

HEPiX CASPUR Meeting Report - Site Reports

Site Reports presented at the Meeting held October 23rd - 25th, 1996

1. Introduction

o Logistics

The Autumn 1996 HEPiX meeting was held in [CASPUR](#), Rome, from October 23rd to October 25th. It was attended by 37 people from 13 different HEP institutes. Local site arrangements were handled by Andrei Maslennikov and Monica Marinucci and Alan Silverman was responsible for the agenda.

o Introduction by Prof. Romano Bizzarri

Prof Bizzarri opened the meeting by welcoming the attendees to CASPUR. CASPUR, more formally the Inter-University Computing Consortium for Central and Southern Italy, is funded by the Italian education administration and various associated universities. The lab is not exclusively a HEP site but had a strong collaboration with INFN. It had started with association with HEPVM and CASPUR has long been a user of HEP software and has maintained its links with INFN and with CERN. More globally, its mission is to provide CPU power and computing expertise to its users in Italian universities, including of course various HEP departments and groups.

2. Site Reports

o [CASPUR](#) by A. Maslennikov

As noted above, CASPUR's mission is to provide services for its distributed members/users. Central services include a 16 node SP2, an Alpha 4100 with 4 nodes of 4 CPUs each, a Quadrics APE parallel computer and a CMOS-based mainframe offering both VM and MVS services (this latter service is used for example by the Italian National Library Service). There is also a 6 drive STK silo and some 40 workstations of various flavours.

CASPUR has a site licence for AFS and provides AFS services for several AFS cells across Italy. There are 5 data servers, some in Rome, some in Bari, with a total of 100GB of disc space and some 200 client nodes. They provide an ASIS mirror site for Italy which is used by 10 other sites and AFS is also used to distribute some site-licensed commercial software packages.

They plan to implement the SHIFT software for their STK silo, augment the data conversion tools available in the Centre and produce work group tools making use of the arc package from R.Toebicke of CERN. CASPUR is putting in a major investment in LINUX, including some training tools using LINUX.

o INFN by A. Maslennikov

Andrei continued with a report from INFN. This currently has some 30 sections spread across Italy with good network links between them, usually 2Mbit, and an 8Mbit line from Bologna to CERN. The Italian universities make heavy use of this network. There are local clusters at individual sites but the links allow data processing tasks to be concentrated at the larger installations, often close to the source of the data, for example CERN in the case of HEP.

o [CCIN2P3](#) by W.Wocjik

Since the last meeting, the VM service at CCIN2P3 has begun to be downgraded and will be stopped during next summer. This is reflected in the gradual decrease in the user base and the corresponding growth in the UNIX user population.

Their AFS cell now consists of 7 AIX servers and 350GB of disc space. There are also 4 AIX tape servers with 400GB of RAID disc spool attached. All the major clusters and servers are interlinked with ATM and a new experimental DLT service has begun to be offered on 2 HP 735 nodes; it consists of 2 model 2000 drives and 2 model 4000. Their local version of CERNLIB has been modified for specific tape access but is still RFIO-compatible.

WDSF for file backup has been replaced by ADSM and a graphical user interface for archiving (ELLIOT) has been written; it will be covered later in the proceedings. There have been improvements to BQS and PIAF has been implemented on their SP2 using IBM's parallel file access method instead of RFIO. CCIN2P3 is now using AFS mail via Pine and Netscape; they are also examining IMAP. Other CCIN2P3 topics will be covered later.

o [CERN](#) by A.Silverman

The inexorable rise in the UNIX workstation and server population at CERN continues but the number of new X terminals shows signs of tailing off, counter-balanced by a dramatically increased rate of purchase of PCs. The major event over the summer at CERN was the switch-off of CERNVM and a lot of energy was spent by the UNIX support groups absorbing much of the CERNVM user base, especially with regard to mail and printing services. Mail in particular has been a delicate issue and a lot of effort has gone into various forms of user education. Currently the mail support team is concerned especially about how their twin-node failsafe solution can cope with the total load as well as how it might scale.

The AFS cell is also steadily expanding but AFS 3.4 has given many problems and CERN has a close collaboration with Transarc to resolve these as they occur. One result of this however has been a postponement of any real advance work on DFS.

There are many PC initiatives at CERN and although the main administration and occasional users are switching to a Windows 95 version of CERN's NICE architecture, there are also many independent Windows NT projects. A few people have expressed interest in LINUX and at least one experiment uses it seriously but there is (so far) no official support for it.

On the network side, the move to structured cabling and routed networks is progressing well and the result is a much more stable network. The CORE batch services are also expanding to take care of user demand and the first HP PA 8000 chip systems were recently successfully installed.

Other things to note include

- The Gnats problem reporting system is in full production use by a number of groups
- The ADSM file backup service has switched to 2 RS/6000 hosts
- The formal closure has been announced of several services, namely DXCERN, ULTRIX and SunOS 4 support and support for SGI desktops.

- The public UNIX services have been augmented by a new service (DXPLUS) based on Digital UNIX.
- CERN's SUE has been ported to Digital-UNIX and the first releases of [certified operating systems](#) (see Vancouver HEPiX) have been announced.

○ **Fermilab by Lisa Giacchetti**

The migration off Fermilab's VMS system is almost complete, there was a big push in June on e-mail in particular and many were helped individually to change to another scheme. The next step will be the removal of the FNALV cluster name, making it a sendmail alias instead. There are a few remaining VMS applications.

The current major activity is the forthcoming fixed target run which will use UNIX as the online system rather than VMS as in the past. Fermilab are also testing PC farms, initially with LINUX then moving to NT. Fermilab's PC group is migrating to NT and are investigating Wincenter. For the future, working groups have been set up to discuss how to deal with the so-called RUN II cycle of experiments.

Fermilab batch is migrating from Loadleveler to LSF and this will be used by both CDF and D0 as well as on the central UNIX farm, which itself is being upgraded with 45 more SGI and 22 more RS/6000 nodes and both CDF and D0 have expanded their batch capacity in disc and/or CPU power. PIAF has been installed on SGI Challenges. On the interactive side, FNALU will receive more AFS space. The Sloan Digital Sky Survey, supported by Fermilab, has added 2 Alpha 8200 systems.

