

HEPiX Rio Meeting Report

Minutes of the Meeting held September 24th-26th, 1995

1. CBPF and LAFEX Intro - Prof Amostope and Prof Santoro

CBP was the first research institute in Brazil, created in 1949. It had a long tradition of hosting famous teachers and visitors.

LAFEX was one branch of CBPF and Prof Santoro presented the structure of the research establishments in Brazil together with a description of the range of physics and cosmology done at LAFEX.

2. Requirements for UNIX from Physics Experiments

o [L3 and UNIX, An Adiabatic Solution](#) - V.Innocente

The speaker described the adiabatic transition of L3 to UNIX. The experiment had begun with VAX/VMS for the online part in 1983 but had decided on a mixture of an IBM mainframe, IBM emulators and various workstations for the offline. The first Apollo had appeared in 1985, running Aegis which was "nearly" UNIX but not quite. Over subsequent years, more and more Apollos arrived, both workstations and the first DN10000s for batch work.

When CN opened the CSF service in 1992, L3, having already ported its code to UNIX, were well placed to take immediate advantage and used a large fraction of the available resources for some time. By this time, Apollos were deemed to have no future and their gradual replacement by HP Series 700 had begun. UNIX replaced Aegis on the remaining Apollos and some RS/6000s appeared. In 1994, a large SGI node was installed in SHIFT for L3 and much more UNIX capacity has since been added.

L3 had always insisted on a heterogenous environment with the result that physicists had been forced to get into contact with UNIX early and this had allowed L3 to take advantage of steadily increasing amount of CPU power at the best price/performance ratio. Since at that time LANs were bottlenecks, a central solution was necessary - such as SHIFT and similar farms. Use of CERNLIB, GEANT and standard FORTRAN eased the transition between UNIX flavours. Graphics was initially GKS-based but has since migrated to X/PHIGS and Motif.

Much of their central production control software had been written in REXX for use on the mainframe and this had needed a very large conversion effort in the move to UNIX production facilities. Code management was CMZ and Patchy and these ran everywhere. They were considering a move to CVS in 1996.

The speaker summarised the architectures as follows -

- VMS was user-friendly but very expensive and now only used for online.
- VM/CMS (the HEPVM variety) was quirky, not user-friendly and had a poor file system. But it was robust and reliable for production running and data base control.
- Apollo/Aegis was nearly UNIX and had a great file system (like UNIX should have been) but when using a cluster as a compute farm the LAN had been a bottleneck at that time
- UNIX (various) were all similar but never the same, at least for system administration; they were all less user-friendly than VM but also less quirky. A major drawback for a large collaboration was that the permission bits in the file system offered insufficient flexibility.

The speaker summarised saying users spend more time inside applications than with UNIX itself and a major source of irritation remained the different keyboards.

o [BaBar's Perspective of UNIX](#) - T.Glanzman

BaBar was a new collaboration at SLAC but many of its members already were familiar with UNIX. Broadly, the statistics of the experiment, due for startup in 1999, were as follows: it was a Beauty experiment where the event rate was expected to be around 100Hz and at an expected 25KB per event, the total data sample would be around 25TB per year. It was a large international collaboration which implied non-trivial project management. This area, project management, plus some of the engineering work, would use NT, but the collaboration had decided to standardise on UNIX for its main computing platform.

A lot of effort was being spent in smoothing the rough edges of UNIX with respect to the user environment and its differences on different architectures in the various labs. Special attention was being paid to keyboard mapping and cross-platform compatibility of tools. The speaker presented several foils of issues raised within the collaboration.

The collaboration, with SLAC Computing Services, had prepared a requirements list to guarantee institutional support. They were interested in the HEPiX Security Working Group and would be in favour of a HEPiX-wide security policy across the labs. They will consider adopting the HEPiX Login Scripts. These and other questions were being addressed in the form of a questionnaire - the so-called "Randy's List" of questions to be answered by the collaboration. Many of the issues had been or were being discussed within the HEPiX community already.

Among vendor concerns were the need for Transarc to provide timely releases of AFS for all the operating

systems in common use and the need for all the vendors to support Fortran 90 properly. There was also a requirement to define a standard data interchange format for the experiment, including the selection of media.

