Annual Report of the Computing and Networking Advisory Panel 1999/2000

June 2000

1. Introduction

There have been a number of major changes to HEP computing in the UK during 1999/2000. The BaBar experiment, the first major experiment to use OO methodology, has come online and large scale computing provision has moved away from commercial Unix workstations to commodity PCs and components plus open source software with the rapidly growing Linux farm taking over from the CSF (HP) farm. In addition to traditional sources of funding, there have been a number of successful JIF and JREI bids to provide computing equipment dedicated to certain experiments, such as the new US experiments. Finally the concept of the computational Grid as a likely solution to the future computing needs of the LHC experiments has grown rapidly over the year and is likely to dominate HEP computing activity in the coming years.

CNAP has been active in all these developments and has continued to consolidate and rationalise the provision of centrally provided services while planning for future experiments, notably those at the LHC. CNAP provided computing input to the PPARC Long Term Science Reviews with the 'LHC Computing Challenge' document, which is provided as an annex to the Particle Physics section.

In November 1999, the PPC asked CNAP to conduct a review into all aspects of the future of HEP computing in the UK, to report by the end of 2000. However, in view of the fact that CERN is currently conducting a review into LHC computing, the outcome of which will have a large influence over what the UK does, and the rapidly changing Grid situation, it was agreed to postpone the CNAP review for a year.

Section 2 of this report deals with centrally provided computing and data store services. Network issues are covered in section 3, a brief summary of the rapidly moving Grid situation is given in section 4, CNAP sub-group reports in section 5, training in section 6 and user representation in section 7. A brief summary and list of recommendations are given in section 8.

2. Central Facilities

2.1 Unix Farms

2.1.1 HP

The HP farm was stable during 1999 apart from security and Y2K changes (see below). It continued at 100% utilisation throughout the year. With the growth of the Linux farm the HPs form a reducing fraction of the power although the effort involved in maintaining them does not reduce so the future of HPs was considered by CNAP. A survey of users was carried out where it was proposed that the service be terminated by October 2000. Most experiments have either moved off, will move off or are happy they can move off when required. One or two are quite keen to have access to HPUX for software releases while it is available. One experiment claim they will have difficulty meeting an October deadline. A termination date for the service will be discussed at the next CNAP meeting.

2.1.2 Linux

An early production service started in June 1999 based on ten dual boxes. The twenty dual processor Linux farm went into full production in August. The Linux systems share a NIS domain and home filesystem with the HPs and run the same batch system and ASIS software. Another 40 dual processors with 600MHz cpus were bought online in June 2000 This represents a 300% increase on the phase 1 farm. Initial take-up has been very good, the Linux farm has already been used by a cross-section of experiments as shown in the figure.

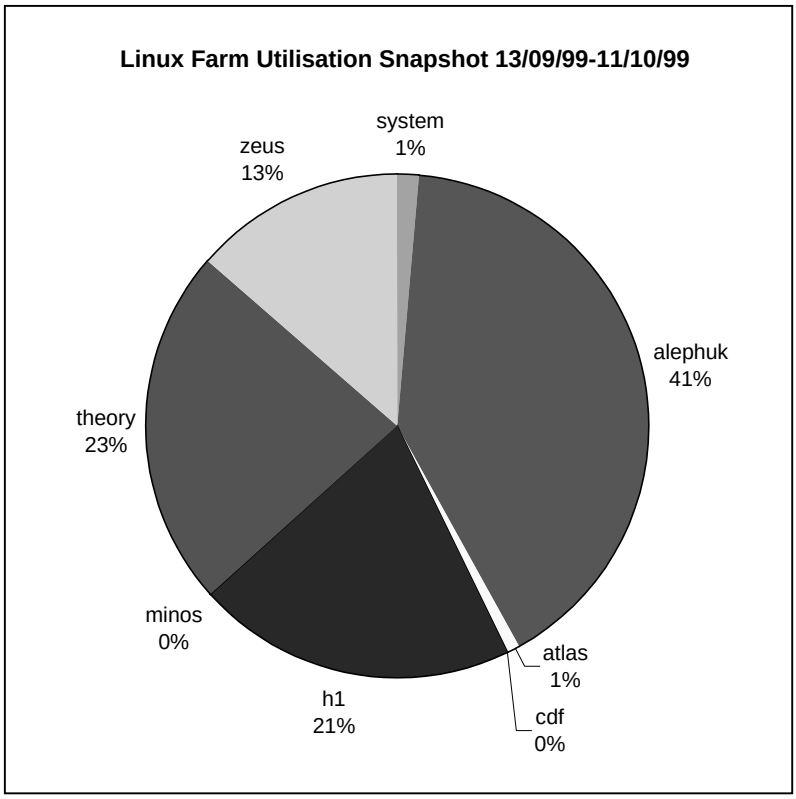

- H1 and the theory group have almost completely moved off the HP farm onto Linux. Given that these were major users of the HP service this has freed up a significant amount of capacity on the HP farm.
- For the first time ZEUS have begun to make significant use of the RAL facilities and CSF is now making a considerable contribution to ZEUS production. A couple of other individual users now use Linux.
- Although the primary ALEPH production remains on the HP farm at least one individual user is now using Linux.
- ATLAS have begun making some use of the service however there are no plans yet for any major production.
- ANTARES, CDF, and MINOS are all in various stages of code porting although no significant volume of work has yet been run. BaBar have ported their Monte-Carlo production to Linux but not yet started using it for production. CMS also, are clearly interested in Linux in the medium term and may well wish to use this service.