○ **LAL by M.Jouvin**

The LAL computer centre supports two major in-house activities, CAD and physics, the latter partly in support of several CERN and DESY experiments. The physics users have a park of 8 Alphas, 4 HPs, a small VMS cluster, 150 X terminals, 150 MACs and a growing population of PCs. LSF runs on the batch systems. There is a file server configured in failsafe mode with 80GB of RAID. The servers sit on FDDI rings interconnected by Gigaswitch. There is an AFS client in the CERN cell but the local file scheme is based on NFS. LAL would be interested to move to DFS but they would prefer to wait until CERN decides its future. Plans for 1997 include a memory channel cluster and the first appearance of ATM 622Mb as a backbone.

They have a mail server based on IMAP (see later) and they are setting up a Wincenter service following the presentations at HEPiX at TRIUMF.

○ **Nikhef by W.van Leeuwen**

Since the last NIKHEF report, both halves of the lab (Nuclear and High Energy Physics) have been amalgamated, including the computer support teams. They have moved to a new building equipped with a new UTP-based network. The work generated by PCs, especially laptops, is growing, running either Windows 95 or NT, but the SUN platform remains dominant along with a significant number of HPs and SGIs. The number of Apollos at NIKHEF is starting to decrease.

For public interactive services, there are two multi-node SUN servers with the home directories spread across several file servers. NIKHEF is participating in the Internet 96 World Exposition.

○ **DESY by K.Kuene**

A story of continuing gradual expansion, in AFS space (to 100GB of home directory space), and in Work Group Servers (68 nodes, 1100 users). Most of the WGS nodes are SUNs. The Computer Centre support group is trying a scheme whereby they install workstations for users ready-to-use; this is now operational for SUN/Solaris and is making progress for AIX, HP and IRIX systems. Similar methods are being developed for ongoing system updates.

Central services have been established for - a mail service based on IMAP but with some POP support; a central registry service based on DCE (DESY's first steps towards a DCE/DFS cell), see later talk; and a rudimentary equipment database for all workstations at DESY.

At the high performance end, the 10 SGI Challenges installed have a total of 186 CPUs of which 32 are R10000s which therefore require IRIX 6. There are 2TB of disc space and the OSM (Lachman) mass storage scheme is used; access to this is possible from all DESY platforms. A new Grau-Abba tape robot is on order for the Hamburg campus with 2100 cartridge slots and 6 Ampex D2 drives. A smaller similar robot is installed at Zeuthen. RFIO is being adapted to trap calls for remote disc access and transfer this to an OSM call.

DESY is starting seriously to investigate PC farms. The Hermes experiment has 10 Pentium Pro systems running LINUX which has been in production for just a few weeks. On the other hand, ZEUS has just ordered 20 similar PCs but will run Windows NT.

The current problem areas at DESY include --

- the diversity of platforms to support, a historical hangover
- AFS bugs and deficiencies and the need for user training
- local WGS disc space management
- the future of UNIX "mainframes"
- what is the future of DFS? What are SGI's plans for DFS?

Current work includes investigating the use of OSM as a data server for medium-sized files and making improvements into the mass storage system.

○ **RAL by J.Gordon**

The computer departments of the Rutherford and Daresbury Laboratories have been merged which brings in a strong theoretical computer science faculty into the RAL sphere, including an SP2 and some SGI systems.

At RAL itself, the CSF farm is now 30 nodes including 4 new HP C Series systems which may later be upgraded with PA8000 CPUs. The service delivers over 5000 CPU hours per week and is not only used for Monte Carlo work but increasingly for physics analysis. It is much accessed from across the UK via SuperJANET.

On the Digital side, there is a 6 node Alphaserver 8400 with 2GB of memory running Digital-UNIX and there is still a VMS-based mainframe which is now becoming expensive to maintain so discussions have started with a target to stop this in October 1997.

RAL are planning a trial NT farm service and this will be discussed at the end of this meeting. Much work has been put into tape service reliability and the efficiency of data throughput. An extra 100 GB of disc space has been added to permit holding more data online and RAL is experimenting with SHIFT software for staging. They are also evaluating various options for file migration and there is the possibility of an update to the ATLAS Datastore (see previous HEPiX, for example [Prague](#)).

A new challenge at RAL may be a plan by the H1 experiment from DESY to install an analysis farm at RAL which would have high output rates as well as the more traditional high input only.

- **GSI by J.Heilmann**

GSI has some 120 UNIX systems installed. mostly used for physics work; 200 VMS nodes used for accelerator control and data acquisition; and about 300 PCs, mostly for administration tasks. Most of the UNIX boxes are IBM RS/6000s running AIX and there is a 16 node SP2 of which 4 nodes offer an interactive login service. There is no AFS on site, only NFS with the amd automounter. They find this slow but cannot yet decide on an alternative.

There is an IBM 3494 cartridge robot with 6 drives and 12TB capacity. ADSM is in use for backup along with an in-house archive/retrieval tool which has been ported to many platforms. They are in the process of migrating their file base (5TB) from MVS to AIX.

They have defined a recommended UNIX desktop environment consisting of fvwmm, gnu utilities and Z-mail; they may move to CDE when they upgrade to AIX 4. There are plans for an NT service and they find their Wincenter service rather popular; it runs on 2 dedicated PC servers with a total of 50 client licences and makes use of Samba to access UNIX files.

They have PIAF running on their SP2 but it uses NFS for disc access and gives poor performance so they are looking at the parallel file system on the SP2 to speed this up. Other plans include a site-wide migration to AIX 4.1 or 4.2; upgrading their login service; and acquiring more nodes for their SP2.

- **CEBAF by Sandy Philpott**

The following report was submitted after the meeting by e-mail.

First of all, we're now The Thomas Jefferson National Accelerator Facility, or Jefferson Lab for short, and soon jlab.org on the Internet. This is our first year of production physics research after completing our construction phase, thus the name change to denote our status change. Of our three experimental halls, we have achieved beam to two (Halls A and C), with the third (Hall B) expected near the first of December.

We are getting ready to use our new Internet domain -- what a project! We've found that our community is concerned about ".org" status, the Department of Energy site office wonders about moving out of the ".gov" domain, and we're scurrying in the Computer Center to make it all work. (Right now jlab.org is simply an alias for cebaf.gov.)

We have implemented our Common UNIX Environment (CUE) across the UNIX platforms. We are implementing it for every single JLAB system managed by the Computer Center -- to the desktop, central computing JLAB systems, IFARM interactive compute farm, FARM batch compute systems, DAQ data acquisition machines, and SWC specialized workgroup clusters (minimum number of special purpose systems). This gives users an identical environment no matter which platform or class of machines they're on. We are still considering providing elements of the HEPiX environment.

