HEPiX Spring 2005 FZK, Karlsruhe

May 9 to 13, 2005

Trip Report

Alan Silverman
18 May 2005

As usual, these notes are taken online and reviewed only once so errors cannot be excluded. I apologise for the length but I hope you can find something of interest and relevance. In response to requests, this time I have also summarised the CERN talks. And the total length is so voluminous that I have added a Table of Contents so readers can skip to those topics of interest. And as usual I have selected what I consider the highlights. I thank Helge Meinhard who shared his notes with me, certainly helping me correct some errors. I apologise to any speaker who feels under-represented or misrepresented. I have added a few URLs where appropriate and the overheads of almost all talks can be consulted online at http://iwrwww1.fzk.de/hepix/hepix.html.

Attendance at this meeting was around 100 (although not all present on all days) including good representation from the usual US sites; clearly more and more HEP institutes are coming to believe in the value of sending people to HEPiX, or perhaps the same sites are sending more people (actually, a check on the attendance list shows it is a mixture of both effects). Also the workshop at the end of the week attracted a number of people who clearly added value to the discussions.

The meeting was rather well organised, the room was functional, sited just outside the FZK main gate and well equipped for portables (it seems that practically every attendee had one). Sessions were videoed and video-conferences were arranged 3 times for particular sessions. Other arrangements included a well-appreciated shuttle bus service from and to the 3 recommended hotels downtown and the ability of the organisers to accept credit cards to pay the meeting fee on entry.

Next meetings – SLAC, October 10th to 14th. The next European meeting will be in CASPUR, Rome in April or May next year. In fact there are tentative offers for European sites until the end of decade!

Table of Contents

Highlights

A number of points which caught my attention, in no particular order included:

- The theme of the week seemed to have been collaboration. On Tuesday morning, Guy Wormser, chairman of IHEPCCC, asked if HEPiX would be willing to act as technical advisors to IHEPCCC on specific questions where it has expertise. Two immediate examples were given, the status of Linux in HEP and the idea of a virtual organisation for HEP physicists. The meeting agreed with an earlier Board decision to accept this role with certain caveats and actions were agreed to respond to IHEPCCC's first two questions.

- On Tuesday afternoon there were reports and discussion on the most recent successful HEP collaboration, the production, distribution and widespread acceptance of Scientific Linux. Discussions focussed on which versions would need to be supported in the years up to the startup of LHC.
- On Wednesday there was a brief discussion of a limited collaboration being setup on the suggestion of Jos van Wezel to discuss the data storage needs of Tier 2 sites. This should happen over the summer, chaired by Roger Jones of Lancaster, and report to the next HEPiX in the autumn.
- On Monday the cat was caste among the pigeons when Alf Wachsmann of SLAC said he was looking into a new monitoring tool for SLAC; he had decided to invest some effort in Nagios, had found some serious drawbacks (scalability limits) and he was wondering what to do next. Since his talk was sandwiched between a most attractive CERN talk on Lemon showing what it could do and a Fermilab talk on the successful experiences of NGOP, each of which monitor in production mode clusters totalling more than 2000 nodes, an open discussion on the possibilities of a HEP collaboration on monitoring tools ensued. Details and some questions are given in the text and discussions were still continuing on the buses back to the hotels.
- We seem to be focussing down on a finite number of tools for system administration tasks, at least in most sites. These include ROCKS (public domain) and Quattor (CERN) for system installation, Ganglia and Nagios (both public domain) for monitoring and management tasks, dCache (DESY/Fermilab) for data caching, SpamAssassin for mail spam fighting, etc
- The workshop at the end of the week, organised by Tony Cass, covered all the batch systems in the major and some not-so-major HEP sites as well as raise some issues concerning the evident gap between local job scheduling and grid scheduling. There were some rather good presentations by people who clearly knew what they were talking about but it was unfortunate that there was no grid scheduler author present to listen to the issues of grid scheduling as expressed by the site administrators who have to make grid scheduling co-exist with local environments. Nevertheless there was a lot of lively discussion and plenty of ideas to reflect upon.
- The recent DoE budget announcements included an end-date for BaBar (2008) and a plan to refocus SLAC away from HEP to Basic Energy Sciences, apart from an interest in the projected ILC (International Linear Collidor). They, and possibly Fermilab, are expecting 5% staff cuts to be announced. JLab is also being financially squeezed. The electrical accident at SLAC last October finally cost 6 months running time but beam has now restarted

Opening Session

The meeting was opened by Prof. Dr. Reinhard Maschow, the Director of the lab. He said that 5 years ago, he would have been surprised to see us because FZK is not known as an HEP site; rather they are an inter-disciplinary centre running research programmes in various fields but not specifically HEP. Their range of interests has in common the need for good computational facilities, hence their recent interest in supporting Grid computing for German groups participating in LHC. They are a Tier 1 Centre for LCG[1], participating in the various LCG Service and Data Challenges. They are very pleased that HEPiX agreed to hold a meeting in FZK and they take it as recognition of their participation in the world of HEP computing. He hoped HEPiX would accept IHEPCCC's invitation to collaborate as a technical body (to be discussed on Tuesday) but he recognised that HEPiX is an independent body.

[1] LCG - LHC Computing Grid

Site Reports

FZK/GridKa – Their computing centre (GridKa) was founded in 2001, not only to participate in LCG but also to supply computing facilities to non-LHC experiments such as BaBar, CDF, D0, Compass. Although their CPU capacity has multiplied by a factor of 3 to 4 over the past 12-18 months, their hardware plan needs to be upgraded in line with the demands of the LHC experiments. The speaker described the computing centre infrastructure, including their flexible power scheme based on power rails suspended from the ceilings and their water-cooled system racks (as described at the last meeting).

They have a mixture of AMD and Intel CPUs, all equipped with CERN Scientific Linux 3.0.4 and LCG 2.3 (going soon to 2.4). There is one exception on a grid file server where they run Redhat Enterprise Linux in order to get IBM support. For a batch system they started with OpenPBS but it did not scale, so they moved to PBSPro with 5000 licences; they use dynamic queue limits and fair share scheduling.

They suffered many problems with IDE discs which were hard to analyse due to poor error analysis software but eventually blamed on the mechanical assemblies. They use LT01 and LT02 tape systems using dCache and Tivoli. They use Ganglia for system monitoring and Nagios for system management, along with some home-built tools. GridKa has helped setup 10 sites in the German/Swiss federation supporting HEP, astrophysics and bio-physics; they act as a Regional Operations Centre (ROC) in LCG and are responsible for setting up the Global Grid User Support Centre for LCG. They also participate in EGEE NA2 (user training).

CASPUR – while their IBM SMP system is fairly stable and unchanging, their HP SMP is being upgraded to a 32 node system with Alpha EV7 CPUs, still running Tru64. They have no further interest in Itanium, rather preferring to expand their Opteron cluster. They offer user single sign-on based on Heimdal Kerberos 5 and offer AFS and NFS access on all platforms. Purchased an 8TB IBM SANFS (StorTank) and 4 SATA/Fibrechannel arrays. They have replaced some LT01 tapes with LT02s and will replace the rest by LT03s in the autumn. They ship their own Linux package (CASPUR BigBox) based on Redhat Enterprise Linux which they claim is fully compatible with Scientific Linux, including also the CERN flavour of this and they have bought some Redhat subscriptions for reference and to get access to critical patches. They are developing a taste for SuSE Linux, preferred by the CASPUR Application Sector as a better fit on Opterons, but they don't have plans to build a SuSE distribution. Working on various collaborations including one with CERN and RZ Garching on Object Shared Devices (OSD) in accordance with the T10 specs and in particular studying the integration of OSD in AFS.

TRIUMF – Although they participate in the ATLAS data challenge, recent budget cuts mean the ATLAS data hub cannot be completed as planned. Their PC base is heavily Dell-based, still largely based on Scientific Linux version 3 but they would like to move to version 4 as soon as possible. Since the last meeting, Opteron servers have largely replaced Xeons although they still have a large variety of configurations and it took time to stabilise them all. There is a suspicion that Opterons may be I/O limited and investigations are in progress. The speaker (Corrie Kost) presented some tables on I/O tests he had performed as well as his reasons for preferring RAID 6 over RAID 5 – see overheads.

SLAC – Beam finally restarted some 3 weeks ago after being off for 6 months since the DoE inspectors were called in to investigate a serious electrical accident last October. There is also a DoE budget cut which, along with greatly increased power bill, will result in a 5% staff reduction. New initiative is the Linear Coherent Light Source (LCLS), to come online in 2009 just after BaBar ends in 2008. Thus the lab will change emphasis. They will maintain their interest in astronomy (the Kavli institute including a Dark Energy Search). Plus plans for an ultra-fast ultra-small science institute based on the LCLS laser. A major lab re-organisation is expected shortly but in the medium term the budget should rise.

All the projected new activities are data-intensive and the goal is to find common solutions. Therefore their large memory project described at CHEP 2004 and at the last HEPiX meeting, renamed the PetaCache Project, is going ahead with a 1TB cluster in place, offering 10GB/s on a single server with 100 micro-sec latency and scaling to 100s of servers. Currently seeking first real users and funding for a 20TB demo model. They are now a node on the US OSG[2] offering 250 CPUs and soon 10TB of discs. Solaris 10 x64 is being used on the PetaCache system and the transfer of Solaris tools and knowledge is straightforward. Supporting SGI Linux for the Kavli Institute (based on Redhat); looking at 64 bit mode Linux on Opteron; and investigating RHEL[3] version 4 based on the Linux 2.6 kernel.