Initial benchmark figures are extremely encouraging - individually the Pentium 450 CPU out performs anything on the HP farm. Following the June 2000 upgrade, the Linux farm now delivers 20 times the HP farm capacity.

2.1.3 Sun

The Sun server bought at the end of 1998 had its memory and disk increased and it has provided a successful software development service for BaBar. When the JREI equipment for BaBar (a six cpu Sun E4500 server and 5.9TB of RAID disk) was delivered in August the original Sun served as a front-end. The JREI server was designed as a database server but for the moment it has sufficient capacity to run batch work as well so the LSF batch system was

installed on both Suns.

2.2 NT

The NT farm has been upgraded with the addition of ten dual 450MHz CPU PCs which increases its capacity by a factor of approximately five. The hardware has been rearranged to separate the functions of the front-end, the file server and the batch processors which makes management simpler. System replication mechanisms have been implemented which make it easier to add batch nodes to the farm which improves scalability, and this also helps with recovery from disk failures. The old front-end has been replaced by one of the new dual processor systems running NT Terminal Server, which makes the system software as uniform as possible with a base of NT 4.0 products throughout.

Prior to the upgrade ALEPH were the major user, but their use dropped off and has finished. However, the farm has now been configured with the LHCb web-based front-end system for submitting simulation work, and some 200k+ events have been produced so far. This work is expected to continue as LHCb finalise a production version of their simulation code and RAL has been allocated a target of 2M events for the forthcoming year. DST production from the simulated data will increase the volume of LHCb work and this will start as soon as the code is released.

The future of the NT farm is being reviewed by CNAP and a decision will be taken once the current LHCb production run has finished.

2.3 Disk

The total of 1.25TB of disk that was achieved in 1998 proved sufficient for 1999 as the BaBar users released space for others.

A dedicated RAID server with 50GB of disk was purchased to act as a common home filesystem server for HP, SUN, and Linux. This provides much better disk and network performance than the HP which previously held the /home data as well as being much more reliable.

2.4 LAN

As most of the machines purchased during the last couple of years have had fast ethernet access the LAN configuration was redesigned to centre on fast ethernet switches and remove dependence on FDDI. The Atlas Datastore now has Fast Ethernet access so almost all of the HEP services at RAL can communicate via switched fast ethernet. The remaining services will be moved during 2000.

2.5 Y2K

Although there were very few problems across the world, ITD put a lot of effort in making the HEP systems Y2K compliant. All operating systems were brought up to the required level and relevant patches applied. Disaster plans were in place and most systems were shut down over the year-end.

2.6 Security

The level of hacking on the Internet is still growing and has resulted in increased work for system managers to avoid problems. CLRC has a firewall in place to protect 'internal' systems and this is effective in protecting most of the batch farm nodes from malicious external attacks but the front-ends and disk servers need to be accessible directly from universities and other sites so cannot be as secure. Great care has been taken with these externally-visible systems but the workload is non-negligible.

RAL is collaborating with other HEP labs (e.g. through HEPiX) and will follow best practice. Adoption of common encryption standards and tools may follow. This will be a continuing topic of

interest within CNAP and the wider HEP computing community.

2.7 Data Storage

The main focus of the group in the past year has been in positioning ourselves to be ready for the large increases in data volumes which will appear from the USA experiments (BaBar, CDF, D0 and MINOS) and from CERN, in particular in the longer term from the LHC. The activities have involved close liaison with the various JIF and JREI bids; as these have significant implications for the area of data storage and data access.

Two formal datastorage group meetings were held on 14/9/99 and 16/12/99. In addition a meeting was organised in May 1999 to review the status and plans of the JIF/JREI bids to try and ensure that they were compatible where possible, and for the purposes of exchange of information.

In the June meeting the needs and plans of the major experiments were reviewed. As a result of the discussions a more detailed survey was conducted and the results reported to the CNAP meeting in October 1999. This report highlighted the very strong growth in storage demands expected in the coming years as shown in the figure and that it was now appropriate to consider expansion of the robotic tape capacity at RAL. The existing robot is running close to capacity and already there is a significant amount of data that have been archived. The CNAP meeting agreed with the recommendations of the storage group and asked this group to formulate plans for expansion.

