

HEPiX Pisa Meeting Minutes

Alan Silverman

11-14 October, 1993

CN/DPG/93/35

1 Introduction

The second European HEPiX meeting of 1993 took place at the Scuola Normale Superiore in Pisa, Italy and was jointly organised by SNS and by INFN Pisa. Once again it was held in conjunction with HEPVMx in the so-called HEPMIX format, starting with HEPiX concentrating on UNIX matters, a joint HEPiX/HEPVMx meeting concentrating on batch and storage issues and finally a meeting devoted to VM. This time the week began with two half-day introductory talks on PERL and TK/TCL respectively, both ably presented by Lionel Cons from CERN.

Some 45 people attended the meeting. A full set of overheads is available at CERN in the Computer Science Library and a partial set is available via WWW under the entries HEPiX - Pisa.

2 Welcome

Alan Silverman opened the meeting by introducing Prof Emilio Picasso, the Director of the Scuola Normale Superiore, who welcomed the meeting to Pisa and thanked the audience for the contribution that HEPVM had made to the success of the LEP experiments. He foresaw that the work of HEPiX in preparing for LHC would be an even harder job. However the saving in effort that comes out of HEPiX standardisation and collaboration was even more important these days when money is tighter.

Prof Giuseppe Pierazzini, the director of INFN Pisa, also welcomed the audience to Pisa. He reminded them that this was the second such

visit to Pisa since HEPVMS had met there in 1998. The problem then had been the increased data rate due to LEP. Now there were two problems, firstly the high data volume foreseen for LHC and secondly the migration from proprietary operating systems to Unix. He stressed the importance of a common Unix interface on big machines and on the desktop and wished us a fruitful weeks work.

3 Site Reports

3.1 IN2P3 - Wojciech Wojcik

The BASTA farm at Lyon, which had consisted of two IBM RS/6000 model 550s, two model 370s and seven Hewlett Packard 9000/735s, had recently been upgraded by the addition of a further three RS/6000 550s, three model 370s and 14 HP 735s. BASTA is now delivering much more cpu than the VM system. A new version of BQS that is independent of NQS and portable is now under beta test. Ztage now has output capability as well as input. Both BQS and ztage are Posix 1003.1 compliant. By December 1993, BQS will also be 1003.15 Posix batch compliant. Future plans will add I/O scheduling and add checks to ztage to prevent duplicate staging.

Lyon have carried out tests driving their STK robot connected via channel cards to an RS/6000 model 550 with 400 MB of SCSI disk. Performance tests had shown bandwidths between memory and tape of 2 Mbytes per second and between disk and tape of 1.2 Mbytes per second for a single transfer. When two tapes were driven the rates halved.

The BASTA project will be enhanced by the addition of two other components, Sioux and Anastasie, to provide interactive and analysis services that require better I/O throughput.

3.2 Fermilab - Judy Nicholls

Fermilab have eight people supporting 650 machines. The bulk of the production physics work has migrated to Unix; the next challenge is to convince the general users to move from VMS.

Both fixed target and collider experiments are making heavy use of Fermilab facilities, both UNIX farms and VAX clusters. The Amdahl is currently being run down and Unix farms are dominating the cpu time delivered, a total of 5000 VUPS delivered today out of 9000 VUPS available.

More general purpose UNIX services are to be provided by the FNALU and CLUBS systems. The CLUBS (Clustered Large UNIX Batch System) service delivers batch using IBM's LoadLeveler software. An analysis server based on an IBM SP1 was now in service. The STK robot driven by the Amdahl at present will soon be switched to two RS/6000 model 580s and data will be delivered to the batch servers by Ultranet.

FNALU is a UNIX interactive service based on AFS. Moving the home directories of CLUBS users to FNALU and serving the home directories to CLUBS using IBM's AFS translator was not reliable initially.