We've purchased a Network Appliance Model 540 NFS fileserver to provide initial file storage and will probably go ahead with the purchase of a second while we investigate high-end RAID solutions. It looks like fibre channel arbitrated loop is 6 months away, so we're planning to delay a little to find what's available in the market then. Our testbed DCE cell is up and running, with cross cell communication between two offsite DCE cells. This is promising as we look to providing wide-area DFS network access to our data files.

In the batch farming arena, we are awaiting the approval of our Jefferson Lab Off-Line Batch System (JOBS) requirements document. The initial implementation of JOBS is scheduled for 3/97. (To refresh your memory, our system uses LSF to manage the farm, and OSM to manage the data.)

We are well into the process of implementing a Fast Ethernet network for both our general backbone LAN as well as for the Experimental Physics segment. We are still planning to dual-attach fibre channel RAID to move the Hall B data directly to our StorageTek silo -- remember it's approximately 10 times the amount of data from Halls A or C. ATM is still under investigation; no production use.

To support this complex new physics environment, we're searching for a Helpdesk Administrator.

3. **OCS by Lauri Loebel Carpenter/FNAL**

OCS stands for Operator Communications Software and it is a UNIX tool to provide tape volume handing across all platforms. FNAL uses the package for tape drive allocation and operator-assisted tape mounts; it also provides logging and statistics. Eleven different project teams use it plus a number of outside sites.

The interface is by means of library calls, shell commands or via a MOTIF interface; the last of these is used exclusively by the operators because it has many features designed specifically for them, especially in the area of problem handling.

A system administrator can define access rights to particular drives by group; a user can select criteria to select drives, to queue a request or not, read-only or read-write access, etc. The speaker showed a number of examples of OCS commands. followed by the structure of the various programs and the command message flows.

Plans for the next version, version 3, include linking it to FTT (see later talk) underneath OCS and beyond that to add robot support and to integrate the currently-separate OCS databases.

4. **FMSS by Lisa Giacchetti/FNAL**

The Fermi Mass Storage System is an interface to a mass storage system based today on an IBM 3494 robot and used with HPSS software (FNAL is a member of the HPSS Consortium). It consists of a set of tools to centrally manage user data files. It offers transparent access to data without the user having to know where the data resides.. It also has special provisions for archiving small files and handling large files in a "bulk area". Group quotas are supported.

Nodes must be registered to use FMSS and it is most-commonly found on central systems. Use can also be restricted to members of a given user table.

The speaker showed a list of the various FMSS calls such as copy, remove, query, etc. The authors had maintained traditional UNIX formats for the commands where possible.

The product was not at that time yet in production status, people were still migrating off the STK robot and Unitree scheme. It was hoped to complete beta testing and begin production use within 6 months.

Future versions should permit direct import/export to tape. More staging options would be added.

5. **Status of the CERN Magnetic Tape Robot Acquisition by Harry Renshall**

A group of CERN users and internal and external tape specialists combined to define CERN's future needs for tape automation and a

market survey based on this was issued late in 1995. As a result of the replies, it was decided to wait for a time, partly in order for more experience to be gained in emerging technologies and partly to allow CERN to bring into full production use its then-current robotics. It was clear that much more information should be learned from other sites. and there were many outside contacts as well as onsite tests.

The objective has been defined as - no more manual tape mounts outside CERN working hours by the end of 1997 and a fully automatic scheme by the time LHC begins. Users should be encouraged to make more use of robotics and in the meantime much effort goes into tuning the tape library such that the busiest tapes are placed in the robots. There is also close cooperation with robotics vendors to improve their hardware and software.

As a result of all these actions, the automatic mount rate does show signs of increasing at the expense of manual mounts and this has permitted the desired reduction of overnight operational support, thus promoting savings for investment in the next stage of automation.

Renshall presented some details of the units under consideration, including performance figures. Tests performed at CERN in general supported the vendors' claims. All tapes gave satisfactory results and differences were in the area of purchase and/or operational costs. For the acquisition exercise currently underway, specific criteria were established such as mounts per hour, the need to demonstrate working sites with the equipment to be offered, the cost of ownership over several years, and so on.

At the time of the meeting, a Call for Tender had been issued and it was hoped to take the result for adjudication to the December Finance Committee meeting and, if confirmed, to have the equipment installed for the 1997 LEP/SPS run.

6. **ZARAH Tape File System by O. Manczak/DESY**

Zarah is the data storage scheme for the ZEUS experiment at DESY. The data is write-once, read-many and it is aimed to store eventually 30TB. The ZEUS computing centre includes 2 SGI Challenge XL systems, tape robotics with 12TB capacity today and 1.2TB of disc space. 200 client workstations access this from on- and off-site and there are more than 250 users per week. Zarah provides a user interface to the hierarchical storage and offers overall a high performance, large capacity file service. The interface permits to hide the file stage process from the user as well as offering distributed access and scalability.

Zarah is implemented via a fake NFS demon using RPC between clients and the servers. On the server side, there is a stub file system which is identical in structure to the real file system on the mass storage and a user's file tree structure is thus mapped.

The speaker explained in detail the mechanism of staging a file. He stated that NFS was very slow and they used instead the RFIO protocol which offered other advantages over NFS as well as better speeds. They used a local implementation of RFIO and he explained how this affected and improved access to files. Speeds achieved had reached 40MB/second from an SGI server via HiPPI.

7. **FTT by Lauri Loebel Carpenter/FNAL**

FTT, Fermi Tape Tools, is a consolidation of various system-dependent tape I/O codes. The package offers support for various operating systems and tape types and the group responsible for it are always ready to add modules for more platforms or tapes if requested.

FTT provides the basis for other data access packages such as OCS (see earlier), where FTT gives the basic tape access routines. It contains modules of library calls, a setuid binary and some test code. Use of the package provides system- and device-independence, statistics gathering and consistent error returns. It is easy to port to new platforms but there is a need to simplify more the procedures for building the required tables of tape characteristics and operating system features and to automate the necessary checking procedures. Another open point is a mechanism to determine which drives are available on a system, made more difficult since each architecture has its own commands. Work is also needed to improve the Porting Guide and it is then planned to include FTT in the Fermi Tools package for public distribution.

8. **A Roundtable Discussion on Tape Interchange Issues led by John Gordon/RAL**

The speaker first summarised the conclusions reached at the [TRIUMF](http://www.cern.ch/hepix/meetings/triumf/minutes.html) meeting of HEPiX (see URL <http://www.cern.ch/hepix/meetings/triumf/minutes.html>). The proposal to standardise on SL labels was already been seen more and more, for example at IN2P3, RAL and DESY/H1; and the maximum file size of 1GB had been set already for example at the ATLAS LHC experiment. However, in Vancouver, the meeting had been unable to arrive at consensus on a recommended transport medium and J.Gordon wished to return to this discussion; had an extra 6 months experience with various media helped us come to a conclusion?