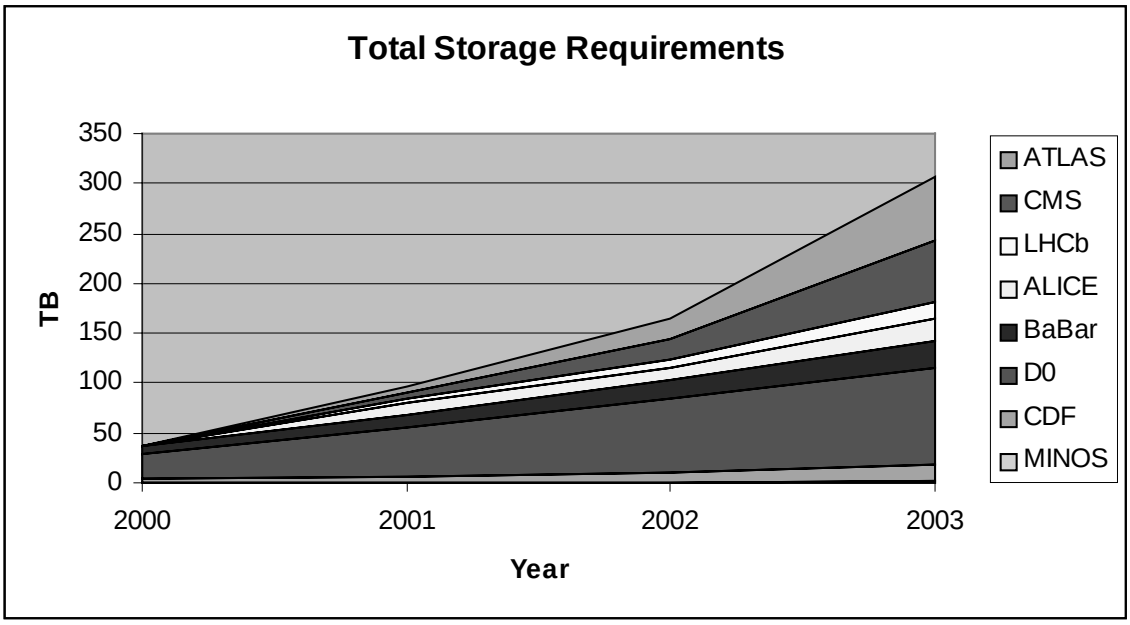

After various consultations the possibilities for expansion were discussed in the December storage group meeting. John Gordon outlined the possible expansion plans, using as a basis the facilities already available. It was agreed that a flexible solution is needed as the tape and drive technology is foreseen to evolve considerably in the coming years. It was agreed that a detailed tender document be drawn up and that the expansion be coordinated with the overall financial planning of central computing.

The various JIF/JREI bids, and their implications, were also discussed. BaBar have now completed installation of their successful JREI storage bid. Another successful JREI bid was from Liverpool, for a large Linux PC Farm (MAP). Although it has only small direct implications for central computing at present, it is important in the context of the LHC bid. Significant increase in the disk capacity of this facility is desirable to exploit the facility more efficiently. The JIF bid by

CDF/MINOS was successful and will involve the installation of large disk capacity at RAL and the collaborating Universities, as well as contributing to the upgrading of the robotic facilities. D0 have also had a successful JREI BID bid involving tape storage. This will be sited in Lancaster. Finally, a bid from the LHC Collaborations to set up a prototype grid system was discussed; in particular the important aspects of data storage. The evaluation of the various possible storage control software projects (HPSS, Eurostore ...) is still underway and a decision from CERN on what they plan to use is awaited.

2.8 Central Computing Resources

2.8.1 FY98-99

For the Financial Year April 1999 to March 2000 the central computing budget was spent as shown in Table 1 and the spend on new equipment in Table 2.

	£k
Maintenance	41
Travel	10
Other	15
OO Training	23
New Equipment	166
Total	255

Table 1 Central Computing Spend 1999-2000

	£k
HP Farm	0.0
Linux Farm	70.0
Sun	5.7
NT Farm	12.7
Atlas Datastore	22.0
Disk Farm	12.0
AFS	5.0
Total	127.4

Table 2 New Equipment

There are a number of points to note:

- A large fraction of this year's spend went on cpu provision. This will also be the case in future years.
- There was little spend required on the Sun service. Unless the LHC experiments show a large increase in demand for Sun services this service together with the BaBar JREI server should meet the demand for next year as well.
- The disk farm expenditure was on an experimental server using commodity disk. This has not yet reported but the first indications look good.
- All of the above were debated and agreed by CNAP.

- In addition to the Datastore spend shown, the remaining capital of roughly £40k was set aside for capital provision to help meet the upgrade plans of the CNAP Data Storage Sub-group.

2.8.2 FY1999-2000

For 2000-2001 the priorities are:

- To meet the increasing demand for cpu on the platforms demanded. It is proposed to buy Intel cpus and implement them as Linux or NT as required and to integrate all the various architectures as closely as possible through a shared infrastructure.
- To meet the target of 20TB of datastore by the millennium that was set by the CNAP review of 1997, as advised by CNAP's Datastorage Subgroup, and create a platform for growth for the LHC era.
- To provide a high-quality infrastructure that can support the addition of other HEP services that may wish to co-locate with the central services. As well as the BaBar JREI service, CDF are also placing servers at RAL. While support of these will not come from the central budgets they will make use of the infrastructure (e.g. networking) and this must be kept up to date.

It is foreseen that the resources should support approximately 9 staff years of effort and that there should be £280K for capital and recurrent spending and £20k for OO training of UK physicists. This represents an overall increase in cash from £700k to £735k. The foreseen breakdown of staff effort is shown in table 3.

Service	**Staff**
Unix Farms	2.2
NT Platform	0.5
Datastore	3.4
Disk Servers	0.8
Network Monitoring	0.6
Video Conferencing	0.3
Development	0.7
Indirect	0.5
Total	9.0

Table 3: Staff Effort for ITD Services

The increase in capital partly represents an injection of capital to fund the datastore upgrade foreseen for 2000, and partly represents the increased costs of supporting a wider central computing bases (BaBar, CDF, D0, LHCb).

2.9 Future Computing Requirements

With the advent of demanding computing requirements for the LHC experiments, it was deemed important to quantify their needs. Given the speculative nature of such a task, the methodology adopted here was to use existing experiments (H1, ZEUS, CDF, D0, BaBar) to also provide data as a guide to the estimates for those in the future (ATLAS, CMS, LCHb, ALICE).

The computing requirements in terms of processing power were evaluated for each of the experiments active in the UK. This was carried out by means of a questionnaire that was then

followed up with discussions and clarifications with the experts concerned.

