

**EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH****CERN-HEPIX-2000-004****November 6th, 2000****Last rev: November 14th, 2000****CERN report from fall 2000  
HEPiX/HEPNT****[Maria Dimou](#) - [CERN/](#) [User Support](#)**

The [HEPiX/HEPNT meeting](#) took place in Jefferson Lab (JLab), Newport News, Virginia during the week of Oct. 30, 2000. Windows2000/HEPNT issues covered the first two days, HEPiX the last three. Approximately 50 people were present from America (BNL, FNAL, JLab, LBNL, SLAC, TRIUMF, University of Wisconsin) and Europe (CEA (Saclay), CERN, Czech republic, DESY, GSI, IN2P3, LAL (Orsay), NIKHEF, Oxford University, RAL). The CERN participants were German Cancio, Maria Dimou, Miguel Marquina, Harry Renshall, Alberto Pace, Alan Silverman and Tim Smith. The full list of [participants](#), the exact [agenda](#) and the talks' material are available from the [relevant JLab site](#).

This report contains my notes from the talks and will be [presented](#) together with more details, on the HEPiX part by [German](#) and on the HEPNT/Windows2000 part by [Alberto](#), in the post-C5 session of November 17th, 2000. Specific issues, like large cluster management ([Alan](#)), afs support and service monitoring systems elsewhere will be commented by the other CERN participants in the meeting. Notes from CERN talks are omitted from this report as the CERN users are regularly informed via periodical presentations. Post-meeting reports produced by other participants will be linked as soon as they are communicated to me, e.g. [Harry's report to PDP](#).

Readers may follow the links, below, to reach the items they are interested most:

[Win2000](#) [WindowsNT](#) [Linux](#) [Netscape](#) [Monitoring](#) [Batch scheduling](#)  
[condor](#) [GRID](#) [afs](#) [LSF](#) [Security](#) [Large cluster](#) [SIG](#) [Curios](#)

**Windows2000**

Christian Caro (previously known as Christian Trachimow) explained the aim of the "Windows2000 Coordination Group" as:

*"Easy and secure access to HEP resources, ideally by the use of one login account across labs."*

Single login is a problem in Windows2000 as we should be restricted to only one schema.

Issues that need coordination include:

- LDAP
- dfs
- kerberos
- HEP applications.

This is a closed coordination group formed by the HTASC sub-committee of HEPCCC. F. Hemmer is the CERN representative.

Other speakers described their experience with:

The Windows2000 **Remote Installation Service (RIS)**: part of the IntelliMirror, useful as it offers action log and does not require a PC reboot but only efficient for simple user installations (DESY and LAL/CNRS). Centrally controlled OS installations are faster by cloning disks (RAL), an accepted Microsoft policy.

RAL chose **DeltaDeploy** for application installation, a 3rd party solution costing less than 10\$ per workstation (too expensive for our number of PCs).

**Group Policies**, originally introduced with NT4 to apply different security policies to various groups of users: still very much used in Windows2000, very difficult to use, except for small scale environments.

DESY complained that the **Installer (MSI)** doesn't take care of the application's life cycle, i.e. doesn't handle well re-installation evolution.

**Overall, the other sites don't have many PCs running Windows2000 so far (the maximum amount of nodes mentioned were 150 in all of INFN sites) and they haven't yet concluded on several migration issues.**

## WindowsNT to Windows2000

The second day (October 31st) was dedicated to WindowsNT operation and migration issues.

LAL runs 130 NT machines with 300 registered users. Two machines are migrated to Active Directory. The users' home directories are on Unix filesystems and accessed with **Samba**. The print server runs on Unix as well. When using SMS, Exceed7 (and other applications) installation doesn't work as an upgrade, one has to de-install Exceed6.

Problems with **SMS** were reported from other sites as well. SMS v.2 behaves better than SMS v.1.2. This was reported by B.Cowles (SLAC) where 50% of their 1600 NT workstations are managed with SMS.

Dual boot is discouraged at SLAC. PC users are recommended to run Linux and to use Windows applications via **Citrix**.

J. Surget (Saclay) reported that they chose the english version of their Windows2000 server configuration but the users customise their own menus using the **Multilanguage User Interface (MUI)**. This leads to an english-french mix-up after IE5.5 installation.

H. Kreiser (GSI) presented the results of a cost per year estimation that a consulting company did for PCs, printers and LAN equipment. The formula is **cost/year = hardware\_price/lifetime** and it can be as high as 9K\$/year.