The speaker declared that in his opinion as a physicist in BaBar, UNIX was "acceptable" but it needed improvements to its user interface and close interworking between the laboratories would be essential.

◦ **BELLE at KEK - Y.Morita**

BELLE was another Beauty Factory being built, this time at KEK in Japan. It was scheduled to start up in 1998 and would be a high-rate asymmetric collider. Total data rates were expected to top 30TB per year and required storage capacity had been estimated at 300TB and computing power at 30K SPECint92. The data flow would be from online farms through data servers to offline farms and users' private workstations and all of these were expected to be running UNIX.

A list of requirements had been put together which included support for large file systems, hierarchical storage facilities, parallel event processing, a distributed batch queuing system and fault tolerance. To this end, several research and evaluation projects had been started, some of which had been reported at CHEP the week before.

User requirements could be summarised as single system images across a distributed environment, easier multi-vendor support, transparent access to data and file sharing (DFS was under test) and single batch queues (they were evaluating LSF).

They felt that they did not want a special HEP UNIX but a better interface to it. Security must also be improved.

◦ **Embedded UNIX at HERA-B by P.Wegner**

HERA-B was due to start in 1998 and involved over 90 physicists from 15 institutes, with more joining over time. It was based on a fixed target and was expected to take some 80GB of data per year. The experimental set up would use a standard data acquisition system with a multi-level trigger in which UNIX would figure in the top levels. The third level involved a read-in throughput of 1 to 2 KHz, 50 to 150 KB per second. Clusters had been considered but rejected because of the desire to avoid administering hundreds of workstations. Instead they had decided for embedded systems, based probably on VME boards.

Simulations had already begun (presented at CHEP). They were looking at PowerPC as the processor to use but the software question was still open - VxWorks, LynxOS, etc? Another option was to develop their own software. One important criterion was not to depend on a single vendor. Evaluation of these various options was continuing.

For the future, they were anxious to see hardware vendors adopting the latest standards and they hoped for a UNIX with better real-time features.

◦ **D0 by A.Parra**

The speaker presented a physicist's view of moving to UNIX. Until now, D0 at Fermilab had been largely based on VMS but UNIX was coming in more and more, especially in SGI and RS/6000 computing farms although the physics analysis and the main data server remained today on VMS. Phase II of D0, starting in 1999, would require to be done in UNIX however.

D0 had several hundred manyears invested in developing and using VMS utilities and there was an understandable reluctance to move away from this. He himself had experienced the adoption of UNIX in three phases:

- Phase I - irritation
getting used to the strange command names and strange and inconsistent behaviour.
- Phase II - confusion
having to make a choice of shell, editor, etc, up to a whole environment
- Phase III - great confusion
customisation questions, different platforms, keyboard differences, different feature sets.

He believed in the end that UNIX was powerful if difficult to get into and that a good Primer was missing. Turning to code porting issues, he said it had been relatively easy and only the use of non-standard features had complicated things. However, code verification against existing code was not easy, the results always varied, and not only simply rounding errors. But certain parts of D0's code suite would be very hard to part - for example they made use of the VMS search list feature and VMS debugger.

The fact of changing the computing paradigm in a running experiment (replacing dumb terminals connecting to a mainframe by workstations, X terminals and PCs) raised many debates, including whether to move directly to Windows NT.

During questions, the example of CEBAF was quoted where there had been a debate for 2 to 3 years between VMS and UNIX protagonists but where the decision had finally been taken recently to adopt UNIX.

◦ **CMS's Requirements for UNIX - V.Innocente**

CMS were currently writing a detailed Computing Technical Proposal which must be published in 1996. Currently they were UNIX-based but it was possible that PCs might appear in strength in the future.

CMS users generally preferred UNIX workstations to X terminals because there was less network disturbance, especially to graphics-related work such as PAW or data visualisation. Apart from their workstations, they made heavy use of central batch (SHIFT and CSF) and used MACs and PCs for office applications. Access to CERN from external sites was essential, including X11 access in both directions.

They had standardised the user environment on the HEPiX recommendation and were promoting the use of AFS

for home directories. They had established their own X terminal boot service with some local tailoring and the standard HEPiX X environment.