Given the relatively long period (in computing terms) between now and the commissioning of many of these experiments, it was decided that the information should be restricted to the performance of simulation and reconstruction programs as they stand at the current time. The intention here is to avoid relying on estimates and expectations for improvements in code design and optimisation.

Experiment	**H1**	**ZEUS**	**BaBar**	**CDF**	**D0**	**ATLAS**	**CMS**	**LHCb**	**ALICE**
Simulated CPU times	169	300	350	1300	3000	9000	1560	2000	19200
(SPECint 95 seconds)									
Reconstruction CPU times	13	15	35	-	1300	1000		800	320
(SPECint 95 seconds)									
Simulated event size (kB)	240	40	150	1000	1000-3000	1000-1500	350	1000	2000
Raw data event size (kB)		100	40	500	250	1000		100	480
Reconstructed DST (kb)	13	150	100	30	150	100	1000		200
Reduced DST (kB)		25	60		50	10			

After the results of the questionnaire were collated, further questions and clarifications were required which led to the results shown in the table. In addition to the cpu demands, an estimate of the expect data storage size per event was also requested in order to compare current experiments with expectations from those planned for the future. These results are also summarised in the table. Note that some estimates are simply unavailable at the current time, or are likely to be updated in the immediate future and so current values are not useful.

It should be noted that the future experiments expressed reservations about the validity of using current simulation timings etc as a guide to future (production) performance. However, the expected improvements in performance quoted, varied widely and no reliable estimate could be determined. Competing factors, such as improvements from optimisation and better algorithms may be balanced by more sophisticated reconstruction analyses and complex algorithms applied at lower stages in the reconstruction hierarchy. As an example of this, the LHCb simulation cpu time per event currently fluctuates within a factor of 2 depending on a combination of the program version, type of events and the generator settings used.

Finally, the above table represents a snap-shot of a dynamically evolving status. The evolution of the processing power needed, even for the currently existing experiments, is changing due to the effects of luminosity upgrades and new sub-detectors etc. Hence, it is expected that such a survey of computing requirements needs to be updated regularly. The CERN MONARC project is likely to cover such ground in the near future through the demands of the CERN LHC review and, as such, should provide an updated estimate.

In order to keep the estimates of future demand up to date, CNAP has introduced an new arrangement with the PPESP. When each experiment is reviewed every two years, it will be asked for its estimates of future central computing requirements and the replies will be forwarded by the PPESP to CNAP. Two experiments have so far undergone this procedure.

3. Networking (PPNCG)

3.1 Current State of the networks

Our major international links (UK-USA and UK-Europe) are providing reasonable overall

performance even at peak times. With firm plans in place to upgrade the London infrastructure through which we connect to the international links and to add a third 155 Mbits/sec link to our US-UK bandwidth it looks as if this good performance could continue.

However, the end-to-end connectivity depends also on the quality of interconnection (peering) between networks. A major improvement in connections to sites such as SLAC and Fermilab occurred on 23rd February 2000 with the implementation of full peering between JANET and ESnet. Peering between JANET and Abilene, the US Internet2 high-speed network, came into effect on 10th January 2000. This improved connectivity to some US Universities.

Connectivity to the SNO experiment in Sudbury has been poor with a major impact on the development of this experiment. This was reported in August 1999 through our normal channels and the necessary peering between JANET and CAnet (the Canadian equivalent of JANET) has just been installed and connections are now good.

Both of the above are major achievements of the PPNCG following monitoring and requests to UKERNA. We would like to emphasise that UKERNA have shown a very helpful and open attitude to requests to change configurations that benefit all.

Isolated problems continue to affect various connections. For example, a protracted problem in the London JANET core router caused problems for some time in late 1999 and early 2000 particularly for sites connected to the router as all external connections were affected. There is still a problem with Oxford's connection to the core router that is technically a simple problem which is proving to be "administratively insoluble".

The PPNCG monitoring tools continue to be widely used, gaining us credit outside of UK HEP. They have recently been consolidated into the ICFA monitoring framework.

There are several aspects of Particle physics network use that do not sit easily on public networks. Network Quality of Service is one solution and dedicated lines are another. PPNCG has investigated using UKERNA's pilot Managed Bandwidth Service over SuperJANET to establish end-to-end connections between particle physics machines as a flexible alternative to private connections. The experimental requirement driving this project was the BaBar collaboration's need to move data around their 10 UK JREI servers. The trial project was to establish dedicated connections between BaBar machines at IC, Manchester and RAL. These sites were chosen for their combinations of network connectivity, BaBar use, and networking expertise. The initial connections will only be at 2Mbit/second but even establishing a connection across all the management domains will be a useful demonstration.

UKERNA had congestion problems that delayed the project but the connections are just about to be established. When the connections have been established a number of tests will be carried out on the performance of a range of basic protocols, middleware applications like database replication and data access from user applications.

3.2 Future Networking Requirements

The computing Grid proposals (see section 4) have implications that warrant specific mention. Bids have been made to JIF/JREI for Tier1/2 centres in the UK that will benefit MINOS, CFD, D0, BaBar and the LHC experiments. These call for the following bandwidth requirements in the years 2001/02/03: From RAL to CERN 4/25/50 and from RAL to Tier2 0/4/25 Mbits/s. There will also requirements for Europe-US bandwidth, although this is not yet specified in UK bids. These numbers are considered conservative, and figures more like 500 Mbits/s are called for at the start of the LHC. The provision across JANET and TEN is likely to be 2.4 Gbits/s in 2001 rising to 10 Mbits/s by 2005. Thus the HEP requirements will account for 2-5% of the total national provision. This is quite high given the breadth of other ac.uk users, thus it will almost certainly be necessary to fund some guaranteed bandwidth connections. This is foreseen in the funding requests that have been made.