## Linux

This seems to be the strategic platform for HEP applications in all sites, although the others didn't mention plans for phasing out other Unix platforms in a time-scale as concrete as ours.

For example, INFN counted 18% of their institutes' boxes run Linux but they have no plan at this moment for a global Linux migration. SLAC currently has 1200 Solaris systems as opposed to very few (89) Linux boxes. However, BNL chose to re-install existing Solaris PCs with Linux. RAL uses **kickstart** to automate installation across multiple Linux PCs. Several sites decided to keep the Unix user homes and make them available on the desktop Windows systems with **Samba**.

## Netscape

As CERN's strategic desktop solution is Windows2000, the relevant services hesitate to continue with Netscape as the recommended browser, mail reader and shared calendar because:

- Internet Explorer (IE) comes up by default during installation
- the MS Installer requires manual configuration to make netscape 4.7x available
- netscape6 (Mozilla) promises to have all problems solved but is not considered mature yet
- the Netscape (company) support contracts are no more valid

As far as the mail is concerned, other sites have fewer users, therefore less material for reliable statistics on Outlook failures. FNAL, LAL, SLAC, JLab and most others said that once imap is used to access the mail on the server, the users are free to use any mail reader they please. The JLab security manager Bob Lukens sees problems with Outlook or Express.

SLAC installs or gives server space to users to install their preferred mailers. FNAL promotes pine and netscape for Unix users and moving from Outlook Express to Outlook (including Calendar) for the Windows ones.

RAL uses Exchange and would be happy for its users to have it available

when they visit CERN as well.

LAL chose **silkymail**, a free mail client written in php, running on apache, LDAP compatible. More info from <http://www.cyrusoft.com/>. INFN reported that they use [imho](#), public domain software to download mail from the web.

The conclusion on the Netscape vs IE/Outlook discussion was to **make both available for a sufficient amount of time and inform the users about the strategic direction.**

## Monitoring

Although the HEPiX meetings take place twice a year there is a surprising number of independent projects in many sites developing software packages that monitor systems/services. We followed presentations on:

- NGOP (FNAL) now in prototype development status,
- Ranger (SLAC) successor of 'nicemon' and 'Patrol', now operational as a Perl toolkit,
- PEM (CERN) prototype planned for next January, faced the need to add in-house developed functionality to JDMK
- and the Global Monitoring system (IN2P3), future Java package, in very early stage, not yet prototyped.

## Batch scheduling

Here again many interesting but unrelated approaches were described like:

- [CODINE](#), installed at DESY/Zeuthen, developed by the company Gridware, now bought by Sun, Linux port imminent.
- [Portable Batch System \(PBS\)](#), used in JLab on alpha and intel nodes, a good free product that requires a server node.
- BQS and its graphical interface Jbqs are used in CCIN2P3.
- Farms Batch System Next Generation (FBSNG) written in Python, a successor to FBS is used in FNAL and no more runs on top of LSF, as did FBS.

## Condor

The [Condor](#)'s purpose is to use idle CPU power. [Peter Couvares](#), now working in the University of Wisconsin, explained the mechanism of moving data to location while other data are being computed. The development team he belongs to, is composed by 5 staff and 20 graduate and PhD students. They collaborate with other research projects such as EMERGE, a high-speed QoS-enabled data network, the GRID Physics Network (GriPhyN) and the Particle Physics Data Grid (PPDG). INFN use Condor for the last 2 years in a pool of 200 machines and reported good results in resource exploitation. Saclay is now testing Globus and Condor.

## GRID

Everyone seems to be have a role in dataGRID or other GRID-related projects. More than 20 FTE people are involved from INFN. Saclay participates as well with 6 Alice PCs. New projects are 'GRID-aware' like Computer Fabric Management (WP4) in the area of distributed computing and JASMine, the mass storage product by JLab.

David Kelsey (RAL) spoke about his role in the Testbed (WP6) in recommending X.509 authentication certificates for use by GLOBUS. The certification authority has to be national. It is not yet decided whether the certificates will only authenticate or, in addition, authorise users. David's view is against authentication and authorisation combined. A meeting will be held in December at CERN to discuss this.

