

Annual Report
of the
Computing and Networking
Advisory Panel
2003/2004

October 2004

1. Introduction

The Computing and Networking Advisory Panel (CNAP) currently acts as a PPARC advisory body to give strategic computing advice across the whole PPARC program. This CNAP annual report gives a very concise summary of the topics considered at the two meetings held in the past year and also includes a section on RAL computing facilities as in previous years.

The report is structured as follows. Network issues are discussed in section 2 and the existing RAL computing facilities are described in section 3. Brief accounts of the current state of the Astrogrid and GridPP projects are given in section 4. Issues raised by the Particle Physics and Astronomy system managers arising from the CNAP and system manager meetings are detailed in section 5. In section 6 some issues concerning the high energy physics computing coordination committee (HEPCCC) are discussed. Finally in the appendices details of committee membership, system manager contacts at each institute and some relevant web sites are listed.

The current PPARC plan for CNAP is that it should be merged with the GRID oversight committee to produce a single advisory committee to advise on computing for PPARC science including high performance computing. This is expected to happen before the end of 2004 and so this is likely to be the last CNAP report. Many thanks to my colleagues on this committee for their valuable input.

2. Networking

2.1 Introduction

The Particle Physics Network Coordinating Group (PPNCG), a sub-committee of CNAP, has not formally met as a face-to-face meeting during the course of this year. It has however continued to be an effective voice for the Particle Physics community in matters of networking through electronic communication, and the involvement of its members in many of the significant areas of development both nationally and internationally. In this way it has enhanced its respected position in the national and international networking sector which gives the community a significant influence in the development of strategic issues. Its contributions have maintained the very considerable worldwide respect which particle physics commands in the area of networking which is so critical to all of its operations.

This has been achieved through a wide range of activities including the continuing direct contact with the network providers and in particular with UKERNA (UK Education and Research Networking Association) and DANTE (Delivery of Advanced Network Technology to Europe Limited) who are responsible for the SuperJANET and Geant networks respectively.

This position results from participation in many joint projects, demonstrations and initiatives including, direct sharing of monitoring information, grid middleware development, leadership roles in the GGF (Global Grid Forum) and through the championing of the UKLIGHT facility.

Work in the areas of advances in high speed data transport, provision of network performance monitoring and diagnostic services provided to the GRID continues with direct PPNCG and member involvement in many significant successes. The list is extensive, but highlights include the EU DataTAG project, the UK MB-NG e-Science project, leadership in the GGF Network Monitoring WG, UKLIGHT and the Global Lambda Integrated Facility (GLIF) developments. Members have also been invited to participate in the SuperJANET 5 development work as representative of a high demand science area.

The Particle Physics Community and the global research network providers continue to have a close and fruitful interaction, and UK PP community continues with its front-line role maintaining this close relationship.

The PPNCG has a formal status as a Special Interest Group (SIG) within the Joint Information Systems Committee (JISC) allowing communications on the provision of the JANET service to be directed through the JNUG structure. The JISC and UKERNA have indicated that the JNUG structure is being reviewed, however UKERNA have made it clear to the PPNCG that it will continue to value the close collaboration with the PP community.

2.2 National and International Networking Status

The SuperJANET4 (SJ4) network provides the 10 Gbits/s backbone upon which PPARC relies and it has continued to perform extremely well in terms of capacity and reliability through the period of this report. There are no existing or foreseeable problems in respect of SJ4 capacity and with the early indications on the likely architecture for SuperJANET5 (SJ5) the future looks equally promising where PPNCG members have been engaged in the design process through the requirements review stage and the more detailed SJ5 planning.

The Regional Partner Academic Networks are similarly well provisioned and they are closely involved with UKERNA to ensure their capability is aligned with the SJ5 developments.

The TVN network which connects SJ4 to RAL runs at 2.5 Gbit/s with the connection from this into the RAL sites provided by 2 * 1 Gbit/s. The new firewall at RAL running at Gbit/s rates has provided a stable operation platform. Advanced planning is now underway for the re-procurement of the TVN network and PPNCG members have been invited by UKERNA to be closely involved in this process

The PPNCG, on behalf of PPARC, would like to formally recognise the excellent service provided by UKERNA, the RPANs and the site networking administrations for provision of the end-to-end service.

GEANT, the European research network runs a 10 Gbits/s backbone with typically 2.5 Gbit/s peering to national research networks and there are no foreseeable problems in this sector. The process leading to the reprocurement of the Geant network is underway, and DANTE have closely involved the PP community through this process.

Commodity and research traffic are now provided through different mechanisms. Commodity traffic continues to be peered in London where a multiplicity of redundancy is provided. Research traffic is now routed onto the Geant backbone, and from there to the US via 2.5 Gbit/s links from London, Frankfurt Amsterdam and Vienna. This has proved to perform well.