Experience had shown that logins were slower than expected, probably caused by the use of the HEPiX login scripts and AFS and there was a perceived lack of user documentation on the HEPiX environment, especially how to tailor user-configurable parameters - the speaker suggested a "man HEPiX" command. And system admins reported that it was not so easy to install the present HEPiX environment. Nevertheless, the fact of having the same environment on all services was very much appreciated, and there was also praise for having a standard X11 for X terminal management.

The speaker suggested that HEPiX strive for

- slimmer, more efficient login scripts
- better documentation, including reference cards and man pages
- having the HEPiX scripts installed in outside labs, at least as an option
- more X applications (in particular an X mailer and not Pine) since their use of dumb terminals was very small and support for dumb terminals was low priority in their opinion

CMS were concerned that AFS was not an official standard and was also late in arriving for new releases of vendor operating systems. On the other hand, they were pleased to have it managed centrally with issues such as file backup taken care of for them. Token expiry was judged to be too short (set for the default of 25 hours), there was additional complication when logging in to remote nodes for long periods in cascade fashion (which token has expired?) and general AFS reliability should be improved.

The speaker acknowledged that security was important but suggested that it was mainly an issue between HEP and non-HEP sites and that inter-site working between HEP sites should be made easier.

3. Site Reports

○ LAFEX/CBPF by M.Miranda

LAFEX had gone through a cycle of replacing an IBM 4381 by a VMS based network and had recently added PCs, X terminals and some UNIX workstations, mostly SUNs running SunOS but a migration to Solaris was planned later that month. Future plans included more SUNs and a farm of RS/6000s. Physicists at LAFEX collaborated in the D0 and Delphi experiments, both today based on VMS but gradually moving to UNIX. The site acted as the distribution centre for the CERN Program Library for Brazil.

○ [Yale University](#) by R.Lauer

The speaker described her goals as being to minimise management tasks; to maximise system uptime and reliability; and to serve multiple communities of users including physicists from experiments in various HEP labs, local physics groups and engineers.

She was in charge of some 30 nodes of which 4 were UNIX, the rest mainly VMS; most of the systems were from Digital, either VAX or Alpha. However, within her UNIX cluster there was a heterogeneous selection of systems - ULTRIX, Digital UNIX and SGI/IRIX. She made use of a common application server for the public domain software where the architecture differences were "hidden" by using a /arch directory pointed to by links from each node. The clients of the cluster were dataless (in discussions this was explained as clients which did not have local disc space except for local page and swap files).

The speaker summarised her current problems as -

- which UNIX batch scheme to choose?
- file system performance and security
- printing
- mail
- various cluster issues

○ CEBAF by R.Chambers

CEBAF were then in the process of acquiring a major UNIX configuration in preparation of a new round of experiments. They had been based largely on VMS and has undergone a long internal study on UNIX, its advantages and its negative points. The first UNIX nodes had been HPs; they had appeared some 5 years before and had increased greatly in number since, to more than 300 today. There were also some 30 ULTRIX nodes, which were currently being retired or re-assigned as X terminals.

CEBAF's current plans included a major acquisition to be decided shortly, with the choice falling between either a large SMP system for central processing or a new farm of processors. Important factors affecting their choice would be integer performance, price of course but also cost-effectiveness, and I/O speeds and the software environment on offer. They had a formal UNIX Procurement Committee which was in the process of reviewing the bids received.

Their next planned project would be in the area of data management and they were looking at both Storage Tek's Redwood and IBM's NTP drives, at storage silos and at HSM systems. Other current concerns included the user environment and security - could DCE be of assistance?

○ CERN by A.Silverman

CERN reported continuing growth in installed systems, especially HP stations and NCD X terminals. The AFS disc space was by then over 300GB spread across 7 servers but they were experiencing some problems with attached RAID controllers. On the batch side, there had been some 12 DLT units installed.

