

# HEPiX Jefferson Lab Trip Report

October 9<sup>th</sup> to 13<sup>th</sup>, 2006

## Introduction

This was the third HEPiX meeting held at [Jefferson Lab](#) (JLab), a nuclear physics laboratory in Newport News, Virginia, some 250 Km south east of Washington. In some ways the meeting could be called HEPiX-lite because of the attendance was around at just under 50 (compared to double that in the past 2 European meetings) and an agenda which could have accommodated more talks. But it was well organised, there were many interesting talks with none (I think) which were repeats from previous meetings and JLab is a pleasant place for a meeting with its onsite hostel (for those of us who remembered in time to book one of the well-appointed rooms), the relaxed atmosphere and the generally clement weather.

As usual this report is mostly a personal set of notes but includes some input kindly supplied by Helge Meinhard and Sebastian Lopienski. Readers interested in particular points are referred to the overheads, all (or almost all) of which are online at

<https://indico.fnal.gov/conferenceDisplay.py?confId=384> .

In the notes which follow I have tried to insert footnotes and links to web pages to explain some of the abbreviations and terms which may not be so common to all readers (including me!).

## Highlights.

- Collaboration was a theme of the week with :-
  - a suggestion to create a HEPiX e-mail whitelist (see SpamCop talk)
  - a suggestion that HEPiX assist LCG to improve and standardise Tier 1 and Tier 2 site monitoring (see LCG presentation on Thursday)
  - a proposal from the IHEPCCC chairman to create joint study groups on benchmarking and global file systems – see Friday’s report.
- As usual security was another hot topic with the introduction of two factor authorisation mentioned by at least two major labs (BNL and JLab), two spamming war stories and a very detailed tale of a particularly serious attack lasting around 12 months (The Stakkato Intrusion)
- An interesting and entertaining talk was given at the Conference Banquet by the JLab CTO on the importance of vacuum tubes over the years. This was (I think) a first such after-dinner talk for HEPiX and was definitely appreciated by the audience.
- A number of sites reported on tests conducted and/or planned on global file systems such as Lustre and GPFS. Fermilab are investigating a new device called BlueArc.
- Yet again Michel Jouvin wins the award for the HEPiX member adopting most CERN software, discovering the CERN-dependencies where they remain and spreading the word that sharing is a “good thing”. Next on his list, when he has completed the switch to Indico, is IT/FIO’s SLS service level monitoring (see SLS report below)
- Next meetings :-
  - April 23<sup>rd</sup> to 27<sup>th</sup> in DESY Hamburg
  - possibly early November in either Berkeley or FNAL, hopefully in the week preceding Supercomputing’07 in Reno
  - Spring 2008 in CERN.

For the DESY meeting, topics suggested were benchmarking, cluster file systems, VoIP and in general discussion topics (as opposed to LCG workshops) likely to attract LCG Tier 2 sites.

## **Welcome – Roy Whitney**

The CIO of the newly-reorganised IT department of Jefferson Lab opened the meeting by describing the mission and activities of the lab which are centred around nuclear physics. He described some of the interesting physics being carried out in the main experimental halls. JLab are planning on an energy upgrade of the accelerator to 12Gev. They also operate a Free Electron Laser as a spin-off of their nuclear physics programme and there is also a healthy QCD research study based round 2 dedicated clusters as part of a 3 site National Lattice QCD Computing Project along with FNAL and BNL.

## **Monday -Site Reports**

**BNL** – information on both the RHIC and US ATLAS computing facilities. BNL is upgrading its AFS storage, retiring lots of NFS storage based on Solaris and replacing it with newer equipment; they are also still using [Panasas](#)<sup>1</sup> which is becoming more stable although much of it is being mounted indirectly via NFS as it is considered as lacking in scalability. On the clusters, older CPUs are being replaced by more recent 1.8GHz dual Opteron-based systems. Their batch system has been upgraded to Condor 6.8.1 with LSF still displayed on 400 nodes in tandem with Condor, which creates its own problems. dCache is primarily being used for distributed storage with some use of xrootd for the STAR experiment. They are moving to two-factor authentication using ssh keys or cryptocards only, no more passwords; ssh keys are stored in LDAP.

**CERN** – Helge reported on various CERN activities, starting with the introduction of the LCG backbone in the Computer Centre for the Tier 0 services and work on CERN's audio and video conference services. He reported that IT management had successfully negotiated an LCG-wide agreement with Oracle for licences on LCG Tier 1 sites. On the other hand, LCG was hit by a recently-found Oracle bug which had caused data corruption. Under the heading "code Management" the CVS service is stable, testified by the decision of CMS to close their own service and move to the central one; but CERN is not currently looking at Subversion. The SURE computer centre alarm scheme has been replaced by the [Lemon Alarm System \(LAS\)](#). CERN will use [Scientific Linux \(CERN\) version 4 \(SLC4\)](#) for LHC startup with support for SLC3 to be stopped by October 2007. He described some of the power interruptions which have happened during this year so far. He listed the hardware acquisitions in progress or planned, including the Computer Centre's first dual-core Opteron servers, 30 of them so far, which have been introduced without fanfare or visible effect on physics services. They are looking at building fatter disk storage servers. He explained the revised acquisition planning based on the recently-published LHC startup schedule. Use of insecure mail protocols has been forbidden (which I met directly when I tried to use Pine on a JLab terminal while my PC refused to see the network). CERN Printer services have been moved to Windows-based servers and there is now significant use of Windows virtual servers. All physics database services are now based on Oracle RAC<sup>2</sup> servers with two new installations planned.