The PPNCG will take responsibility for procurement, implementation and testing of Grid related networking work packages, as well as the provision of networking between sites.

As stated clearly in the previous report, many HEP applications will benefit from the introduction of "Quality of Service" production services, which allow prioritisation of time critical protocols. This certainly includes all interactive work and video conferencing, and possibly some bulk data transfer. QoS is also vital for the Grid. We are currently working with SLAC and UKERNA to measure the effect of introducing CAR and WRED onto the transatlantic link as a part of the management of traffic flows across the network. The PPNCG recommends that HEP UK supports strongly the development and deployment of QoS.

3.3 Video-conferencing

Video-conferencing is now an essential requirement of the HEP community. The PPNCG has published detailed recommendations on what equipment should be purchased. These recommendations have been forwarded by CNAP to the Particle Physics Grants Sub-committee and distributed to the wider community. Further details can be found in Appendix A.

3.4 HEPTEL

The HEPTEL service continues to offer a private telnet connection between the UK HEP sites and CERN and thence to DESY. It runs on a 64 Kbits/s RAL-CERN leased line and may be used as fallback circuit when interactive response over TEN-155 is poor. HEPTEL is also configured to route transatlantic calls across a private CERN/US link, of which we have negotiated a small share. This route to the USA, via CERN, has been essential during long periods of poor network performance over the normal transatlantic link. The main development in the last year is the installation of secure SSH capability on HEPTEL. The demand for the HEPTEL link will, as always, be reviewed annually.

4. Grid

The computational ‘Grid’ concept grew within the year, and will prove to be central to UK particle physics computing in coming years. The Grid provides seamless connection between distributed computing power, high-speed bandwidth interconnection and data storage. The ‘LHC Computing Challenge’ sets requirements which are fundamentally different from those in previous generations of experiments, and the Grid is seen to be the best solution to meet the distributed computing demands.

Substantial Grid activities are already underway in the UK. A community meeting was held in March at RAL, and an email distribution list established. A workshop given by the Ian Foster team is planned for June. A JIF bid has been submitted to fund a prototype Tier 1 (a national computing centre) and the effort required to put in place a pilot Grid which will be connected to CERN. Similarly a request for resources has been made against the next Spending Review.

A European consortium has been established named the ‘DataGrid’, and centred at CERN. The consortium has made a bid to the European Commission to fund activities in setting up (primarily) a Grid for LHC computing. There are 6 primary contractors in the bid, and the UK (in the form of PPARC) is one of these. A number of ‘Work Packages’ have been formed and the UK is leading two, and making substantial contributions to a further two. It is anticipated that forthcoming Grid activities in the UK will be formed around the structure established for the DataGrid project.

The experiments with most immediate potential need for Grid functionality are BaBar and the Fermilab experiments. Collaboration with Ian Foster is being set up to develop Grid connections between the US and UK for CDF.

The Grid is an extremely exciting initiative, with strong support from Research Council management, and will provide a crucial tool to facilitate particle physics analysis.

5. Other Group Reports

5.1 Computer User Forum (CUF)

The Computer User Forum (CUF) was set up at the end of 1997 by CNAP to improve information flow between the computing service providers, policy makers, users and experiments. Meetings at RAL and UCL, were held in the CUF series during 1999 and were chaired by Tim Adye (RAL) with representatives from most of the experiments with UK involvement.

These meetings had the dual role of discussing the configuration and operation of the central computing and networking services; and of informing the community on computing technologies that are becoming important for particle physicists. On the central services, the main issue this year was the potential rundown of the HP Farm, and useful feedback was obtained on how this could be managed with minimal impact. The main informational themes this year have been Networking, Linux, and Object Oriented technologies.

The CUF meetings attracted an attendance of 20-30 people from the particle physics community. Feedback on the concept, topics, and venues of the meetings has been very positive. Based on experience from this year, at least one CUF meeting each year will be held outside RAL.

All CUF meetings are publicised in advance using Hi-Phi, a dedicated LISTSERV mailing list, and a set of Web pages. All the meeting transparencies are available from these Web pages.

5.2 System Managers Group (HEPSYSMAN)

Lee Thompson (Sheffield University) has been appointed to CNAP, doubling institute representation as proposed previously by the HEPSYSMAN Group. With two CNAP members now responsible for institute-related matters it is hoped that the wide spectrum of local computing environments can be fully represented. The two appointments are staggered for continuity.

The third meeting of system managers was held last November at UCL. Attendees at this meeting confirmed their support for the two to three meetings a year suggested in the original proposal and the next meeting was consequently scheduled, as a 2-day RAL event, for March, 2000. It is planned to follow this up with a one-day meeting hosted at one of the other institutes later in the year. Meanwhile the mailing list set up with the re-formation of the HEPSYSMAN Group continues to prove invaluable for the exchange of tips and queries and dissemination of information.