In the previous period we reported on two outstanding issues,

1. The connection to Jodrell Bank being inadequate; and
2. The connection to the Boulby mine to be very rudimentary.

We are pleased to report that both these problems areas are being addressed and for both sites the quality of connectivity will shortly be substantially improved.

2.3 Network Monitoring

An important duty of the PPNCG has always been to monitor national and international connections (see <http://ppncg.rl.ac.uk>, <http://av9.physics.ox.ac.uk>). These tools continue to be used extensively to keep abreast of the state of the connections relevant to all PP experiments and remain an essential and very powerful part of the PPNCG armoury. These are now extended through the availability of the network performance monitoring across e-Science site within the UK (<http://gridmon.dl.ac.uk/gridmon/>)

Work continues building on the successes of the EU DataGRID, and UK e-Science network monitoring with work now actively underway for the UK PP community, for the EGEE project and more widely within the GGF. This work is being led by members of the PPNCG who are very conscious of the need to coordinate this activity on a world scale to ensure that the delivered Grid Services are of general applicability.

The PPNCG continues to believe that it is in the interests of the PPARC community to deploy these tools at every UK PPARC supported site as they provide access to the diagnostic information needed to tackle residual throughput problems.

2.4 UKLIGHT & UK development network

The UKLIGHT facility is now in place with 10 Gbits/s connectivity to Amsterdam and to Chicago (StarLight) through the London C-PoP. Connectivity to this infrastructure is provided using spare capacity within the SJ4 core with access initially provided to a small number of Institutes. The RAL site will be connected in November 2004. The UKLIGHT facility is being managed and operated by UKERNA although through the various committees of governance, the PP community is well represented.

There have already been "early users" of this facility with the CDF community establishing a 1 Gbits/s "light path" between UCL and FNAL. Early results from this connectivity were reported at a recent UKLIGHT Town Meeting at the NeSC.

The first funded project to make use of the UKLIGHT facility - the ESLEA project - has prominent PP and Astronomy involvement and the experiences gained will feed both the SJ5 planning process but also the PP community's understanding of who best to exploit such capacity. Related to this PPNCG members have organised and run a Workshop (<http://grid.ucl.ac.uk/NFNN.html>) aimed at end users to highlight the technical issues and practical solutions to make use of this considerable end-to-end capacity.

2.5 High Performance Data Transport

Members have been heavily involved in the EU DataTAG project and the MB-NG project where high performance data transport has been a key area of research and development. Both projects have been great successes both for understanding and progressing the issues but also in cementing the PP community's relationship with the computer science and network research community throughout the work. The Workshop referred to above resulted directly from work carried out in these projects, and will be taken forward to develop "cookbook" style instruction and education to roll this capability into the community.

2.6 Related projects and affiliations

Members of the PPNCG have been key to several high profile international network projects under the e-Science/GRID banner. These include:

1. The EU-EGEE project, to build a European wide production Grid. PPNCG members have been responsible for formulating the network sector of this proposal.
2. The EU-DataTAG project. Several PPNCG members were prominent in this project to demonstrate high throughput and QoS with US partners. In particular a PPNCG member led the main network work-package.
3. The Managed Bandwidth – Next Generation (MB-NG) project which was initiated at a PPNCG meeting, and run in collaboration with UKERNA and CISCO. It has involved several PPNCG members and many Network Researchers from the UK Computer Science Community. It has resulted in the first very high performance QoS enabled network in the UK NASA Research and Education Network sector

Members of the PPNCG serve/have participated in the following:

1. The JCN (JISC Committee for Networking)
2. HEPCCC/HTASC (High Energy Physics Computing Coordinating Committee / Technical Advisory Sub-Committee)
3. ICFA (International Committee for Future Accelerators) Network Committee
4. The UKERNA Monitoring and Measurement Working Group
5. TERENA (Trans-European Research and Education Networking Association) Conference Programme Committee
6. The Global Grid forum steering Committee (responsible for network issues) and chair of the Network Measurements Group.
7. The JISC SJ5 Planning Group
8. The Global Lambda Integrated Facility (GLIF), (<http://www.glif.is/>)

3. RAL Computing Facilities

3.1 Unix Farms

3.1.1 Linux

The Tier1A Linux service in July 2004 consists of:

- 256 dual processor 2.8GHz Xeon systems (523 KSI2000)
- 80 dual processor 2.66GHz Xeon systems (160 KSI2000)
- 156 dual processor 1.4GHz PIII systems (200 KSI2000)
- 90 older dual processor systems of various ages (72 KSI2000)

The 2.8GHz systems are just coming on-stream and are not reflected in the CPU plots or descriptions below. The service now consists of:

- The Tier-A service - a “non-grid” service providing front end hosts for job preparation followed by submission to the PBS batch service. This is hosted on the 1.4GHz PIII systems and the old legacy systems. It mainly provides CPU to Babar and other non Grid enabled experiments.
- The LCG Tier-1 service based on the 80 2.66GHz Xeon servers. The LCG service is beginning to become a real production service within the Tier1 and substantial additional

resource will soon be added. Access to the service is solely by Grid Certificates via an LCG resource broker.

