

HEPiX Prague Meeting Report

Alan Silverman

September 6, 1995

The meeting was opened by Prof.V.Dvorak, director of the Institute of Physics in Prague who welcomed the participants. The main direction of their Institute was solid state physics but there was a sub-group concerned with high energy physics led by M.Lokajicek. The Institute was the largest of some 60 similar Institutes in the Academy of Sciences of the Czech Republic with some 200 researchers.

Many of the overheads presented at the meeting are available on the World-Wide Web at URL – <http://www-hep.fzu.cz/computing/HEPiX/HEPiX95/Overhead.html>.

1 Site Reports

1.1 Prague Institutes - M.Lokajicek, J.Hrivnac

The user community of the high energy physics section in the Institute included also physicists working in Charles University and in the Czech Technical University, both also in Prague; this came to some 70 physicists in all. They began using small mainframes in the 80s, moving gradually to PCs and workstations of which they now had about 15 in total. They also had some access to a Cray Y/MP locally and a 4 CPU Power Challenge belonging to the Academy of Sciences but these were already approaching full capacity. They were trying to develop a distributed farm in collaboration with HP.

They had been using Internet since 1992 with some 28 sites and 1200 nodes on the country's internal CESNET network. They had international links to Europanet in Amsterdam, to EBONE in Vienna and to Slovakia. There was an ATM backbone pilot setup in Prague.

The physicists were members of DELPHI, OMEGA and H1 as well ATLAS and ALICE for LHC. There were also some theorists.

In developing further all of the above, funding was a most serious problem, closely followed by a scarcity of skilled system administrators.

Since the Saclay HEPiX meeting, they had added 2 more HP nodes and some disc space. They still used local login scripts but were trying to merge in the use of the HEPiX login scripts; this was

complicated by the need to share the nodes with non-HEP users. They will follow the HEPiX working group on X11.

There had been more use of MBONE but bandwidth was still limited. They hoped to acquire an SGI station to send video images. They were planning to migrate file backup to using OMNIBACK. They mirrored ASIS from CERN and hoped to implement AFS soon; they also made use of CERN's ZEPHYR service.

1.2 DESY - K.Kuene

DESY had decided to move to an interactive UNIX-based computer model similar to that adopted at CERN, based on AFS and with specialised servers for different tasks. There will be originally 3 home directory servers, 3 binary servers and 9 work group servers with at first a total of some 60 registered users.

For permitting sudo use they will use an extension of the AFS arc procedure developed by R.Tobricke and Loadleveler with his token-extension scheme. They have seen an AFS crash which resulted in some data corruption. The new HERMES experiment will be the first at DESY to be uniquely UNIX-based.

File backup will be a mixture of ADSM to their mainframe and AFS volume backup. STK access was limited to their Challenge systems and Legent OSM was under investigation for file management.

Parallel computing was being concentrated at Zeuthen which had 2 Quadrics, a 10 node SP2 and a 22 node Challenge. A 2 Mbit connection has been installed.

Future plans included installing a data server based on Legent OSM, using DCE for user registration, deploying a UNIX-based mail server and investigating the centralisation of UNIX management.

1.3 FNAL - M.Wicks

Later in the conference, J.Nicholls would summarise the newly-announced Fermilab plans to migrate away from VMS, to UNIX and PCs. Major issues arising out of this included

- a UNIX mail strategy (see later talk)
- information distribution (news, WWW, etc)
- distributed file system technology, including DCE/DFS
- access to PC applications for the casual user
- diversity of licence management software

Digital UNIX (OSF/1) has been officially added to the list of supported architectures and although HP-UX 9 is not, they will review the situation for HP-UX 10. AFS-controlled disc space was expanding quickly. Fermilab were looking at possible successors to NSL Unitree. CDF were planning to expand their SGI Challenge to 16 processors and D0 were also moving more to Challenge.

1.4 SLAC - C.Boeheim

For their tape servers, SLAC had installed SHIFT software for the staging. For future tape technology, they were investigating the so-called Redwood drive from STK.

AFS use has continued to grow, with 84 registered users of whom over 30 used it for their home directories. They offered systcl for users to set up their own directories.