DLT 2000 units are still very popular, can this be the default? Or should we wait for DLT 4000 to become more widespread? Are the various DLT models (7000, 4000, 2000) at least downwards compatible in both read (claimed) and write (less certain)?

H.Renshall stated the CERN position - they are user-driven and try to guarantee support for (almost) any reasonable media requested by CERN users. However, he agreed with the suggestion that if there was a HEP-wide recommendation, it might help to reduce diversity.

It was reported by one delegate that Quantum had stopped manufacture of DLT 2000 models.

After a long discussion, it was agreed to declare a recommended default for any site who sought advice while of course allowing any site who wished to support additional types. On this basis the default should be DLT 2000 media format although any new drive should be DLT 4000 or DLT 7000 (by this time the meeting was convinced that both these units could write DLT 2000 formats). This recommendation would not only provide the lowest common denominator for data interchange between the labs but also offer small sites the possibility to use the same drive (e.g. the DLT 4000) with sufficient local capacity for file backup.

9. **DANTE: Data Analysis and MC production cluster for Aleph by Silvia Arezzini and Maurizio Davini/INFN Pisa**

DANTE is a cluster made up of 7 HP 735 nodes and an HP D Series 350 twin-CPU server. There is also an HP Magneto-optical robot and a DLT robot. DANTE is a member of the Pisa AFS cell (all users have AFS home directories) and in fact it hosts the Pisa AFS cell. The tape cartridge robots run under the control of HP Openview Omnistorage. This provides a form of HSM (hierarchical storage management), it supports various magnetic media types and makes the storage areas appear as a single large file system. It is controlled via a Motif GUI.

The cluster is used for ALEPH Monte Carlo and data analysis work. The standard ALEPH software modules are mirrored nightly from CERN and the results of the jobs on DANTE can be stored locally or written back via the network to SHIFT at CERN. One node of the cluster is used for interactive work and the others for NQS batch queues, the CERN flavour of NQS with the gettoken scheme for AFS tokens. The HEPiX login and X scripts are in use and there is an ASIS mirror.

Among the problems seen have been occasional Ethernet congestion and a network upgrade is planned using HP's AnyLAN plus a major problem with NFS for which a patch was found. In the discussion, it was noted from the audience that NFS under the latest HP-UX releases showed significant performance improvements.

10. Parallel Processing at CERN by Eric McIntosh/CERN

The speaker gave first a review of the concept of parallel programming - there was clearly an opportunity to use many processors in parallel, how are we to split the job and then recombine the results. We must also be careful with the communications overhead.

Parallelism at CERN began with the use of the late Seymour Cray's CDC 6600 and its various functional units. From there we moved to vector processors, to RISC computers, MPP systems, the GP-MIMD2 (a European Commission-funded project) and currently to the Meiko CS-2. This also is EU-funded with 3 industrial partners and 3 user sites including CERN. Meiko also sold two large systems to Lawrence Livermore Laboratory, which had the unfortunate side-effect of diverting vital support staff personal away from our project although this sale did assist in the debugging of the overall system. Recently, Quadrics has taken over Meiko, and its best experts, and more developments of the architecture are planned. So far, Quadrics have given excellent support and Eric also acknowledged here the support and help of Bernd Panzer of CERN in the work around the CS-2.

The CS-2 is a SPARC-based machine with a proprietary internal network rated up to 50 MBps. It has a vector option but this is very expensive and only really useful with long vectors (not our situation). The initial SPARC chips used suffered poor floating point performance and recently Ross 100 MHz HyperSPARC chips have been installed. Memory is physically distributed but logically shared among the 64 nodes of 2 processors each. Each node also has a disc for local swapping. For data space, there is some 100GB of internal disc and 400GB externally. The Meiko parallel file system uses striping but the user sees only improved performance as more nodes are brought into use to access the discs.

The systems are connected to the CERN network by FDDI but HiPPI will be tried with NA48. The operating system is Solaris 2.3 with extensions from Meiko, including an attractive system management tool permitting control of remote nodes. As set up, four nodes are designated control nodes for the whole system. Users' home directories are local, NFS-exported to all nodes in the system. Eight nodes are dedicated to PIAF, 5 for interactive logins and other nodes are used for batch or parallel tasks. AFS access is via the NFS exporter.

There are three programming models available:

- o full shared memory - this is hard to program and hard to debug
- o a Meiko feature known as "libatom" which exploits global memory access (any node can access the memory of any other node); this is not interesting for HEP uses since it makes the code totally non-transportable to other platforms
- o message-passing using MPI, PVM, etc
- o sockets and TCP/IP - the lowest common denominator

Eric gave some examples of applications in operation

- o NA48 central data recording (up to 6 MBps has already been achieved and it is planned to reach 20 MBps) and real-time reconstruction and calibration; some event-parallel Monte Carlo work has taken place
- o Parallel GEANT, used by CMS for detector design; this uses MPI message passing. (Some of this work is also performed on the CERN SP2 system.)
- o PIAF for NA48
- o Various PVM applications for smaller experiments and for LHC design.
- o Some classical batch processing

11. The CERNVM Migration Experience by Harry Renshall/CERN

When it was decided in December 1993 to close the CERNVM system, a special Task Force was established to concentrate and motivate the effort needed to effect this major change. At its peak, this task force involved some 30 people. The largest task was to build up a UNIX environment for physics and engineering users but it also required a specific marketing exercise to really get the message to users in a simple and timely fashion. Key technical areas included mail, editing and migration tools. In each case, various alternatives were studied and recommendations were made. Solutions independent of a particular UNIX flavour were given priority. The final list of [recommended utilities](#) was published (URL <http://consult.cern.ch/umtf/1996/tools4certification>). On the marketing side, a series of user seminars and open user meetings were held and much of the [material presented](#) is freely available on the web at URL <http://wwwcn1.cern.ch/umtf/>. There were also lots of articles in the CERN Computer Newsletter and extra copies of this were printed and distributed.

After a slow start, the so-called User Migration Task Force (UMTF) was re-organised in October 1994 after the HEPiX meeting in Saclay/DAPNIA to include more input from other labs, notably DESY. The result was CUTE (Common UNIX and X Terminal Environment, see for example [HEPiX Rio](#) at URL http://wwwcn.cern.ch/hepix/meetings/rio/Rio_cute.ps). The first incarnation of CUTE was the interactive service opened on CERNSP, CERN's SP2.