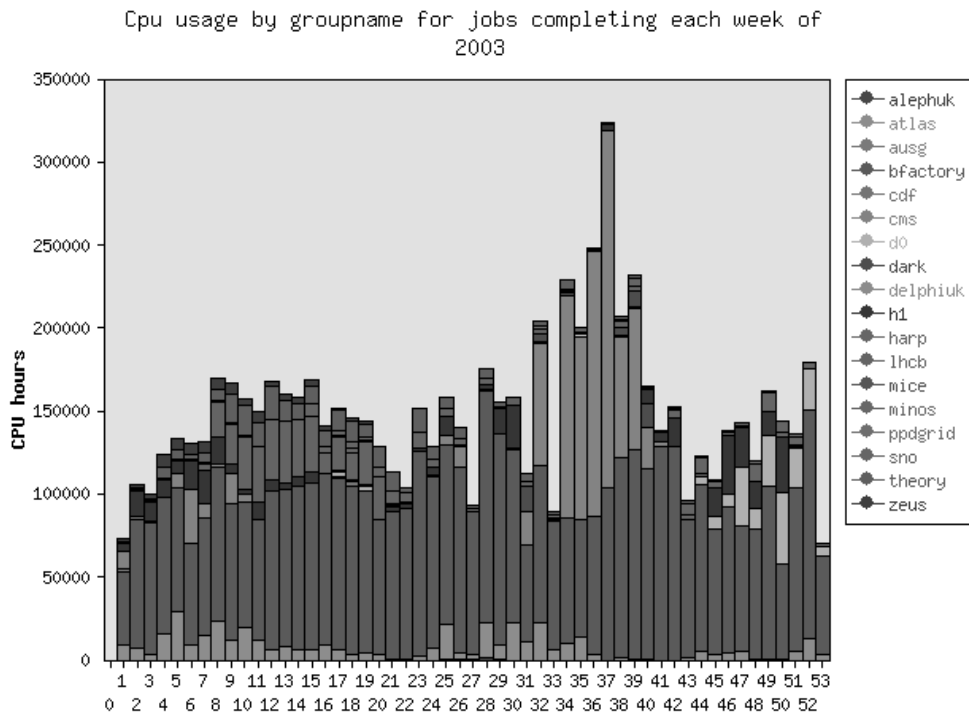


Figure 1: CPU Utilisation (all farms) 2003

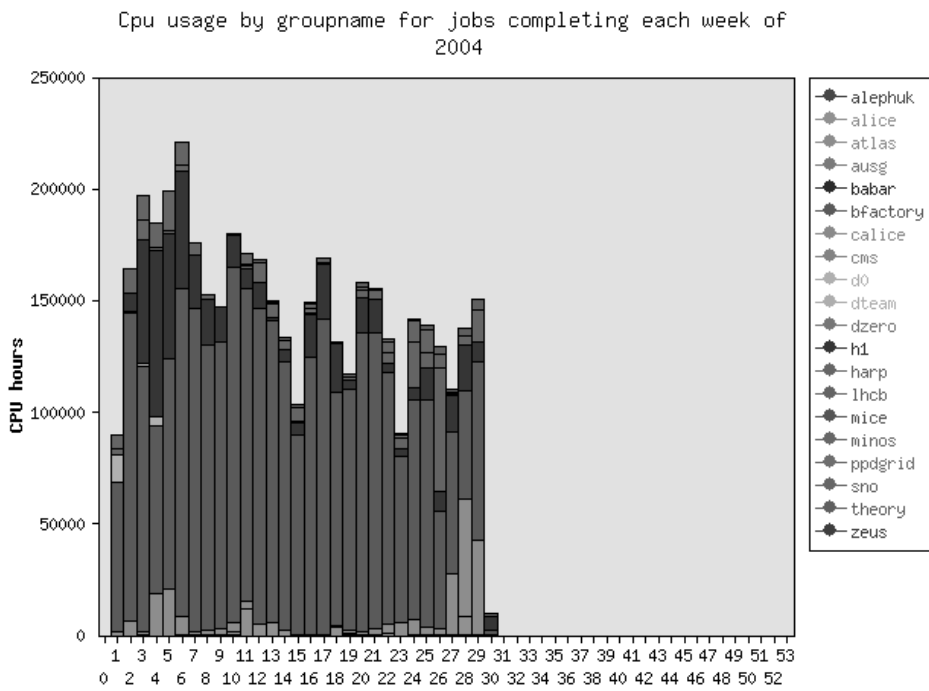


Figure 2: Tier1 CPU Utilisation (all farms) 2004

3.1.2 Sun

The CSF Sun service is run for the BaBar collaboration and most of the equipment has been purchased by BaBar with funds from JREI and JIF. It consists of 4 quad processor SUN 420R CPU nodes and a 6 CPU E4500 disk server hosting 6TB of disk. This service is now over 5 years old, and will be terminated in September 2004.