**Fermilab** – first FNAL site report in 18 months! There is a new director, Pier Oddone, since July 2005 and there is a change due in the management of the lab where the results of an outsourcing tender are expected to be announced by the end of the year. Lisa Giachetti described the multiple 10 GigE dark fibres constituting the Fermi Lightpath connections to the STARLIGHT international network exchange point in Chicago. Much work has gone on to automate computer security, moving to

---

<sup>1</sup> A commercial disc file system described at various previous HEPiX meetings.

<sup>2</sup> Real Application Cluster

commercial tools in many areas and open source in others, replacing many home-grown tools in the process. Given that the Feynman Center is running out of capacity, a second 2000 sq.ft computer centre was opened with space for 3080 CPUs and ancillary equipment and another room is under construction, re-suing an old experiment counting room. There is also a separate “Robot Room” with so far 1 Sun/STK SLA8500 silo with another on order. And there is yet another computer room for the Fermilab contribution to the National QCD collaboration described by Roy Whitney. She described the various dedicated and general purpose compute farms, mostly operating under the Condor batch scheme with some legacy LSF and a little PBS. The general purpose farm is the central point for the Fermi Grid, part of the OSG (US Open Science Grid) although some of the dedicated farms (US-CMS for example) also have direct access. US-CMS has 700 TB of dCache space, expected to grow to 2.5PB by Fall 2007 and their batch farm is expected to grow from 700 nodes today to 1600 in the same timeframe. They use [BlueArc](#) NAS for online storage systems and they are so pleased with its performance that there are plans to expand this facility although she agreed with John Gordon that it is expensive and one requires a strong negotiator!

**GridKa** – since the last meeting various hardware upgrades of CPU and disc capacity has been installed or ordered, including new water chillers. Use dCache on a GPFS file system. There is also an expansion of tape capacity due later this week, a Grau Datasystems robot with IBM LT03 tape units. Their recent Opteron server installation was the first with significant problems and required a series of BIOS, BMC and NIC firmware upgrades. They had also some recent problems with their water cooling, traced to leaks in the water chillers. The speaker, Manfred Alef, then turned to his “hobby” as he described it, reporting his latest benchmarking results, this time on Xeon Woodcrest systems. They show a 60% jump in performance compared to his previous Opteron tests but he has not yet tested the newer Opterons. His results are listed on the new HEPiX [benchmarking web page](#) at CASPUR. He will soon move to using the SPEC CPU2006 benchmarks and he noted that the so-called “rate” results based on the 2006 release of SPEC are now close to the raw speed tests.

**Jefferson** – as mentioned earlier, there is a newly-constituted Information Technology Division; this results in a new (outsourced) management contract for the lab; the new Division merges the previous computing department with the previous high performance computing group and the previous Physics Computer Centre group. JLab now has an OC-192 connection but it is not yet connected to OSG<sup>3</sup> although this is under consideration, especially due to the fact that a group have made this connection part of their proposal for an experiment in the new experimental hall. JLab had a recently so-called “friendly” visit of the DoE security team (officially called a Site Assurance Visit) to help improve security ahead of a possibly less-friendly visit to test the results. There had been considerable activity in preparation for the security visit and Sandy Philpott listed a significant number of issues raised during the visit and leading to ongoing effort. A problem without an obvious solution is how to change the attitude to security of the users but one concrete result is more emphasis on a long-standing plan to move to two-factor authentication, at least for accounts with privilege.

**TRIUMF** – Corrie Kost started by listing the network connections to LHCOPN and to other Canadian sites and noted that upgrades to 10Gb links are due soon, once some local difficulties are ironed out with their connections to downtown Vancouver. A recent mail migration exercise caused considerable disruption, exacerbated by a network failure. They have plans to move to SL4<sup>4</sup> in the near future. Future expansion plans still include use of blade systems which Corrie believes offer better environmental advantages over pizza box configurations. A recent test of their air conditioning units almost ended in disaster when the single unit left running shutdown due to overload and stand-alone

---

<sup>3</sup> The US Open Science Grid

<sup>4</sup> Scientific Linux version 4 from FNAL (see later talk)

fans had to be quickly installed. He also showed how ATLAS is building up their Tier 1 facility at TRIUMF.