Fermilab – like other sites, budgets are tight and like SLAC they are preparing for a 5% staff cut. The DoE in their latest review unexpectedly cancelled the Btev project. Two new computer rooms have been installed in existing experimental halls, one for LQCD and the other for farms for CDF, D0 and US/CMS. Various production farms are being expanded with more nodes, the one for CMS with an extra 300 nodes on order. Average growth per year is around 1000 nodes across all farms, with concomitant growth needed in infrastructure (not only buildings but also network – WAN traffic doubles every year; for example), Various clusters now use or are migrating to the Condor batch system. 58TB of dCache space is used for CMS with another 38TB on order. Counting the ENSTORE Mass Storage cluster, there is a total of 225TB of dCache front end disc storage. ROCKS is used for installation. They have leased dark fibre links to Starlight and hope to offer 10Gb/s to experiments soon. As well as FBSng use dropping (described in more detail in the batch workshop), so is LSF (down to 158 licences from a peak of around 700). NGOP now monitors 2362 hosts. On the Windows side, migration to Windows server 2003 is underway as well as updates to various Windows and third-part utilities.

CERN – Most central farm nodes are now running SCL3[4] with only a few remaining on RH 7.3. The refurbishment of the main computer room is well advanced and almost able to accept the predicted influx of new servers..Setting up a test cluster for architectural studies and tests are about to start on various Intel and AMD processors. Serial console cabling is almost complete and in use by the insourced system administration team, now numbering 9 people. Acceptance tests based on va-ctcs no longer seem correct and the presenter (Helge Meinhard) wondered what other sites were using. CERN has recently installed 75 disc servers with 360 TB SATA discs and is awaiting 225 Nocona 2.8 GHz farm nodes. Currently tendering for 2 times 350TB of disc space, concentrating on SATA discs as these give more stable performance, thought to be due to better assemblies (fewer cables). CERN expects to issue even larger disc (1 PB) and CPU server (1000 nodes) tenders for delivery next year and Helge has estimated a price of 1.40 Euro per GB of usable installed RAID 5 disc capacity including server support up to the Gbit interface but excluding infrastructure (racks, network switches, etc).

Openlab's Itanium cluster was recently able to achieve 660 MB/sec to Caltech. CASTOR achieved over 450 MB/s in a recent data challenge, sustained for a week. There was a serious of spurious crashes of AFS, eventually traced to a system update but AFS migration to Kerberos 5 was finally completed. The public SUNDEV Solaris service is being closed in the summer. InDiCo has become a very popular product for supporting conferencing. Linux and Windows service updates will be reported in full in separate sessions although it is worth pointing out recent successes in reducing SPAM mails and the recent drop in these, thought to be linked to the successful prosecution in the US of a well-known Spammer. He described the dramatic rise in the number and seriousness of security intrusions, including one making use of an unknown weakness in the LSF protocol. LCG software is now deployed at more than 120 sites, on more than 10,000 CPUs.

---

[2] OSG – Open Science Grid
[3] RHEL – Redhat Enterprise Linux – the distribution from Redhat
[4] SCL – CERN's packaging and distribution of the Scientific Linux distribution from Fermilab

RAL – The RAL PPD group is part of the UK Southern Tier 2 network. The 550 node RAL Tier 1 centre is a separate group and they are replacing old PCs and SUNs with more modern CPUs over the coming summer; some of these old PCs are being recycled in outlying outposts such as the Liverpool and Glasgow HEP sites. A major Tier 1 upgrade is expected next year so investigations are in progress to decide which architecture this should be. Except for NFS servers, they run SL[5] 3.0.3 and LCG 2.4, public domain tools such as Ganglia and Nagios and the Maui scheduler. They plan to move from NFS to SRMs for data access; there are load problems with their Objectivity server (appears to be a memory issues) and more servers are being added. They have a project to look at SL 4 for servers. The ATLAS data store is up 16TB total disc space and 6000 tape silo slots with 200GB per tape. They are considering what should be the next generation of tape drives. The Tier 1 WAN should be upgraded shortly, aggregating 1 Gbps links to form a 2Gbps, later 4Gbps, backbone and a later update is planned to 10Gbs. They were connected to UKLight via 2 1Gb/s links for the recent LCG Service Challenge and achieved 80 Mb/s sec sustained for 2 weeks, short of the target but they have understood why. Now preparing for the next Service Challenge which will bring in Tier 2 sites.

GSI – They have updated their IMAP server from Exchange 5.5 to 2003 and there is more effort being added to the SPAM war on the Linux mail server (MTA); they have introduced greylists which contains domains from which spams appear to be more frequent and this seems to reduce spam by some 30%. They have established a Wiki service (talk later in the week). The Theory group has adopted 64 bit Opterons running Debian in 64 bit mode and no problems have been found as existing programmes were migrated over. Further use of their SATA-based file system has reinforced the reliability message of the last meeting.

LAL – 6TB of SATA disc is on order and a tender is out for ~20 new dual Opteron servers. They are moving to Scientific Linux, and preparing to do this with Quattor. Further upgrades will be driven by the needs of LAL as a Tier 2 LCG site in a partnership with DAPNIA and LPNHE (see later talk). Moving to Subversion, a replacement for CVS, and also to using certificates for Wiki authorisation. They are switching their anti-virus tool from F-Secure to McAfee. Had planned to install CERN's InDiCo for conference scheduling but delayed it because of problems so they have installed Agenda from the same CERN team in the meantime.

BNL – They support RHIC experiments and US/ATLAS. They are currently ramping up the ATLAS farm, for example they recently added 288 compute nodes with 144TB of local disc storage. There is also 39 NFS Solaris servers with 220TB of RAID5 storage available. They suffered NFS scalability issues but have now moved to Panasas software[6], offering 81TB of storage and they are very happy with this product; it allows load balancing, uses inexpensive SATA discs and offers NFS access where the Panasas client is not available although there are occasional hiccups and user quotas are not currently supported. Still running HPSS on their 4 STK Powderhorn tape silos and planning an upgrade to HPSS version 5.1 after RHIC run 5. They plan to buy a new generation STK tape silo later this year, probably an SL8500. The RHIC central analysis and reconstruction farm consists of some 2000 dual CPU Intel systems with 714 TB of local disc storage in a variety of architectures. The CAS/CRS farm now runs Scientific Linux 3 and this is currently being installed on the ATLAS farm. In the middle of replacing their LSF with Condor. Kickstart is used for system installation, Ganglia for monitoring. A complete migration from NIS to Kerberos 5/LDAP is underway. To support ATLAS, their external network was upgraded from OC12 (622 Mbps) to OC48 (2.488 Gbps). ATLAS in BNL is a member of OSG. dCache has been deployed for both ATLAS and the PHENIX experiment.

CNAF – this was a first site report from CNAF, a multi-experiment Tier 1 site. They have a large (underground) computer centre but currently insufficient power to fill it. They have some 700 dual CPU

[5] SL – Scientific Linux – Fermilab's packaging of the source corresponding to RHEL
[6] Panasas is described in more detail in Maslennikov's talk on Tuesday.

servers of various speeds and a variety of disc architectures. Networking is based in the GARR-G backbone and the current WAN is 2.5 Gbps, going soon to 10. The INFN Tier 1 site has 1 Gbps of this and uses an additional 1 Gbps during SCs. They use Quattor to install nodes and most recently adopted SLC 3.0.4 and LCG 2.4.0 (successfully installed over 500 nodes in one day) with LSF as the batch manager. CASTOR is used as the storage HSM. The speaker presented the results of some NFS stress and disc storage hardware tests; the speaker stated that "the NFS protocol may scale to over 200 nodes without aggregate performance degradation". They compared the performance of GPFS from IBM and the Lustre file system; the graphs (see overheads) would appear to show that Lustre won fairly easily. For the recent LCG SC[7], they used 11 SUNFire V20 systems and successfully achieved the goal of this test – 100MBps disc to disc sustained for 2 weeks.

Lemon Monitoring Task

Lemon is used for monitoring performance and status in a distributed system, part of the ELFms tool suite originally developed under the banner of the EDG project[8]; this talk, presented by Harry Renshall, was an update on recent developments. Various levels of data are available to different communities from users and managers to system administrators. He described the different Lemon components. It currently handles some 260 metrics with only a few more expected to be added. It currently monitors some 2200 systems in 100 clusters gathering some 70 metrics per node, a total of 1.5GB per day. It integrates into a legacy alarm display scheme (SURE) but a new alarm display tool is in development. He showed some typical screen shots and described some use cases, for example checking kernel version on system reboot and scheduling an update if the kernel version is too low. In another example, he showed how easy it is to track down one node in a cluster with extra high load.

A web interface via plug-ins has been added to search for related information, for example in Oracle or Remedy (CERN's chosen trouble ticket tool) databases. Results can be graphed to identify easily particular problems or instances. Lemon has a correlation engine to generate recovery actions, for example to restart a dead daemon; there are currently 70 recovery actions defined today. When this feature was first added, the number of trouble tickets halved. Metrics can be written by and for specific groups, eg. to measure various ORACLE events for Database Suppoprt. Virtual clusters can be defined or redefined inside Lemon, for example to take a node out of a cluster for an intervention and put it back afterwards.

Host and Service Monitoring at SLAC

Although they already have Ranger (written by Chuck Boeheim and described at various previous meetings) Alf Wachsmann stated that this is not a monitoring tool as generally understood. BaBar has started using Ganglia but this has no alarms. "SLAC is late in the game of monitoring" he admitted. Alf has sounded out system administrators at different conferences and not found a consensus. He has asked users what they would like to see monitored. After collecting the various inputs, he started looking seriously at Nagios. This has modules for data collection, alarms and status display. There are plug-ins for probes, storage and graphing.