3.1.3 CDF

As a result of the TQUARC JIF award to CDF, 6TB (nominal - 5.2Tb available as data space) of RAID disk, hosted by 2 server nodes, and a 16 node batch farm, both purchased from IBM, were installed during August 2001. In August 2004 the maintenance contract will have ended and it will be no longer financially viable to host this relatively small amount (in available disk capacity) of specialist equipment at RAL. It is planned to distribute the hardware between the 4 universities with CDF groups.

3.2 Disk

The Tier 1 and Tier A disk services are based on commodity hardware. Linux PC servers, each hosting a pair of external SCSI/IDE RAID arrays. The farm is configured as follows:

- 26 Linux dual 1.26GHz servers each hosting 2 Fibrenetix SCSI/IDE disk arrays. These provide 1.6TB available disk capacity per server.
- 11 Linux dual 2.2GHz servers each hosting 2 Infortrend 6300 SCSI/IDE disk arrays. These provide 3.8TB available disk capacity per server.
- 20 Linux dual 2.8GHz servers each hosting 2 Infortrend Eonstor A16U-G1A3-M1 SCSI/SATA disk arrays. These provide approximately 7TB of available disk capacity per server. These servers are just (July 2004) being brought into service.

In total the service provides approximately 198TB of file system after all system overheads.

Considerable experience in the operation of commodity disk servers has been built up at RAL. The operation of this equipment in a production environment continues to be challenging. The Tier1 sees drive failure rates of about 2% per annum (an MTBF of about 400K hours), this is coupled with a range of complex problems in almost all parts of the I/O chain from the disk drives themselves to the Linux operating system. This experience is consistent with other large sites such as CERN and CNAF.

3.3 Network

The local area network is based on a commodity switched network with gigabit backbone switches. This continues to provide excellent reliability and sufficient performance to meet the services requirements. However it is likely that over FY 2004/2005 this network will become saturated and require some re-engineering.

The RAL connection to SuperJANET 4 via TVnet (Thames Valley MAN) provides 2.5Gbit/sec link.

The main site firewall was upgraded in late 2003 to a Netscreen 5400, capable of delivering up to 8Gps of throughput. Significant problems have encountered with the interoperation of the GRID middleware and the firewall. These had a severe impact on LCG farm utilisation for several months until resolved in May 2004.

In September RAL will have a connection to UKLIGHT. UKLight will allow direct tunnels across UK and internationally bypassing the regular networks so hopefully giving improved bandwidth, latency, packet loss etc. <http://www.uklight.ac.uk/>.

3.4 Security

There were no security incidents. Patching frequencies were substantially improved on the service after the deployment of Yumit, a locally written tool for monitoring and displaying the patching status of all Tier1A hosts. This was coupled with a regular monthly security review meeting and the deployment of network auditing tools such as Nessus.

3.5 EU DataGrid Testbed

The Tier1A maintained sustained membership in the development and production European DataGrid (EDG) testbeds, In particular, the development testbed at RAL contributed significantly and continuously to the development of EDG software. Along with two other European sites the GRIDPP Tier1 became one of the major contributions to this development resource. EDG received a successful review by the EEC in February 2004. EDG services are now being wound down – although similar work will continue as part of EGEE JRA1.

3.6 Atlas Datastore

The Atlas Datastore consists of an STK Powderhorn Tape Silo and associated drives and servers. A major hardware upgrade was completed during FY 2002/2003 which was described in last years report. This hardware was fully deployed by June 2003. Storage Resource Broker (SRB) and EDG Storage Element Grid interfaces were provided into the Datastore and both have subsequently been used in LCG Data Challenges.

Particle Physics currently occupies approximately 68TB (about 67% of total utilisation).