**NERSC** – also a first report from [NERSC](#)<sup>5</sup> (or [PDSF](#)<sup>6</sup>) for a number of years, this time due to absence of a site representative. The report dealt mostly with PDSF of which the speaker, Cary Whitney, has been appointed project head after the departure of Shane Cannon to Oak Ridge. The next biggest event has been the introduction of the NERSC Global File system (NGF) in production mode with 70TB project space (see later talk on NGF). NERSC has acquired a Cray Hood system with more than 19000 processors in 102 cabinets and 70TB of disc space; the speaker described it as basically a blade system for general purpose use. NAGIOS is their preferred standard monitoring system and this will be extended to the Cray. Ganglia and Cacti are being looked at for some functions although the latter has a poor user interface. The 19,000 sq.ft computer room is now full and any expansion will need to move up the hill in Berkeley although there are plans for significantly more electrical power for the existing centre in the expectation that more power will be needed as storage systems are replaced in place by more power-hungry units. And as may be expected in a Californian computer installation, extra care needs to be taken for seismic risks.

**NIKHEF** – they have installed a new storage device from EMC NS502 for home directories. It has a nice snapshot feature although this is not used for backup. Migration from the previous scheme was a lot of work, especially for users with restrictive ACLs on their files. There have also been some severe administration problems for which it took some time to get fixes from EMC. On the desktops and servers, there is a gradual migration to SLC4 and new systems are equipped with that. With SLC4 comes the [CUPS](#)<sup>7</sup> protocol for printing so they have introduced a CUPS server for printing from Linux only. For their latest CPU farm acquisition, Dell won the order, based on supplying a given number of SPECmarks using Intel 2.66MHz systems and future acquisitions will be based on the same principle – vendors challenged to supply a given number of SPECS.

**RAL** – RAL has recently been re-organised into business units and the LCG Tier 1 service is now in the e-Science Department which is in the Facilities Business Unit. RAL has submitted a bid to PPARC for the GridPP3 project for the exploitation phase of LHC. Aside from that, they have approved funding for a new 3 floor computer building with multiple computer rooms on the ground floor; overall there should be space for over 300 racks plus robots and telecom equipment; the power rating will be 2700 kVA initially and 5000 kVA by 2012, including air conditioning. As for all sites, there have been various modest additions of CPU and disc and tape storage and more are planned, including their first Woodcrest-based servers. There were multiple problems with their March disc acquisition traced to an un-notified change of electronics in the drives since the initial testing. There was also a major air conditioning concern on a particularly hot period (in England?) during July. They have experienced some performance issues with dCache which appear to have been stabilised only recently. CASTOR2 is expected eventually to replace dCache. Nagios is gradually replacing RAL's implementation of CERN's SURE.

**INFN** – network news include the new connection of CNAF by a 10GigE link to the GARR network backbone with a second one due shortly dedicated to LCG. GARR is predicting 40 to 100 GigE links by the beginning of 2008. GPFS has shown enough robustness and reliability that it is being adopted by more and more sites for shared file systems, despite tests of Lustre at a few sites. Also more sites are taking advantage of a shared INFN LSF licence. INFN has a development plan to build-up more Tier 2 sites and the first 5 sites have been selected, supporting three LHC experiments; all Tier 2 sites

---

<sup>5</sup> National Energy Research Scientific Computing Centre in Berkeley, California

<sup>6</sup> Parallel Distributed Systems Facility at NERSC

<sup>7</sup> Common UNIX Printing Service

are served by a 1GigE link to the GARR backbone and current work is concentrating in making sure they have a suitable power and air conditioning environment to cope with the equipment build-up needed for LHC running. Rather than this being a case of “catching up, the overheads stated that “[the] new LHC start-up schedule well matches INFN prudent approach to Tier2 initial growth”. There were lots of details on tests on iSCSI with the conclusion that “throughput is satisfactory but not completely satisfying”. EMC is offering an iSCSI device which may resolve their performance issues but they are also looking at home-grown solutions which may permit more flexibility for future acquisitions. There is a project called TRIP to provide a common infrastructure for travelling INFN users without needing specific site authentication. It is based on the [EAP<sup>8</sup>-TTLS](#) protocol and a distributed architecture of radius servers.

**GSI** – major problem with new large disc storage farm when 100 nodes out of 120 did not come back up after a planned power cycle; the vendor traced it to a vibration resonance problem of the cooling fans!. Also they had what the speaker called nightmare problems with new Foundry network switches.

**DAPNIA** – there is a new site security scheme in Saclay where DAPNIA falls into the second-tightest zone but the DAPNIA-GRID systems are in the outside of the 4 zones. Windows servers are arranged so that admin systems are in the innermost zone. Quattor is used via LAL at Orsay on the GRIF nodes at Saclay, recently upgraded by the addition of 23 new worker nodes. Lustre is in test in a new non-HEP collaboration with EDF Research, LAL and some CEA sites to test the use of fast networks (10Gb going to 40Gb) in applications; Lustre tests will be included.