Load on SLACVM was expected to decrease as future experiments moved their users off towards VMS and UNIX, starting in the autumn.

1.5 CERN - J.Richards

The steady growth in workstation population had continued, except for Apollos and ULTRIX systems which were now being gradually phased out, with particularly fast growth in the X terminal numbers. For X terminals, two central boot and font servers had been established and the first remote instance of this service also. CERN was investigating a suitable tool to provide casual access to PC applications from UNIX but so far without finding anything with sufficient power which offered access to its installed NETWARE-based PC servers.

The AFS service had also increased but they had been seriously hindered by the late ports to the latest operating system releases. Nevertheless, there was now more than 300GB of AFS disc space, more than 1800 users and 570 client nodes.

Another evaluation was of GUIs with parallel investigations into COSE/CDE from Triton for various platforms and a home-built simpler interface called ADAM (to be covered later by M.Marquina). The HEPiX login scripts were heavily used as was ISS, the interactive load balancing scheme from Loadleveler.

There were two public login UNIX services (PLUS), the SP2 and one based on HP; and a SUN-based work group server had been established for CMS, adding to the RS/6000 and HP based WGS already running for other experiments.

On the batch side, CERNVM had been halved in size at Christmas and would be further reduced next January, much of the load transferring to the SP2 and the CS-2 (later talk by F.Gagliardi). Disc capacity for batch was increasing rapidly and on the tape service, the DLT service was expanding and CERN was looking at IBM's NTP device.

1.6 DAPNIA - P.Micout

They made use of various workstation farms at CCIN2P3 in Lyon, including BASTA, ANASTASIE and DATASERV. Locally they have seen a large rise in the use of X Terminals. They were evaluating EPOCH on SUN with a DLT stacker. Their main VMS node has been replaced by a DEC ALPHA system running VMS.

DAPNIA had created a DEC OSF/1 cluster consisting of 1 file server and 4 CPU servers using a home-built load balancing scheme based on the use of Xdm, plus 1 node for Internet integration. However, they had discovered that Digital's ADvFS (Advanced File System) was not as reliable as might be hoped and made a plea for AFS for OSF version 3.

1.7 RAL - J.Gordon

Twenty more nodes, model 712/80s, had been added to the HP farm, making a total of 26 nodes; it was being used for physics analysis and Monte Carlo work. The evolution of the configuration was shown demonstrating the heavy use made of NFS.

They now had their own AFS server with clients in UK universities and they were moving their mail name resolver from VM to UNIX, adding more functionality in passing such as a WWW interface and a socket interface for remote sendmail.

RAL had completed trials of Digital's new Turbolaser processor and had been very impressed, a view shared by other labs represented who had had access to similar systems. Performance figures were presented comparing it to a Cray Y/MP.

1.8 NIKHEF - W.van Leeuwen

The central group were about to take over the running of the systems for NIKHEF-K. The team were currently revising their public domain multi-media tools for most of the major platforms since use of multi-media was growing inside NIKHEF and abroad. They were using a video server to replay recorded sessions.

1.9 CASPUR - A.Maslennikov

The CASPUR configuration was described; LSF was the batch scheme in use on most systems but Loadleveler was also present on some. AFS use continued to expand with CASPUR acting as server for many clients across Italy spread over multiple AFS cells. More and more home directories were migrating to AFS. Their Quadrics systems were used for SIMD work, their SP2 for MIMD. For databases, they relied on DB2.

They collaborated with the CHORUS and NA48 experiments in the realm of high energy physics.

Current work included development of distributed management tools, further downsizing of their mainframe and revising their network configuration.

1.10 CCIN2P3 - J.O'Neall

Their latest workstation farm, DATASERV, to be used for storage and access of data, was in the process of being built up and production analysis was getting under way on the ANASTASIE farm described in previous HEPiX meetings. All the farms were linked by ATM and the AFS-based disc farm by two ATM paths, effectively multi-homed AFS.