3.3 DESY - Karsten Kuenne

Current services include an Apollo cluster (frozen and gradually being phased out), 17 Hewlett Packard 730s, seven of which will upgrade to 735s soon, with 70 Gbytes of disk, a Silicon Graphics farm containing seven new challenge machines - a total of 84 cpus and 4.5 Gbytes of memory. The SGI machines were chosen because of their ability to pack several cpus in a box and because of their large number of SCSI controllers. This year, DESY has installed two Ampex D-2 tape library robots each with three drives and 256 25 Gbyte tapes. The robots have an NFS front end. They are now working satisfactorily after initial problems. During trials the head design was changed, before that they needed cleaning every day. The goal is to write raw data directly to Ampex tapes. It is planned to include an interactive service using part of this system.

Two hundred X terminals had been added making a total of 600. All X terminals are controlled by XDM. Most use a chooser, usually a built-in Tektronix or NCD one. Authorisation is by magic-cookie with special scripts xrsh and xrlogin to pass this onto remote machines. Hamburg has more NCD than Tektronix X terminals but Zeuthen has Tektronix only.

The so-called HEPiX scripts are in use at both Hamburg and Zeuthen and although they differ, it is hoped to resolve this soon and use a single set. Versions for tcsh and csh have just been finished and are under test. A set of supported tools have been defined, they include vi, pico, emacs, pine, elm, mh, www, news, xrn, xmosaic, tin. Reference cards are provided for all of these. Backup on SGI currently uses Legato Networker. The software has been fine but the Exabytes have failed, especially the 8500 stackers. Backup will soon be changed to use ADSM (formerly DFDSM) in MVS to write to the STK robots.

Problems encountered have included the performance of the Hewlett Packard systems, NFS and FDDI are not much better than ethernet

and HFS. Vendor supplied automounters have given difficulties and the management of the increasing number of X terminals can cause problems.

NFS is used to cross mount users' files but AFS is being considered for next year. It is planned to replace the mainframe by a smaller machine next year. DESY currently use Ultrinet but are considering moving away from it because of doubts about future vendor support after the recent take-over of the original supplier.

3.4 DAPNIA - P.Micout

It is planned to use IN2P3's central computer services next year. The Cray X/MP-128 will be stopped in February 1994, the 116 (in Marseille) in January 1994. The Cray II in Grenoble will be replaced by a Cray C94 with a T3D (128 Alphas). Central CPU server at DAPNIA will be based on an IBM SP1 (16 RS/6000 CPUs) + a file server.

The nuclear physics community has 100 Suns + 1 Sparccenter 2000 for Nuclear Physics. Particle physics has a mixture of DECstations 5000, RS/6000s, Alphas running VMS, Vaxes and X terminals.

3.5 CERN - A.Silverman

AFS is becoming strategic for home directories and ASIS. One experiment, Chorus, is trying out the workgroup concept based on AFS. The increased use of X terminals, contractor assistance for updates and plans for standard configurations should reduce the work required for systems administration. In the discussion, several sites were interested in a possible offer to install a single AFS client licence as part of the CERN "cell" in order to access the new ASIS (in which the master copy will be on AFS). What side effects will this cause?

Currently, there are 1400 workstations on site with central support for 7 platforms. WWW, Unix user guides and installation guides are being used to reduce the number of queries.

Central services include Crack for passwords, backup using ADSM and AFS home directories services. DFS tests have been made with IBM; HP and DEC are expected soon. For systems management, Tivoli, Patrol and FullSail have also been evaluated, with Tivoli the most promising but it is very expensive.

The challenges are providing support for a large range of platforms,

buying in contract assistance if necessary. Standard profiles and utilities must be developed to reduce the variety of systems. New SunOS server licensing rules are also causing concern.

3.6 KEK - T.Sasaki

Central computing based around a Hitachi mainframe, a farm of 11 Sparc stations and a Sony DD-1 tape drive (1Tb/volume!) plus a super-computer, several MPP systems and lots of VMS and UNIX stations. Projects include Root, a support service for workstation administrators, and Kiwi, an NFS based file server from Auspex.