Production deployment of the HEPiX login scripts was more and more widespread and work was starting on the X11 part of the proposed HEPiX environment. The recently-installed Mail Server was being gradually brought into general use.

o **FNAL by J.Nicholls**

Since the last meeting, both CDF and D0 had made major purchases of large UNIX systems (both SGI Challenges) and in general UNIX use was rising sharply. Work was continuing on further developments to FUE - the Fermilab UNIX Environment - with regular weekly meetings. They had agreed so-called levels of FUE compliance, as explained in a CHEP session the previous week. This allowed the support group to define the level of support a user could expect.

They were also establishing a Central Software Distribution point for the encrypted distribution of commercial software as well as a server for the provision of public domain packages, today some 40 packages with 4 to 6 architectures for each. A certification scheme for Operating Systems had been drawn up in order to maximise both consistency of the delivered product and also the use of their limited resources for support.

Over the past few months the group had been asked to perform a steady stream of software installations and were looking at ways to automate this, including the handling of operating system patches.

One of the principal new projects was Nirvana, a scheme to package various FNAL-written utilities for use by offsite groups. Among these utilities were included nedit, histoscope, the OCS (Operator Communications Software) which was in use on FNAL's production clusters and farms, and several others. Another new development in the offing was DART - UNIX-based data acquisition - for whom the first customers would go into operation next Spring.

o **DESY by K.Kuene**

DESY's AFS Work Group Server population was growing steadily, over 150 users at the time of the meeting. The users' home directories were now stored on RAID discs. There were a total of some 20 servers of different architectures and some replicated applications binaries servers.

There was central support for system installations but these were found to be labour-intensive so some automation was being sought. An AFS installation script was provided for users and the HEPiX environment was installed via the Salad procedure. They planned to migrate their X environment to the one being proposed by the HEPiX X11 Working Group.

There was a long list of planned developments:

- Group-based administration of users' AFS space, including group quotas
- A central mail and registry service, possibly DCE-based
- Work on batch for AFS token extension; Loadleveler was being used for the HP platforms but NQS on the SGIs
- Access to the STK robots was possible only from the SGI Challenges at that time and HP-UX access would be added. 16 drives were connected via 3 RS/6000s. OSM was under study for hierarchical access to data.
- A questionnaire had been circulated on the needs for applications software and the result was a list of supported products made available in a standard place and accessed via AFS.
- Other future plans included a data server for medium-sized files; centralised UNIX management; DCE user registration; and DFS.

DESY Zeuthen were still largely based on NFS with AFS only being used for applications software. There were plans to add a cartridge robot, possibly a DLT system.

o **SLAC by C.Boeheim**

A number of central file servers had been installed since the last HEPiX meeting, including a RAID file server. Further server additions would include one for news as well as possibly some NFS servers. Experiment E154 was about to start operation, including data logging to the main file servers.

SLAC were evaluating some technologies in preparation for the BaBar experiment, including SSA-connected discs on RS/6000s.

Plans were being made to migrate SPIRES to UNIX, by recoding most of it in C with the remaining parts using a 370 emulation scheme under UNIX.

o **NIKHEF by W.van Leeuwen**

The speaker described briefly the divisions at NIKHEF: Section H (HEP Physics) run a wide range of UNIX architectures; Section K (nuclear physics) use only SUNs. There were several work group clusters, based variously on HP, SUN and SGI platforms; some were shared by different groups, some dedicated to one group. In addition, two large SUNs were used as public UNIX servers and there was a tape server acting as an interface to VM.

NIKHEF have laid plans to reconfigure their network into a star topology using UTP with subnets. They will copy CERN's ASIS scheme for public domain software and introduce AFS.

They maintain a very active group working on audio- and video- conferencing and have established a Media On Demand service making use of many freely-available tools. They will participate in the 1996 World Internet Fair.

o **Prague HEP Institutes by J.Hrivnac**

The support team were currently trying to create a common structure across the four HEP-associated institutes in Prague, an environment based on CERN tools and the use of ASIS and AFS. An agreement has been signed with

Transarc for AFS. The plan also included ATM connections between the sites. So far the HEPiX login scripts were installed and enabled for some 40 users at the Institute of Physics; no serious complaints thus far.

They were looking at an AFS/WWW project whereby URL links would benefit from the presence of AFS if available. ASIS was mirrored from CERN in the Institute of Physics and other Prague institutes were mirroring the software from there.