He described how Nagios works, pulling monitoring data centrally to the master. Probes work locally on the client nodes and can decide themselves if the result is ok, serious, critical and so on and can optionally return performance information. Probes exist for most common operating systems and more can easily be created. Data can be stored in simple or structured databases. It comes with a web interface

[7] SC – Service Challenge; LCG is performing a series of these, most recently SC2
[8] EDG – European DataGrid project, CERN's first Grid project

to display current state and a histogram tool to display up/down history. More facilities are available with plug-ins.

Problems with Nagios include

- Poor scaling of the web interface in the current version; the beta version of the next version is better but only somewhat.
- Lots of probes do not scale
- Need to write your own Windows probes
- No recovery actions
- And database support seems to have been removed in the new version.

In summary, currently at SLAC they have Ranger on all UNIX nodes, Ganglia on BaBar nodes, Nagios on SCS and students working on a Ganglia interface for Nagios data. Where to go next? [There was a somewhat heated discussion in answer to this but a full discussion was postponed to after the NGOP talk.]

NGOP Experiences

Marc Mengel, one of the original authors of NGOP has been considering its future and possible alternatives or upgrades. According to the published documentation, Ganglia is designed for up to 512 systems, Nagios for 650; like Lemon, NGOP currently monitors over 2000 systems. He said that although it was designed for up to 10,000 systems, he suspects it might saturate at around 5000. Several other, relatively minor, problems have been found (listed in the overheads) and he now sees several places where it could be better designed, handling configuration data for example. He also believes that the tool should be expandable to a grid scale for some user communities although site operators need not to be distracted by offsite views. He then proposed a new architecture where hosts would be grouped in a hierarchy which would scale much better. These groups could of course be different sites. At the lowest level, an “agency” using the SNMP protocol and a log file watcher could watch up to 200 nodes and a site event server could monitor up to 5000 agencies. In conclusion, he invited anyone thinking about a new monitoring tool to get in touch. This led naturally into …

Monitoring Tools Discussion

After the talks on Lemon and NGOP and the question from SLAC on which tool to move forward with, the question was asked - why invest some effort in a tool such as Nagios having found some serious drawbacks (scalability limits). Why not investigate collaborating with either Fermilab on NGOP or with CERN on Lemon. Among the points made were

- How dependent is Lemon on Quattor? Harry Renshall thinks it is but probably performs better with the rest of the ELFms suite (which includes Quattor); he will find out.
- How easy would it be to tailor Lemon for another site and package it for installation elsewhere than CERN? Again Harry will investigate.
- It would be interesting if we could get a practical example of a non-CERN instance of Lemon – any volunteers?
- Lemon, or NGOP, will never gain wide acceptance, even by other HEP sites, until they become truly open source and are offered to sourceforge.org. But would CERN or Fermilab management condone this, even if support was not promised? There may be no “legal” obligation for authors of sourceforge software to maintain their submissions but there may well be “moral” pressure and it may reflect badly on an institute that did not in fact offer any support.

- Nagios remains attractive for small sites with a limited number of nodes and limited available manpower. But how long will “small” HEP sites, especially those linked to major experiments, remain small compared with suspected scaling limits of such tools?

Day 2

Security – Bob Cowles

Day 2 commenced with the bi-annual Bob Cowles show. And as usual he started by displaying the various passwords he had trapped and/or decrypted the previous day sent by attendees in the room! This led into an explanation of the various tools and protocols which expose user passwords and how to protect these. He then moved on to some of the more recent major incidents, starting with the so-called Teragrid Incident which started with an ssh-based keylogger and ended with session hi-jacking.

Phishing, spyware and Google hacking are all showing an increase. New technologies to be aware of (afraid of?) include Bluetooth, RFID, VoIP and iPODs. Finally he listed some well-known vulnerabilities in different operating systems and web browsers.

ENSTORE, dCache and SRM at FNAL

Enstore manages access to data files and volumes; dCache provides a user interface to reading or writing cached files and SRM provides a consistent interface to the underlying storage systems There are 3 instances of each of ENSTORE and dCache, separating the major user groups from each other and from the rest of the world. ENSTORE manages 6 STK tape silos with a variety of drives, controlled by 127 Linux boxes mostly running Scientific Linux. ENSTORE management and monitoring is performed using web pages, cron jobs, resource usage plots and a graphical image of the hardware components with their current status displayed. There is currently some 2.6 PB of user data on tape in 10.8M files over 25,000 volumes and with traffic peaking at 27TB of data in a single day. In all of this, limits are pushed and constant checking of error logs is essential.

dCache works over ENSTORE to provide a buffer between the user and tape. [It also works in standalone mode.] Users can access it via SRM, Gridftp, kerberised ftp, http or the native dCache protocol (dcap). At FNAL there are 100 dCache pool nodes interfacing to 225TB of disc. The transfer record for dCache is 60GB in a single day. They have found that XFS is the best file system for dCache discs and direct I/O is the best manner to access local dCache discs

SRM is the storage resource broker that sits on top of other storage systems such as dCache or ENSTORE (the ENSTORE interface is now in development). CMS Service Challenges have achieved 50MB/s sustained rate from CERN through SRM and dCache disc to ENSTORE tape and up to 700 MB/s when the ENSTORE tape staging is dropped.

Overall they are happy with performance, scalability, availability and data integrity.

CASPUR Storage Tests and Evaluations

Next was Andrei Maslennikov's annual update on his testing activities in Rome. As usual, there is a lot of detail in his overheads. For most tests they used Infortrend controllers, SATA discs, dual CPU 3.4GHz PC servers and non-blocking GigE switches. He went through the results for the various tests – for example, comparing different controllers showed that Data Direct is faster and more performant under most loads; it is also more expensive but you get more flexibility. He reported on a file system test comparing NFS, AFS, Lustre and Panasas. Panasas is a hardware/software appliance; physically it is a blade server with a combination of storage blades and director blades. You must install a Panasas module on client nodes to use their native protocol. Lustre performs best, as expected, but Panasas compares well. AFS is, of course, well behind for very well-known reasons. Stress tests hung up NFS for reasons not yet understood. Lustre can be tuned but has undocumented parameters!

IBRIX at FNAL

Reported on by Lisa Giacchetti. US/CMS were looking for a global file system and storage solution with the same features as AFS but better performance. IBRIX chosen for evaluation in front of other commercial solutions, including Panasas (described in the previous session) because it required no specialised hardware required, enabling them to use their choice of hardware (which could then be re-used if the tests failed) and this also reduced the price relative to other vendors' solutions.

Lisa described in some detail the main features of IBRIX. The test setup consisted of 10 RAID5 arrays of 5 discs each; configured as so-called IBRIX segments: 5 segments are combined to make file systems. NFS was chosen as the preferred access protocol. The first installation in December 2003 was followed by several months of high-level system configuring and testing and the first users were given access in Spring 2004. By August, serious problems were recognised with NFS access and despite good IBRIX response, these were considered a block to moving forward.

IBRIX agreed and proposed waiting for the next release which duly arrived early in 2005 and users were re-connected. The system is currently stable, more users are being gradually added and there are plans to expand the configuration as we move towards LHC data taking. Unfortunately the speaker was unable to present performance figures due to their emphasis thus far in making the system stable and usable.

IHEPCCC/HEPiX Collaboration

Guy Wormser, chairman of IHEPCCC (a committee mostly of the directors of the major HEP institutes throughout the world and which reports to ICFA), repeated in a talk the request he had made by e-mail to the HEPiX Board last October, namely that HEPiX act as a sort of technical consultative body to IHEPCCC, somehow replacing the role filled by the previous HTASC sub-committee of HEPCCC, the predecessor to IHEPCCC. Guy first presented a description of IHEPCCC's mandate, membership, objectives and working groups.

IHEPCCC would occasionally formulate questions on technical subjects which they hoped HEPiX could respond to. The first two questions are:

1. What is the status and future of Linux in HEP?
2. Could HEPiX discuss the idea of a HEP virtual organisation, a subject for which IHEPCCC has created a high-level working group chaired by Manuel Delfino? Such a VO should cover all bona-fide HEP members, address the problem of the travelling physicist and so on. Could be based on certificate technology but would require support by the institutes to keep the membership list up to date. Any solutions would need to be discussed with appropriate technical experts.

An open discussion was then chaired by Alan Silverman. HEPiX is generally positively disposed towards the proposal. Alan himself offered to write, in cooperation with the Fermilab and CERN experts, an answer to the first question. For the second, HEPiX will contact Manuel to see where it may be able to help and perhaps schedule talks on this subject at the next meeting. Further topics would indeed be considered but finding volunteers who could add value is essential and support by lab management for this activity is vital. This condition, backing my management, feeds back directly to IHEDPCCC members because frequently these are the same people. The point was also made that for some of the listed topics, working groups already existed or were postulated in another for a such as EGEE or GGF and HEPiX should collaborate with such groups when they exist, not try to compete.

Xrootd Data Serving

Chris Brew described his experiences deploying this at RAL for use by BaBar. Xrootd stands for extended root daemon and is a suite of tools for serving data. It was written at SLAC and Padova during the migration from Objectivity to Root I/O. At RAL it deployed painlessly across all 26 file servers with 75TB of disc and first use seems to work as expected. For users, the load balancer also appears to work correctly so users do not suffer if a server goes down and queues do not get stopped in such cases. This benefits users and obviously eases the life of the system administrator as well. In SLAC there is over 300TB behind Xrootd and IN2P3 has over 100TB.

SATA Storage Experience at FNAL

In her second talk of the day, Lisa Giacchetti described Fermilab's experiences with SATA storage. SATA is found in commodity or mid-level storage configurations (as opposed to enterprise-level) and cannot be expected to give the same performance as more expensive architectures. SATA controllers can be FC or PCI/PCI-X. She then listed a number of pitfalls which may be met on the way to purchase a SATA configuration – imperfect firmware, misleading claims by vendors, untested configurations, known but unfixed bugs – all summarised as "you get what you pay for". A number of suggestions were made for improving the prospects of a successful SATA purchase –

- careful selection of vendor
- firmware upgrades
- consider paying more if you can be sure of reduced ongoing maintenance costs (including human costs)
- it is vital to understand properly your needs and estimate the cost and effect of disc loss balanced against investment and maintenance costs (disc loss may be acceptable, data loss is not).