**SLAC** – a BaBar run has just completed; SLAC has been made officially an ATLAS Tier 2 site and equipment acquisitions will start soon. And the new non-HEP physics developments are gradually building-up; for example [GLAST<sup>9</sup>](#) launch is only a year away. There is a new 153 Opteron cluster, space being made by removing 900 old cluster nodes. They are planning a 150-200 node Infiniband cluster for MPI work for large-scale modelling with 8-16GB memory per node – checkpointing is a challenge. Lustre is in test on a 5 node cluster for a high performance application. Although SLAC is still based on native Redhat linux (RHEL), there is some work on Scientific Linux (SL) which BaBar uses to build their code so that it runs on both RHEL and SL worldwide. They are testing ssh with Kerberos 5 to see if they can reduce AFS token passing. As already adopted at other sites, they have introduced a wireless registration scheme with redirection to a registration server on first contact. Chuck Boehm reported on work he did at CERN on his Console Server during his visit in the summer including client authentication and ACLs on individual console lines.

**INFN-CNAF** – similar to the report from RAL, CNAF also had a cooling problem in the summer (more understandable in Bologna than Didcot!) and this has delayed a scheduled hardware tender while the cooling of the centre is upgraded. Quattor is used as the installation tool and the speaker said some nice words about it. Lemon also will come into production use soon for monitoring. On the other hand, he noted that some CERN customisation did not always make them happy (SLC was quoted as an example) and some CERN products (e.g. CASTOR2) proved hard to install – he pleaded for better documentation. Their (conservative) users are remaining on SLC3 for the time being, including moving to V3.0.8 which supports Woodcrest processors but an upgrade to SLC4 will happen if the EGEE community decides for that. Development continues on StoRM, described as a candidate SRM<sup>10</sup> solution if CASTOR proves not suitable.

---

<sup>8</sup> Extensible Authentication Protocol

<sup>9</sup> Gamma Ray Large Array Space Telescope

<sup>10</sup> Storage Resource Manager

**LAL** – concerning EGEE participation (GRIF) there has been little change since the last meeting beyond normal capacity expansion although they still hope for major investment in the near to medium future. A PhD student is building a BOINC-like cluster using desktops for the XtremWeb grid to provide power for the Auger project. Elsewhere, web services are (finally) being moved off their True64 cluster to 2 dedicated Linux nodes although the cluster is still relied on for files and databases. They have current calls for tender open and the arrival of the Intel Woodcrest chip on the market has the side-effect that Dell is back in the bidding game. Sun appears still the winner for acquiring grid disc storage. They have installed SL4 and gLite3 ran without modification. Quattor is in heavy use and there are plans to implement Lemon monitoring. The migration from CERN Agenda to Indico is still in progress “with good support from CERN” but they are short of resources to complete it. Trac/Subversion use is growing (see later talk).

## Tuesday

### **SpamCop** – Jim Fromm (FNAL)

Earlier this year Fermilab landed in the SpamCop blacklist and this talk described why and how they dealt with it. Fermilab’s e-mail policy is to deliver all e-mail except those containing a virus or creating denial of service although spam mails are tagged on delivery. One of their 2 mail gateway servers was blacklisted because of an attempt to deliver a mail to an invalid address causing a bounced mail back to CERN which triggered a report to SpamCop by the CERN gateway. The first fix was to stop the bouncing of invalidly addressed mail and secondly they have returned to the situation where obvious spam is no longer delivered but quarantined in a gateway for a week. They also propose to create a HEPiX community of SMTP servers which would mutually not blacklist each other. Attempts to talk to spamcop.net got no replies despite repeated calls over several days. Finally, they will prepare alternative IP addresses for their mailservers in case of future incidents. They were eventually removed when their e-mail “score” for good mails against bad dropped below a threshold which happened approximately 24 hours after the incident but Michel Jouvin reported that LAL had a similar incident and they stayed on for several days; he also confirmed the unresponsiveness of spamcop.net. To set up the proposed HEPiX whitelist, the speaker suggested that interested people contact Kevin Hill at FNAL.

### **Trac** – Michel Jouvin (LAL)

[Trac](#), an open source project, includes a wiki, an issue tracker, a milestone roadmap manager and a source code browser. An instance consists of independent projects although project cross-reference is possible and projects can contain multiple components. The architecture is based on plug-ins, of which there are currently over 100, with an SQL backend database. It is written 100% in Python, so highly portable. There is a central http Trac server for multiple projects which can run inside an Apache server or as an independent web server. The only client is a web browser. Default authentication is via http and certificates can be used and there are plug-ins for Kerberos and LDAP. Authorisation is managed by privileges for one or several groups of components. TracWiki is a standard, simple but powerful wiki available to all components. In the milestone roadmap, milestones can be attached to issue tickets. The tickets are organised by components. Tickets have standard and customisable fields. There are standard reports but creating new ones requires knowledge of SQL although there is a newly-released plug-in to simplify this. The overheads include links to the [LAL Trac pages](#) and Michel showed a publicly-readable wiki on the LCG Quattor Working Group project

thread. The Trac community is fairly active but Michel did not know if any commercial support is available.

