers to recently used documents)
```

Windows 2000:

```
Menu: Start - Settings - Taskbar & Start Menu
- Advanced
- Clear
```

Internet Explorer:

```
Menu:  Tools - Internet Options
- Clear history, ok
- Delete files
- Tick "Delete all offline content", ok
```

Note that this might last quite a few seconds, depending on the number of web pages you have accessed.

Question [3311] – Who to contact to (re)install Windows2000

Who can I contact to install / reinstall Windows 2000 on my PC?

Answer

Concerning re-installation, users can:

- contact the Helpdesk - call (x78888) / email (`Helpdesk@cern.ch`) - and ask that somebody come for (re)installation

- or (for users who feel comfortable in doing it themselves) follow the instructions in the Winservices/docs web page (see related links).

Please note that installations of obsolete operating systems like W95 or NT are not done since the beginning of 2003.

Related links:

- Installation page from the Winservices/docs web site at: `http://winservices.web.cern.ch/WinServices/docs/Installation/`.
- End user documentation on Windows 2000 at CERN at: `http://winservices.web.cern.ch/WinServices/docs/Win2000Intro/`.

Question [3339] – Popup window to update programs - be CAREFUL that it comes from CERN

I have a PC connected to NICE (Windows 2000 or XP). I received a popup window telling me to update the Norton Antivirus program. Shall I accept it (and reply "Yes")?

Answer

ATTENTION please!

Users assumed the pop-up message came from a CERN service provider. When they followed the upgrade recommendation, it messed up the antivirus installation, requiring specialist intervention for the clean-up.

Displaying a window on your screen (rather than sending an email) can be another method of intrusion in your system (and a way to infect it with a virus).

If you are sure the message is provided by the NICE team at CERN then you can trust it. In case of doubts send an email to `Nice.Support@cern.ch` with the exact message that you got, and they will confirm if this is valid and issued by them.

Question [3313] – Windows - Integrated login failed: user doesn't exist

Each time I log in on my Nice 2000 machine, I get the following message:

"Integrated login failed: user doesn't exist".

I press OK and can log in normally. What is this all about?

Answer

This pop-up window's header line shows `"AFS logon"`.

Check the AFS client config and untick the option "Obtain AFS token when login into windows". The message will then disappear at login time.

Related links:

- Information about AFS on Windows (`http://consult/qa/0731`).

Question [3329] – Merge old / new Word versions prompt on Office XP (or 2002)

Whenever I open an updated version (e-mail attachment, for example) of an old Word document, a dialog box appears asking whether I want to merge the two versions. If I say yes, the old version is automatically merged and modified. I use Office XP.

Can you please explain why this happens and how to revert to the previous functioning, that is: to open the new version without having any questions appearing and without modifying the old version?

Answer

If you use Outlook XP and Word XP you may see the message *"Do you want to merge changes?"* when opening a document attached to an email.

One of the new "features" of Outlook XP is to automatically add Reviewing data to every Office document you attach to an email. This is only explicitly visible on a Word document, and it is added to the File Properties without your knowledge and is saved with the document you are sending.

Possible replies:

```
If you choose to merge changes and answer Yes, Word will merge the
original document into the copy of the document you are opening. This is
similar to using the Compare and Merge Documents command found under the
Tools menu. All modifications between the two documents are marked as
```

```
revisions and the Reviewing toolbar is turned on.

If you choose not to merge changes and answer No, the documents are not
merged; however, the Custom File Properties will remain and you will be
prompted to merge changes each time you open the document.