Most major batch work on CERNVM was stopped at the end of 1995 and transferred to CORE (see for example [HEPiX Fermilab](#) at URL <http://dcdsv0.fnal.gov:8000/hepix/1094/talks/batch/shift.ps>). At that time, the CERNVM configuration was halved in size. That left an interactive user base of some 11,000 accounts.

The planned stop of CERNVM was originally end-96 but this was advanced by 6 months to end-June '96. Cleanup of the unused interactive accounts started early in 1996 and this removed some one-quarter of the remaining accounts. Gradually the CERNVM load was reduced until it was formally closed at the end of June. Migration of the VMarchive database (300 GB, of which some 105 GB was "active"; this was compressed to 35 GB) and users' minidisks (50 GB) to the RS/6000-based ADSM service was completed by the end of August. Tools were written to permit users to retrieve their own files from these archives but such occurrences have proved very rare indeed. Also by this time the client side of the ADSM backup service had been changed on user nodes to point to the RS/6000 version of this service.

To cope with the expected problems, staff on duty in the Computer User Consultancy Office (UCO) was doubled throughout June and July and there was a great deal of individual "hand-holding" of users to help with their migration. During the last month of use, some 2500 users were mailed individually to encourage them to migrate either to UNIX/CUTE or to PC/NICE.

After the official shutdown at midnight on June 30th, the system was run throughout July with no users except for 2 particular applications. After this, and once all the archive and minidisk files had been archived on ADSM, the system was finally shutdown, except for a brief restart in mid-August to retrieve via ADSM a disc for a UNIX client.

Harry drew only a few lessons from the process, the main one being the amount of work needed to help users migrate and feel happy with their new environment. He was concerned that this might not always be possible in the future as support services were gradually scaled down. Also he stated that users often felt happier with fewer choices being presented to them, especially in the UNIX environment where there are always several ways of tackling any problem.

12. Fermilab Farms by Lisa Giacchetti/FNAL

At the time of the meeting, there were some 158 SGI nodes and 144 IBM systems in the Fermilab central processing farm, a total of

about 8000 MIPS. All nodes have local disc, often very small capacity, and often minimal memory. There are also 10 larger nodes designated as I/O nodes with some 250 GB of disc space and some 90 Exabyte drives. Worker nodes are split by experiment into different Ethernet segments with one I/O node and one file system per segment and appropriate disc and tape allocation; great use is made of NFS for disc file access although bulk data moves are via TCP. This arrangement makes it rather cumbersome to effect changes in configuration, for example when an experiment's node allocation changes.

Parallel processing is typically effected by CPS, a locally-produced socket-based scheme which distributes events to worker nodes. The base CPS package has been extended to provide a batch queuing scheme.

Problems with the above arrangement include :-

- the system management load of over 300 individual nodes
- failures of I/O nodes are severely felt (Fermilab report an average of 7 broken Exabyte drives per week)
- the load of system reconfiguration among the user groups as needs change, especially file system moves and retuning
- obsolete technology in some places, for example on many nodes they cannot afford to run the latest releases of AIX or IRIX respectively
- segmentation of the nodes is felt to be less than optimal

For the future there is known to be a need for much more capacity, both data and CPU, better inter-node communication links and more memory for physics analysis jobs. To this end, Fermilab issued a tender to the suppliers of the 4 platforms already present on site and as a result they are currently purchasing a further 45 SGI Challenge S nodes, 22 IBM 45P nodes, 2 more I/O nodes, one each from SGI and IBM, plus more disc and tape capacity.

On the new farm, the configuration will be modified to use an Ethernet switch to permit a much more flexible allocation of worker nodes to experiments, including node sharing where appropriate. The new farm will have a single file system, possibly AFS, and this should eliminate the need for bulk file or data moves. All the equipment had arrived by the time of the meeting and the target was to have the new farm in production by January 1997.

Speculating on possible futures, Lisa mentioned the suggestion of PC/Linux farms, multi-CPU systems, perhaps a Windows-NT farm. For the moment however, Fermilab intends to stay with Exabyte drives, largely for compatibility with the running experiments.

13. [Fixed Target Experiment Support under UNIX](#) by Lauri Loebel Carpenter/FNAL

The challenge facing the FNAL support group was how to support many UNIX systems coming from different sources staffed by teams of different skill levels. A prime requirement was obviously the adoption of standards, this especially in the area of operating systems and patches. It was also felt that centralised support would be most efficient, including call logging and support mailing lists. Also, make available plenty of documentation.

A script was written to be run locally on target systems to gather information about their configuration; this information is then posted on a web page for each system. Another useful tool was found to be a Suggestions Document where one can easily find a list of Do's and Don't's.

So far, there has been a good level of acceptance by the user community although concerns have been expressed about too-frequent SGI patch releases. Another currently-open question is system level backup: when should it happen, how often, and what to do if there is no local tape drive.

Among the plans of the support group is the idea of creating a more automated method of installing the operating system on client nodes.

14. [CERN WGS and PLUS Services, an Update](#) by Tony Cass/CERN

This was a review of the Public Login UNIX Servers and dedicated UNIX Work Group Servers relative to what had been presented at previous HEPiX meetings, see for example [the talk at HEPiX Rio on WGS](#), URL http://www.wcn.cern.ch/hepix/meetings/rio/Rio_cute.ps.

- AIX: 24 nodes of the CERNSP SP2 system were dedicated to interactive use and known as the CERNSP PLUS service. There were over 1500 active users per week, 100 active at any one moment with a peak of about 180. Current investigations include the consideration of twin-CPU 43P model 240 nodes for additional capacity.
- HP: the HPPLUS service consisted of 21 model 735/125 systems although an upgrade was due shortly. Its usage was about half that of CERNSP PLUS. There is also an ATLAS WGS cluster based on HP and this will move shortly to twin-CPU model D250s. HP-UX 10.10 was run on the HP nodes without the so-called transition links, thereby moving to a more UNIX System V layout. Few problems had been seen. The Logical Volume Manager (LVM) found on UX 10.10 was thought to be very useful.
- Digital UNIX: two ALPHA model 250 had been installed recently. The feeling was that D-UNIX changed too frequently. Use of this PLUS, DXPLUS, remained low for the moment but DELPHI were expected to convert to it soon.
- SUN: no SUN-based PLUS service yet but several SUN WGS clusters, all under-utilised.