Fermilab Scientific Linux Update

Marc Mengel presented an update of SL. SLF 3.0.4 (F to denote a release from Fermilab) was released in February; it supports 64 bit Linux. Development has started for SLF 4.0 but the basic SL 4.0 package was made available in April. So far only Fermilab and CERN among HEP sites perform their own customisation but PSI has also done this.

In the near future they expect bug fixes to RHEL 3 to be released and these will be packaged into an SL 3.0.5 release but this will probably be a fast rebuild without many new features. In due course they would expect to release a 4.1 release. For people planning to perform their own packaging, he laid out various rules or guidelines they would like to enforce to maintain compatibility between the different packages. They would like to build a tool to test compatibility and would welcome possible collaborators on this or other support aspects such as the web pages, FAQs, etc. He put forward the idea

of a Scientific Linux workshop in Fermilab in the autumn. The FNAL team welcomes all contributions, errata, suggestions, etc.

CERN Scientific Linux Update

Jan Iven then described the status and plans for the CERN distribution. SLC3 was finally certified on Nov 1st 2004 and declared the default version. A number of 7.3 instances will be around for a while yet but support for this and other older versions will stop in December 2005. Since certification, the main batch and interactive clusters have been steadily switching more and more nodes to SLC3.

The CERN experience with the Redhat TAM support contract has met with mixed success, the person is responsive and knowledgeable but internal Redhat procedures are less responsive than we would like or expect, for example sometimes long delays to get bugs fixed in the source (as opposed to a providing a workaround).

Why was Scientific Linux a success? Jan put forward a number of possible reasons

- Last year many sites were in a position ready to move forward and all were affected by changes in Redhat's licence policy
- Sites preferred to remain if possible with the Redhat flavour
- Slow progress in negotiating an acceptable contract with Redhat (apart from some labs who decided to take up an offer to DoE sites – only SLAC?)
- The potential for SL (and consequently SLC) to be compatible with RHEL was attractive along with built-in flexibility for local tailoring
- No forced commercial contracts needed
- The example of major sites adopting it with the underlying implication that they would fix problems (although neither FNAL or CERN was in a position to guarantee support)

There is still a risk of divergence down the line but since this was always the case, it is not a show-stopper. CERN in fact inserted such a divergence, forced by the need to interface to the LCG POOL utility, but it was accepted and will be fixed shortly. Is there however a problem of dependency on particular point releases, or on particular packages in the release (for example the chosen browser)? We may require formally to define a minimal core package but hopefully that will not be necessary.

How to handle support for different combinations of SL release versions, 32 or 64 bit, different compiler versions? Do we need ongoing IA64 releases? What about AMD or Intel flavours? CERN is returning responsibility for the compiler to the experiments; the system compiler is the compiler of choice but groups may diverge to a different one if they prefer.

Should we skip SLC4 and go to SLC5 when RHEL 5 is released? Due to the constraints of LHC startup (scheduled for summer 2007) and the need to lock down to the then-current version in the autumn of 2006 at the latest, we may prefer to move to SLC5 quickly rather than stay on SLC4 for up to 3 or even 4 years from mid-2006 (will new hardware in 2007-8 even run on SLC3 or SLC4?). There are risks in either solution and lots of feedback is required, from CERN users and other sites but CERN expects to decide only much later on which version we will freeze in autumn 2006 ahead of LHC start in mid 2007.

Scientific Linux Discussion (chaired by Alan Silverman)

In the ensuing open discussion, in which the Scientific Linux development team in Fermilab participated by video link, the following points were made:

- RedHat is aware of SL, they ignore it. They even knew about it before they made their last offer.

- The LHC experiments (through their CERN representatives) have requested to stay on SL3 this year and have agreed with the proposed twofold certification of SLC4 and SLC5 and the ensuing "late" decision (formal agreement from CMS still outstanding but initial positive comments). It was clarified that it is mostly the LHC online groups (and accelerator support) that require "locking" onto a given operating system version in 2006; Kors pointed out that offline code should not have strong dependencies on the OS.
- Other experiments, such as CERN Fixed Target and BaBar, are likely to accept skipping release SL4 although the Italian part of BaBar has asked for SL4 to fix a compiler bug. It was noted that the Linux 2.6 kernel in SL4 supports laptops better than SL3.
- Corrie Kost expressed the fear that waiting for SLC5 risked, in the LHC startup timescale, not having a supported Linux release able to run on new hardware. He was somewhat relieved when both Fermilab and CERN stated that SL4 is, and SLC4 soon will be, available but not at this time certified nor guaranteed to run with a given experiment's software.
- No site represented expressed interest in IA64 (Itanium) support. But almost all expressed significant interest in using Opterons. In this respect, it was confirmed that both the Fermilab and CERN releases support both Intel and Opterons, in both 64 and 32 bit mode. How fast experiments recompile to take full advantage of 64 bits depends on the incentive of faster execution.
- Kors Bos stated that in his opinion, the experiments should be taking all steps to avoid dependencies on any part of the operating system release; further, he believed this goal was shared by the various experiment development groups.
- There was lukewarm support for the idea of a video-conference to discuss SL directions or a workshop on site tailoring and it was suggested to test these ideas via the mailing list
- A request for a security-only update mailing list was favourably received.

Day 3

Opteron Evaluation at Fermilab

The intention was to evaluate Opteron chips and in particular to compare them to the Nocona 64 bit Xeon chip and to get experience of Linux running in 64 bit mode. Although the standard SL 3.0.3 release ran well, Opteron chips have NUMA memory and this requires a Linux 2.6 kernel to take full advantage and a special release with this worked well also. They were able to run almost all of their 32-bit applications under the 64-bit kernel distribution in compatibility mode with little trouble. Standard Fermilab acceptance tests ran mostly very well and finally they performed a variety of benchmarks; the results are shown in the overheads and more details are available in their published paper. They indicate that Opteron 2.4GHz and Xeon 3.4GHz chips are roughly equivalent but the Opteron show a larger gain in performance (up to 40% in the best case) when running in 64 bit mode when codes are recompiled and optimized for 64 bit running. Further, their tests showed a clearly reduced power requirement for Opterons. And lastly, there appears to be no problems running on Linux in 64 bit mode.

Anti-Spam Measures at GARR

GARR is a research network in Italy. Their main anti-spam tool is SpamAssassin and he described how this works, how it is configured and how it uses Bayesian filters to maximize spam rejection and minimize false positives (legitimate mails wrongly identified as spam). They use 3 DCC: servers: mails

which appear with the same signature in several sites are tagged as suspect. They looked at SPF – Sender Policy Framework – where each DNS server must publish a reverse MX record to identify the authorized smtp server for that domain. This has major disadvantages for the roaming user who must then always use his “home” smtp server. A test at Salerno University showed that only 32% of mails were sent from SPF-compliant domains and they are currently unsure how to pursue this. He listed a few “best practices” concerning mail server setup. He closed the talk with a (long) list of possible tools and methods which could be investigated in order to further reduce spasm.

SPAM Fighting in DESY Zeuthen

Wolfgang Friebel has set himself the ultimate goal of 100% rejection of spam mails and 100% acceptance of good mails - while realising that it is certainly not a fight he is going to win, in the short or probably longer term. He advocates legal action against spammers since we can identify the most prolific of them but in the meantime, he listed a few measures to reduce e-mail address harvesting, how to use blacklists of known spammers and where the best blacklists are to be found. He described how not to become a mail relay by accident or deign. He proposes rejecting mail identified as spam rather than deleting or quarantining it as then the real sender gets a rejection message. He has seen significant reduction in spam as seen by the user when he implemented this but he noted that such behaviour (rejection of mails) may violate some national laws and spam identification must not rely on only simple checks. He suggested a form of greylisting, rejecting the first mail from an unknown sender with an smtp message to try again later: a real sender will usually retry whereas most spammers only send once. But he admitted there are some drawbacks in this scheme (delivery delays, some mails not re-sent, a few of them rather serious. There were some members of the audience advocating allowing only authenticated mails between HEP sites but it was noted by Alberto Pace that such measures would not be possible until all HEP sites decided on a common policy for mail, which ports to allow or block, which authentication scheme to use and so on.

FNAL E-Mail Scheme

The speaker described their configuration, how the various tasks are distributed to different nodes and how they use commercial tools at each stage (IMAP and POP mail servers, Listserve, anti-spam fighting and virus checking). He described how incoming mails are restricted to certain systems so they can be properly checked, first for spam and then for virus content, before eventual delivery to end users. They measure a throughput of around 1,000,000 mails delivered per week. They permit relaying mails to permit authenticated users to send mails while travelling. IMAP is the preferred mail protocol (3300 registered users); there are a dwindling number of POP users but no new POP account is allowed. The user is free to choose his preferred mail client and there is a web interface to the IMAP service. The Listserve server manages some 2000 lists with a web interface to access archived mail.

Wiki at GSI

Initially implemented at GSI to help produce system documentation for the system administration team. He explained and demonstrated how Wiki works and its main features. Of course he used Wiki as the presentation tool. Of the various implementations, he had selected Twiki because it seemed well suited to a cooperative intranet and he explained some of its main features (revision control and rollback, integrated access control, management of attachments, a wealth of plug-ins). He described some of the weaker aspects of Twiki that he did not like (scalability fear if Twiki webs become too large, not well-defined interface between code and data making upgrades more difficult, and file system backend performance) but overall he was very pleased with it and it is heavily used throughout the GSI IT departments and being evaluated by other GSI experiments. Its current uses include

- Knowledge base
- Installation log
- Hardware Database
- Meeting minutes
- Trouble ticket system
- Group calendar

Other uses are planned, as well as hardware and software upgrades and more user training.