At the November meeting it became clear from institute updates that there had been considerable evolution since the last “round-table” session the previous year. Most groups were already well down the road to replacing their desktop miscellany with PCs, as recently recommended in HEPCCC guidelines. Simultaneously with the rise in popularity of PCs confidence has been growing in the adoption of Linux as a mainstream operating system. Almost all groups now run Linux as well as NT, and only a few sites are persevering with a rigorous NT desktop policy. For many the issue has become how to provide auxiliary Windows compatibility in a basic Unix working environment; products such as Windows Terminal Server, Vmware, VNC, and StarOffice are being investigated as solutions. From another perspective, however, Windows 2000 might provide the kind of single platform panacea originally hoped for in NT, but there are by now many Linux devotees who will be hard to convince. A summary of the computer equipment currently installed in the Institutes is given in Appendix B, taken from the Particle Physics Grants Sub-committee Form Gs.

In the networking area there is an evolution too, to faster LANs and higher-speed WAN connections. Many groups have installed 100 MB/s equipment, with some connectivity already at Gigabit speeds. These faster LANs are generally switch based, providing not only better performance but also increased security against “sniffers”. Most installations and plans assume

Ethernet, though ATM is not completely unrepresented. Areas in which interest is growing strongly include shared high-capacity file-servers, desktop video-conferencing, and firewalls.

Looking to the future, groups express many concerns. Management overload does not go away in spite of efforts to standardise. This is partly because “software” personnel are occupied with experimental commitments, but also because of the new demands made by security measures and new software and peripherals. The number of server types is increasing fast as well, whilst experiments are creating their own diversifications with different choices of, for example, Linux distributions and tape media.

Another concern is the shortness of the desktop replacement cycle. Not only must new equipment be funded, but it must be suitably phased in, resulting in inevitable hardware incompatibilities. Groups are also worried about unbudgeted expenses such as LHC++ software and network charges.

Security continues to be a major headache. The distribution of information by CNAP last year has helped to improve the general appreciation of security issues, but system managers still have difficulty finding time to set up the requisite monitoring tools and recovery mechanisms.

6. Training

The estimate made this time last year that the demand for IT training was falling off has turned out to be quite wrong. As a consequence, even though a small reallocation of funds to IT training was made, it was only possible to offer only the most basic courses this year. The course appraisal forms which attendies complete indicate that most people are very satisfied with the training they are receiving.

44 people have received training this year - most of them have followed 4 days OO Analysis and Design followed by 3 days of introductory C++. Most of the students have been Ph.D. students this year as the more senior people who want to discover the delights of OO have already done so. Requests for other courses have had to be refused due to lack of money. It seems probable that with the unsatisfied demand for the extra courses, the continued demand for the OO A&D and C++ courses that a 50% increase in spending on IT training is necessary.

Since a large number of those now being trained are PhD students, CNAP is in discussion with the Education and Training Section of PPARC to see how this increase can best be financed.

7. Geant4

Geant4 is the new version of Geant which has been completely rewritten using OO methodology and replaces the Fortran version that has been in use for many years. Geant4 is an international project by constitution. It is self governing and is unique for such a big software project. Each participating organisation - institute, laboratory or experiment - gets a seat on the boards. So far the following laboratories have joined - CERN, ESA, LEBEDEV, KEK, SLAC, TRIUMF and the following experiments - ATLAS, BaBar, CMS, LHCb. IN2P3 has just joined as a national institute and there is a substantial Italian group, already with board places, who hope to come in as INFN. The existing organisations cover many of our needs but not all in that at present only members of the experiments listed above are entitled to support. A proposal was made, which CNAP supported and encouraged, to get a group together who can join the Geant4 Collaboration as an identifiable UK contingent under the auspices of PPARC to get a seat on the governing and technical boards of Geant4. That way the UK will get a say in how Geant4 is run and get a right to support effort on behalf of all PPARC activities, including the astronomy community where there are already some UK members designing spacecraft shielding for instance. To qualify the UK needs to contribute 2 CUs (Contribution Units), effectively 2 fte of effort from those not already contributing to other groups.

Sufficient manpower has now been identified and a draft proposal has been submitted to the

Geant4 Collaboration and has been warmly received. The next stage is the signature of an MoU between the new PPARC Group and the Geant4 Collaboration, agreeing to fulfil our contribution. A small amount of money is required to allow the members to travel to Geant4 Collaboration meetings.

8. Representation

Group representation on CNAP was strengthened this year with the appointment of a new Institutes Representative as suggested by HEPSYSMAN. The needs of the University Groups are now directly addressed by the Institute Representative and the Chair of HEPSYSMAN while the needs of the experiments are addressed by the Chair of CUF.

Andy Halley was recently appointed to CNAP to oversee future computing needs and monitor JIF and JREI computing bids. Since then Andy has been appointed Chair of a committee (the 'Halley Committee') that provides Particle Physics input to PPARCs Grid planning committees and so this role will naturally include reporting on Grid activities to CNAP. At the moment UK Grid activities are well covered by the CNAP membership but the whole relationship between CNAP and the UK Grid committees will have to be reviewed periodically.

9. Summary & Recommendations

The last year has seen the start of the new generation of experiments in the USA and the emergence of the Grid as a likely solution to the longer term computing needs of the LHC experiments. In addition there are several additional sources of funding available which present many opportunities but also make central management of resources more complicated as some facilities are funded for particular experiments.