3.7 RAL - B.Saunders

The central Unix service is based on 5 Alphas running OSF/1 today and providing interactive and batch services for the whole community and a simulation facility based on 6 HP 735s. RS/6000s provide tape and file services and act as a front-end to the Cray. It is planned to cease VM by about April 1994, replacing it by a large VMS system and a CSF-type farm. The file server is an RS/6000 with 27Gb of disk. Each file system has its own name defined in the Name Server which should allow easier migration in future.

A clone of ASIS is used but RAL misses the EPIP function under OSF/1 to install public domain software on client stations.

The central tape services are based on the RAL virtual tape protocol, 6.5Tb in total, staging through the 3090 and RS/6000s. The HEPiX profiles and a standard user environment have been established. Networking with SuperJANET has given 140Mbps between 6 sites and an ATM pilot using real-time video has been established.

3.8 NIKHEF - W.van Leeuwen

Central computing facilities are based around a Sun 4/690. There are Apollos, DEC, HP, IBM, Next, SGIs and Suns on site. At the end of 1995, the SARA VM mainframe service will be stopped and the Robot will be managed by the Cray or RS/6000. Future projects include video conferencing and AFS.

3.9 CASPUR - N.Sanna

An IBM 3090 and an 8 way Alpha Farm using OSF/1 and a Gigswitch provide the central computing. The Utopia batch system will be used for scheduling batch and interactive work around the farm. Some parallel applications using PVM are being developed. The APE project, a parallel system, is being developed at CASPUR. It currently has 8 CPUs moving to 128 CPUs next year. The 3090 is being used a front end to ADSM, nstage, and the accounting database.

3.10 GSI - M.Dahlinger

Main CPU systems are an IBM 3090 600J, RS/6000s, HP 720s and a Memorex robot. The standard workplace uses PCs and X terminals to access these resources. The MVS service will be downsized, but not completely removed, and the VMS Alpha and IBM RS/6000 clusters will be increased to compensate. Standardising on KornShell, LoadLeveler for batch and backup using ADSM.

Channel connections between the RS/6000s and the mainframe, running MVS, give problems and poor performance.

3.11 CINECA - C.Bassi

There is an HP cluster based on the CSF configuration from CERN. NQS, RFIO and RTCOPY from the SHIFT software package are all used. Currently tape serving is performed by the Cray but there is a concern about future software support now that there is no Cray installed at CERN. This is a new field for CINECA and is currently in test; they hope to start production soon.

4 X11 at DESY (Administration aspects) - Thomas Finnern, DESY

Th. Finnern gave numbers for the supported systems at the DESY computer centre. Out of 600 Xterminals, 435 are supported centrally and of 250 Workstations, 65 are supported centrally. X support has to be provided for a mixture of X Terminals from NCD and Tektronix and various workstation consoles. For an X terminal the following resources should be present (depending on the usage 50 to 1000 percent of these numbers might be needed):

- 3 MB of graphics memory
- 2 MB of I/O memory
- 1 MB of local client memory
- 5 MB of session server memory
- 5 MIPS of the session server CPU
- 100 kbit/s net bandwidth

X stations use the following protocols: bootp, tftp/NFS, xdmcp and snmp. To establish an application on an X terminal several servers come into action, which all can reside on the same host or be duplicated for enhanced reliability (boot, configuration font login chooser, login session and application server(s)). For security reasons the MIT magic cookie protocol should be used where possible (not available on DEC/VMS) instead of the xhost-based authorisation. Experience with the different vendors has shown that not all tasks can be performed on all terminals (example: xdm host list is not configurable for login chooser on XP 20 terminals from Tektronix).