◦ **CCIN2P3 by H.Hue**

The BASTA farm at IN2P3 had been expanded with more IBM CPU power, ANASTASIE had received a boost with more HPs and ATM on most of its nodes, and the DATASERV cluster had more disc, including a RAID 7 array on test, and a tape stacker. On the software side, there had been improvements to ANASTASIE. to XTAGE and to BQS.

Development was starting on operation tools for the clusters to manage users, batch and disc space management. This would include remote management tools.

◦ **DAPNIA by J.Le Foll**

DAPNIA was now using CCIN2P3's computer centre in Lyon, including its IBM mainframe VM service which however they would like to migrate away from as soon as possible, encouraging their users to move instead to ANASTASIE. Some new workstations had been installed at DAPNIA and used as work group servers, connected to RAID arrays for home directory space.

Their SP2 had arrived; it would be shared by many experiments and was attached to an IBM 3494 cartridge robot as a data repository.

Recabling of the main building was underway, installing twisted pair cables and Cisco switches with 100Mbit Ethernet as the backbone between the switches.

◦ **TRIUMF by C.Kost**

TRIUMF relies on user groups to share system administration tasks. They have a large variety of equipment on site but with a concentration on Digital Equipment. There was a VMS-based Computing File Server, a UNIX File Server and also an overloaded UNIX-based Compute Server. There were workstations running ULTRIX and Digital UNIX, the former to be phased out soon.

Among the problems under current investigation were

- Tape allocation - they were looking at FNAL's solution
- Batch systems - nothing was used today but user pressure was rising
- VMS - there was a plan for a gradual phase-out, over some 4 years

◦ **KEK by Y.Morita**

In January of 1995, KEK had installed a Japanese MPP Supercomputer. Currently out to tender was a 15,500 SPECfp UNIX central computing work group server to be installed in January '96. They had plans for some 10 Work Group Servers and as part of this had developed "PURE" - a Pilot UNIX RISC Environment. It used an Auspex file server.

They had implemented an SMTP mail server and made use of Majordomo as a mail list server. There were also several information servers in operation. For tape devices on their tape servers, there was a mixture of Exabyte and DLT stackers.

Future plans around the Work Group Servers included investigating DCE (with one cell per cluster), LSF for batch work, DFS for home directories and adding POP3 support to the mail service.

◦ **RAL and Oxford**

Although unable to be present, J.Gordon had submitted a short report. RAL was by then using ASIS on their 6 node OSF/1 farm and on the 26 systems in their HP farm. AFS introduction was progressing slowly. An Alphaserver 8400 with 6 CPUs had been purchased and the last use of VM was scheduled for the end of September '95.

The Oxford Uni Particle Physics Department was a largely Digital shop but recently some Pentiums running Windows NT had begun to make an appearance, utilising the Exceed X windows interface. There were plans to evaluate the HEPiX login scripts. Elsewhere in the Lab there was heavy use of SUNs.

4. **Security Issues on Digital UNIX by R.Lauer**

Ms Lauer described why security auditing was important, password protection, X access, etc, and the features which were present in ULTRIX to assist this if the so-called "enhanced security" mode was enabled. Some of these were missing in D-UNIX, at least in the early releases. She then described the various tools which were present in D-UNIX 3.2 and how they were controlled by the system administrator. The use of NIS complicated the security tasks.

She believed that auditing could be a useful, perhaps even one day essential, tool in case of troubles such as a serious security incident. D-UNIX offers to build an audit demon and to provide reports. Many parameters were configurable. She had found D-UNIX enhanced security easy to use and considered it to be a good starting point but stated that it required better documentation and a more consistent user interface.

5. **Lights-Out UNIX - C.Boeheim**

The speaker described plans at SLAC for cutting back on overnight operator cover in their computer centre. SLAC had adopted the Control Console Server from FNAL; it was an RS/6000 with a patch panel connecting the console lines of up to 128 nodes with a console package from Purdue and a log watch package. SLAC had developed a control program

with ideas from CASPUR as presented at HEPiX in Paris the previous year. Messages were passed to a control host module where they were filtered by the log watch program and the "interesting" ones forwarded to specific destinations where appropriate actions could be taken.