The systems were installed with the appropriate certified version of the operating system (see [HEPiX Vancouver](#), URL <http://www.wcn.cern.ch/hepix/meetings/triumf/cern-certify.ps>), then SUE tailoring (see for example [HEPiX Prague](#), URL http://www.wcn.cern.ch/hepix/meetings/prague/CERN_sue_update.ps), then the WGS/PLUS customisation scripts and then tailored for each cluster according to use. Of course, AFS was vital for all of this. Several tools had been developed for performance monitoring, to identify trends, predict problems and so on. Users had requested some remote monitoring which would then include any networking effects. This will be added shortly.

For the future ...

- Where, if anywhere, does Windows-NT fit into this (interactive) scheme? In particular, what about access from a PC to AFS files? For the existing experiments, this is not likely to be an issue as they seem to wish to continue with UNIX into the new century.
- There have been repeated requests for investigations into LSF at CERN, including its ability to "soak up" unused capacity on interactive nodes. It was agreed at a meeting with the physics users (FOCUS) to purchase trial licences and perform some more tests before finally deciding on this direction.

Other topics under more or less active consideration are CDE, DFS and NT.

15. [FNALU, Revising the Vision](#) by Lisa Giacchetti/FNAL

FNALU is the multi-purpose central UNIX cluster at Fermilab. In its original version it was targetted at physics applications as well as some CPU-intensive batch work. It is AFS-based and multi-platform. Over time, the AFS servers themselves have migrated from AIX systems to SunOS, currently on 6 SUNs with 600 GB of disc space. Use of AFS has permitted a physical move of the cluster to

take place without interruption of the user service. Overall, usage of FNALU has risen gradually, especially with the rundown of the FNAL VMS service.

On the negative side of AFS, its performance was found to be poor for large files and this has led to the installation of 2 times 60 GB of local, non-AFS, non-backed-up disc space on the cluster nodes, dedicated to certain applications. Also, AFS proves to be less than optimal on multi-CPU systems, for example 20 node machines, giving system cache problems.

There is now a move to merge CLUBS (the heavily I/O batch-oriented cluster) with FNALU in order to reduce the central support load and to provide a unique user interface. One implication would be a migration from Loadleveler to LSF. This would unify the user interface to services and offer a single queuing scheme, attractive to both users and system administrators

16. [Interactive Benchmarks](#) by [Claudia Bondila/CERN](#)

An attempt has been made to model and benchmark a typical CERN PLUS (see above) interactive load. It is hoped that this would help in fine-tuning nodes of the PLUS and WGS clusters and predict trends for future investment. Today's standard CERN benchmarks are batch only. Claudia has developed a Monte Carlo scheme to run simulated UNIX accounting and attempted to model a diverse applications mix. She uses accounting records because these are easier to generate and handle although it is realised that they only provide information on the end-state of a job with no details of the resources used over the life of the job. Also there is no information on network (including AFS) use and we cannot distinguish between shared and "normal" memory.

Claudia explained the methodology she used: accounting data produces a huge amount of information and this needs to be compressed by splitting it into groups and categories. Users are grouped into "clusters", typifying their use of the systems and different combinations of these clusters are used to produce different benchmarks. By varying the resources used by a member of each cluster type and the total number of users, a Monte Carlo method can be used to generate system loading and hence statistics. She presented in graphical form some of the results.

So far, she had modelled the ATLAS WGS load and she will now try this for other services, refining the command definitions and comparing different platforms.

17. [Interactive Services at CCIN2P3](#) by [W.Wojcik/CCIN2P3](#)

This talk was concerned only with the SIOUX and BAHIA clusters at CCIN2P3 although they share file and tape services with the batch clusters also. BAHIA consists of 4 HP J200 series nodes and 4 RS/6000 model 390s in the CCIN2P3 SP2. It is intended for batch job preparation and general interactive work. They investigated response times by running test scripts and they found that HP model 735s become overloaded at 15 users logged-in but moving to J200 systems has moved this limit up to at least 50 users without degradation.

SIOUX has 2 HP 735s and 4 RS/6000s, again from the SP2 configuration, and is intended for users with heavier CPU tasks. In fact BAHIA uses nodes of SIOUX for large CPU jobs. Like BAHIA it uses load sharing but based on a simple round-robin sharing algorithm where the user gives a generic name (BAHIA or SIOUX) and the name server allocates a particular node name to the connection request.

PIAF has been ported to run on the SP2 using IBM's parallel file system for data access. They needed to develop a group quota scheme to use this parallel file system. It is in use by 6 user groups. The conclusion so far is that the PIAF/SP2 arrangement is well suited to data mining tasks and that the parallel file system offers an interesting speed-up in a completely transparent way.

18. [Roundtable Discussion on Security Around HEP Sites led by L.Cons/CERN](#)

This was a discussion on computer security around the different HEP sites led by Lionel Cons, convener of the [HEPiX Security Working Group](#) which was established after the HEPiX Prague meeting. Since then there has been little activity however. The session started with some site reports.

- [CERN](#)

CERN has a Computer Security document but it is more a set of rules than a real policy. By default, most incoming access is denied to a machine on the CERN site unless requested by the administrator; outgoing access is open. Local system managers have a few tools: SUE which sets up some standard security checks as one of its features; a CRACK service for password checking; a new password program to improve the quality of AFS passwords. There is a screening router between CERN and the outside world which blocks certain traffic and a firewall is planned.

All IP traffic is monitored and this helps not only to produce statistics but also provides a means of tracking back when security incidents occur. Since the CRACK service started there has been a dramatic fall in the number of weak passwords (passwords failing the tests have fallen from 25% to 4%). Nevertheless, there is about one major security incident on site per month on average and it is estimated to cost some 4 person-months per year to clean up and repair the effects. Plus there is a small (3 man) team working on security of whom one is full time on security.

- [DESY Zeuthen](#)

DESY in Zeuthen were working on defining a more official policy than the unofficial one which then existed, including rules for users and for administrators and what to do in cases of violation of these rules. Passwords are set for expiry, there are no guest accounts permitted and there is a scheme to check for good passwords as they are changed. Weak passwords are forbidden. They also check for world-readable sensitive files. Satan is used and a TCP wrapper allows them to monitor the traffic. This means they are quickly able to react to an attack.

- [DESY Hamburg](#)

The Hamburg site of DESY has suffered some incidents recently. They have established a local security mailing list and the site now required AFS home directories to be closed to public access although this was in conflict with physicists' desire for openness.