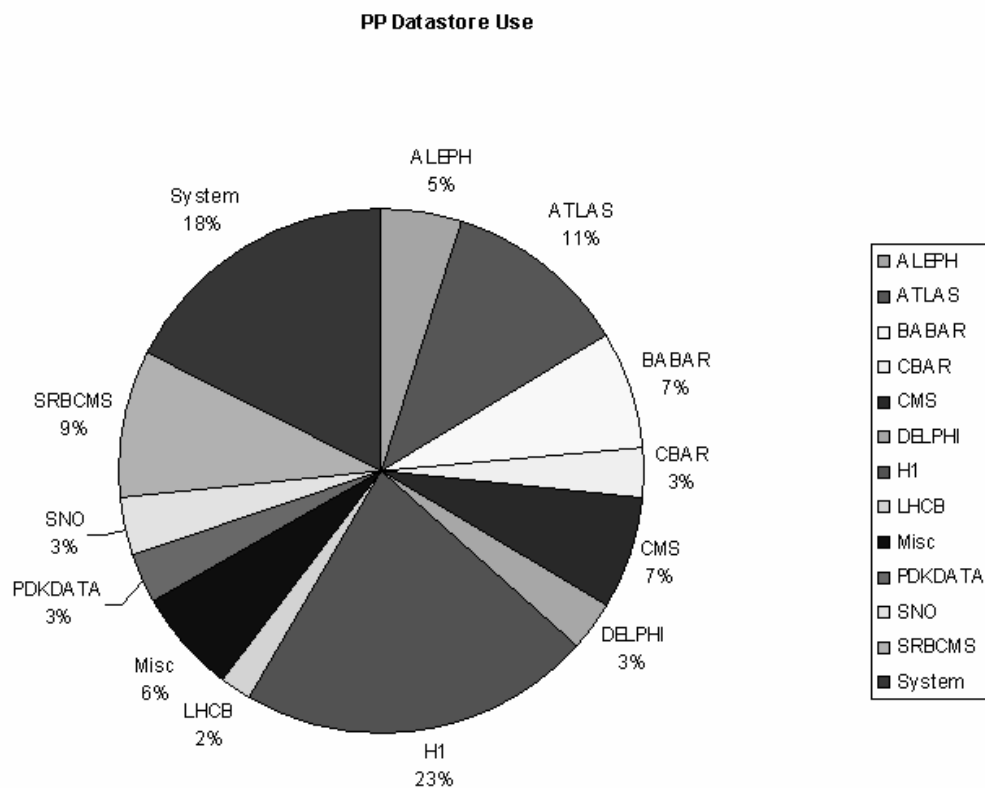


Figure 3: Atlas Datastore Use by Experiment

4. Grid

4.1 AstroGrid

Software development within AstroGrid adopted the principles of agile development, with the intention of releasing software at the end of each 3-month iteration. For various reasons, however, the first usable release to the Science Advisory Group (AGSAG) came at the end of iteration 5, which at the time of writing (October 2004) has not yet provided any significant feedback. The Project Manager is confident, however, that version 1.0 of AstroGrid software will be released to the astronomical community before the formal end of the project in December 2004.

The consortium was expanded by another four universities to bid for the continuation of the work in the years 2005-7. While the bid amounted to nearly £9M, the award from PPARC was only £3.7M, which has forced the project to de-scope and cut staffing considerably. Among those whose contracts are not being renewed is the AstroGrid representative on CNAP.

Since everyone in the project who had significant expertise in deploying large astronomical databases and in data federation will have departed from the project by early 2005, this area of the work will be among those to be severely curtailed.

4.2 GridPP

In 2003 GridPP successfully applied for funding for a second phase of the project, extending the existing project for a further three years up to September 2007. A total of £15.9m was awarded to develop the prototype grid produced in GridPP1 into a full scale production Grid for use by the LHC and non-LHC experiments with UK collaborators. In preparation for this more mature phase of the project, GridPP has already employed a full time Dissemination Officer and a Production Manager who will lead the production team. The current GridPP Experiments Board will become a more generic User Board that will set requirements and target shares, and provide feedback on the facilities provided. The existing Technical Board will become a Deployment Board which will coordinate the technical and deployment activities.

During 2003, GridPP continued to make excellent progress with three quarters of the high level tasks now complete. The deployment of LCG-1 in September, and the recent deployment of LCG-2 marked significant progress. The EDG project, partly funded through GridPP, was completed at Christmas and was reviewed as a great success in the spring. All the LHC experiments made extensive use of resources in the UK to provide leading contributions to their international data challenges. The running experiments in the US continued to make excellent progress: BaBar dominated the use of the Tier-1/a centre at RAL and the Fermilab experiments developed SAM into a real Grid system that was used to reprocess a portion of the D0 data. Over the last year, the University-based resources have been organized into four Tier-2 collaborations and their relationship with GridPP and the Tier-1 resources is currently being formalized.

5. System Managers

5.1 HepSysman

The HEP system managers had a one day meeting in Birmingham in December and a two day meeting at RAL in July 2004. The most important topic discussed at these meetings has been Operating Systems choice. The background to this is outlined next.

In October 2003 Red Hat announced that there would be no Red Hat version 10. They were discontinuing the free major releases, which the Particle Physics community had come to rely on. In stead there would be the option of paying for support and stability by buying the Enterprise Linux version, or trying to use the free 'Fedora' version. Fedora is not suited to our environment as new releases are scheduled for 6 month intervals and support for the previous release is only available for 3 months. Security patches for Red Hat 7.3 which was the predominant version used in PP would be dropped at the end of 2003.