5 Control Host, A UNIX-based central commander/monitor host - A. Maslennikov, INFN Roma/CASPUR

A. Maslennikov reported on a tool that allows to control and monitor processes on remote hosts. A typical application is in the field of slow control of experiments and interaction with real time systems such as for example OS/9. The tool is free of commercial software, easily configurable and can be tailored to further HEP applications.

The GUI is based on Tk/Tcl/Xf. Graphics is based on CERN standards like HIGZ and PAW. The interprocess communication goes via tcp sockets and shared memory (local communication). The package consists of a main dispatcher process that handles communications and data transfers, several data driven client processes and a library of action programs, utilities to access and transfer data and to configure the dispatcher.

The achieved rate of handling events with a setup consisting of an RS6000 and an OS/9 system was up to 5 Hz. A beta version will be available from CONHOST@ITCASPUR.CASPUR.IT by mid november. A final version including a user manual is expected after Feb '94 at CERN.

6 Experiences with automounters in a multi-vendor environment - A.Koehler, DESY

Automounters are used mainly to simplify the system administration, to avoid stability problems with crossmounted filesystems in the net and to introduce naming schemes of mounted filesystems. Sysadmins can benefit from additional tools provided with the automounters and from the built-in NIS support.

In this talk an introduction to the principles of automounter operation and configuration was given. A discussion of vendor-supplied automounters was followed by a summary of their disadvantages:

- different sets of supported automounter maps
- replicated filesystems not supported by all vendors
- update of the maps causes problems (reboot)
- overall stability in a heterogeneous environment is poor.

The decision at Zeuthen was therefore to use the unofficial (enhanced, public domain) version amd of the automounter which is available from ftp.cs.columbia.edu. Advantages include the many additional features such as an extended syntax for the automounter maps and the support of additional filesystem types.

Examples for the usages of amd were given including

- home directories mounted locally, via Ethernet and FDDI
- replicated servers for a filesystem
- architecture sharing of parts of a filesystem (e.g.fonts)

The speaker also gave some examples of amd administrations and he concluded that the scheme used at Zeuthen was stable and the administration effort for mounted filesystems could be reduced. Some of the functionality of AFS could already be achieved by properly using amd.

7 Experiences and Conclusions with the CHORUS Group Server - Tony Cass, CERN

The CHORUS group server, a concept for a modern, Unix- and X-based environment for a new collaboration, was presented. It is based on the

assumption that different functions can be devoted to different, relatively small dedicated workstations. It was realized with 2 RS/6000 which work as process and file server respectively. Tony emphasized the crucial role of AFS as an institutional file system for the project.

Still open questions are the lack of a well developed user setup, especially in the AFS environment, poor X-support and other configuration issues requiring manual intervention on individual machines.

At the end he touched on some problems with mail agents currently in use and the lack of useful desktop utilities.

As a conclusion he rated the concept as generally viewed successful with the consequence that interactive services could be moved off CERNVM in 1995 with sufficient investment in manpower and resources in 1994.

8 Product management at DESY R2 - Thomas Finnern, DESY

Thomas presented the work of DESY on a tool for distributing products easily on various workstations and platforms, called SALAD. It allows to copy the product from a reference machine or distribution via tape or ftp. It is available for all major Unix flavors.

SALAD recognizes automatically the appropriate binary type. Binary classes with different levels of subdivisions depending on operating system, release and version levels as well as additional hardware requirements are taken into account.

The product description is documented in a short and easy customizable form and is distributed together with the product itself. This so-called salad-card controls then the automated installation procedure.

9 Unix at Pisa - Maurizio Davini, SNS Pisa

In 1992 the Physics Department purchased an RS/6000 model 950 to replace an IBM 9370 system running VM/CMS. This system supports individual logins (from Xterminals, PCs and Macs) and also acts as a central server for the clusters in the various physics groups (theory and astrophysics for example). Overall, the Physics Department has about 25 workstations, 15 Xterminals 70 Macs and 30 PCs.