### **Twiki at CERN** (Hege Hansbakk)

Hege started by describing Twiki and the features which make it attractive and flexible. The first public Twiki service at CERN started in 2003 for the software development tool service and it has grown steadily ever since to over 1100 users last month across over 400K topics which were updated 4600 times last month. User groups include 4 CERN departments, ATLAS and EGEE. The service is based on 3 production servers using DNS-based load-balanced SLC3 servers with the data in AFS. There is also a stand-alone (non-AFS) read-only Twiki service which is effectively a backup version to cover for loss of data or loss of connection to AFS; it will soon be moved to CERN's Preveessin site as part of IT/DES group's disaster recovery planning. The team is currently fighting performance problems trying to find a particular bottleneck. They will also compare AFS to NAS storage. They will migrate to using NICE logon rather than AFS today.

### **Update on Scientific Linux** – Troy Dawson (FNAL)

Presented by Jim Fromm with Troy back in FNAL along with Connie Sieh, co-authors of SL. The steady growth in usage has continued since the last update; statistics show nearly 20,000 nodes and 25 public mirror sites. On the release front, SL 4.4 is due for release later this month with a 3.0.8 release candidate also in the works for this month – this will be the last 3.0 release although support for 3.x will continue until Fall 2007 as agreed at HEPiX last year. As follow-up on the Spring HEPiX, they have created a bugfix repository and a fastbug repository but XFS has still not made it to the main release. Work has started on SL5 with installer changes and removal of trademarks but they will not create a release based on the RedHat alpha release; the SL pre-alpha release will be based on RedHat beta release.

### **SLS Project at CERN** – Sebastian Lopienski

Sebastian described the [Service Level Status service](#) he has developed at CERN. As explained at recent talks in IT, this tool is used to show the availability of IT services (as opposed to individual nodes) and inter-dependencies between services (for example which services may be affected by problems or scheduled maintenance on other services). He explained its main features and how services are defined by their managers in terms of dependencies, what levels of service mean for a given services and how this is measured. Dependencies can be strong (required) or weak (used but not dependent on); services can be grouped into meta-services such as services for physics, for engineering, for an experiment. Both services and meta-services and dependencies are displayed on SLS web pages depending on the level you descend to. The web page shows services from a user's perspective where the level is expressed as a percentage, not in number of nodes for example. The service manager defines a threshold (which need not necessarily be 100%) above which a service is declared "fully available" and similar thresholds for "service affected", "degraded" and "not available". SLS also measures and displays Key Performance Indicators (KPIs) of performance levels compared to target values.

Sebastian gave an online demo of the service and he ended with how other sites could install their own instances. He claimed no CERN dependencies but experience tells me to be reassured of this only when Michel Jouvin of LAL confirms it although in Sebastian's defence he did a quick installation of SLS itself on the PC being used for display in the meeting room at JLab along with the addition of a simple service (a coffee machine). Based on the first 6 months experience, he noted that it is important

to trust service managers to define correctly levels of service; users must accept that it is not real-time in the instantaneous timeframe – it takes a short but finite time for problems to be reported on the SLS screen. He said it would be easy for sites to replace the Oracle database by SQL if they wanted.

### **System history and issue tracking with SVN/TRAC - Michel Jouvin (LAL)**

At LAL, they needed to track configuration changes and wanted a way to share service configuration across machines. In particular, versioning of configuration information and tracking open issues was desired. Subversion (SVN) was chosen for the versioning; it allows for atomic commits and supports renaming files and symlinks as well as reverts, branches and merges. TRAC was not a requirement, but was found useful as it provides at the same time a Web interface to SVN, issue tracking, and the possibility to link documentation, configuration and issue tracking. Michel showed live examples.

### **Using RT to manage installation workflow - Chuck Boenheim (SLAC)**

The SLAC computer centre is facing a rise of installation workload due to an increasing number of small orders for different stakeholders, a consequence of the diversification of the lab. In the recent past, this has caused confusion, dropped work, and stalls. They looked into project management software and found it too heavyweight, and not well adapted to administering multiple projects. [RequestTracker](#), on the other hand, in use at SLAC since some time, met most, of not all, of the requirements. Installations are now a single queue in RT, within which each batch of hardware creates one ticket. Each ticket has fields to be ticked off during installation such as order status, host name, DNS, power, rack, network. Child tickets are created for each required step, and assigned to the respective responsible party; tasks can proceed in parallel. RT supports attachments (e.g. PDF, spreadsheets) to tickets; it is possible to equip tickets with links to the purchasing queue (or any other bookmark). An installation dashboard provides a vertical view (by step) to all incomplete installation tickets. The initial experience is very positive, although some work needs to be done to convince everybody involved to enter the relevant information into RT. However, as RT was in use at SLAC already, the introduction was less difficult than it would have been with a completely new system.