Managing CERN Desktops with SMS 2003

Michel Christaller described how CERN uses Microsoft SMS to manage some 6000 desktops, mostly Windows XP. He showed its architecture and the interactions with client nodes. He described the method used to install a new system via a CERN-produced CD and how this calls for a network installation after establishing a minimum bootable system on the desktop (the Windows Pre-installation Environment, WinPE). Once installed, during normal operations, SMS Package can be used to prompt the user to install new system tools or patches. For example via pop-ups which appeared every day for up to a month, 90% of XP SP1 users were encouraged to launch the upgrade to SP2, eventually forcing unresponsive users. This allowed the team to roll out SP2 department by department. SMS is also used for system patches, first “advertising” the patch to the desktop user but eventually forcing him to install it after a certain time.

A subset of SMS is SUS – Software Update Service. This is run regularly to check that all required tools and patches are installed. In summary, SMS greatly eases the management of the Windows desktops at CERN although it has a few drawbacks – heavy inventory phase on low-powered clients and the need to package non-Microsoft products in VB scripts.

Web Content Management at FNAL

This is embodied in a tool called Plone, a content management system designed for publishing documents, tracking workflows, serving portals and for group collaborations. It is open source but supported by some 100 developers (and commercial support is also available). It was easy to install and is easy to use. It inter-operates with relational databases and has security features such as user authentication and document access permissions. Supports ftp, http and webdav and it is platform-independent (Windows, Linux, Solaris, etc).

It used by the OSG team at FNAL for documentation management, wiki support and Grid certificate authentication. It is used extensively by the Scientific Linux team and in the area of Windows policy discussions, for example for workflows for document approval. FNAL have added Kerberos 5 certificates to the package and modified the workflow implementation. The speaker described the configuration they had implemented and what user groups have to do to obtain a web site managed by Plone. In the long term, they may buy-in commercial support and development.

Lessons learned include the need for a lot of memory and the training of site maintainers is important but customisation should be pushed only so far (the law of diminishing returns). Although Plone pages are slower to access than static web pages, they can be compiled into static pages although this may remove the advantages of content management. Overall, they find it very user-friendly and hence popular and they feel they have only just scratched the surface of its potential.

Single Sign-on at CERN

Alberto Pace presented how CERN intends to make Kerberos and X.509 certificates (PKI) interoperate. Today at CERN we use PKI in all Grid-related projects and Kerberos is used in the Windows and AFS domains. He gave some examples of the use of certificates. CERN will establish a certification authority to create and sign X.509 certificates pre-matched to existing Kerberos accounts. We will also publish a web interface allowing users to create and download such certificates to their computers or map their existing certificates to their existing Kerberos accounts. In this mode, we would have a master/slave state with Kerberos the master and PKI the slave so the Kerberos password establishes the mapping.

In a second phase, the requesting and installation of certificates should be automated for "managed" computers. And we could re-discuss which should be the master. For example, in order to improve security, we could move to stored certificates, for example in a smartcard protected by a pin code, and move away from typed-in passwords. However, this would have significant consequences and will require a lot of discussion.

A few CERN web sites are piloting client certification authentication. A CERN Root Certification Authority is being implemented, we will investigate automatic mapping to Active Directory and various desktop authentication methods and smartcard technology is being investigated.

DESY User Registry

DESY has had a UNIX-centric user registry for many years but the integration of other systems was only rudimentary and many services have their own account registry, passwords and policies. This, among other problems, creates a heavy load on User Support. A series of workshops and discussions resulted in some 300 user requirements for a single unified registry. A market survey did not discover a suitable commercial product, although Tivoli came close, so a home-built scheme was devised.

The DESY-Registry should be a single account registry, have a portal for user administration with delegation to non-IT administrators. The speaker showed the various roles of people who will access the registry and what access each role has to given objects. The main entities are namespaces, persons, accounts, passwords and groups and he showed the relationships between these. The service went into production in February and first reaction has been good from both users and administrators. In the future they will add more quota handling, X.509 certificates and coupling to grid Virtual Organisations.

A PBS Graphical Management Tool at GridKa

GridKa has 640 worker nodes, 1280 CPUs and is running PBSPro, a commercial implementation of OpenPBS. It would be nice to make it easier to manager PBSPro for the cluster. A graphical interface was created and this was presented. The top view shows the systems per rack but you can drill-down to particular nodes, for example to detect system failures. The internal architecture is an object-oriented model implemented in Java with a PBS library implemented also in Java. The current version is described as "alpha". Several sites represented expressed interest to export it to their sites. GridKa has no plans for ports to other batch systems but believes this should be straightforward.

US ATLAS LCG Tier 1 Facility

This was about what the BNL ATLAS Tier 1 site learned from LCG Service Challenge 2 and what are their aims for SC3. In SC2, their 1Gbps WAN network connection was able to support their 4 file transfer servers and they met the target of 70-80MBps disc to disc transfers between CERN and BNL. But many components need to be tuned to go further. The SC3 target is to double the transfer rate, to

also achieve 60MBps to tape and sustain these rates for several weeks. WAN bandwidth has already more than doubled to OC-48 and the LAN will be upgraded from 1 to 10 GigE in June. Another imminent change is the addition of 288 new dCache nodes, each with 0.5TB of disc space. .Still missing is the required tape capacity for SC3 but they hope to borrow this from other BNL teams. After their experiences with Gridftp in SC2, they see advantages in an early adoption of gLite FTS (described later). They will link to 2 Tier 2 sites in the US, each of which should receive half the data volume received by Tier 1.

LAL LCG Tier 2 Proposal

Need to start getting experience in setting up Tier 2 centres to satisfy expected needs of HEP physicists in the Paris region. They would offer facilities for simulation and analysis for all 4 LHC experiments plus some capacity for EGEE and local use. They target 350 TB of discs, 10Gbps internal network capacity, either a 1 or 10 Gbps WAN connection and 1500 CPUs. It should be a collaboration between LAL, DAPNIA and Saclay, all of which are in the same region of Paris. They estimate they need a budget over 3 years of 1.6M Euros and have identified possible sources.

Initially they will concentrate on storage and they have established a collaboration with HP on Lustre in a Grid context and this year they will try to build a small prototype multi-site configuration and a technical support team, gaining experience with gLite. They would also hope to participate in SC3. Further out, they would build a mini-Tier 2 site for the 2006 Service Challenge and a full size production site for LHC startup.

LCG Service Challenges

Sophie Lemaitre described the various services challenges conducted and planned for LCG. SC2 achieved its goal with more than 600 MB/s daily sustained for 10 days, with more sites being active than planned. But the service infrastructure is far from complete and this is an area to be worked on for SC3 which will be starting setup in July this year. Two major components will be the addition of gLite and a more widely-deployed SRM. Also tests between Tier 1 and Tier 2 sites will be included in the so-called service phase which should be scheduled in September. She then listed the individual SC3 milestones. To quote Jamie Shiers – “the Service is the Challenge”.

There has been a lot of discussion of Tier 2 issues, their roles and the services expected of them – and those required by them. Tutorials will be held on how to install LCG software at Tier 2 sites and Tier 1 sites are expected to help their associated Tier 2 sites. The various LCG modules to be installed by Tier 2 sites were listed along with the Tier 2 targets for SC3. Only a limited number of Tier 2 sites have been identified for participation in SC3 in order to have enough resources to support them properly.

LCG File Catalogue, File Transfer Service, Disc Pool Manager

Jean-Philippe Baud described the new LFC being developed at CERN. SC2 showed that the file catalogue delivered by EDG has several performance limitations so a new one has been developed to fix in particular scalability and performance problems. There is also now the concept of a transaction with rollback in case of failures. Other new features include hierarchical namespaces, integrated GSI authentication and authorisation, access control lists and checksums. There is support for both Oracle and MySQL database back-ends. Performance tests show a major improvement over the previous catalogue but they now need real use cases to verify this. Future developments will support bulk operations and integration with VOMS/LCAS/LCMAPS, to AUTHZ which seems to be more common in US sites and to the ARDA Metadata Catalogue.

Sophie then returned to describe the work on FTS. This is a fairly simple service providing point to point file movements on request from their VOs. A single server can handle multiple channels. An ORACLE database holds the database (a MySQL version is on the list of updates but has low priority) The models supported today are Tier 0 to Tier 1 distribution via a push server on Tier 1; Tier 1 to Tier 2 distribution via a push server on Tier 1; and Tier 2 to Tier 1 upload via a pull server on Tier 1. Currently under test at Tier 0. FTS will be distributed with LCG but it will be available also standalone.

Jean-Philippe then returned to describe the new light-weight Disc Pool Manager (DPM). It performs a similar function to dCache but it is much lighter, easier to install and configure. It must be noted however that it is by no means as optimised or powerful as dCache and possibly better suited to use at Tier 2 sites. There is no central configuration file, all states are kept in a database. It is easy to remove and add discs and partitions. There is a range of interfaces including direct sockets and interfaces to various releases of SRM. Data access is via Gridftp and RFIO; Rootd and XRootd access could be added. The new DPM now supports pinning of files to keep them in place if required. It has been tested by means of 120 sub-tests and stress-tested with 40 suites running parallel, all the tests written by a separate testing team. It will be released in LCG 2.5.0, due end of June, but a pre-release will be available later this month. The target audience is to replace 'Classic SEs' at Tier 2 sites as it satisfies the gLite requirement for an SRM interface at these sites.