Besides the use of (NSCA) telnet, Mac users can also use CAP software, developed at Columbia University, to access the 950. This provides a Mac-like front end access to Unix through the System 7 AppleShare facility and also allows the RS/6000 to act as a repository for Mac software.

Elm and Pine are the supported mail agents for Unix users. POP is supported for PC and Mac users with Eudora as the Mac mail agent.

Gopher and especially WWW are widely used for information access with the latter providing access to CERN manuals. MacMosaic is a good WWW browser, but really requires a powerful Mac.

Future Unix development at the Physics Department will focus on AFS (in collaboration with SNS and INFN Pisa), IBM's LoadLeveler product, the introduction of a CSF service on an HP cluster and investigations into parallel processing, especially using PVM.

There was a brief presentation of the University of Pisa Network Service (SIERRA). The hardware aspects (a chief goal being the linking of all University departments and ISDN testing) were covered briefly, the focus was more on the services provided. Here the major work has been carried out in the information access area with News, Gopher, WWW, Archie and Veronica servers all installed - the latter being the first Veronica server in Europe. WWW is seen as a key technology for the future and there is a project together with CNUCE-CNR to develop WWW servers for museum guides and manuscript access.

Finally Italo Lisi covered the situation at the SNS. The current VM and VMS users will be moved to an upgraded Central Unix service as quickly as possible. Together with the University and the Department of Physics they will be exploring MAN technology to improve the communication infrastructure between the various sites.

10 A new model for CERN computing services in the post-mainframe era - Chris Jones, CERN

This presentation fell into two distinct parts. Firstly an overview of the post-mainframe CERN computing strategy based on foils provided by David Williams and secondly an overview of the strategy for the provision of interactive (desktop) services. Batch CPU capacity will be provided through an extension of the existing CORE services at CERN. It was assumed that most people present knew about the CORE architecture and Les Robertson would cover CORE status in the afternoon

as part of the joint HEPiX/HEPVMx batch meeting.

a) Overall Strategy

Basically the message is that CERN has to move away from the current situation, developed over the past 20 or so years, where a central mainframe provides a natural integration of all computing services. In particular it is clear that the CERNVM service cannot continue for more than 5 years or so.

The migration strategy planned is, unsurprisingly, to move to a distributed computing model which, when fully developed, should provide better services for the same money. The move will not save money and will even involve greater expenditure in the short term. CN division is aware of the problems of switching the computing model for LEP experiments (it is assumed that other existing experiments will largely end before VM goes and that new experiments will build up their computing environment in the new structure) and it is clear that CERNVM must be maintained until replacement services are available.

A replacement interactive service is a key element - CORE already provides a solid foundation for future batch services - but will not be easy as so much depends on the "look and feel" of the environment rather than the basic system hardware. Fortunately a lot of the required software technology is appearing at the moment (AFS, DCE, COSE) but there will also be a much increased reliance on the CERN internal network. In fact, CERN will need to rebuild the internal network, moving to a structured wiring based system as the existing cheapernet network is at the limits of its capacity and manageability.

A particular problem for CERN in trying to provide a unified interactive environment will be the lack of ultimate control over the hardware and software platforms to be supported - it is assumed that CERN will have to support systems chosen by institutions in the member states. As such, it is hoped that HEPiX could perhaps provide some help to reduce the number of hardware/software combinations.

Between now and the end of 1994 (when the current IBM 9021-900 lease expires) CN plans to move most of the batch load off CERNVM onto an expanded CORE service and develop an attractive Unix interactive service. In 1995/1996 CERN will acquire a 390 mainframe of sufficient capacity to cope with the interactive load which will then be moved to other services. If necessary a "rump" VM service will be provided in 1997. CN division will be reorganised to cope with the focus on batch and interactive computing by creating two new groups - led by Les Robertson and Chris Jones respectively - out of the old System Software and Consultancy and Operations groups.