### **High-availability methods at GSI - Karin Miers (GSI)**

Karin described several approaches in use at GSI to provide high-availability services: software solutions (both commercial and open-source) as well as hardware implementations. All software solutions imply some sort of heartbeat checks between the servers involved, an IP address for the virtual service, configurations in active/active or active/passive mode, an area for shared data, and monitoring tools. For their Exchange, file and print servers, they use Windows clustering in active/active or active/passive mode and SAN storage or local shared disks. They also run an Oracle Real Application Cluster, which they have very good experience with. Linux services requiring high availability use the packages heartbeat (for the communication between nodes), drbd for file sharing (creating a common block device of a local and a remote disk device), and mon for the monitoring. Currently they use these techniques for the central file services (nfs), the central Web server, and the entry server for their mass storage system. Finally Karin described a hardware solution (Coyote equalizer traffic management); it provides failover and monitoring. So far, the experience with all these solutions is good, although it is rather complex to configure the software stack under Linux correctly.

### **Managing a Grid Fabric with Quattor – Michel Jouvin**

In his third talk of the day, Michel started by describing Quattor, a machine installation and configuration tool developed as part of EDG. Designed to handle thousands of nodes; it uses a central configuration database, originally stored in CVS. There is an extensive framework of components to

manage different systems. Originally the EGEE middleware was targeted to be YAIM but it threw up a number of serious and less serious issues, for example being node-oriented rather than service-oriented and with thousands of nodes at LCG Tier 1 sites this is important. Eventually, a Tier 1 working group led by Cal Loomis at LAL decided on Quattor. The general architecture is to generate as many templates as possible automatically with no modifications to standard templates and taking advantage of the layered nature of templates for low-level services building up into a single template for each highest-level service. There is also a collection of single templates for each major machine type. VO configuration, especially difficult under YAIM, is currently under development.

Michel appreciated especially the ease with which the Quattor Working Group-produced templates for O/S and Middleware updates can be downloaded. He estimated that Quattor is being used at around 20 LCG sites around Europe on many thousands of nodes (4100+ at CERN alone as of October 2006). German Cancio, head of the CERN Quattor development team, estimates from March 2006 show that it was used then on over 7000 nodes worldwide. It has a definite learning curve and documentation could still be improved. However he considers it an effective tool, even for small sites with as few as 10-20 nodes although the templates from the QWG are not used everywhere and these are using Quattor via a YAIM interface. He hopes the use of templates will increase and he reported that the second Quattor user conference is scheduled for DESY next week.

### **Spam Fight at GSI – Walter Schoen**

Tools used include postfix, [amavisd-new](#)<sup>11</sup> and [Spamassassin](#). In Germany, sites are forbidden to reject spam – the actual rule is that once mail is opened at, that mail must be delivered although there are exceptions for e-mails with viruses but even here the user must be informed. Each user can set their own level of spam rejection. In Spamassassin, GSI use standard configuration and rules as far as possible with few local modifications. They have switched off the user learning feature in the Bayes filter because it was being misused and effectively decreasing the filtering efficiency. They use “smart” greylisting of sites, temporarily blacklisting domains which trigger the spam signal more than 10 times but “freeing” them if a valid mail with the same sender, source IP and recipient arrives within 10 minutes. The greylisting on a particular day in September caught 14000 mails of which 1500 were delivered and 12500 rejected compared to 13000 detected spam and 14000 non-spam mails in the same day. He described the mechanics of a recent event whereby SpamCop set up a honeypot trap into which GSI fell because of the legal obligation to accept all mail, including forwarding it on request, and to inform the sender (fictitious in this case) when the forward site, CERN as it happens, correctly rejects the spam mail which of course triggered SpamCop to blacklist the GSI mail server (which had performed the forward) for a while.

### **Scientific Linux Inventory Project (SLIP) – Troy Dawson and Connie Sieh (presented by Jim Fromm)**

The need for such an inventory was mandated by the US DoE and also wanted by the developers to enable security queries and hardware and software queries, all to see where given software is being used. SLIP is a set of programs to build a software inventory built on [OCSInventory](#) from sourceforge. The tool uses MySQL which was also essential as it should be bundled with SL which therefore prohibits the use of a commercial package. The inventory contains much information on when an RPM was built, where it is installed and so on. Interestingly, the feature list of OCSInventory reads almost exactly to the list of features listed by Michel Jouvein earlier in the afternoon when

---

<sup>11</sup> Described on its web pages as a high-performance interface between mailer (MTA) and content checkers: virus scanners, and/or SpamAssassin

describing Quattor ... The current status is that it has passed initial tests and is being rolled out workgroup by workgroup. Better reporting tools are under development.