- [RAL](#)

RAL was relatively open in the past but is changing nowadays, partly because of recent attacks and partly because of pornography on the Internet. There are now a couple of security committees, a CRACK service, a COPS and SATAN service and CRC checks of system files. There were plans for a firewall but the level of access was not yet decided. There were also plans for a sniffer and a central web cache to catch pornography and other illegal material. A part-time security officer had been appointed.

- [CASPUR](#)

CASPUR had suffered a major incident in April caused by malicious use of a sniffer on the internal network. This had caused an important rethink of security but little manpower was available to implement this. Hence CASPUR was eager to see an active HEPiX Security working group with which they could collaborate. Meanwhile, a certain number of rules had been enforced for selecting AFS passwords and all known security holes had been closed.

- **FNAL**

A security policy exists but seems to be little known or used; there is a part-time security officer and some monitoring is performed. There are rules for the unauthorised use of CRACK. But the presence of many visiting teams and particular national and state laws make computer security sometimes difficult to implement at an individual or corporate level in some cases.

From the above it was clear that things are indeed changing for the better vis-a-vis security with much more awareness now of the risks and possibilities, at least in the larger labs. AFS in particular makes security that much more important. For example how to protect sensitive files normally stored in the home directory, should they be partially open to HEP sites or completely protected.

During the discussion it was proposed that each site appoint a named security contact whose name could be included on a mailing list. The question was raised if HEPHX should or even could define a recommendation for a default security policy and if it did how that could be implemented. Could one achieve a secure HEP collaboration over the Internet?

Other subjects raised included

- AFS cross-cell authorisation, did it still work?
- the use of ssh and arc to protect passwords and AFS tokens
- how to balance complete security with physicists' desire for openness
- how to get individual lab managements to agree on the need for a common security policy

19. **Managing a local copy of the ASIS repository by Philippe Defert/CERN**

For people who do not already know, [ASIS](#) is CERN's central repository of public domain UNIX utilities and CERN-provided tools and libraries. Programs and packages are "cared for" by a local product "maintainer" according to specific levels of support ranging from full support to acting merely as a channel to the original developer. There are 600 products/packages contained in some 14GB of disc space and 10 architectures are represented although some such as ULTRIX and SunOS will be frozen at the end of 1996. Clients can access ASIS via AFS, NFS or DFS (although external DFS access was not yet available).

Products "appear" to clients as if they sit in /usr/local/... although in fact many are actually links and there is a regular update procedure (ASISwsm) which normally runs nightly on client nodes to update these links. This procedure is tailorable by a local system administrator. Another utility, ASISsetup, declares which version of which products should be used on a particular node, similar in function to FNAL's UPS/UPD procedures. However, enabling access to multiple versions at a time requires extra effort from the maintainer and this is not possible for all packages.

Product maintainers have a certain number of tools at their disposal: to build a new version across all supported architectures with one command; for testing; and for public release of new versions. Products passed through various states - under test, under certification, release - and Philippe illustrated the flow of modules and the directory hierarchy used.

To offer ASIS at remote sites, a replication scheme has been developed which was usually run daily. It checked for the presence on the CERN master repository of new products or new versions of products, thus minimising file transfers. This extended to checking when the change was only a state change (for example, only moving from certified to release state) in which case the files on the remote site do not need to be updated, only their state and/or directory placement. The ASIS local copy manager tool was currently being updated to take account of the latest ASIS structure changes. At the time of the meeting some 6 remote sites were mirroring ASIS locally.

Product state changes and releases are defined as "transactions" such that if a transaction does not complete fully (removal of an old version for example, followed by release of a new one), then some recovery is attempted and warning messages issued.

20. **ELLIOT: An archiving System for Unix by Herve Hue/CCIN2P3**

CCIN2P3 felt the need for a UNIX file archiving scheme to replace VMarchive on their VM system, to be ready by mid-97 when the mainframe is due to be stopped. It should have both a command line and graphical user interface and be based on ADSM, the agreed file backup utility. ADSM in its "natural" state was rejected as the archive tool after some tests because not enough information about the file was stored, there was poor authorisation control and the command line and GUI actions stored different items! There was also a suspicion that ADSM, at least the then-current version, suffered from a Year 2000 problem - it could not accept an expiry date beyond 2000.

The chosen solution was to add a client part "above" ADSM using ORACLE to store the required extra meta-data about the files and to build a MOTIF interface to this. XDesigner was used to build this interface. The client code on each node communicates to an ELLIOT server demon which itself communicates to ADSM and ORACLE. The ELLIOT server must reside on a node with access to the files it is to archive, normally the file server itself.

ELLIOT offers powerful query facilities and if only meta-data is requested then only ORACLE queries are issued, ADSM is only called if the file itself must be accessed. Further, the query calls are contained in a library such that eventually another database could be substituted, even another backup product to replace ADSM.

The speaker showed examples of both the graphical and line mode interface to ELLIOT. Man pages and online help from both the GUI and line mode versions already exist. The program was then in beta-test. In answer to a question, Herve stated that the files could be fully accessible even without the ORACLE meta-data since sufficient archive data was stored in ADSM as well.

21. **The DESY UNIX User registry System by S. Koulikov/DESY**

There are many clusters at DESY, AFS and non-AFS, UNIX and non-UNIX (e.g. Novell). A scheme has been developed to have a single user registry for at least the UNIX clusters based on DCE registry and QDDB, a public domain package from the University of Kentucky. Both databases contain almost the same information, DCE being the master, and nightly consistency checks are performed. QDDB is normally used for quick information searches.

The normal registry administrators and the DESY User Consultancy Office can register and modify users' information and group administrators can perform certain maintenance tasks, for example modify a password or space quota. The user can modify his/her own password, choice of shell, mail address and default printer.

The speaker described the 2 databases in some detail and how the information flows when a registry entry is made or altered. He also showed the graphical interface.

This is seen as a first step in the migration to DCE when the plan is currently seen as every user having a DCE account from which accounts would be generated on particular clusters or services.

22. **A Scalable Mail Service - LAL Experience with IMAP by M.Jouvin/LAL**

Since some time LAL has been largely based on VMS mail, accessed via dumb terminals or a Motif interface. PC and MAC users were told to use terminal emulators to the VAXes and mail from there. UNIX mail was not supported. As the use of VMS has decreased and

that of UNIX has risen, a good UNIX solution was needed, including a native GUI for PC and MAC users and support for PPP access from home. Other requirements included

- a unique mail message database
- support for non-text mails (e.g. MSoffice) using the MIME protocol
- independence from the actual user mail agent chosen.