Integrating Tier 2 Sites in LCG

The day, and the main part of the HEPiX meeting, ended with an open discussion chaired by Kors Bos following on from a proposal by Jos van Wezel of GridKa to discuss the resources needed by Tier 2 sites. The previous paper covering this subject (contained in the 2001 Hoffmann Report) is no longer in line with the current software proposals of the LHC experiments, in particular the disc component of a Tier 2 configuration. Jos had suggested a small task force of experts to discuss this. Although not directly a HEPiX topic, the HEPiX community contains many people who could add value to such a discussion and is also a valid forum in which to present the result. Roger Jones of Lancaster has been nominated to chair this group and the names of some people who could contribute will be forwarded to him.

Day 4

Batch Workshop

The fourth and fifth days contained a series of talks and discussions about the various batch scheduling systems in use at or of interest to HEP sites. It was run under the banner of HEPiX's Large System SIG and organised and chaired by Tony Cass. The topics cover the range from local to grid schedulers. First the main sites represented presented their setups. Then the various vendors invited presented their products. Then finally there was an open discussion of the various issues raised.

LSF at SLAC

- Running LSF 5.1 now, soon going to 6.1
- They use Platform HPC for parallel processing plus cross-queue and queue level Fairshares but no use of the LSF job starter
- Solaris 480R 4 CPU master with a clone for failover, running Solaris 9

- 2500 mostly dual-CPU batch servers on a variety of platforms but mostly RHEL 3.0 with some Mac and one Windows NT; 80 of these servers are connected via Myrinet
- Some queues are dedicated to specific groups, mostly BaBar, some are public (or "general").with a mixture of execution time limits including one with a 10 day limit and one "idle" with no limit.
- There a few MPI parallel processing queues, some for groups who have funded their own servers.
- The 10 node Mac server cluster is dedicated for the astrophysics group.
- Queue level administrators manage group queues.
- Batch jobs are usually scheduled by queue name, rarely by execution time, but sometimes by host name to get a particular resource, for example O/S version, hardware architecture, memory (rare) , etc.
- They have added a local scratch space resource, available on some servers (not all) and this resource parameter can be used to select a host although this is rare.
- Various locally-written scripts exist to manage queues and jobs, for example a script which intercepts jobs whose output file is too large to send by mail..

Torque/Maui at RAL

- The batch hardware includes 980 CPUs on dual-CPU PCs some rather old and some new.
- In 2003 there were 2 farms, CSF and LCG, with OpenPBS/Maui on the former and OpenPBS with native FIFO on the other. These differences increased management overhead and reduced the flexibility to fully use all CPU resources, especially when RAL is contractually bound to guarantee a certain percentage of capacity to particular experiments.
- So in 2004, they upgraded CSF to Scientific Linux 3 and took the opportunity to move the LCG worker nodes to the CSF farm, keeping a CE as an LCG gatekeeper. Now CSF is running Torque 1.2.0p1, a variation of OpenPBS with lots and lots of stability patches and enhancements, plus the Maui scheduler, version V3.6.2p11 with 3 local patches to increase the number of job classes and the maximum numbers of jobs queued and running; and to implement fairshares
- For non-grid users, there is standard OpenPBS; these users are able to specify CPU and/or memory needs and jobs are queued accordingly.
- For Grid users there is a shared short job queue there plus one queue for each VO which appears to becoming more popular in LCG; the advantages include independent ETTs[9] for jobs in different VOs and queue lengths can be easily customised. CPU scaling can be solved by normalising the speeds of the different CPUs but it makes wall time more confusing for users so RAL appears less attractive to run jobs at. Therefore they "interpret" the ETTs. This should be fixed when they move to a new normalisation unit.
- The Maui default scheduling method is FIFO, as is PBS, but it supports fairshare scheduling based on wall times; the administrator can give targets to Maui (allocate 20% of capacity to an experiment over a 7 day period for example) and Maui will strive to achieve the given target. They plan to introduce Maui to manage the VO queues also.
- Emerging issues – the amount of available RAM on a node (ATLAS is approaching 1GB which may block the job on the second CPU)
- There is no facility to use scratch space as a resource or scheduling parameter,

SUN Grid Engine (SGE)[10] at London E-Science Centre (LESC)

LESC must support a variety of users from different disciplines so a wide mix of job types including parallel jobs, large memory jobs, long running jobs, and so on.

[9] ETT – Estimated Traversal Time) for the user, estimated start to completion times.

[10] Confusion between Grid Engine and SGE; LESC are using the open source (free) version so in fact they are probably using Grid Engine, not the SUN productised version of this, SGE. Sun claim there is no difference.

- They have a 24 way SUN E6800 with Solaris 8 with a large memory and disc configuration and a tape library; a 260 dual-node P4 Xeon cluster with RH 7.2; and a 204 dual CPU 64 bit Opteron cluster with RHEL 3.
- There are two logical SGE installations: one of the PC clusters runs SGE 5.3 and one runs 6.0, both mastered by the SUN but only one queue with SGE sorting out all queuing.
- Issues are:
  - SGE has no queue checkpoint-and-pre-empt facility (the speaker claimed)[11], so a short job cannot run if a number of long-running jobs are executing. To get round this, they reserve some nodes for short jobs where the user specifies the job is indeed short.
  - Internally, they have heterogeneous clusters (variations in memory sizes, processor speed, disc sizes). So sometimes a job with low resource needs gets queued to a large configuration node and a job with large resource requirements has to wait. They plan to change the SGE configuration such that by default a job gets queued to the least capable node available, putting the onus on the user to specify the resources he needs. But they have not implemented this yet because of fears that users will exaggerate their requirements to be sure of getting enough, putting the situation back to square 1.
  - Sometimes a job needs multiple cluster nodes and this can cause such jobs to be queued for a long time if not enough nodes are available. They then intervene manually to manipulate queues in the 5.3 cluster to schedule the job.

They wish to participate in LCG as a Tier 2 site but need to make SGE a valid scheduler in LCG software. They have developed Globus Job Manager and Information Reporter interfaces for SGE, and are starting to expose some PC nodes to LCG and setting up the various interface services.

LSF and OpenPBS/Torque at JLab

JLab has two batch farms, one of 200 dual Xeons for experimental physics and a Lattice QCD farm consisting of 3 clusters for parallel jobs for theorists operated by the HPC (High Performance Computing) group. The physics farm is LSF-driven, mastered by a locally-written front end (Auger) and coupled closely to their locally-written JASMine mass storage system. The HPC group uses Torque with a locally-written front end. With 2 different sets of queues, management overhead increases and they are.looking at developing an LSF-like interface to Torque.

There are LSF queues with different characteristics but most jobs run in the Production queue with short ones going into the so-called Priority queue. The Production and Low Priority queues use Fairshare, the Priority and Idle use FIFO, but only Production and Low Priority have time limits.

Jobs in PBS queues are scheduled according to resource needs, including memory, scratch space, number and type of nodes needed, etc. They have implemented tight integration of JASMine and LSF to improve resource utilisation, for example to pre-stage data before jobs are admitted to the farm. They feel they need a better way to specify resources needed for a job, currently they have a local workaround inside Auger.

They would like common job submission across both environments and then to add grid job submission seamlessly as JLab becomes more involved with Grid. For this, they are collaborating with STAR at BNL on RDL[12] – Request Description Language.

[11] A.Maslennikov claims SGE does indeed support checkpointing!

[12] RDL is used to write Java description files for submitting STAR batch jobs. This sounds a most interesting tool and we will try to schedule a talk at the next meeting.

BQS at IN2P3

BQS is locally-written in IN2P3, supports heterogeneous groups of users, job classes, AFS token management and it is grid-enabled. Classes are defined a combination of memory sizes, disc space, CPU type and so on. They can define virtual resources, either public or private, and use these to schedule jobs, for example, one for CMS jobs. A workpoint is a virtual machine execution point and only one job can run per workpoint but several workpoints can be scheduled on a worker node, each with different profiles (which themselves are built out of classes). Before job submission, BQS checks if resources are available on the assigned worker node and resource consumption is monitored during job execution. Jobs are submitted by class name, or by specifying real or virtual resources. Jobs can be tracked via remote job queries and users can be informed by e-mail when a job starts or when it completes. Scheduling policy criteria take account of the type or amount of CPU, memory and disc needed, O/S version, virtual resources, job priority (there are 3 levels), etc. Jobs are automatically re-queued if worker nodes crash and are not impacted by a crash of the master server. There are 680 dual CPU worker nodes, all dual CPU, 2500 workpoints (3 or 4 per node) and a mix of mostly Linux and Solaris nodes. BQS supports now parallel jobs (PVM or MPI) on a dedicated farm and they plan to support quasi-interactive jobs (which seems to mean jobs being scheduling and executing almost immediately).

Batch at FNAL

- The big experiments each have their own dedicated reconstruction farms on which scheduling is trivial – large numbers of very similar jobs.
- Plus there is a smaller general purpose farm where things more complicated – a mix of jobs of different execution times. Each experiment on the general purpose farm had a maximum number of CPUs they could use in parallel and a maximum share of the cluster (because the total of maximum number of CPUs was oversubscribed compared to the cluster size). Most queues do not have fixed time limits and priority is calculated by recent previous use by that group plus there is an overall guarantee of execution within a maximum time.
- Before “the Grid” arrived, they used to use FBSng, a locally-written package which had a kind of fairshare model. It had advantages such as light resource consumption by the batch daemons, being free and being reliable. But it was designed for a very homogeneous environment (Run II) and is not grid-friendly. So FNAL is moving towards Condor which seems to match their needs – free but well supported, widely used in our coounity (and elsewhere), Kerberos 5 and X.509 authentication, Grid interfaces, etc. US/CMS has switched and CDF is in the process of switching.
- There is a growth of analysis farms for the experiments which support a wide mix of jobs, replacing large SMPs mainly. For example CDF went from a multi-CPU SGI to a PC-based farm running under Condor. D0 use PBS/Torque on their analysis farm; US/CMS still use FBSng but will switch to the Cisco load balancer (?) soon. CDF have written, with INFN, software to monitor FBSng in these environments and are modifying this for monitoring Condor jobs.
- D0 uses SAMgrid (see previous HEPiX reports) to manage data for Monte Carlo production worldwide.
- For the Open Science Grid (OSG), there is a Globus gatekeeper which distributes jobs to a Condor pool.