Hans Klein raised the issue of the staff needed to help move LEP computing environments - he felt that such people would need a good knowledge of CERNVM and not just of Unix. Chris replied that this was true, but that it was also likely that many "production control" systems would be rewritten rather than ported - L3 are already planning to do just this to be ready for the start of LEP200.

b) Interactive Strategy

The basic strategy for interactive services, on whatever platform, is to minimise the "personality" of any individual platform. Both PCs and Workstations should have as few as possible locally-altered files. Standard configurations and networked services provide the best means for a relatively small number of people to manage a large number of boxes; CN staff will manage services (home directory services, software distribution services) and not systems.

The main focus for interactive services will be the basic Mail, News access and document preparation type of facilities common to all users. Today nearly 3/4 of CERNVM users do little else and use CERNVM only because there is no CN provided alternative. However, CN also has to take account of the needs of physicist users which include the requirement for rapid turnround for program testing. Thus CN must provide sufficient CPU capacity and also data access as part of the interactive environment.

It is imagined that there will be two interactive services - one Windows and Novell-based for PC (and Macintosh) users and one Unix-based. The PC service will be the natural further development of the NICE service which provides an integrated environment for at least 1000 PCs. In particular, NICE provides today a standard installation facility (DIANE) with which a newly-purchased PC can be setup by booting from a single diskette. However, NICE does not yet provide a particularly solid home directory service and this will be the focus for improvement.

The Unix-based interactive service is not in such an advanced state, but will be developed from the experience gained in providing a Unix environment for the CHORUS experiment. Here CN started to build a server for CHORUS but, through the availability of AFS and the early realisation of the problems of non-standard system configurations, ended up concentrating on the services that were needed. Thus the aim is to move towards a NICE-like system with "shrink-wrapped" installation procedures and standardised environments. To do this, we will build on AFS (and later DFS) as the fundamental technology which makes distributed home directory and software services a practical proposition. In addition to CN-provided, experiment-based and publically-available Unix servers, we hope this architecture will extend to privately-owned workstations and to X terminals for which CN will provide the necessary

support services, e.g. boot servers.

In moving towards a standardised Unix environment, CN will rely heavily on external sources of standardised software, notably the HEPiX work on standard user profiles and the COSE work on a common desktop environment.

Given that VM-based interactive capacity will be available until 1996, there is less immediate pressure for a fully-fledged interactive service than for a batch service. However, CN has to be in a position to start moving VM users to a better interactive environment in 1995 and there will be much work required in 1994 to ensure this will be possible.

c) Discussion

John Gordon - ATLAS want an HP-based Unix system at CERN and an OSF/Ultrix-based system at RAL - what do they really want?

Chris Jones: They want a good interactive service rather than a service on the same architecture as the CPU service. There is some concern about whether or not HP systems can provide a good interactive service but this has to be tested.

Hans Klein - Likes the description of the Unix service but wonders why it is being restricted to ATLAS and CMS.

Chris Jones - Basically we need to stop these new experiments getting entrenched in a (soon to be) outdated environment and the limited number of people available prevents us from doing more.

HK - But Delphi are currently moving to Unix, we don't want to go off in the wrong direction here either.

CJ - Point taken.

John Gordon - If batch jobs are moved off CERNVM in 1994 will the CPU be idle at night? And RAL experience is that it is not trivial to move off batch work from individuals as opposed to physics production jobs.

Chris Jones - Exact position still unclear but we do imagine that there will be unused capacity and we also see that moving production jobs will be much easier than moving private analysis jobs.

11 Status of The HEPIX common user environment project - W. Friebel, DESY

The talk focused on 5 items:

1. Components of the environment:

- Startup
- Keybindings for the various shells and utilities, e.g. emacs, the less browser
- X11-environment; X11 usage was standardised by providing various scripts and setup files and a standard chooser
- A Common pool of programs such as emacs 19, elm and pine, less, etc.
- SysAdmin-Support.