## Afs

[Gary Gerchak](#), the IBM Marketing Manager based in Austin, Texas is responsible for the afs support since IBM acquired Transarc in July 1999. Country support lines should be contacted first. IBM is re-writing the support contracts to adapt previous Transarc terms and conditions to standard IBM ones. The official versions available are:

Afs v.3.6 since March 2000, running on Linux Redhat 6.0 and promised for Solaris 8 and Windows2000 at the end of this year. They don't perform the official Microsoft certification procedure. They classify the product as "Windows tolerant". What they call "end of service" was announced for Year End 2002. Such end dates are normal business practice for IBM and they assured us it does not mean end of support but we regard it as an important warning signal.

DCE v.3.1 since November 1999, running on Solaris7 or higher and AIX 4.3.x. User data under DCE will be (optionally) put in LDAP in the 1st quarter 2001. DFS v.3.1 is available now. No details were mentioned on this.

Afs development continues in the IBM/Transarc lab and in IBM India. A board is now being formed to review OpenAfs code and possibly incorporate it in the IBM tree. IBM, CMU, MIT, Univ. of Michigan.

## LSF

Rich Hall, from the company Platform computing gave a sales talk on the future of LSF. He announced that LSF 4.1 will run in parallel on Linux and that LSF development is continuing in close collaboration with SGI.

## Security

Americans take computer security very seriously. They are all about to shut down 'telnet' and 'ftp' in favour of 'ssh' and 'scp'. They run special monitors and consult related web sites continuously. JLab security manager Bob Lukens explained that they monitor their servers (30 hosts) every hour and a selected group of 350 systems (out of 1800 in DNS) 3 times/day. Certain types of messages call Bob's pager for immediate action. He mentioned information sources like the [Computer Security Technology Center](#), located at the Lawrence Livermore National Laboratory and tools like '**shadow**', a Perl handler of 'tcpdump' output presenting it in HTML, possibly generating alarms in its most recent version. It also offers a web page for searches per day or other time period. He also mentioned home-made scripts in Perl that cut off external systems attempting to connect many times or attempting to connect to more than 10 JLab systems simultaneously.

The SLAC security manager, Bob Cowles, explained that hackers, themselves, publish a list of sites compromised/defaced every day. Common programs like **xlock**, **bsd telnet**, **php logging**, **ftpd**, etc are vulnerable and need very recent patches to become safe. **tcpdump** was recently found to cause a buffer overflow, if attacked. Web servers today allow cross-site scripting with no known fix so far. SLAC has port 80 blocked by default from the outside.

## Large Cluster SIG

Alan Silverman explained that the LHC Regional centres need a single computer environment. As HEPiX is the forum for sharing tools he announced that he visited FNAL (Tier1 centre for CMS) and BNL (Tier1 centre for Atlas) and plans to organise a workshop to adopt common solutions in selected concrete areas:

- OS certification (CERN, FNAL, others?)
- distribution of **jpasswd** (the brilliant common passwd tool across platforms by Marty Wise, JLab) for CLASP use together with relevant FNAL Kerberos5 implementations
- kickstart for large clusters
- instructions for building clusters, e.g.
  - qualify PCs
  - choose hardware configuration
  - packaging on racks
  - evaluate software products' scaling capabilities

The outcome of this SIG (Special Interest Group) workshop will be reported at the next HEPiX.

## Curios

FNAL decentralised computing support and re-designed group structure separating Science Soft from Office Computing.

The OS count across INFN sites gave as many MacOS boxes as the Linux ones (18% of the total).

JLab has many open posts in computing but a great difficulty to find any interested candidates. The CEBAF upgrade will give 2/3rds of CMS data volume. They are seeking advice for thin client X-terminal to replace their NCDs. They use **Myrinet** between 40 Linux alphas for Lattice QCD.

SLAC uses Myrinet too between 16 nodes. IBM announced that most HPSS data are at SLAC. HPSS 4.2 will be tested there for the Solaris port.

IN2P3 being Babar users developed **bbftp**, a parallel, secure ftp used in Babar. They also evolve the rfio code for HPSS and Castor use.

BNL speaker Stratos Efstratiades showed us photos of their multiprocessor computer, whose chips are made by Texas Instruments.

There was a discussion dedicated to LDAP. Although, LDAP will be used in GRID to allocate resources, and the HEP tree structure has to be decided, the Special Interest Group is inactive and no interest is shown by anyone.

## **Finally...**

This was the first time I was present in a complete session of a HEPNT/HEPiX meeting. I found that there is a good team spirit amongst all participants. Our hosts from JLab worked hard to make this event a success.

---