SGE at DESY

- They started in 1994 with CODINE which became SGE; they are currently on version 5.3 but moving soon to 6.0. SGE offers Fairshares. Hosts are grouped by host type then they define queues with short or long execution time limits but there is no checkpointing.

- There are 130 dual CPU PCs, mostly Intel but more recent purchases are Opterons. Currently they feel the need to replace older, smaller memory systems which are under-used while other systems are overloaded.
- There is a mixture of SuSE and Scientific Linux. They have developed tools for auto-installation and regular administration tasks.
- There are 2 databases for host configurations, one for host information (whether it is a master node, submit node, worker node, etc) and a second for hardware configurations. Users use these databases to specify resources rather than rely on the SGE-detected information as SGE does not always give true results (it does not properly detect hyper-threaded nodes) or it does not give results which are wanted (SGE quotes memory including swap space which DESY does not want to use to specify useful memory).
- Unlike LESC, they have created a number of queues although users do not submit to a queue by name, rather SGE matches job resource needs and queue characteristics.

LSF at CERN

Users login to the 44 node LXPLUS service and submit jobs to LXBATCH, currently around 950 CPU servers. 400 disc servers, 80 tape servers with a very heterogeneous LSF environment. There is a mixture of PCs of different ages, mostly running SLC3 but still some with legacy RH 7.3. There are many queues with different execution time limits including one per LCG VO. They quote normalised CPU figures based on kilo-SPECint2000 currently (SPECfp is not used in this respect). There are Lemon agents (see talk on Monday) on the LSF nodes to show many load values.

Issues include –

- the queue definition model does not match well to the grid model of VOs;
- need to allocate shares between the experiments; LSF tries to deliver the defined share of each experiment and experiments can pay to for more nodes to expand their share. A small percentage of the total share is withheld to permit emergency priority jobs for execution but it is put back into the main system pool when not used in emergency mode.
- LSF supports hierarchies of user groups and each level of a group may use a different scheduling policy – Fairshare, FIFO or pre-empt..
- Users have a very wide range of parameters which they can use to specify needed resources but most users know (or quickly learn) to which queue they wish to submit. Beyond queue name, typical parameters used are CPU type, memory and scratch space.

Batch at BNL

- They have a total of 2750 CPUs, growing to 4400 later this year. The operating system is mostly SL 3.0.2 with a few RH 7.3.
- There are two central batch masters, one as backup and 3 machines dedicated for job submission, each managing ~600 jobs simultaneously; some other execution nodes can also submit jobs.
- They have been running LSF version 6.0. LSF policies are defined per experiment with Fairshare and multi-level job pre-emption (Phobos experiment only).
- But they are moving to Condor (ATLAS has already moved), mostly Condor 6.6.6 with some 6.7.0 for RHIC and 6.6.8 for US/ATLAS. There are Globus gateways and they allow, also for LSF, jobs to be submitted via Globus as well as locally-submitted, There are 6 Condor pools with 2 configuration models for each, restricting which job can run in each part of the pool, reserving one part for high-priority jobs. This restricts resource sharing of course. This is a more complex version of the Bologna model to improve sharing.

- The Bologna model creates virtual machines and allocates multiple VMs to a CPU in a multi-CPU system; then allocates different priority jobs to different VMs and since a CPU can have a number of VMs, this increases the resource sharing. Pre-emption is based on user priority; CPU time limits are assigned to different VMs. ATLAS uses a version of this model; machines are defined as having three sections with several VMs per section. Grid jobs have higher priority in two sections and local jobs higher in the third but unused capacity can be utilised by the "other" type of job (grid jobs in the local third when it is not used for example).
- Problems include Panasas (see earlier in the week) creating badly formed process names in some kernel threads which breaks Condor but Panasas fixed their bug. And there have been scalability issues with the number of jobs which had to be worked around.
- Unresolved issues include trashed jobs not making available their output files; Kerberos support not fully working; and some missing features in Condor. They will move to 6.8.x everywhere in the autumn but they are concerned about scalability and high availability as usage rises. They would like to integrate BNL Condor pools with outside pools but are concerned about security.

Interface between Local and Grid Resource Management

Jeff Templon of NIKHEF discussed the gap between LRMS and GRMS. Typically a grid site will guarantee X% of capacity to the grid and the remainder to local jobs but you must not allow unused capacity to go to waste, for example by sharing unused shares and also perhaps allowing a group of users to scavenge any left-over cycles. He defined three types of users which formed the basis of much discussion over the rest of the week:

- The polite user just submits his jobs as documented in the grid user guide and let's the GRMS sort out scheduling. Rule 1: polite users almost always lose!
- The high-throughput (HT) user blasts jobs at a site until he sees queue pile-up and then moves to the next site.
- The sneaky high-throughput user has X jobs to run but submits N times X "dummy" jobs to many sites, ignoring any queue pile up; when these dummy jobs are queued, they select the real job from a linear queue somewhere and when the remaining (N-1)X dummy jobs are scheduled they simply fail to find a real job to run and immediately exit, potentially quickly draining job queues! Sneaky HT users are good for filling sites but this hits fair-share commitments.
- The fast turn-around user needs to negotiate priority to get fast access.

The question of the meeting is - how to match these users to this typical site? In particular, faced with sneaky HT users, how to somehow throttle series of jobs which repeatedly "fail".

A user's key question is - how long to run my job? A site which seems 100% "full" to an ATLAS user may be only be full because ATLAS are fully using their share. So the answer to the key question actually depends on who is asking and whether the person is local or "on the grid"! And you need normalised CPU times, and a normalised way to express the length of your job and therefore a solution taking account of heterogeneous PCs in a cluster.

Another problem is that grid scheduling has a very high overhead – typically 250 second life cycle for a 20 second job. High priority users need dedicated CPUs (which is expensive) or some form of virtual machines where low priorities can be suspended and restarted afterwards, a form of checkpointing. Some batch systems indeed permit this.

Little information is passed from the grid to the LRMS such as estimated job length, and this makes it hard for the local scheduler to queue the job correctly. And how to properly account for MPI jobs?

BQS and the Grid

Yves Fouilhe presented what improvements were necessary to make BQS work in a Grid environment. The introduction of grids should produce no partitioning BQS farms. Need to authenticate Grid jobs via certificates and VOs are mapped to a unique local UNIX group. BQS was modified to use the mappings from the CE as job characteristics. The VOs are thus treated as UNIX groups and objectives are given by group (VO). Users with no local account are given a nickname. LCG and EGEE middleware are considered as worker node resources used for scheduling (jobs mathed to particular versions for example). Changes were also required for the job accounting and tracking routine and for monitoring.

During the introduction of Grids at CCIN2P3, a number of problems had to be met and solved, among them the sharing of the installed capacity. At the moment of submission, the resource broker does not pass all the required information, for example expected CPU time, memory needed. One solution is for the gateway to set these numbers to some maximum but this leads to under-used resources. Another solution is to open more queues.

Need a good and clean way to map a user in a VO to a unique generic user; it was suggested to use the Grid Map Directory mounted in NFS. BQS real priorities were not being passed on by the Grid Information Service so virtual queues were created reflecting the number of CPUs used by a VO to better reflect the status of the site. On the other hand, the Grid requires a site to publish a lot of other information about jobs and this results to many requests to BQS. Also BQS needs to cope with jobs whose proxies expires and need to be renewed.

Loadleveler and the Grid (IBM)

Recently released for Linux, 32 and 64 bit chips supported. Very little used in HEP (only one site present uses it). No fair-share scheduler, only back-filling (filling empty gaps in the queues) New feature is advanced reservation, most frequently asked for by Grid computing sites and sites which have to perform tasks at a given time or as soon as possible: a time slot of a certain length is reserved for a given time. To actually get the job into execution, need to have pre-emption to suspend running jobs on the target node(s). Advance reservations can be “as soon as possible” or can be for a specific time if submitted at least 24 hours in advance.

Loadleveler can control how many jobs per user or per group, but Loadleveler groups do not correspond to UNIX groups. LL supports pre-emption on its backfill scheduler except that suspend does not work today on Linux. The new version includes the Globus toolkit for integration with Grids and it contains toolkits, examples and documentation on how to integrate it in a Grid environment.

LSF and the Grid (Platform)

Started with a quick review of the features of version 5 and the recently-released version 6 but the main part concerns interfacing to Grids. Some of the issues we face, Platform claim to have seen in their so-called multi-cluster LSF configurations. He listed a number of issues as seen by him, many already brought up by previous speakers. One of these is the visibility of what is going on inside the middleware. Another, already mentioned is that LSF’s notion of a user is very close to that of a UNIX user while Grid middleware uses PKI, mapping Grid users to generic local users which results in Grid users becoming somehow “second class” citizens. Perhaps we can re-define both local and grid users as “consumers”.

Where are applications and how are they configured? Users get used to their home configuration. We really need some form of standardisation, for example a standardised configuration description language; does this tie in with DRMAA job categories? Could provisioning services provide an answer?

How to set scheduling policy, are they different for Grid and local users? Why? Suggest setting policies at a higher level and less at the local level. Do all users not belong to a VO, even local users? Set policies at the VO level.