If you choose not to merge changes and answer No, and don't ask again,
the documents are not merged and the Custom File Properties are removed.
```

A solution within OUTLOOK (not Word) is to: Click `"Tools"` in the menu bar. Then select `"Options"`. On the Preferences tab click `"E-mail options"`. Then click `"Advanced E-mail options"`. At the bottom: disable `"Add properties to attachments to enable Reply with Changes"`

Please note: Office XP (Outlook XP, Word XP) is also called Office 2002 (Outlook 2002, Word 2002). It can be installed under Windows 2000.

Question [3330] – Modified NICE access rights within a group don't seem to be valid

Via the NICE group manager, I have modified folder access rights for an individual user within a group. But although all looks well, the new user doesn't seem to have the rights as per the new definition.

Answer

When one adds a user to a NICE security group, (s)he has to logout / login in order to validate his/her membership for it. Otherwise, the user will only be member of the group after the next expiration of his/her Kerberos ticket which might last up to 12 hours.

Tthis explains many of the "access denied" problems when user permissions are granted via a user group.

Question [3345] – Winzip at CERN on Windows

What is the status of WINZIP at CERN on Windows (NICE)?

Answer

As of May 20th, 2003, Winzip 8.0 will be automatically replaced by version Winzip 8.1 Service Release 1.

The update of Winzip on NICE 2000 will start on the first reboot of your computer after that date.

When the computer is restarted, you will see a message saying: "Removing managed software Winzip 8.0", followed by the message "Installing managed software Winzip 8.1 Service Release 1"

You must wait until this message disappears before being able to use your computer again.

IMPORTANT NOTE CONCERNING NICE XP USERS:

Winzip 8.1 Service Release 1 will not be deployed on NICE XP. In fact, Windows XP has a built-in functionality for creating compressed folders and opening ZIP files.

Related links:

More information about "Using WinZip at CERN" is in the documentation at: `http://cern.ch/Win/docs/winzip/`.

Question [3347] – No VPN certificate on official CERN PC at home

I have a CERN PC that I use at home.

When I tried to use the CERN VPN service I had an error message saying I did not have a certificate. Is the registration procedure bound to a machine or a user?

Answer

The problem came from an incorrect network configuration. Taking home a PC from CERN, one must change the network configuration because the user is no more physically connected to the CERN network.

Go to:

```
Start Menu -> Settings -> Network -> Local Area Connection ->
TCP/IP properties
```

and check if the IP address is a CERN one (137.138.*.*).

If it is a CERN address then this has to be changed: choose "`Obtain an IP address automatically`" or use a private address like 10.0.0.1, depending on the Internet Service Provider (ISP) settings.

Unix Support

Question [973] – Exceed connection to Lxplus

What shall I do to connect to Lxplus using Exceed on my PC?

Answer

The simplest is to open an X session for a CERN pre-defined connection from the menu:

```
Start - Programs - Hummingbird Connectivity... -
CERN Settings - Predefined X Connections - ...
```

and then select the machine you want to connect to (e.g. `Lxplus`).

You can select (left button) Lxplus to start an X session, or "drag & drop" the icon on your desktop to retrieve it next time (without going through the menu).

Related links:

Exceed documentation on Windows 2000 is available at: `https://cern.ch/Winservices/docs/Exceed/`.

Question [3346] – Batch job - status PEND and option -R

I try to run a job in batch, but it does no execute and seems to stay for ever in pending status (PEND).

```
JOBID   USER    STAT  QUEUE   FROM_HOST   EXEC_HOST  JOB_NAME    SUBMIT_TIME
423626  lechows PEND  8nm     lxplus047              *ast_prod2  May 26 11:55
```

Answer

When a job seems stuck you can execute the command `bjobs -l <jobid>` which will give detailed information about it, including pending reasons.

In your case:

```
[lxplus092] ~ > bjobs -l 423626

Job <423626> Job Name <atlfast_prod2> User <lechowsk> Project <default> Sta
                     tus <PEND> Queue <8nm> Command <atlfast.sh prod2>
Mon May 26 11:55:09: Submitted from host <lxplus047> CWD <HOME/public/LittleH
                     iggs/batch> Output File <atlfast_prod2>
 PENDING REASONS:
 Not the same type as the submission host: 363 hosts;
 Job slot limit reached: 353 hosts;
 Closed by LSF administrator: 82 hosts;
 Load information unavailable: 47 hosts;
 Queue's resource requirements not satisfied: 1 host;
```

It seems that you have submitted your job from a node on `lxplus6` (with Red Hat 6). However LSF batch jobs go by default to machines of the same type, so a simple `bsub myprog` will go to RedHat7 nodes. To cross-submit jobs between the different Linux versions you need to specify the option `bsub -R "type==linux?"` where `"linux?"` can be LINUX6 or LINUX7.

Related Links:

- Known RedHat 7 issues at: `http://cern.ch/plus/issues.html`.
- Redhat Linux version 7 (7.2) at: `http://consult.cern.ch/qa/0898`.

Question [3344] – gcc on Linux - version numbers and settings

What is the C / C++ compiler available and recommended at CERN on Linux?

Answer

On Linux (e.g. Lxplus) at CERN you should use the GNU project C and C++ compiler `gcc` or `g++`.

Version 2.95.2 is still recommended but 3.2 is being increasingly used.

To access it you may have to change your `PATH` (and `LD_LIBRARY_PATH`) variables (e.g. with `csh` or `tcsh`):

- g++ 2.95.2:

```
setenv PATH /usr/local/gcc-alt-2.95.2/bin:{$PATH}
setenv LD_LIBRARY_PATH /usr/local/gcc-alt-2.95.2/lib
```

- g++ 3.2:

```
setenv PATH /usr/local/gcc-alt-3.2/bin:{$PATH}
setenv LD_LIBRARY_PATH /usr/local/gcc-alt-3.2/lib
```

Related links

- gcc on Lxplus - cannot open shared object file (`http://consult.cern.ch/qa/3261`).

Question [3271] – NCD X terminals - problem login on Lxplus

I am trying to connect from an NCD X terminal to Lxplus. I was using "New Terminal" from the "Terminal" menu to open a telnet connection. This works for lxplus6, but does not work now for Lxplus (i.e. lxplus7).

What's the problem?

Answer

After the migration of Lxplus to Redhat 7 (lxplus7), `telnet` is no longer supported for security reasons. Therefore, on all the X terminals still available at CERN, the `"New Terminal"` from the `"Terminal"` menu (which was opening a `telnet` connection) no longer works. N.B. It is unfortunately impossible, for manpower issue, to suppress on all X terminals this item in the existing menu.

Now you have to start an X session (`"Login new X Session"`) from the `"Login"` menu. This method which is more secure was already recommended in the past.

Please note: If you try to open an X session with `"Login new X session"` from the `"Login"` menu and your terminal is blocked, you have to switch off the Window Manager running inside the NCD terminal in order to start the KDE or GNOME Window Manager. Otherwise it complains that another Window Manager is already running. To switch it off:

```
Select WindowMgr on the NCD service console
Uncheck "NCD Window Manager"
```

Related links:

- Known RedHat 7 issues at: `http://cern.ch/plus/issues.html`.
- "Exceed on Lxplus - login panel re-appears" at `http://consult.cern.ch/qa/3269`.
- "Problem to start X session (X terminal or Exceed)" at: `http://consult.cern.ch/qa/1456`.
- Public Interactive Logon Service - Security Web page, at: `http://cern.ch/plus/security.html`.

Question [3341] – ssh - login window in inverse video

I am logged in on a Sun / Solaris host (or HP-UX) and I open a remote connection to a Linux machine (e.g. Lxplus) via ssh.

The window where I open the connection (the login window) is in inverse video (colors are swapped such that the prompt is in foreground black and the words are white).

However when I open another window with the command "xterm" this new window is OK.

What shall I do such that colors in the first (login) window are not swapped?

Answer

A possible solution is to change the definition of the environment variable `TERM` in your login file. Try to a add a line (e.g. in `.tcsrc`):

```
setenv TERM vt100
```

and logout / login again.

Related links:

- with a Google search: look at a similar problem
  "Linux terminfo or xterm problem" (reported in the IT resource center forums from HP)

Question [3232] – Problem with Putty connection to Lxplus - colors swapping

I am using `Putty` on my PC to connect to Lxplus and since the migration to lxplus7 (RedHat 7) I have problem with colors. For instance, using `pine` to read my mail, the background is black and I have white background below black text.

Answer

Please try to re-install PuTTY on your PC (from `Start - Add/Remove Program` you should first remove it and them add it again). When done, please do `'Reset CERN Putty default settings'`. After that, remove your old session and create new ones.

The settings should be the following:

- black text on white background
- blue block non-blinking cursor
- non-proportional easy to read on a screen font
- correct vt100 emulation

You can of course modify the settings later yourself.

Reference link:

- The problem is most probably linked to a wrong terminal emulation. For PuTTY the command `echo $TERM` must give "`vt100`" (and not "`xterm`"). If you do not want to get all the "CERN settings" but change only the terminal emulation, a recipe is provided at: `http://juslin.home.cern.ch/juslin/putty/`

Question [3337] – bjobs command - bjobs -u group

Can you please explain how to use the LSF command `bjobs -u group` in order to get information of the batch jobs in the group ZP (for Atlas)?

I have tried `bjobs -u zp` but the output is not correct. I get:

```
[lxplus008] ~ > bjobs -u zp
No unfinished job found
```

which is wrong (jobs are still pending).

Answer

Group `ZP` is the registration group and the command `bjobs` works with LSF (or Unix) groups. Unix groups are implemented as `u_XX` names (to distinguish them from real user names that never contain "_"), so for Atlas you should specify: `u_zp`, and execute:

```
bjobs -u u_zp
```

to display information of all jobs in group `ZP`.

Question [3336] – scp / ssh on Lxbatch - Connection timed out - lost connection

I am running a batch of jobs on Lxplus. At the end the job copies the output back to an external machine in Lancaster via the command `scp`. However the jobs do not finish properly, and, when looking at the job standard error output produced by unfinished jobs with the command `bpeek 39980`, I get the following message:

```
++ scp simdet_v4.hist ajf@lapw.lancs.ac.uk:simdet4c/lxplus2/qq_209.hist
ssh: connect to address 148.88.89.197 port 22: Connection timed out
lost connection
```

What is the problem? Please note that this worked a few weeks ago when I last ran this production.

Answer

Some of the LXBATCH nodes are on CERN-internal networks. In this case, it looks like your job 39980 is being executed on `lxb0003`, which is neither reachable from the outside, nor can connect to there (`External TCP/IP Connectivity set to NONE`, as you can see in the `http://network.cern.ch/register` database).

N.B. This is a fairly recent change in the Lxbatch cluster.

It is possible to specify that your jobs need WAN-wide connectivity by adding `'-R wan'` to the job submission (command `bsub -R wan`).

If you need more details on this change you should contact directly the LSF administrators at `Lsf.Administrator@cern.ch`.

Mail and Web Support

Question [3343] – Cannot open URL (HTML link) in mail

I have received a mail which contains an URL, but when I try to open the URL directly form my mail client (e.g. Outlook) I get an error:

`The page cannot be found`

However I have checked the URL is valid, e.g.:

`http://info.web.cern.ch/info/Announcements/CERN/2003/05-22G8Info/Message-en.html`

Answer

It seems that some mail clients (or programs) sometimes "cut" the URL and then when you open it the URL is not complete (which means you open an invalid link).

You should first check, in your browser, that the URL which is open is complete: if you notice the end is missing, you will need to add it manually (for instance with a "cut & paste").

In the example you give, the URL has probably been truncated after the `"-"` character and you probably need to add `"en.html"` at the end of the URL in the browser line for the "Address".

Question [156] – Read mail from outside CERN

When I travel for work, how can I read my mails? (i.e. access my CERN mailbox from outside).

Answer

The *recommended way* to access your CERN mailbox from outside is to migrate to the new Exchange 2000 infrastructure which has been approved by the Desktop Forum to be the new platform for CERN's Mail Services. This contains a supported product, Outlook Web Access (OWA in short), which allows to read/send email from any web browser.

You can easily migrate your CERN mail account transparently to the new infrastructure by using the appended link to migration instructions, then begin to use the new OWA interface. More detailed technical information follows.

For reading email remotely:

- One option is to authenticate by some way to the CERN network (e.g.: connect by ACB or VPN) and use the same mail clients as when at CERN.
- Another option is to connect through an ISP without VPN then:
  - Use OWA as stated above.
  - Use Outlook XP configured with ISA, do the same as when at CERN (see related links).
  - ssh to lxplus.cern.ch then use pine or Mozilla through a tunneled X11 connection.
  - Configure your local mail client to use IMAP/SMTP protocol to access the mailbox. For Outlook XP this is not recommended (use the ISA config), but for netscape / pine / others you can do it. Users migrated to the new infrastructure should activate SSL encryption.

For sending email remotely: The only tricky issue is when the connection is through an ISP without VPN and the local mail client is configured to use IMAP/SMTP. The SMTP server must be changed from `smtp.cern.ch` to:

- `mmm.cern.ch` (for users migrated to the new infrastructure). Activating encrypted connection through SSL is also recommended.
- `maildev.cern.ch` (for users still on the old mail infrastructure).

This must be done because `smtp.cern.ch` does not authenticate users, and we don't allow un-authenticated relaying.

Related links:

- Instructions for migration to mmm (Exchange infrastructure) at: `http://mmmservices.web.cern.ch/mmmservices/Services/Migration/`.
- Instructions to configure Outlook XP with ISA server at: `http://mmmservices.web.cern.ch/mmmservices/Help/contents/0_4_isaconnect.asp`.
- Connecting from home via ACB at: `http://it-div-cs.web.cern.ch/it-div-cs/public/services/acb/acb_contents.html`.

Question [3290] – mp: command not found on Lxplus

In order to print a mail from Pine on Lxplus I was using the command "mp" (POSTSCRIPT pretty printer) for a nice formatting. The command I used from Pine was:

```
mp | xprint -P30-6043-hp -h
```

This command now gives an error message:

```
sh: mp: command not found
not a readable, nonzero length file - /tmp/lpr20790
nothing to print
```

Is it possible to re-install this nice formatting program?

Answer

This program is not included in the standard RedHat distribution and this is why it is not installed on Lxplus today (after the migration to RedHat 7). It used to be included in ASIS, but now, there is no-one to maintain this product.

So, basically, the 2 options provided by the Lxplus system administrators (at mailing list `it-div-fio-lxplus7@cern.ch`) are:

1. Install and use `mp` yourself only.
2. Help us as well as other people: Become a maintainer of `mp` and deliver us a RPM that you are happy with and we can install on LXPLUS.

Related links:

- for MP installation:
  - `http://rpmfind.net/linux/RPM/contrib/libc6/i386/mpdist-3.5.2-2.i386.html`
  - `ftp://sunsite.cnlab-switch.ch/mirror/redhat-contrib/libc6/i386/mpdist-3.5.2-2.i386.rpm`
  - `http://sourceforge.net/projects/mpdist`
- "Tools (or software) missing on Lxplus (compared to lxplus6)" at: `http://consult.cern.ch/qa/3197`.
- "Print with Pine on Lxplus - mailprint: command not found" at: `http://consult.cern.ch/qa/3281`.

Miscellaneous

Question [3334] – Computer accounts for ATLAS members

I am working at an external institute (outside CERN) but I am a member of the ATLAS collaboration. I need computer accounts and login at CERN. What shall I do?

Answer

All people who wish to use an ATLAS account (in group ZP) must be officially registered as ATLAS user in the CERN Human Resource database. Then, the first thing to do is to fill-in the Atlas registration web form at:

`http://cern.ch/Atlas/registration_form.html`

and to send it to the ATLAS secretariat:

```
Atlas SECRETARIAT  74279 71259 163293  EP      40   4-D06
```

After that the ATLAS secretariat (or ZP group administrators) can register you and create accounts for you on any CERN central computer.

Related link:

List of group administrators for ATLAS (group ZP): enter the group-code ZP in the "Search for Administrators" web page (`http://consult.cern.ch/xwho/people/form/admins`).

About the author(s):

Editing and revising the "Question and Answers" database is a daily task of the members of the User Assistance Team in the User Services group.

4.2 User Services Book Catalogue

Jutta Megies and Roger Woolnough, IT/User Services

Situated in Building 513 1-022 (tel. 74050) is the Computing Bookshop, provided by the *IT / User Services* group, where CERN users can find computer books and CDs at discount prices. The service is open weekdays from 8.30 to 12.30 or contactable by e-mail to `Bookshop@cern.ch`. Books are purchased from some 15 publishing houses and the catalogue offers a selection of documentation aimed at the range of computing utilities available at CERN. The service welcomes suggestions from the user community for new acquisitions. Purchasing may be done internally via EDH or TID or alternatively cash payments via the CERN Bank.

The list, with all the relevant information, is compiled regularly at URL: `http://cern.ch/consult/books`.

If you need Help

FRONT LINE Services

	Phone	Location	E-mail Address
Computing Help Desk Opening hours: 8:30-17:30 (Monday-Friday)	**78888**		`Helpdesk@cern.ch`
Administrative Applications (EDH, CET, etc.) Opening hours: 8.30-12:30 and 13:30-17.30	**79933**		`Ais.Support@cern.ch`
Central Computer Operators 24/24 hours on call	**75011**	**513/R-049**	`Computer.Operations@cern.ch`
General Network **(Operational Problems)**	**74927**		`Network.Support@cern.ch`

Contacts at CERN

The "Computing Help Desk" is the call center for help on IT services related issues. Following is a non-exhaustive list of services that are provided by the IT divison, and the people responsible for it.

Please note:

- In the Web version of this page the "Service Definition" can be a link to the Welcome page of this service: we invite you to read carefully all the information given in this page before submitting a question/problem.
- In addition you might find it useful to search in the "Questions&Answers" (`http://cern.ch/consult/qas`) database for possible answers to your query.
- Many generic e-mail addresses for support lines are just an automatic re-direction to the Computing Helpdesk (this is the case, for instance, for `Mail.Support` or `Www.Support`). Those addresses have been kept mainly for backward compatibility, and also to help the Helpdesk to escalate the query towards the right specialist whenever this is needed.

USER SERVICES

Service Definition	Name	E-mail Address
User Relations and Service Manager	L.Pregernig	`Ludwig.Pregernig@cern.ch`
Help Desk Managers	R.Woolnough	`User.Relations@cern.ch`
	N. Crémel	`User.Relations@cern.ch`
User Registration		`User.Registration@cern.ch`
Computing Bookshop	J.Megies	`Bookshop@cern.ch`

CENTRAL SERVICES

Service Definition	Name	E-mail Address
Distributed File Services (incl. AFS)		`Afs.Support@cern.ch`
CASTOR and Tape support		`Castor.Support@cern.ch`
Central Data Recording (CDR)		`Cdr.Support@cern.ch`
Print Service	J.L.Vosdey	`Printer.Support@cern.ch`
NAP and Parallel Applications	E.Mcintosh	`Eric.Mcintosh@cern.ch`

IT-Experiment Coordination

Information about IT/experiments/accelerator coordination is at the URL:
`http://cern.ch/it-exp-coordination`

DESKTOP COMPUTING

Service Definition	**Name**	**E-mail Address**
NICE (CERN Windows Services)		`Nice.Support@cern.ch`
PC Desktop (Repairs and Sales)	C.Ball	`PC.Admin@cern.ch`
MACintosh Support	F. Ovett	`Mac.Support@cern.ch`
MACintosh Sales	W.Hug	`Waltraud.Hug@cern.ch`
X Terminals Support		`Xterminal.Support@cern.ch`
ASIS Support		`Asis.Support@cern.ch`

Operating System Support:

Service Definition	**Name**	**E-mail Address**
HP 700 s/w		`Hp.Support@cern.ch`
SUN s/w		`Sun.Support@cern.ch`
Solaris s/w		`Solaris.Support@cern.ch`
RS/6000 s/w		`Aix.Support@cern.ch`
Digital Unix s/w		`Dunix.Support@cern.ch`
Linux		`Linux.Support@cern.ch`

SCIENTIFIC APPLICATIONS & SOFTWARE ENGINEERING

Service Definition	**Name**	**E-mail Address**
CERN Program Library	L.Besson	`Heplib.Support@cern.ch`
PAW Support	O.Couet	`Paw.Support@cern.ch`
Anaphe		`Heplib.Support@cern.ch`
Oracle Databases	N.Segura	`Oracle.Support@cern.ch`
Objectivity DB Service		`Objectivity.Support@cern.ch`
Software Development Tools		`SDT.Support@cern.ch`
Adobe FrameMaker		`Docsys.Support@cern.ch`
IT Controls		`ITControls.Support@cern.ch`

COMPUTING FOR ENGINEERING

These activities are now in the Product Support (PS) group.