The speaker made some remarks on realization and the advantages versus disadvantages of using commercial tools or locally-provided tools or a widely-available set of common tools.

2. Benefits for users of the DESY scripts:

- Same environment everywhere
- not necessarily supported locally, but present and available

3. Usage summary:

- Used at DESY Hamburg + Zeuthen
- Tested in: RAL, CERN, GSI Darmstadt, Dortmund
- 7 reports of experience, about 40 fetches of "profiles" + "key-bindings" during May - September '93

4. Known problems:

- scripts too long
- overhead due to conditional code needed for multi-platform coding
- missing documentation
- incomplete installation instructions
- no complete description explaining certain choices
- overall packaging could be improved

5. Future of the project

- reduction in size of scripts by removing site-specific parts; perhaps a method to produce production scripts for one architecture/site from a large master script.
- binaries for e.g. tcsh, zsh, elm, emacs, ...

- more documentation
- improved packaging and installation scripts; perhaps a make file.

In his summary, Wolfgang pointed out:

- Benefit for current users is small, but huge for newcomers
- The installation is still difficult, but will be improved
- The number of installations using the common environment will increase by having simpler scripts, better documentation, better packaging and closer cooperation with other sites.

12 To a LHC computing model - Willem van Leeuwen, NIKHEF

In his talk, Willem, drew attention to the differences between LHC versus LEP in the amount of data, the required computing power and the greater number of people and institutions involved. He commented on:

- Basic concepts such as distributed computing, transparent access to data from the desktop, the need for a uniform environment and platform independence. It is be important to use industry standard h/w and s/w and to provide adequate network bandwidth.
- ATLAS strategy and ATLAS computing infrastructure
- Activities outside CERN: Participation in pilot projects, Monte Carlo simulations, Mail services and WWW services
- Global Data Access using WWW

A final comment was on "some change in the routing of IP network traffic" which made interactive work between NIKHEF and DESY almost impossible; the careful handling of such undertaking is very important.

13 Are You worried about the DATA flow and storage for LHC - Hans Joerg Klein, CERN

In a reply to the previous talk, Hans Joerg identified some of the daily problems with data and data processing in and outside CERN concerning LEP.

14 Wrap-Up of HEPIX - Alan Silverman, CERN

Alan covered the following points -

- availability of minutes (scheduled for publication in November if possible); they would be made available via anonymous ftp at Pisa, HEPiX news and mail, and WWW.
- ideas for further seminars along the lines of the PERL and TK/TCL seminars given this week. Future topics could include AFS or Emacs 19. Suggestions and offers to present something should be addressed to him.
- a document directory currently being established by Judy Richards. She was expected to announce this shortly. The primary access would be via WWW. She will gladly accept references to documents which people wish to see publically available.
- a Tools Data Base: Alan would be discussing this further at the forthcoming US HEPiX meeting.
- next meetings. These include US HEPIX on Oct 27-29 at SLAC, a HEPiX "World" meeting to run adjacent to CHEP 94 scheduled on April 20-27 at San Francisco, HEPIX "Europe" autumn '94 together with HEPMiX. [After the meeting, DAPNIA kindly offered to host the next European meeting at Saclay sometime in October 1994, exact date to be fixed later.]

15 Acknowledgements

We are greatly indebted to the staff and management of SNS Pisa for welcoming us and organising the meeting with such efficiency. Thanks are especially due to Prof Emilio Picasso, Director of SNS and Mario Soldi and his staff.

We must also thank the speakers of all the sessions throughout the week for agreeing to present their work and share their experiences. We especially thank Lionel Cons who spent the whole of the first day on his feet presenting the two half-day seminars.

On a personal note, I would like to thank the attendees who agreed to be volunteered to be chairmen (chairwoman in one case) and minute takers. The details above are the results of their efforts. Any errors are probably the result of my editing.

Alan Silverman (Editor)

2 December 1993