There were problems however with the RS/6000 ATM driver and IN2P3 were concerned by the lack of standards regarding broadcasts in ATM. Nevertheless, they had measured up to 10 MBps per circuit in memory to memory transfers.

They had developed a scheme called BAHIA for spreading logins among members of farms, a home-built equivalent of the functions offered by LSF. A 16 node SP2 had been delivered of which 2 nodes would be for login, 6 for X servers and the rest dedicated to other uses.

1.11 KEK - T.Sasaki

Plans had been made to replace their existing Hitachi mainframe with distributed systems by next January, their Fujitsu supercomputer having already been replaced this year. As part of the migration, a UNIX farm had been installed, a mixture of HP and Hitachi workstations. The Tristan experiment had ended and their mainframe would be replaced by a system for the forthcoming KEK-B system in 1-2 years time.

KEK had over 400 UNIX workstations now and a site-wide support contract for them with Fujitsu and HP. Software installation made use of central file servers and user training was a major issue. Plans for the rest of this year included tests with ATM, including an ATM wide-area network linking KEK with some of its collaborating institutes in other parts of Japan.

2 Security

2.1 Security and Worldwide Collaborations - C.Boeheim

It was a fact of life these days that Internet-based attacks were becoming ever more frequent and more sophisticated. There were more tools available for hackers and some even contained "cloaks" such that once accessed was gained detection remained difficult. HEP Labs were particularly vulnerable to attack because of their requirement to be open to worldwide collaborations with large shifting user populations.

Security was labour-intensive with few good vendor-supplier tools. However the speaker had come - the hard way - to learn that break-ins were more expensive in the long run. It is vital to have supportive users but these users require a working atmosphere which allows them to get their work done with minimal overhead.

The speaker gave a list of tools which could be called upon including AFS, Kerberos and X Proxy. The most vital areas of security were X-11 and file access.

2.2 Security on X - L.Cons

The speaker noted that both UNIX and X11 were intrinsically vulnerable, there was a long list of potential threats and it was hard to trace the source of any attack in many cases. Among the most common X11 protection schemes were host-control methods (for example using xhost) and three varieties of authorisation protocols (xauth) -

1. the so-called MIT Magic Cookie
2. XDM authorisation using secure DES security - available in the USA only
3. the SUN-DES-1 scheme, based on SUN's RPC protocol

However, all had their own drawbacks and the speaker compared them.

He presented mxconns, a public domain tool used to protect an X display from intrusion. It could be integrated into xdm or xrsh to increase its user friendliness. It solved a number of problems, was extensible and was able to work in the presence of firewalls to pass X11 traffic; however, it was not a universal cure and some problems remained.

2.3 Security Interest Group

In the discussion following these two sessions, the suggestion was raised to inaugurate a working group on security matters. Its initial goals would be to provide a forum where security managers, and system administrators in general, from the different HEP sites could share experiences. Another spinoff could be to generate security presentations at future HEPiX meetings. Later in the course of the conference, the idea of a working group was refined more to an interest group since it was realised that such a group would not in general have as a goal the production of software or products but rather a sharing of information. It was left to the HEPiX Board to promote the formation of such a group.

3 Tools

3.1 Fermitools - J.Nicholls

This was a scheme under which Fermilab distributes software packages to outside collaborators while still keeping to rules established by the US DOE for such software distribution. Packages available under this scheme must adhere to certain rules; there was no commitment by FNAL for support to outside users; and there was a heavy emphasis on providing documentation. The list of software available included CPS (Cooperative Processing Software), Nedit, juke box robotic control software and also some PC software; more tools were in preparation. Significant interest had already been shown in the scheme, especially for Nedit and further information was available from Ruth Pordes, e-mail Ruth@FNAL.GOV.

3.2 NETSCAPE

This session took the form of an open discussion led by M.Marquina of CERN and concentrated on the new popular web browser NETSCAPE. While early versions were freely available, the latest 1.1 release seemed to have subtly changed the licensing conditions and it was no longer clear if it could be used freely in HEP sites like CERN, SLAC, etc. A poll was taken of the different sites represented and a variety of opinions presented. Some sites would purchase a few licences while others would prefer to negotiate with the supplier on terms for the current and future versions while trying to restrict their users to accessing version 1.0. No clear consensus emerged.