For reporting problems, they had considered several tools, including direct alerts via pagers, Zephyr-grams and eventually they had purchased a commercial tool, TelAlert.

For system monitoring they used Patrol which checked for healthy status (for example adequate swap space, free memory, a list of required processes). They used home-grown tools for some of these checks and they had a policy of reducing priority on CPU-bound jobs on their interactive servers.

Other automated tasks included scripts to clean up the disc space periodically, removing temporary and core dump files, producing audit summaries and so on. Another important area handled by such tools was of course security and some of the tools used in security checking were described at the previous HEPiX and they had now been merged into this automated scheme.

6. **User Training for Migration by A.Silverman**

As part of CERN's plan to migrate several thousand users off its IBM mainframe, the migration team was setting up a series of seminars on very specific subjects. These would be typically one to two hours long and would include online demonstrations. For example, a course on emacs would be illustrated by editing a file via an X Terminal connected to an overhead projector. Some subjects would be treated in more depth by arranging courses of several half days.

In choosing speakers, it was hoped that experts from different labs could be used where possible but in all cases it was necessary to allocate sufficient time for the speaker to prepare the material, a non-trivial time often. And it was also the case that the author or person responsible for a product was not always the best teacher.

7. **Work Group Servers at CERN by T.Cass**

At CERN, the speaker is responsible for a team supporting both Work Group Servers (WGS) for particular experiments or groups and Public Login UNIX Servers (PLUS) for general interactive use. Both offer a similar user environment and this is also offered to users with their own desktops. The support team has developed management tools to be independent of platforms since they must currently support a range of systems from IBM, HP and SUN. They aim at one systems specialist per platform and so adding a new cluster of the same architecture as is already present is relatively cheap in resources. Next year however they plan to add a Digital Alpha PLUS system. One tool they use is the SUE scheme described in the Saclay and Prague HEPiX meetings.

The decision to have a cluster of nodes in a given WGS or PLUS rather than a single large configuration was partly in order to minimise the single point of failure risk although every cluster did have one node which had special roles (account registration and accounting for example). Nevertheless, jobs were shared across all nodes in a cluster using a tool called acrontab developed by CERN's AFS team. Similarly, incoming accesses from users were shared across the nodes by the ISS scheme of crude load balancing.

In the past months, a lot of effort had gone into understanding the system behaviour and many monitoring and analysis scripts had been developed. From the results, they could now plot the rise in use of CERNSP compared with the gradual rundown of CERNVM.

The next question was system modelling - what effect would be gained by adding extra memory for example for the same number of logged-in users? Other tools under study or in some cases in development were --

- o control tasks - for example, to lower the priority of looping or CPU-bound jobs on interactive service nodes
- o user tools to interface to tasks which normally require a system manager's privilege - to change the user's shell for example, or to perform a command on all nodes of a cluster
- o system management - they were evaluating Digital's Polycenter Console Manager.

Concentration this year was on consolidation of existing clusters and services plus some help to be given to the LEP experiments to assist them to migrate off CERNVM,

8. **Plans for a Central Mail Server at CERN by T.Cass**

For some years, DXMINT had been the main central gateway for mail both within CERN and between CERN and the outside world, linking non-similar mail schemes and providing the Internet access. Now however, CERN believed it required a machine-independent mail scheme which permitted users moving between hosts, even between sites, to be able to access their mails from anywhere. Also, mail addresses should be independent of particular user hosts.

The first attempt to satisfy this demand had been afsmail - storing the incoming mail spool in AFS which could then be read from any AFS client. However, use of Berkeley mail's flat file scheme for the inbox meant that the AFS file was read afresh with every incoming mail, hence continuous file transmission as the user's mail spool was updated; if the inbox was not cleared regularly this could, and did, generate a lot of AFS traffic. It did not scale.

Use of IMAP as a mail protocol was much more efficient than AFS, only the headers were transferred as they changed until the user actually requested a specific mail. CERN had therefore installed a Digital UNIX Alpha twin host server with RAID discs as an IMAP-based mail server. The configuration included several fail-safe options and automatic failover of the discs. Incoming mails were stored on the server but once saved into a user's mail folder, they were transferred to files on the user's normal home directory, usually AFS. Today the service was IMAP2-based but there were plans for IMAP4. There was also some support for POP.