We note the following:

- That the demand for data storage and management continues to grow and the likely demand in the future will far exceed what is available or can be provided through normal funding routes.
- That the adoption of the Grid as the solution to the LHC Computing Challenge will have a profound affect on HEP computing in the UK and that the whole role of CNAP may have to reviewed in light of this development.
- That the management of ITD resources will change in the coming financial year and CNAP will review this during the next year.
- The UK and European bandwidth provision has been by and large adequate over the last reporting period. SuperJanet4 and GEANT (European Backbone) should offer sufficient bandwidth over the next few years.
- The transatlantic link has remained a great problem area. This was noted in the previous report, and accordingly we have investigated possible solutions with UKERNA. Networking changes have been implemented as a result and we are hopeful that these will at least alleviate the problem. We note that in the meantime CDF/D0 have been awarded funds for a direct link which will may solve the interactive problem for those experiments
- The use of video conferencing is growing, but slowly. ISDN should be considered as a mature production tool and could be used more extensively than it is. H.323 is emerging as the future standard.

We recommend:

- That the RAL ITD budget to support central computing facilities and datastore be increased from £700k to £735k including 9 SY of staff effort. The increase is primarily to pay for increased data storage and only includes £20k out of an estimated £30k required for OO

training.

- That a suitable funding mechanism be found to pay for OO training for PhD students. Training is likely to be vital if the UK is to develop the necessary Grid technology for the future and to successfully export this knowledge to other disciplines.
- That the HEPTEL service be maintained for the time being as this is still an essential lifeline for many users.
- That PPARC support the UK joining the Geant4 Collaboration.

Appendix A Video-conferencing

Simply put there are currently three options:

1. Dial-up ISDN circuits using the public ISDN network, and a "multipoint switch". This form of VC is in regular use by the LHC, CDF, D0 and MINOS, with several regular meetings occurring each week. The quality is high (and of course completely dependable) and the system is easy to use. Multipoint switches exist throughout the world, thus the only issue is the equipment needed at the institute. Until recently this has normally been a central institute facility that needed booking. The PPNCG has researched simple desk top equipment and produced a document recommending several low end desk top systems. Such systems have already become invaluable in some institutes. The disadvantages are (i) high cost of phone calls and (ii) need for someone to pre-book the multipoint switch (iii) need for ISDN connection near equipment. Apart from institute equipment there are no other issues which need addressing on behalf of HEP UK. We recommend this as the de-facto solution for VC at present.
2. IP packet based VC using the older "MBONE" applications. The technical quality, both internationally and within the ac.uk domain, is just about adequate (for some uses) , and this method of VC has the great advantage of requiring almost no booking, and being free at the point of use. The equipment costs are minimal (£200-300). Several groups within UK HEP have successfully tested this technique for various small meetings, although it is limited to one or two conferees per site. The disadvantages are (i) much time must be invested in session setup each time, (ii) all video compression is carried out in software(iii) there is generally no echo cancellation available, leading to "push to talk" operation. These three disadvantages, particularly (i) are sufficiently severe that, sadly, this method is unlikely to find favour with the majority of the UK HEP community. We do not recommend that this is pursued as a primary VC facility, and indeed there is little action that the PPNCG can take to alleviate the situation since the problems are mainly at the application level.
3. IP packet based VC using the new H.323 standards based equipment. This is clearly the preferred method of the future. The potential of this is substantial. Advantages are (i) free at point of use (ii) much more powerful compression algorithms implemented in dedicated hardware, (iii) echo cancellation built in to most equipment. Thus the quality promises to be much better, although at present it is not very tolerant of packet loss. Such hardware also supports NetMeeting, a freely downloadable videoconferencing application from Microsoft which has proved to be very popular because of its accessible user interface and its support for sharing Microsoft documents (e.g. Word, PowerPoint). This system is not yet mature (at present no multipoint switch infrastructure exists, there are problems with firewalls). However, pilot work is in progress are and we forecast that shortly this may become a useable service at least within ac.uk. We recommend that UK HEP support both the development and deployment of H.323 VC. The PPNCG will actively promote H.323.

It should be noted that ISDN and H.323 can interoperate easily, and that most equipment is dual standard, therefore equipment purchased now for ISDN can be used for H.323.

The PPNCG has actively advertised options (e.g. in documents published on WWW and advertised in Hi-Phi, as well as a demonstration at CUF). Nevertheless we suspect that the use of at ISDN as an alternative to travel both within the UK and internationally is not as prevalent in UK HEP institutes as it could be.

Appendix B Summary of Form Gs

	Desk Top			Server			Farm		Disk (TB)			PPARC	Support		Software		LAN	SJanet
Group	**< 3yrs**	**> 3yrs**	**SI95**	**< 3yrs**	**> 3yrs**	**SI95**	**cpus**	**SI95**	**Desk**	**Server**	**Farm**	**PPARC**	**FTE**	**Other**	**System**	**Desktop**	**Mbps**	**Mbps**
Birmingham	16	58	370	7	4	380			0.22	0.63		100%	1.0		£750	£1,500	10	34
Bristol	19	17	380	2	1	88			0.08	1.10		100%	0.7	0.4	£1,600	£1,500	100	40/15
Brunel	3	6	65	1		17			0.03	0.26		80%		0.3	£1,000	£1,000	100	20
Cambridge	29	23	380	1	1	21			0.15	0.45		70%	0.7	0.3	£2,000	£15,300	10	155/10
Durham	1	1	22						0.02			100%					10	34
Glasgow	28	8	400	1	1	60	3	36	0.10	0.75		95%	1.3	0.2	£6,158	£8,655	100	24
Edinburgh	11	15	230	1		34			0.25	0.50		80%	0.2	0.2		£1,000	100	20
Imperial	37	43	394	3	6	74			0.07	0.11		83%	1.0	0.2	£7,800	£3,700	100	34
Lancaster	15	3	260	2	0	25			0.15	0.05		100%	0.5		£3,000	£56	155	155
Liverpool	25	24	466	1	10	140	300	4500	0.15	0.20	0.02	100%	0.5		£8,000	£2,000	10	28
Manchester	18	12	255				3	52	0.20		1.00	100%	1.5		£5,000	£2,000	100	2x155
Oxford	66	65	988	6	2	126			0.66	0.50		57%	2.0	1.6	£3,000	£14,000	100	8
QMW	19	11	458	3	2	73			0.30	0.56		90%	1.0	0.0	£910	£7,007	10	8
UCL	28	4	210	2	1	82			0.11	0.12		95%	0.7	1.0	£9,000	£5,000	10	80
RAL	90	26	1430	5	9	203			0.58	0.34		100%	3.0		£5,000	£17,000	100	20/34
RHUL	20	7	270	3	1	60			0.08	0.60		50%	1.0	0.1	£10,000	£1,000	100	34
Sheffield	27	27	388		2	8			0.50	0.10		40%	0.1	0.0	£3,000	£8,000	100	33
Sussex	13	4	350						0.10			90%	0.1	0.1		£1,000	10	4
Totals	465	354	7316	38	40	1391	306	4588	3.75	6.27	1.02		15.3	4.4	£66,218	£89,718		