3.3 Controlhost - A.Maslenikov

This had been described at the Pisa meeting; its goals were to provide a method of organising data and message traffic in a complex multi-host environment. The objects and data flows were presented. After production tests with the CHORUS experiment at CERN on a VME/OS9-based data acquisition system during 1994 and 1995, the product had now been released. NA48 were interested in the product and at least one group was evaluating it as a possible tool to build a GUI to control their distributed computing network. For anyone interested, it was freely available from the CASPUR centre and limited setup advice was available.

3.4 gman - a Unified Interface to man Pages - D.Pop

The need for such a tool was justified by the inconsistencies in options and arguments of the man command and in the search order of the man pages as distributed by the different vendors. The gman command from the man_db package currently under development had been adopted at CERN and a number of observed problems fixed (and reported back to the man_db authors). The result could be found on the ASIS server.

4 Data Storage

4.1 The ATLAS Data Server - D.Rigby

The speaker began by explaining that ATLAS in the title referred to the name of the computing laboratory at RAL, not to the LHC collaboration.

The store was created to satisfy the need for distributed access to central data according to the new model of distributed computing recently drawn up at RAL. It was felt important to decouple the users' applications from the data storage mechanism with a well-defined interface. This would allow applications to remain the same while benefitting from any new interesting storage technology which became available. The original plan called for 30TB of managed storage with 30 MBps access speeds to be developed over three years. Today there was some 15TB.

The user interface was defined via a single catalogue of some 272K datasets which all have associated checksums for added security; the speaker pointed out that adding checksums did not in itself guarantee improved integrity but at least changes would be spotted and they had found several instances of such changes which would have gone unnoticed without this extra protection.

As well as experimental and Monte Carlo tape data, the store was used for file backup, VM minidisc store and the migration of files from the mainframes. There was a uniform hierarchy consisting of robotic devices, manual devices and offsite storage. Since establishing the store, they had implemented in turn IBM 3494 and STK 4400 robots and DAT units and were planning a move to IBM's NTP technology (announced as "Magstar"), all without changing the users' interface.

The datastore contained 250K datasets totalling 15TB, with some 4TB held in 2 robots, an IBM 3494 and an STK 4400, and the rest on some 30K DATs. Weekly throughput was around 5000 files, 500 GB. Future plans called for a move to Magstar, to develop an FTP interface and to permit easy import/export of data, for example DLTs from CERN, including exploding the data to the original CERN datasets conserving their original tape IDs. The major challenge facing the project was to provide intelligent indexing and retrieval.

4.2 Storage Strategies at CERN - L.Robertson

The speaker pointed out that the quantity of data to be dealt with was awesome; LEP experiments were producing 1TB per year, NA48 would produce 10TB this year alone and NA49 would add 10TB per year from 1997. The favoured medium until recently had been 3480/3490 cartridges, stored in robots and linked by a variety of ESCON channel, IBM parallel channels and SCSI. DLTs were currently gaining in popularity, again in robots, and Exabyte was still present. The central CORE facility currently mounted some 29K volumes per week while the effective maximum was expected to be around 36K mounts with the equipment and funding currently available.

The long term strategy was

- to keep active data online, a situation they were steadily approaching
- to move to system-managed storage; they were evaluating the Legent OSM package and hoped to begin a pilot test this year
- to introduce high capacity tapes to cut media costs
- to minimise manual operations since robot mounting was much cheaper.

They were looking at

- Digital DLTs: good experience but still only low statistics; a few broken tape incidents
- IBM's 3590 (also known as NTP and Magstar): now announced and in test in many labs which all report good experience
- STK Redwood: not yet officially announced

Although DLTs were cheap devices, 3590 drives in a 3494 or 3495 robot were cheaper per GB of storage; the costs per mount were less clear, it depended on how much data would be read per mount, for example a 200 MB file or a 1-2 GB file.

Priorities were to reduce costs (more use of robots and fewer manual mounts) and to replace Exabytes. It was important to distinguish between open shop and closed shop tapes, i.e. those used to import/export raw data and those used for DST storage.