### **From a Spark in a Vacuum to Sparking the Vacuum – Fred Dylla**

The day ended with an entertaining after-dinner talk at the Conference Banquet by the CTO of JLab in which, as he said, he covered 200 years of the history of vacuum tubes in 20 minutes.

## **Wednesday**

### **PXE Installation Management of RACF at BNL – Chris Hollowell**

RACF is the RHIC/US-ATLAS computing farm. The initial installation scheme based on Redhat's Kickstart was found to have limited scaling capability largely because of the need for boot media to initiate the installation and a PXE<sup>12</sup>-based scheme has been adopted. He explained briefly how PXE works and why it was chosen. Among the alternatives studied is the [Scyld Clusterware](#) tool but it is not available free of charge and even the vendor admitted it would not scale to thousands of nodes. [ROCKS](#) is another alternative but it had seemed too complex for their more modest requirements. He then described the server configuration in some detail and the actual installation procedure.

### **Support of a Kerberos 5 Authenticated Environment at DESY – Bogdan Lobodzinski**

He started by telling the HEPiX audience where DESY is in relation to a nearby World Cup venue, what it does and offered a brief view of the computing environment at DESY – a site report by any other name. He then explained how DESY relied heavily on AFS and that handling of Kerberos 5 tickets is essential for long jobs in their Torque batch system. For this they use a mixture of the Arcv2 tool written locally where the arc server must run on the KDC<sup>13</sup> and local client/server RPC connections which is independent of the KDC. He explained why they had re-written the arc tool and how it operates and then showed the same procedures using the RPC method. Finally he compared the alternatives in terms of complication and security implications and noted that in the end they had decided to use local RPC client/server solution.

### **Plans for Hall D at JLab – Bryan Hess**

This talk was subtitled “Hazards of Fast Tape Drives”. He reminded the audience how JLab had developed Jasmine<sup>14</sup> to make the most efficient use possible of their tape drives and how it works at a high level. Experiments in JLab's planned fourth experimental hall are expected to require significant tape capacity expansion. He analysed their expected needs in terms of what Sun/STK T10000 drives can achieve and also what would be the effects on their networking infrastructure. While noting that small files have a high overhead, he questioned what “small” means as experiments grow in size and complexity. Another problem is to avoid tape stop/start – the “shoe shining” effect as he called it.

### **Porting to 64 bit Platforms – Carl Timmer (JLab)**

---

<sup>12</sup> Pre-Boot eXecution Environment

<sup>13</sup> Key Distribution Centre

<sup>14</sup> See [previous HEPiX](#) for a report on Jasmine

He emphasised that this would not be a tutorial but a report on his experiences. He noted that while Java code should not create significant issues moving from 32 to 64 bits, C and C++ code often could, especially concerning data models and he listed a number of data model dependencies in particular to look out for and some hints to avoid them. He listed switches needed to force 64 bit compilation on various platforms and showed examples of source-level Makefiles that he found useful. On performance, he reported some results with various chips running under Linux, one on shared memory data sharing, the other on a message passing programme. He showed graphs of the results – see overheads; the 64 bit shared memory application did better in every case although only just on the Xeon. However, for the message passing application, 64 bit mode was actually slower in every case but he could not explain why.

### **NERSC Global File (NGF) System – Cary Whitney**

The need was to manage very large scale file stores among an audience of very mobile users moving between NERSC services. [NGF](#) is based on GPFS, uses 24 I/O servers, 70TB of end-user data and 3GB of network bandwidth for streaming I/O. A global /project filesystem, mounted on all major NERSC services covering a variety of UNIX and Linux flavours; it is used to share data between project participants. NGF is currently deployed on over 1200 user nodes. Backups are fortnightly to HPSS but it is intended to move to nightly incremental backups. Users report good feedback with respect to ease of use and performance. There have been outages but in general performance is judged as acceptable. Some of the outages are blamed on older hardware and the architecture has in-built redundancy. They are concerned with possible future requests regarding compatibility with current local file systems, including local GPFS instances (there are 12 such instances on the site). They would have liked to expand this to users' home directories but the use on different UNIX/Linux flavours complicates this. There is no concept of the AFS @SYS variable so users would need to somehow distinguish filenames which may have different significance on different operating systems.

### **Storage Classes – Jonathan Schaeffer (IN2P3)**

This is an attempt to improve the usage quality of the storage. It covers storage on tape only, disk only and stored on both. When a file makes a transition from one storage class to another, it can create confusion within the namespace and it would be ideal to maintain a single storage class within a given namespace. In dCache, storage characteristics of a file, known as Storage Groups, are determined by its namespace and the physical location of the file is determined by its namespace. The namespace can change without affecting the location thus dCache features support storage classes quite well. This is not the case in HPSS at IN2P3 where files are stored according to their PNFS namespace and HPSS has the concept of Class of Service (COS) which is not linked to namespace. He then showed the relationship between PNFS namespaces, dCache storage groups and HPSS namespaces and storage classes and how the various transitions could be handled.