These announcements caused some concern among the System Managers. Would this decision by Red Hat cause a divergence into many different distributions of Linux?

At the Hepix meeting in October 2003 Fermilab and CERN explained that they were entering negotiations with Red Hat for a HEP wide price for Enterprise Linux (RHEL) and that this would be there preferred option. At the same time they would work on recompiling the source RPM's for RHEL (which are still available without cost) and produce a HEP Linux with these.

By the time of the next HEPIX in May 2004 a HEP wide deal had been struck with RH which offers RHEL at a price which is similar to the educational price available to individual Universities.

However it had taken some time for this deal to be made and both Fermilab and CERN had produced their own versions of Linux from the RHEL sources. These were 'Scientific Linux' and 'CEL3' (CERN Enterprise Linux). During the HEPIX meeting it was agreed that the CERN release would be reworked as a "site" of Scientific Linux and be known as SLC3. This means that the whole of the HEP community will be running essentially the same version, more details are available on <http://linux.web.cern.ch/linux/scientific3/CEL3-vs-SL.shtml>. It was also agreed that Scientific Linux would be used for the LCG (LHC Computing Grid). The next major release will be ported from RH7.3 to SL as a priority.

The position is much clearer now, LCG2 continues to use RH7.3 for most of this year but will migrate to SL. Institutes are likely to adopt SL for there local systems and desktops in the later half of this year.

Other topics discussed at the meetings include installation of GRID software such as LCG2 as the majority of sites are participating in GRIDPP. The agenda's for the meetings are followed loosely allowing for discussion and questions, and the exchange of ideas outside of the talks is extremely valuable. The Chair would like to thank the sites, Birmingham and RAL, for hosting these meetings and Gareth Smith for his secretarial efforts.

The HEPSYSMAN group continues to provide an essential focus for sharing of information and helping to avoid duplication of effort.

5.2 Astronomy System Manager's Report

The past year has seen major changes in the organisation of the astronomy system manager's group. The lack of central support for system managers has lead to an increased workload on system managers, a lack of general cohesion and common forward direction, and unnecessary duplication of effort.

After a short hiatus, Astronomy System Managers Meetings returned in September 2003 with a meeting held at UCLan. This one-day meeting proved very useful and was reasonably well attended given that it was arranged at relatively short notice during the summer. The next meeting took place in April 2004 in Bristol. This meeting returned to the traditional and favoured, two-day lunch-to-lunch format, since this enables most attendees reasonable travelling time to and from the meeting. The local organiser also acts as meeting secretary. The secretary and the Chair prepare a report on the meeting for circulation to System Managers. The next meeting has been scheduled for September in Leicester. The Chair expresses his gratitude to Stephen Higgins (UCLan) and Rhys Morris (Bristol) for their organisation of meetings.

The Chair of the Astronomy System Managers Meetings (Barry Smalley, Keele) is now their representative on CNAP.

During the year, a private JISC mailing list (STARMAN) replaced the now unmaintained, but popular, Starlink Forum. The mailing list has proved a very valuable means of information exchange and continues to be widely used.

Two major issues have arisen during the year, 1) Choice of Linux Distributions and 2) lack of central support for software licenses. Other issues discussed during the year have including experiences with Opteron and Itanium 64bit-processors, security incidents, experiences with wireless networking, the need for a System Managers' FAQ for new and recently appointed managers, and backup strategies for the ever increasing amounts of astronomical data.

Linux Distributions

Following RedHat's decisions in 2003 to drop support of the current versions of RedHat Linux (i.e. RH 7.3, 9) and only support their commercial Enterprise Linux (RHEL), there has been uncertainty in the best way forward. There has been considerable discussion and debate during the System Managers Meetings and on the STARMAN mailing list.

Support for the widely used RedHat 7.3 ended in 2003, leaving many sites running machines without any operating system support and no clear upgrade strategy. RedHat 9 has not favoured, due to some broken features. Fortunately, www.fedoralegacy.org is providing some essential (security) updates for these unsupported releases, allowing sites to remain at RH 7.3 until the way forward becomes clearer.

At the Bristol meeting, there was a considered technical debate on the choice of Linux Distribution. The meeting was also attended by Pete Gronbech and Tim Naylor. After excluding all other distributions, discussion concentrating on the merits of the preferred distributions: RHEL and Debian.

- **Debian** is a free distribution and always will be. It is cross-platform, which enables sites to continue to run legacy Alpha and Sparc Workstations, with same operating system on all machines. Against Debian, however, was the fact that the current stable distribution (Woody) dated from Dec 2002, and so it is not good on some newer hardware. There is a steep learning curve to get competent in administration.
- **RHEL** has support for new hardware, long release cycles, and no learning curve for most managers. Against RHEL was the cost, especially since many sites do not have the budgets to retro-fit RHEL into older machines. Some concerns were expressed over the long-term future of RedHat now they are a commercial entity open to take-over, etc.