The speaker was in favour of the 10GB TZ87 DLT variety but he thought the DLT 4000 was too large and should be discouraged. His philosophy was to support any reasonable SCSI device for import and export of raw data.

4.3 HEPiX Working Group on Tape Interchange

J.Gordon asked if it made sense to establish a HEPiX working group on tape interchange formats between the sites; such a group might be able to arrive at a finite number of formats for writing tapes which were to be exchanged between the labs. The HEPiX Board were asked to investigate this and a few names were suggested already who could participate in such a group.

5 Central Services

5.1 The CS-2, a European Supercomputer - F.Gagliardi

The Meiko CS-2 at CERN was a European Commission-supported supercomputer installed at CERN as part of a collaboration under the ESPRIT III project to investigate general purpose uses of a MIMD

(multiple instruction, multiple data) system. The collaboration included 3 vendors and 2 labs, CERN being the leading partner. The target of the project was to demonstrate scalable supercomputing in scientific research and industrial applications.

The CERN node contained 64 HyperSPARC processors with a very fast internal switching network. The configuration was scalable. The system software was Solaris 2.3 with Meiko additions and CERN had installed the SHIFT package, various FORTRAN 90 compilers and AFS and ASIS software. The applications under study included parallel GEANT, NA48 Monte Carlo, parallelised versions of LHC accelerator design codes and a port of PIAF; principle clients were the NA48 and DELPHI experiments. There was also some second level trigger simulation for the RD11 collaboration. The speaker presented many performance figures and graphs.

To exploit the parallel options, various message passing schemes had been evaluated and a parallel file system was used. The Meiko job queuing system had been enhanced by the addition of CERN's NQS++. The system was fully integrated in the central CERN computing infrastructure and offered reliable, high quality production service. Possible future uses included online data recording and analysis for NA48 and more use of PIAF.

5.2 CERN's SP2 - H.Renshall

Since the Saclay meeting, the system had been installed, on its own Class C subnet with a special configuration to ease X terminal access. Some nodes were sited on their own FDDI ring to permit fast access to some AFS servers. The installed software included standard AIX 3.2.5 with some parallel extensions and CERN's various software packages. Login load balancing was implemented using IBM's ISS product and the batch queuing system was Loadleveler with many extensions added as part of a CERN/IBM collaboration, for example support for parallel systems and AFS token lifetime extension for long-lived batch jobs.

The main data access methods were AFS to get to some users's home directories and ASIS software, NFS to access CORE home directories and RFIO to access CORE data files. The user environment was set by the HEPiX login scripts.

The SP2 was steadily taking over more and more of the CERNVM batch load and also some of the interactive load. The speaker presented a number of graphs to illustrate this trend. They would soon start tests on parallel applications, in particular parallel GEANT. In this, they were helped by the flexible scheme for assigning nodes to tasks.

However, there had been some problems, mostly minor, and the speaker listed a few. Some products, for example the FORTRAN compiler, were controlled by a licence manager and this was rather inflexible, it had no failover for example when a licence manager crashed. Another problem was disc errors, made more difficult to trace because of the mixed vendor configuration chosen.

Nevertheless, at this early stage in its career, the SP2 was judged to be giving very reliable performance and offering a good service to users migrating off CERNVM.

5.3 LSF at SLAC - R.Melen

The central UNIX facility at SLAC now had 20 CPUs for batch work and some 5-6 interactive nodes. It was considered to be in prototype mode and would be expanded during 1995 and put into full production. When selecting a batch queuing system for UNIX, their IBM-based SLAC Batch scheme weighed heavily with regards to desirable and necessary features. The speaker presented a long list of such criteria with some indication of their relative importance. A few of them included –

- it must support a number of UNIX architectures including AIX and SunOS going to Solaris; other architectures may be introduced over time
- some form of load balancing was essential among the nodes in the cluster dedicated to batch
- it must be able to use scavenged cycles from distributed CPUs across the laboratory with no interference to users of these systems
- it must support also interactive users, with some load sharing if possible
- it should have facilities for use in parallel applications
- it must permit fair shares and stop individual users monopolising queues.
- it must support NFS, AFS and amd automounter
- etc, etc, etc.