Another feature of the Mail Server was that it generated so-called user-friendly addresses of the form Tony.Cass@cern.ch in the sender field automatically so that the recipient could reply to a host-independent address and the mail server would be able to deliver the reply wherever the originator accessed the server from.

9. Mail List Server at CERN - M.Marquina

Together with the above central mail server, CERN had recently introduced an AFS-based mail list distribution service based on the Majordomo package. This was to replace the popular LISTSERV service on the CERNVM IBM mainframe. The speaker explained the file layout of the new service including how AFS ACLs were used to protect the lists and permit list administrators to manipulate their own lists.

A user interface package had been developed and was made independent of Majordomo, just in case the latter was later replaced by something newer and better. With this package, a list administrator could use his or her choice of editing or updating scheme, based for example on ORACLE, Filemaker Pro, a simple text editor, etc. Other tools were planned, including some using WWW.

The service had been in production for some months and many users had already successfully migrated to it, including the LHC experiments.

10. FNAL VMS Migration update by J.Nicholls

Since the Prague meeting of HEPiX, UNIX use in Fermilab had increased rapidly with most major software packages being ported now. But use of VMS had remained fairly constant. Among the impediments to migration were cited

- o user comfort level with VMS
- o VMS offered a reliable and stable service
- o many familiar tools were missing in UNIX; examples included the CMS code management system, VAX News and NOTES, RDB, file archiver and others.
- o UNIX was difficult to get started on, offered too many confusing choices and too complex an environment.

The goal of the migration team was now to provide a home, either on UNIX or on a PC which was as comfortable as VMS. The migration task force included representatives from the major user groups and experiments. Communication within the community was considered very important and to this end UNIX User Meetings were held fortnightly although they had not yet attracted a lot of user interest. On the other hand, the first VMS to UNIX public meeting held in August had attracted a large audience.

A VMS migration questionnaire had been issued to help identify the hot topics. People had been appointed to specific responsibilities such as mail, batch, editors, information schemes, etc. A training schedule was in preparation and they hoped to share material with other labs. The UNIX at FNAL Guide was being updated and a UNIXHelp scheme developed at Edinburgh University seemed to be interesting. All this was pulled together within the Web.

For electronic mail, they would prefer something more advanced than Pine, perhaps based on mh, for example exmh. VMS to mh folder conversion tools existed.

They expected most general VMS users to have migrated by the end of summer 1996 although some specific groups, (e.g. CMS and D0) were setting their own timetables.

11. Report from the IEEE Mass Storage Meeting - C.Boeheim

It had been a smaller meeting than usual with little new material since the previous occasion. Minor refinements were made to the Mass Storage Reference Model which was planned to be released to the vendors in Spring 1996. The question for HEP was, do we have any interest in this?

Between the use of ATM and Fibrechannel switches, there was still no clear consensus. Both were featured but they had different target markets, different features, and different problems.

At CHEP the previous week to the HEPiX meeting, J.Shiers and R.Melen had run a BOF on Mass Storage where it had been agreed that

- o HEP sites needed to share information
- o Concern was expressed to understand the implications of the model on HEP and to encourage its completion
- o A number of ad-hoc tasks were defined
- o A mail list would be established and a [web page](http://wwwcn1.cern.ch/~hepmss/) at URL <http://wwwcn1.cern.ch/~hepmss/>
- o The current draft specifications would be circulated and comments collected
- o A meeting would be organised to discuss this subject [subsequently dropped because of lack of interest].

12. WABI - A Solution to Windows Access to UNIX - D.Fagan

This was a report on an evaluation performed at FNAL. They had started with Softwindows to provide access to desktop office products from UNIX but found it to be incompatible with AFS on all platforms. Also, support for some specific fonts was missing, although the speaker agreed that use of a specialised font server might have helped here. Lastly, Softwindows had been found to be a major consumer of CPU cycles.

The evaluation had then switched to WABI which appeared to have fewer limitations and offered better application performance. Licensing of WABI was free on SUNs but not on some other platforms.

[Alan Silverman](#), 2 April 1996