A few sites had already committed themselves to a distribution. Many others sites were still waiting at RH7.3 or 9. Several Site Managers have also expressed concern over the support for applications software. It is unclear which way ESO will decide to go.

The Linux distributions supported by applications software often lags more than a year behind. This creates a tension between the needs of the users and system security. System managers want a distribution that will be supported by a significantly long period of time to enable users to have a stable working environment without constant upgrades, which invariably cause vital software to break.

After narrowing the field down to just RHEL and Debian, no consensus could be reached regarding which of the two distributions was more suitable for the astronomy community's requirements. Both distributions had different advantages for astronomy. So it was agreed that the System Managers would recommend that we have *two, and only two*, supported Linux distributions: RHEL and Debian. Keeping a diversity of operating systems was deemed a good idea for reasons of code portability, as well as giving a fallback in case something happened to one of the supported distributions.

Subsequent to the Bristol meeting, *Scientific Linux* has become available. This is a free distribution based on the RHEL source code, which RedHat has to make freely available under the terms of the open source GPL. This distribution could provide an alternative upgrade strategy and will be a topic for discussion at the next System Managers Meeting.

Central Support for Software Licences

As a direct result of the funding reductions to the Starlink programme there is no longer anyone officially responsible for administering the astronomy-wide software license for IDL. This has been causing some delays and frustration amongst system managers and users alike. Dave Giarretta at RAL had been acting as the unofficial contact and attempting to obtain information from RSI. Other centrally provided software licences include NAG numerical libraries and the NAG Fortran 95 compiler.

On a related matter, System Managers have expressed concern that the demise of Starlink could cause problems with their supporting of applications software, with consequential effects on users.

6. High Energy Physics Computing Coordinating Committee

The chair of CNAP represents the UK at meetings of the High Energy Physics Computing Coordinating Committee (HEP-CCC). HEP-CCC is a forum for the directors of computing at major HEP institutes where members analyse potential problems and develop common or complementary strategies to solve them in the most effective way for the HEP community. This committee includes observers from the USA and frequently invites speakers from around the world but is dominantly European in membership. Topics for discussion at HEP-CCC are very wide ranging and include security, computing for the major HEP experiments, grid computing and networking performance. Most of these are now world-wide issues and so this year a new international committee (IHEP-CCC) has been set up under the auspices of the International Committee for Future Accelerators (ICFA) and had its inaugural meeting in May 2004. This has world-wide membership but the majority of members are from America, Asia and Europe. The chair of CNAP has continued to represent the UK on this new committee which has now replaced HEP-CCC. There will normally be an annual open meeting in Europe to report on the work of the committee and receive input from the Particle Physics community. However in 2004 the Computing in High Energy Physics (CHEP) meeting, recently held in Interlaken, has fulfilled this role.

HEP-CCC has a Technical Sub-Committee (HTASC) and currently David Bailey is the UK member as well as being a member of CNAP.

All HTASC documents are available on its website at <http://htasc.pi.infn.it>, including archived minutes of old meetings and technical reports. Details of the most recent meetings are now archived in the CERN agenda system at <http://agenda.cern.ch> under Committees→Technical Committees→HTASC.

Recent HTASC discussions have focussed on the developments in Grid computing and interoperability issues surrounding the various toolkits at tier 1/2 sites, the perennial problems of IT security at large sites, operating system support – in particular the cost of Linux support contracts, and electronic meeting technologies – with particular emphasis on videoconferencing. A special videoconferencing working group is assessing the available technologies with a view to writing a report with recommendations in the near future. This work is being carried out in conjunction with the LCG RTAG group which has a very similar remit.

Contributions and/or suggestions for future HTASC discussions are always welcome. If you would like to contribute, or would like further information, please contact the UK HTASC representative either directly or through CNAP.

Appendix A CNAP Terms of Reference

- To advise PPARC on all aspects of computing and networking in support of PPARC's Science programme, monitoring and reporting on the existing status of and future trends and requirements in, CPU provision, data storage, networking, software/tools/environment, the needs of the theory communities and of particular projects or experiments, and any related areas which may be specified by PPARC or CNAP;
- To act as interface between users in the Institutes and the service providers of networking, data storage and computing power. It will liaise with accelerator laboratories and observatories on these matters;
- To report at least once a year to PPARC via an annual report and to the wider community, for example through reports at the Town Meetings.