LSF from Platform Computing was selected after an evaluation of several such packages in which some 13,000 jobs were run during 638 CPU days. It has been found robust in use and the development team were very cooperative and provided good support.

6 User Environment

6.1 SUE at CERN - R.Tobbicke

SUE, a Standard (or Shrink-wrapped) User Environment scheme for the installation of UNIX workstations at CERN, had been described already at the Saclay meeting. Version 1 had been done for HP-UX 9 and AIX 3.2.5 and was used in production to help install new workstations.

A revision which had been in progress at that time had now been completed and the speaker reviewed how the major features and environment profiles were defined. He also described the various scripts available for initial installation of the environment and subsequent updates and maintenance requirements. A default CERN profile environment had been defined containing those features judged to be necessary on all user systems. A formal definition paper was available (URL - <http://wwwcn.cern.ch/dci/sue/>). The main changes were to gain from the experience of the first two implementations, to provide more structure and more modularisation, to simplify the design and to make it more easily extensible.

This new definition was currently being implemented on HP-UX 9, AIX 4 and Solaris 2 and was already used to install HP systems. Other ports would follow.

6.2 Distributed Printing at DESY - Z.Jakubowski

The requirements for this had included the need to support a range of printers accessed from a range of systems including some at outside labs plus a nice GUI (graphical user interface) to the print command on various platforms. They had decided to standardise on postscript as the printing language and considered BSD 4.3 printing (the lpr suite) to be superior to that of System V despite its major drawback of having a fixed set of options. A number of high availability features were included in both hardware and software.

They had to deal with some 120 printers and include print submission from VMS, MVS, MACs and PCs. They have built their own non-queuing client lpr interface for both BSD and System V. They provided a web page showing the geographic spread of printers on the site, permitting the user to select the most convenient device. Another web page displayed printer status and queues. Their GUI used TK/TCL but they found that the latest TCL release produces compatibility problems.

Other problems met include

- postscript was far from being 100% compatible across vendors' implementations
- files were often incomplete when sent for printing
- they required page accounting
- they required a lot of buffer space

Future plans included developing a possible fax service and generally working towards a lowering of paper costs at DESY.

6.3 Printing Discussion

In a discussion about printing led by W.Wojcik of IN2P3 a number of topics were covered including

1. printing from inside applications: M.Marquina presented some ideas which had been developed at CERN to make use of the xprint command written there.
2. wider use of xprint in HEP: R.Tobricke described very briefly xprint, a scheme based on remote printing via a BSD host using a simple home-made protocol since he considered that the lpr protocol had too many major defects. It already had provisions for outside access (in both directions) but the current version had the name of the central host in CERN hard-coded in and he agreed to look into making a package of xprint which could be used outside CERN to make access to printers both local to that site and at CERN (or elsewhere).
3. access to printers at CERN and DESY from abroad (linked to the previous topic).

6.4 A Generic HEP Interface to FORTRAN 77 - M.Marquina

CERN had developed a generic interface to FORTRAN 77 in order to hide the various differences of the vendors' compilers. It defined standard file extensions for source and object files and included some special switches for debugging and optimisation. The interface handled all the standard options valid for the native compilers on each platform.

Startup files could be installed at various levels (user, group, system, etc) to set preferred defaults for the switches. The interface was available today on RS/6000 and was being ported to other platforms by the CERN Program Library.

It was noted that something similar existed at IN2P3 and from their experience, CERN was advised as soon as possible to convert the interface from a C program to a shell script.

It was remarked in passing that there were no plans today for a 64 bit version of the CERN Program Library.

6.5 Thoughts on a UNIX Mail Strategy - M.Wicks

The broad goal of this was to "solve the UNIX mail problem" for FNAL users migrating off VMS to UNIX. FNAL currently had a mail server/router, based on VMS and offering various protocol conversions and mail list services. Also, administration users make use of Quickmail. On the UNIX side, various user agents were in use, most commonly xmh, but there was no support for POP or IMAP.