X.400 was considered and rejected - largely because the address format has no fit with the sendmail protocol, because X.400 is difficult to manage and because there are too few clients.

POP was rejected, despite its nice user agents, because

- it is an "offline" protocol (store and forward rather than server/client)
- only the Inbox is accessible on the server
- due to the above, it is not suitable for mobile use or use by a floating population
- it is considered by many to be difficult to manage
- it offers only limited privacy

LAL finally selected IMAP 4 (based on RFC 1730) as the basic mail protocol. Messages are stored on a server and accessed by clients. The so-called IMAP consortium drives the standard and its extension ICAP (the Internet Calendar Access Protocol). For IMAP, more and more clients and servers are becoming available, both public domain and commercial. Among the promised servers are Netscape mail server, Microsoft Exchange (both due end '96) and SUN Solstice.

IMAP supports MIME, it has an efficient header-only search scheme and it is optimised for low speed links. Many mailbox management schemes are possible including ACLs and quotas. The protocol (ICMP) will evolve to ACAP (Application Configuration Access Protocol), which is expected to be ready at the end of 1996. This will help with the configuration of address books and should permit common and similar access for address books as for mailboxes. ACAP should also extend IMAP to multiple mail servers with load sharing and replication.

The CMU/Cyrus project suite (also known as the Andrew II Project) is effectively IMAP+IMSP. LAL runs this on a pair of high-availability Alphaservers with automatic failover. The Cyrus suite supports Kerberos login authentication and has a TCL-based management tool.

As clients they offer one or more on every platform, even on VMS, PC and MAC although on PC the choice is commercial (Netscape 4). At that time, only one tool was available on all platforms, Simeon (from a small Canadian company). This happens to be the only one today supporting ACAP although Pine is expected to do so shortly and others later.

The talk closed with a review of the features of Pine and Simeon and a list of URL references.

23. [X11 problems: how the HEPiX X11 scripts can help you](#) by Lionel Cons/CERN

Lionel was convinced that in general debugging X11 is hard: he illustrated this by showing examples of the rich variety of client/server combinations, the many inter-linked libraries used by the applications, some of the many environment variables, keyboard mappings, etc. Then there is a window manager, a session manager, and so on. Often even performing a "simple" task can be non-trivial - changing the cursor size or selecting a particular font for one application were just two examples.

X11 is resource-hungry, both in CPU, memory and network traffic, especially if it is badly configured or contains "unfriendly" X11 application (Xeyes is a trivial example, or a pretty pattern screen saver if not run on the X server itself). Lastly, X11 has several inherent security concerns.

From all of this we can say that X11 needs

- good defaults
- easy customisation
- usage resource control

This was the driving force for the development of the HEPiX X11 scripts. The project started as a joint effort by DESY, CERN and RAL with most of the actual development being done at CERN. They are now widely deployed at CERN and gradually elsewhere (e.g. DESY, SLAC).

For the future, Lionel and his X11 team are looking at

- connectivity tools such as ssh
- further work on the scripts
- wider deployment (help can be made available)
- Wincenter for access to PC applications
- CDE (needs significant manpower - and system resources)
- X11R6 transition - when to plan this? Should we wait for R7?

24. [HEPiX X11 Working Group Report](#) by Lionel Cons/CERN

The meeting, the day before this HEPiX conference, had consisted of various site reports, technical reviews of the current state and open problems and the planning of future work. A list of proposed enhancements was agreed as well as a list of topics to be investigated in the future by the working group. The [minutes](#) from this meeting can be found on the web at URL <http://www.wcn.cern.ch/hepix/wg/X11/www/Welcome.html>.

25. [HEPiX and NT, a Roundtable Discussion](#) led by John Gordon/RAL

John introduced the subject by describing RAL's Windows plans - nothing on Windows 95, go directly to NT; it is expected over time to become the desktop system of choice. At that moment work was just beginning on an NT infrastructure at RAL where, although the UNIX base was rising, the interest in NT was rising even faster. The HEP processing model indicated a potential wide use of PC/NT in Monte Carlo production where the price as well as the price/performance could possibly solve the need for massive CPU power for LHC experiments at a price they could afford.

He posed the question - should these PCs run NT, or Linux or Solaris - and does it matter? HEP needs good Fortran, C and C++ and these are all already there with more coming. For NT there are several options for NFS access (including Samba), AFS and DFS are announced, there are several NQS ports and a beta of LSF. RAL is porting their tape software (VTP) and Sysreq to NT. And remote access to NT is possible with telnet, WinDD, Wincenter, Ntrigue, etc.

WOMBAT was a trial NT farm for HEP. It consists of 5 Pentium Pro systems with a local file server, connection to the ATLAS file store and Ntrigue for remote access. It was used to demonstrate several ports and provide a platform for experiments to try, including

benchmarking. It offered an alternative to UNIX for groups now on VMS which was being rundown. It was hoped to have the first results by the end of 1996.

Areas which John felt need to be addressed include

- integration to NFS and AFS
- integration to magnetic tapes
- submitting a job from UNIX to NT and vice versa (integrated batch)
- CERNLIB
- ASIS

RAL would be interested in involving other labs, perhaps a HEP-wide effort? They believed that portability was more important than uniformity; and sharing experiences, and mistakes, was very important.

CERN had an approved project (RD47) evaluating the use of PC/NT. The LHC experiments were interested in NT and there were a certain number of individual initiatives happening around the lab.

DESY has established an NT Interest Group including representatives of the Computer Centre, major user groups, and Zeuthen. The largest effort was that of the ZEUS experiment which plans to base itself on NT. They have ordered 20 Pentium Pro systems. On the other hand, there are no current plans for an NT-based desktop.

CASPUR is heavily into PCs but rather under Linux as has been reported at this conference and elsewhere. Windows is only used for secretarial work. CASPUR will investigate the NT/AFS client (as will CERN).

FNAL, in the progress of migrating away from VMS, is recommending administration workers to consider PC and MAC rather than UNIX which is the recommended platform for physicists. Some NT servers are beginning to appear on the site.

NIKHEF have bought some PCs for evaluation. LAL have one NT server, dedicated to a Wincenter service and GSI have several of these. Prague Uni have one PC for work on Objectivity.

The session closed with a discussion on how to keep up-to-date with this area and the suggestions included

- a news group for HEP NT activities
 - extend HEPiX, either generally or for particular meetings (for example, schedule NT sessions at a forthcoming meeting)
 - create HEP-NT
 - ask HEP-CCC for advice
- and various people undertook to investigate these.

[Alan Silverman](#), 27 February 1997