Index of tools: `http://cern.ch/product-support/support_index.html`

COMMUNICATIONS AND NETWORKS

See the "Communications Systems Group" web pages, at URL `http://cern.ch/it-div-cs/public/help/` and `http://cern.ch/it-div-cs/`

INTERNET SERVICES

Service Definition	**Name**	**E-mail Address**
Web Services		`Www.Support@cern.ch`
Mail Support		`Mail.Support@cern.ch`
MMM (Exchange 2000)		`Mail.Support@cern.ch`

In addition to this generic address, the normal support lines are listed at URL:
`http://cern.ch/mailservices/Support/`

COMPUTER SECURITY

Computer Security Information	`Computer.Security@cern.ch`
Security Alert Reports	`cert@cern.ch`

Other Newsletters

Administrative Information Services: CERN/AS AIS Newsletter (`http://ais.cern.ch/newsletter`)
Contact: `AIS.Newsletter@cern.ch`

Some Useful Web Pages

IT Home Page	`http://cern.ch/it-div`
TvScreen (Systems Status)	`http://tvscreen.cern.ch`
Network Connection Request Form	`http://network.cern.ch/register`
Remote/Mobile Computing	`http://cern.ch/it-div-cs/public/services/mobile/`
User-Oriented Catalogue	`http://cern.ch/consult`
User-Oriented Service Directory (*)	`http://cern.ch/services`
FrameMaker Support	`http://cern.ch/docsys/framemaker/`
LaTeX at CERN	`http://wwwinfo.cern.ch/asdoc/textproc.html`
NICE (CERN Windows Services)	`http://cern.ch/winservices/`
PC Hardware	`http://pcwww.cern.ch/`
CASTOR Query Report	`http://wwwinfo.cern.ch/htbin/pdp/castor/castor_report.pl`
GEANT4	`http://cern.ch/geant4/`
MAD	`http://cern.ch/mad/`
Software Development Tools	`http://cern.ch/sdt`
Joint Controls Project (LHC Experiments)	`http://itcowww.cern.ch/ITCO/jcop/`
PVSS Service	`http://cern.ch/itcobe/Services/Pvss/`
Mail Services	`http://cern.ch/mailservices/`
MMM Services (Exchange)	`http://cern.ch/MMM/`
Web Services	`http://cern.ch/WebServices`
Computer Security Information	`http://cern.ch/security`

(*) Useful service entry points at this page include the Computing Helpdesk, User Registration, Computing Bookshop, Mail, Netscape, Listbox, etc.

Seminars and Presentations

Designation	**Usual Location**	**Organizer(s)**	**E-mail Address**
Computing Colloquia	Auditorium / bldg 500 see `http://cern.ch/computing-colloquia`	J.Shiers, IT	`Jamie.Shiers@cern.ch`
Computing Seminars	IT Auditorium 31/3-005 see `http://wwwinfo.cern.ch/seminars/`	J. Blake, IT	`computing.seminar@cern.ch`
Technical Presentations	IT Auditorium 31/3-005	S. Cannon, IT	`stanley.cannon@cern.ch`

Computer Resources Allocation and Divisional Contacts

Computer Resources Allocation Committee (COCOTIME):

Chairman:	H. Hoffmann/DG	60/5-006	75458	`Hans.Falk.Hoffmann@cern.ch`
Secretary:	A. Morsch/EP	160/R-018	78617	`Andreas.Morsch@cern.ch`

Divisional Representatives for Computing (Desktop Forum)

Contact person for your division for computing issues (resource allocation, divisional servers, licenses matters).

The list of contact persons is at the URL: `http://cern.ch/mgt-desktop-forum/Members.htm`.

Divisional Representatives for Networking

For questions about new installations, testing, local coordination, etc., please ask the contact person for your division (or, if you are an experimental physicist, the networking contact person within your experimental group).

See list of Divisional contacts at URL: `http://cern.ch/it-div-cs/public/di/divreps.html`

Editorial Information

Editorial Board

General Information

The CNL Web home page is at URL: `http://cern.ch/cnl`

The policy is to have between two and four CNL editions per year.

Issue number	Deadline
CERN-CNL-2003-003	Friday 7th November 2003

Information for contributors

The CNL is processed using XML technology and templates. Articles marked up in HTML are preferred, but plain text is also acceptable. A Web form to be used for article submission (contributions) is provided at the CNL Home page (`http://cern.ch/cnl`) as a link “Submitting an Article”.

Distribution

- An electronic mail is sent to the “CNL Announcement Distribution” list (`http://cern.ch/cnl/announcement`) as soon as a new issue of the CNL is available on the Web.
- A paper version is sent to the users who have subscribed to the “CNL Paper Version Registration” list (`http://cern.ch/cnl/register`).

You can cancel a subscription by filling in the appropriate form:

- cancel your subscription to the CNL announcement distribution list (`http://cern.ch/cnl/cancel_announcement`);
- cancel your subscription to receive the paper version (`http://cern.ch/cnl/cancel_register`)

The compressed *PostScript file* containing the complete printable version of this CNL is available via the Web at the URL `http://cern.ch/ref/cnl-2003-002/main.ps.gz`. It is possible that, in order to optimize typesetting, the PostScript printable version differs slightly from the master HTML version. In any case the HTML version on the Web must be considered as the reference.