Appendix B CNAP Committees and Membership

B.1 CNAP 2003-2004

Name	Function/Representative	Institute
P.M. Watkins	Chairman	Birmingham
D. Bailey	Community Representative: Particle Physics	Bristol
D. Britton	GridPP	IC
P. Clarke	PPNCG Chairman	UCL
P. Gronbech	HEP System Managers chair	Oxford
T. Naylor	Community Representative: Astronomy	Exeter
C. Page	AstroGrid	Leicester
B.Smalley	Astro System Managers chair	Keele

In attendance:

Name	Function/Representative	Institute
B.M.Barnett	Secretary	CCLRC
D.L.Giaretta	Starlink	CCLRC
N.Geddes, G.Rickett	PPARC	PPARC

B2 PPNCG

Name	Function/Representative
P. Clarke	Chair
R. Hughes-Jones	Secretary, Video Conferencing / End-to-end monitoring
G. Fayers	End-to-end monitoring
A. Flavell	Mbone & Multi-media / Monitoring International Links
J. Gordon	Central Computing
J. Hart	Video Conferencing
R. Henderson	DESY
P. Jeffreys	International Links
J. Macallister	Monitoring Local and International Links
D. Sankey	Monitoring International Links – DESY
R. Tasker	Monitoring Infrastructure / Special Projects
D. Terrett	Astronomy Representative
F. Wickens	WWW, HiPhi, US

B3 System Managers Group (HEPSYSMAN)

Name	Institute
P. Gronbech - Chair (from 2002)	Oxford
*L. Lowe	Birmingham
*R. Morris, J.-P. Melot, R. Tapper	Bristol
*P. Hobson	Brunel
*J. Hill, S. Wotton	Cambridge
*S. Thorn, A. Khan, P. Mitchell	Edinburgh
*D. Martin, A. Flavell	Glasgow
*R. Beuselinck, K. Georgiou	IC
*R. Henderson, A. Finch	Lancaster
*A. Moreton, M George	Liverpool
*A. McNab, R. Hughes-Jones, S Salih	Manchester
*P. Gronbech, I. McArthur, J. MacAllister	Oxford
*C. Williams, T. Hartin	QMUL
*G. Smith, C. Brew, D. Kelsey, B. Saunders, K. Dunford	RAL

*S. George, B. Green, S. Johal	RHUL
* P. Hodgson, M. Robinson	Sheffield
*R. Cranfield, G. Crone	UCL
* principal contact	

B 4 Astronomy, STP and Solar System Managers

M. Murphy	Armagh
R. Vallance	Birmingham
R. Morris	Bristol
P. Bunclark, S. Cowell, H. Sanghera	Cambridge, IoA
H. Brimmer, D. Titterton	Cambridge, MRAO
R. Smith	Cardiff
A. Lotts	Durham
J. Barrow, H. Meyerdieks	Edinburgh
J. Rowe	Exeter
G. Stewart	Glasgow
T. Gledhill	Hertfordshire
I. Valtchanov	ICSTM
R. Dickson, A. Holloway	Jodrell Bank
B. Smalley – Chair (from 2003)	Keele
M. Price	Kent
M. D'Vali	Leeds
M. Mahabir, N. Wade	Leicester
A. Scott, M. Tomlinson	Liverpool, John Moores
M. Carter	Nottingham
G. Bradshaw	OU
N. Clifford	Oxford
R. Frewin	QMW
R. Ryans	QUB
D. Russell	RAL, Starlink & astro
R. Balthazor, P. Kerry	Sheffield
M. Hill	Southampton
T. Lister, R. Stapleton	St Andrews
S. Oliver	Sussex

J. Deacon, G. Milward	UCL
S. Higgins	UCLancs
G. Fuller	UMIST

References & Contacts

C.1 CNAP web sites

CNAP	http://hepwww.rl.ac.uk/CNAP/
PPNCG	http://icfamon.rl.ac.uk/
HEPSYSMAN	http://hepwww.rl.ac.uk/sysman/

C.2 ITD web sites

CSF Farm	http://www.e-science.clrc.ac.uk/web/services/csf
DataStore	http://www.e-science.clrc.ac.uk/web/services/datastore

C.3 Other web sites

AstroGrid	http://www.astrogrid.org/
GEANT4	http://wwwinfo.cern.ch/asd/geant4/geant4.html
GridPP	http://www.gridpp.ac.uk/
HEPCCC	http://tilde-djacobs.home.cern.ch/~djacobs/Hepcccw3/HEPhome.htm
HEPiX	http://wwwinfo.cern.ch/hepix/
HEPNT	http://hepwww.rl.ac.uk/hepnt/
HEPpc	http://hepwww.ph.qmw.ac.uk/HEPpc/
HTASC	http://htasc.pi.infn.it
IHEP-CCC	http://www.fnal.gov/directorate/icfa/Int'l_HEPCCC.html
IVOA	http://www.ivoa.net/
National e-Science Centre	http://www.nesc.ac.uk/
RD-45	http://wwwinfo.cern.ch/asd/cernlib/rd45/index.html
Starlink	http://star-www.rl.ac.uk/