Appendix C . CNAP Committees and Membership

C.1 CNAP

Name	Institute	Function	Rotate Off
S.L.Lloyd	QMW	Chairman	2001
T.Adye	RAL	Central Facilities/CUF	2002
P.Clarke	UCL	Chair PPNCG	2000
R.Cranfield	UCL	Chair System Managers	2000
D.Newbold	Bristol	Coopted (CMS)	2002
P.Renton	Oxford	Chair Data Store	2000
L.Thompson	Sheffield	Institutes Representative	2002
A.Halley	Glasgow	Computing Coordinator	2002

In attendance:

Name	Institute	Function
A.Coates	PPARC	PPARC
S.Fisher	RAL	Training
A.Flavell	Glasgow	UK HTASC Rep
J.Gordon	RAL	ITD
P.Jeffreys	RAL	PPARC SLA
D.Kelsey	RAL	Chair HTASC
R.Middleton	RAL	Secretary

C2 Data Storage Sub-group

P.Renton (Chairman), T.Adye, T.Folkes, N.Geddes, J.Gordon J.Hart, P.Jeffreys, G.Patrick, D.Sankey, H.Wittig, S.Wotton.

C3 PPNCG

P.Clarke (Chairman), R.Hughes-Jones (Secretary), G.Fayers, A.Flavell, J.Hart, P.Jeffreys, D.Kelsey, J.Macallister, D.Sankey, D.Terrett (Astronomy representative), F.Wickens.

C.4 Computer User Forum (CUF)

Chairman	T.J.Adye (RAL)	
ALEPH	R. Edgecock(RAL)	
ALICE	I. Bloodworth (Birmingham)	
ATLAS	J. Baines (RAL)	S.L Lloyd (QMW)
BaBar	J.Allison (Manchester)	
CMS	G. Heath (Bristol)	G.N. Patrick (RAL)
Crystal Barrel	D.V.Bugg (QMW)	
DELPHI	D.Crennell (RAL)	
H1	D.P.C. Sankey (RAL)	J.V. Morris (RAL)
LHC-B	N. Brook (Glasgow)	C.P. Ward (Cambridge)
MINOS	R. Edgecock (RAL)	D.Petyt (RAL)
NA48	B. Hay (Cambridge)	
OPAL	D. Charlton (Birmingham)	G.N. Patrick (RAL)
SNO	N. West (Oxford)	
SOUDAN	D.Petyt (RAL)	
WA94	O. Villalobos-Baillie (Birmingham)	
ZEUS	J.C. Hart (RAL)	

C5 System Managers Group (HEPSYSMAN)

Chairman	R.Cranfield
Birmingham	L Lowe
Bristol	J-P Melot
Brunel	P.Hobson
Cambridge	J.Hill
Edinburgh	S.Gowdy
Glasgow	D.Martin
Imperial	G.Fayers
Lancaster	R.Henderson
Liverpool	A.Moreton
Manchester	A.McNab
Oxford	I.McArther
QMW	M.Landon
RAL	D.Kelsey
RHUL	S.George
Sheffield	L.Thompson
UCL	R.Cranfield

Appendix D References & Contacts

D1 CNAP web sites

CNAP	http://hepwww.rl.ac.uk/CNAP/
Data Storage Sub-Group	http://hepwww.rl.ac.uk/CNAPDS/cnapdata.htm
PPNCG	http://icfamon.rl.ac.uk/
CUF	http://hepwww.rl.ac.uk/CUF/
SYSMAN	http://hepwww.rl.ac.uk/SYSMAN/

D.2 ITD web sites

CSF Farm	http://www.itd.clrc.ac.uk/Activity/CSF
WNT Farm	http://hepwww.rl.ac.uk/HEPNT/
DataStore	http://www.itd.clrc.ac.uk/Activity/DataStore

D.3 Other web sites

HEPCCC	http://nicewww.cern.ch/~Djacobs/Hepcccw3/HEPhome.htm
HTASC	http://wwwcn.cern.ch/~eauge/htasc/public/
HEPiX	http://wwwinfo.cern.ch/hepix
HEPNT	http://hepnts1.rl.ac.uk/HEPNT/
HEPpc	http://hepwww.ph.qmw.ac.uk/HEPpc/