A set of requirements were presented, including the need to support non-X devices, easy learning curve, tailorable, consistent across architectures, etc. The first conclusion was that mh seemed a good base, perhaps selecting one of its derivatives. However, reaction from the central UNIX support team and a few users was not altogether positive and more study was being undertaken including the need or not for POP or IMAP, the need for more tools and better documentation, and so on. A major problem was the mailbox format inconsistency between the mh family and other user agents.

A review of chosen mail strategies from the other labs gave

- CERN recommends pine; a central mail server was coming online soon based on IMAP to user agents
- DESY - pine also, no support for mh, plans for a central mail server also using IMAP; IMAP had been added to elm by W.Friebe
- SLAC had tried ORACLE mail but too heavy. Currently using elm and a central mail server. Used Eudora on MACs and PCs with POP.
- RAL had a name server for mail distribution using POP protocol. Most common user agent was pine although they had looked at zmail. Would like to use IMAP for the future; used Eudora on PCs.

- IN2P3 had no mail server yet but discussions were on-going.
- Saclay/CEA had 2 central nodes in Saclay; used MSmail on PCs mainly with a little use of Eudora; had a central POP server
- KEK was a largely VMS mail shop but were looking at POP3 support for PCs and MACs.
- NIKHEF had a central UNIX server and Eudora on MACs.
- The Prague institutes recommended pine.

7 X11

7.1 HEPiX X11 Working Group Report - T.Finnern, A.Taddei

This group was formed in February and had now met twice, once immediately before this conference. Their work was aimed at "average" users and the intention was to produce a reasonable working default X11 environment while offering facilities for "above average" users to tailor this. It should also offer system administrators tools to customise the environment locally at a site or a group level.

From the initial meeting at DESY in February, and from parallel meetings at CERN throughout the Spring a proposal for a standard X session has emerged which has now reached about 90% acceptance by members of the working group. It is hoped to publish this for comment in September but it requires more feedback and also volunteers to help implement it.

The speakers, joint coordinators of the group, presented details of the topics covered by the group including those points already agreed. These can be found on the World Wide Web at address <http://wwwcn.cern.ch/hepixon/wg/X11/www/meeting-2/talk.html>. There remained a number of still-open points and an action list for the group members to move ahead. Future activities would be published in the HEPiX news group and via WWW at the address given above.

7.2 Central X terminal boot servers - L.Cons

In order to use an X terminal, a number of servers are required, including a boot server, a font server, a login server (to provide session management, typically xdm) and a work server. CERN has established a model for boot and font servers and installed two such servers in the Computer Centre as well as one instance remotely near a concentration of X terminal users. Other remote instances are expected to be installed later according to traffic patterns.

Support was provided for a large number of X terminals (over 1000) from a predetermined range of suppliers but mostly from NCD and HP. Some problems met and solved were presented as well as some tricks used to ease the life of the system administrator and the user. A lot of work had been done on offering a wide choice of fonts across the various X terminals and UNIX workstations at CERN. Because of the need to cope with remote instances of these servers, some work had been done to centralise system monitoring of them via SNMP.

8 AFS and DFS

8.1 arc - a sysctl replacement - R.Tobbicke

arc stood for Authenticated Remote Control and was intended to be a tool to permit the delegation of privileges to permit sharing of system management tasks. Authentication was most important. Uses found so far included remote shell execution (xrsh), disk space management at a user group level and batch job AFS token extension.

Whereas sysctl from IBM processed commands through an interpreter running on a server, arc used a client/server Kerberised socket protocol. The commands permitted included any PERL command. Of the applications listed, the speaker explained in particular the batch job token lifetime scheme in great detail. His scheme was starting to be used in the Loadleveler installation on CERN's SP2.

8.2 sudo with AFS - T.Finnern

The aim of this attempt to modify sudo to be AFS-aware was so that root authority could be delegated across the AFS network at the DESY site. It should provide a simple interface for system admin tasks and also logging of the commands. After a number of iterations, it was finally decided to merge the well-known sudo command with arc, described by the previous speaker. Examples of its use were given, in particular by DESY UCO staff for dealing with users' AFS accounts and by the Applications Software Group to release AFS volumes.