Should data go to the job or the job to the data? Current schedulers do not consider this. Ideally you need to know the pattern of data access and minimise overall data flow, not just for one job but for a work load.

CSF is Platform's Community Scheduler Framework to permit experimenting different scheduling concepts and exploring issues such as those in Grid environments. A nice quote on the difference between Grid and local schedulers - "Grid Schedulers are just matchmakers, local schedulers are resource brokers." Grid-level schedulers are still in their infancy.

PBSPro plans for Grid (Altair Consulting)

Altair market a version of OpenPBS as PBS Professional. New version, V7, supports Redhat Linux. New features include advanced reservations to reserve a set of resources which can be used, as described earlier, in a Grid environment. It also allows peer scheduling – coupling multiple clusters, this allows sharing of resources and queues, the scheduler pulling across from a remote site a job which can run immediately in a local site which has spare capacity at that moment. The peer scheduler needs to be aware of data needed by a job; the user needs to specify how much data is needed and where it is stored. It also needs to take account of user authentication and all clusters taking part in such a scheme must share the same user space; mapping could be added but Altair do not support it. They do make sources available free to customers with a valid support contract.

SGE and Grid (SUN)

N1 Grid Engine 6 (N1GE6) is the latest release of SGE. Features include some for scalability such as a high-throughput scheduler. It supports cluster queues. One definite difference in the product version is the provision of a database (Oracle or PostgreSQL) for spooling and accounting information. Resource reservation is supported and there is a GGF DRMAA 1.0 implementation for submitting, controlling and monitoring jobs and Globus/GT4 integration. The open source project (Open Source Grid Engine) consists of half a million lines of code and binaries for the most common platforms but he was not sure if it is up-to-date with respect to N1GE6. The speaker suggested that we need to group jobs by categories based on actual properties; jobs are then interchangeable and there could then be one Grid per interchangeable job category. [But for him a Grid is a logical construct and a sub-cluster is a site.]

Condor (Uni of Wisconsin-Madison)

Total Condor package is now some 9 million lines of source. The scheduler service sets job policy, the matchmaker acts as the resource broker and the Start daemon implements the resource policy on target nodes. All these are linked by ClassAds. A machine ClassAd, issued by the Startd daemon describes the main attributes of a system, a requirements expression declaring what type of job it prefers to accept with a ranking of the degree of match to the requirements. The job ClassAdd, issued by the Sched daemon after job submission is similar but from the job perspective and the Matchmaker makes the best match between the two. Jobs can be pre-empted by startd, e.g. when the client owner takes over his own

system; by the matchmaker, e.g. to assign the node to a higher priority job; and schedd, for example to stop runaway jobs.

Condor-G has interfaces through Globus to various Grid implementations via the middleware. Flocking is similar to Peer Scheduling in SGE, sharing jobs across Condor pools; sched has lists of possibly compatible matchmakers.

Problems: possible lack of information if Grid middleware does not fully advertise available resources or job information during execution. Could lead to a loss of job semantic guarantees (for example, run this job once and only once). Solution - startd and schedd are simple non-privileged jobs so start them via Condor creating a Condor pool out of grid resources (so-called Glide-in); thus you effectively replace the Grid middleware by a Condor environment restoring all the benefits of Condor scheduling which were “washed away” by the grid middleware (see overheads)

They are looking into a tree of matchmakers in future versions. Also looking at Stork, a scheduler which takes account of data placement; it understands characteristics of the job and how and where to place to data.

Lessons: Grid scheduler needs a clean separation between resource allocation and work delegation (scheduling). Scheduling policy should be at the VO level and should deal with resources first; nevertheless, data handling should be treated as a first class citizen.

EGEE Batch Helper (BLAHP)

Batch Local ASCII Helper Protocol – provides a set of plain ASCII commands to Condor-C (and CREAM). It is a set of scripts, 3 for each batch scheme, to submit a job, monitor its status and cancel it. The BLAHP daemon is a gLite component. Decided not to make a set of API calls or to include the code of the scripts directly in the daemon to remove dependencies, make version upgrade easier and make it easier to add an interface to a new batch scheduler. Because PBS can lose jobs during submission, they developed a status cache server to keep track of BLAHP-managed jobs by monitoring the system log files. For proxy renewal, they have a proxy receiver server which is started on the worker node listening on certain ports. Still open are the issues of two-phase commit in job submission (several ideas, no decision yet); how to use job requirements to steer jobs to better-matched hosts, in particular how to express such requirements to make sense to the different batch systems running in a given Grid; how to inspect job outputs at runtime.

GLUE Schema

Laurence Field presented this, explaining first of all what it is. GLUE = Grid Laboratory Uniform Environment, a common schema for grid computing. It is used to describe attributes and groupings published in the information system and used by grid components such as middleware, resource brokers and replica managers in both LCG2 and Grid3. For example its information is used to decide where to send jobs. Early versions were good at service discovery but did not handle dynamic information well so it was (is) poor at service optimisation; first version had no real concept of a VO and was missing descriptions of an SRM or web services. And early versions lacked practical experience and scalability. .Finally, job submissions are according to ETT but early versions had only one ETT value per queue and this does not match reality at a grid scale.

New version aims to solve these deficiencies, interfacing to more grids while being backwards compatible. But it still assumes a homogeneous cluster (which is rare – see previous talks) and the scheme still depends on the grid model. The next version, which will enter development late this year

will look at this without the need to be backwards compatible. It is hoped to maintain one ETT per VO but it needs, for example, a new definition of what an ETT means which is consistent with all batch systems. To really produce a better scheme, the development team requires contribution from experts from each of the batch systems with which the various grids need to interoperate. He asked the audience to become candidate CE satellite testbeds, explaining what that entailed; Harry Renshall volunteered to contribute an LSF testbed. All suggestions on possible improvements in the scheme should be submitted via a mail list.

When asked about gLite, he was unsure of the interaction of the two, he thought “someone is writing” an interface. Jeff Templon claimed that we know do what ETT is (time to start execution) but it is impossible to calculate accurately, only to estimate it statistically.

Discussion on Batch Scheduling

Chaired by Tony Cass. Q1: is there a problem? Yes, there is a mismatch between local and grid level scheduling, we don’t know how to match the attributes of particular VOs to site policies, or cleanly map user identities at sites.

Can we (HEP site administrators) help find a solution? Will vendors soon or eventually solve our problems? Worse, are we going in an orthogonal direction to the rest of the non-HEP world? The Platform representative, who is on GGF, advises us not to wait for the rest of the (grid) world but move forward as we are, although watching what GGF and other grid communities are doing. Jeff Templon agreed, “waiting for (grid) standards is bad for your health”; he hoped for more progress from GGF working groups. The Altair representative said we share many problems with other grids (security, scheduling).

Should we optimise for sneaky HT user submitting a massive number of jobs or the polite user issuing one job at a time? Expressed in another way, should we loan resources for a user to run his jobs or run the jobs for the user? Templon – the real problem is to guarantee a given share for all users; Marc Mengel – are we over-complicating the issue? Can we not devise a method to handle both types of user by managing jobs as they are scheduled – too many jobs from a user would be discovered and could be throttled? Or penalise the user whose job does not run for the advertised time. It is suggested to use credits to run jobs, this would penalise the sneaky user – why use 10,000 credits to run 1000 jobs? SGE, Condor and some other batch schemes have an option for maximum number of jobs per user and the grid scheduling broker could take account of this; for example the batch system could “pull” jobs for execution when the resources were available but Templon suggests the pull model may well create other problems, for example flooding the grid by polls of batch systems trying to find jobs to run. What is the problem we are trying to solve? Is it the sneaky user submitting thousands of fake jobs to run only a few faster than submitting them serially? If real jobs are run within published commitments, what is the problem?

GLUE Schema update – ClassAds for a site? For example, they would declare which VOs can run jobs there, what resources are available (CMS software available for example) and with what priorities. A ClassAd for a site would publish site policies, not site status. Laurance Field – sounds similar to how BOINC (SETI) works – what VOs are possible to run on a given site. Condor representative – some Condor projects effectively do this, SAMgrid is an example.

Should we continue or follow-up this discussion at the SLAC HEPiX? Is this too soon? We should certainly solicit a presentation on the RDL project at STAR/BNL, as mentioned by Sandy Philpott the previous day. There was a suggestion to invite middleware developers to meet with site administrators in

this (HEPiX) environment[13] - does this ever happen? Those present seemed to welcome the idea. Perhaps schedule a half day at SLAC and perhaps partially via video-conference.

LCG/EGEE Standard Working Environment Proposal

This was provoked by an incident when D0 jobs encountered problems in CCIN2P3 because of a mismatch of environment variables – there is no standard definition of an environment of grid batch jobs today. The speaker had discovered that not all LCG sites define all suggested LCG environment variables as described in the LCG User Guide. So this is a proposal for a minimal standard for this which should be based on the corresponding POSIX standard. The proposal is given in the overheads. One feature is to alias all the three temp directories (EDG, LCG and GLITE) into one. The naming convention for such variables must be independent of particular grids, hence independent of particular middleware. The speaker would like feedback on the proposal and a fast decision to implement, distribute and deploy the template.

There were some detailed questions on particular entries in the template, what particular variables and directories are for, but noone was against the proposal. The authors will now write a short document on the proposal, distribute it to the sites and submit it to a workshop in Bologna. It was noted that the released documentation for the templates should contain the semantics – how the variables and directories are meant to be used. Should each site implement this, translating the template for their batch scheme, or should it be translated once per batch scheme? The speaker thought per site but some in the audience thought a version should be produced for each scheduler by an expert in that scheduler and these then installed per site. This will be discussed at the Bologna workshop.

[13] In fact, the most relevant EGEE developer had planned to attend this meeting but had to cancel at the last minute.