8.3 HEPiX AFS Working Group - M.Wicks

Since the Saclay meeting last October, a working group had been established to investigate many aspects of AFS. They had had a two day meeting in February and drawn up a long list of action points (see URL <http://www-oss.fnal.gov:8000/hepix/afs/>).

M.Wicks, Chairman of the group, listed the most important of these action points with their status. These included a software repository for AFS administrators, feedback and input to Transarc, suppliers of AFS, and the sharing of various system admin and user tools written at different labs. In particular, a formal letter had been sent to Transarc and the various operating system vendors to voice concern about the lateness of AFS ports for the latest releases of the different operating systems in use. A reply had recently been received from Transarc with some explanations and a promise to improve this situation in the future.

The group had met again during this conference and revised the list and work would continue and be reported in HEPiX news and at future conferences.

8.4 DFS Status at KEK- T.Sasaki

This was an update on the plans announced at the Saclay meeting. It had been found non-trivial to configure the DCE cell needed for DFS although the vendors claimed the procedures would become simpler in future. The DFS server had been moved from an HP node to a SUN SPARCcenter 1000 and the DCE/DFS ports on HP-UX from HP and SUN from Transarc were under test. Performance figures were given with reservations about the observed results on FDDI. ATM tests will be performed later in the year.

9 UNIX Migration

9.1 UNIX Migration at FNAL - J.Nicholls

Fermilab had recently issued a Statement of Direction announcing a gradual rundown of their VMS service over three years, the users expected to move to a variety of UNIX, PC and MAC. The three largest VMS clusters at Fermilab hosted some 1000 users although VMS formed only a small percentage of the total installed capacity in pure CPU terms.

A group had been established to handle this migration and they would be starting to draw up tasks and possible problems, develop plans for user training and collect user feedback, both formal and informal. From these, a plan and timetable would be established, highlighting major issues. It was noted that the major experiments already had codes which ran under UNIX because of where the major part of the batch capacity was.

9.2 User Migration at CERN - M.Marquina

This was an update from the session given in Saclay. Among the tools built since then was ADAM - a simple graphical user interface to X-capable devices. The speaker described why it had been produced, compared it with other GUIs such as COSE/CDE and showed some examples of its use. It was lighter on resources than other GUIs, used Incremental TCL and was available on PCs via Xvision.

Another major effort was being put into mail with work to adapt pine as the recommended mailer, a mail list server had been established using majordomo and work was currently in progress on a centralised mail server based primarily on the IMAP protocol.

9.3 User Training

This was a discussion on the various UNIX user training schemes adopted by different sites. RAL organised occasional formal courses by an outside consultant tailored to the users needs. DESY offered short lab-style courses to specific groups, presented by User Support and devoted to subjects such as a UNIX Introduction or an introduction to emacs. These introductions were popular; "how to" classes

were also requested. However, they laid great stress on good documentation, including short flyers and cards on single products.

FNAL noted that outside commercial courses were not popular with their users who preferred tutorials tailored to their perceived needs. Their plans included lunch-time seminars. CERN had a formal Technical Training programme which included several UNIX courses but they were looking at more specialised, less formal seminars, especially in the CERNVM migration exercise.

It was agreed that it could be interesting to establish a database of courses given at the different labs. Sites might then be able to use this database to benefit from already-prepared courses either by inviting the speaker or by sending their staff to another site for a course. A.Silverman offered to establish this.

10 Windows NT and X11 - T.Finnern

This was a short presentation of some tests performed at DESY with the WinDD tool from Tektronix to provide a Windows interface on X devices, in particular on Tektronix X terminals (although ports were now appearing for various UNIX workstation flavours). The goal was to offer a single terminal for users for all their work. The results of the evaluation were presented; performance was described as satisfactory, although it provoked a fairly serious network overhead, especially on non-Tektronix devices, and it used propriety protocols. Its major worry however was how well it scaled in multi-user mode. As more users accessed it, how would performance from the NT host degrade?