### **CPU Benchmarking in CERN – Helge Meinhard**

Benchmarking has been a recent hot topic on the HEPiX mailing list and Helge reported on recent work in this area in CERN, mostly on CPU benchmarking but also including some I/O performance results and (electrical) power efficiency. He reported that although SPEC benchmark codes get close to CERN's requirements and there appears to be no need to develop our own benchmarks, the results presented in spec.org are meaningless because the environment under which these are produced does not match the operating conditions in CERN, nor probably in other HEP sites. CERN therefore has a fixed and reproducible environment and performs or requires vendors to perform the SPECint2000 V1.3 default source benchmarks. CERN also requires multiple streams to be run corresponding to the

number of cores in the processor. The last two CPU tenders have been based on this scheme, asking for a given number of SPECints, requesting the vendors to reply with the number of processors they needed to fulfil that target. These two tenders have also included a measure of the power consumption in relation to the CPU power and Helge described how this was done. As mentioned in another site report, CERN is considering moving to the 2006 SPEC revision but Helge feels he would like to understand better how closely these correspond to typical CERN usage as well as to refine the power criteria in the calls for tender.

### **Fabric Management at CERN – Tony Cass**

Tony covered some items concerning batch queue management and system monitoring. CERN is still using the LCG CE<sup>15</sup> and considering when to implement the gLite CE, partially at least depending on when its use becomes more widespread across LCG. He reported that the LSF information provider has been completely rewritten, in particular to support multiple CEs. [SMART](#) is a tool for monitoring wear and tear problems on hard disks, originally written by Compaq and now available from sourceforge. It supports various disc architectures but not yet all RAID cards. Its use permits CERN to pro-actively replace discs which SMART reports as being about to fail. CERN is “encouraging” its suppliers to use this data to replace a disc before failure and several have already accepted this policy. CERN proposes to extend its use, for example to trigger the automatic draining of a failing disc. They are also considering linking messages from SMART to the Lemon general monitoring service, for example to avoid failing discs affecting the CASTOR service.

### **The Stakkato Intrusion – Leif Nixon (NSC, Sweden)**

This was the story of a security incident at the Swedish National Supercomputer Centre which had started in 2004 and culminated in the summer of 2005 with the arrest of a suspect. It began with an intrusion which had led to a succession of Trojan horses, use of the Suckit root kit and the theft of many user and system administrator passwords before detection. The story was told in a beautifully-controlled calm and flat tone despite the trauma and stress it must have caused the presenter and his colleagues at the time. He estimated that over a thousand sites all over the world were affected at one time or another during the attack. He described a number of clues to look for to see if your node has been hacked and some (painful) lessons learned. Interested hackers (sorry, security experts only) should consult the overheads for more detail or talk to Sebastian who took detailed notes.

### **Network Security Monitoring – David Bianco (BNL)**

Network Security Monitoring is a procedure or set of procedures rather than a particular tool and it includes activities such as intrusion detection, event log and system log scanning. The tool BNL uses is called [Sguil](#), an open source tool based on servers, clients and sensors. Sensors are deployed around the network and they send data to the (single) central server and its MySQL database and network/security administrators monitor the results from clients. Sguil sensor components include [Snort](#) for intrusion detection, [SANCP](#) for collecting session information and Snort again for full packet capture. On the server, the sensor data is gathered by a daemon and stored in the database. Squil is not an alert browser, all alerts must be dealt with. He showed some example alerts and described how they are typically dealt with and finally walked through an example case study from last December.

---

<sup>15</sup> Compute Element – interface between the Grid queues and local batch nodes.

## Thursday - Grid Talks

**GridX1: A Canadian Computational Grid for HEP applications** - Ian Gable (Victoria University)  
GridX1 is a collaboration of eight Canadian sites (2500 CPUs, 100 TB disk, 400 TB tape) that aims at providing services to HEP and non-HEP applications alike. The CPUs are running a variety of operating systems; VDT (a packaged version of the Globus Toolkit 2.4) is used as Grid middleware. The Grid has contributed significantly to Atlas production, having completed 20,000 jobs. For BaBar, they run Monte Carlo production as well and they account for 15% of the total BaBar production. They are investigating using Xen virtual machine scheme in order to overcome the mismatch between the operating systems the worker nodes run under and what the jobs require, in particular for non-HEP applications (see their contribution to CHEP2006). They intend to move the Grid middleware to Web services, using the GT4 Condor-G JobManager, the MDS ClassAd extraction tool, and potentially the GridWay metascheduler.

### **GridPP** – John Gordon

John presented this as a talk on experiences running a production grid. He started by situating GridPP in relationship to EGEE and LCG. GridPP is the UK and Ireland part of EGEE but each UK region has its own grid and they hope eventually to unify them all. They have funding for a second three year period and a third funding request has just been submitted for the exploitation phase of LHC. Where RAL is the UK Tier 1 site, there are 4 virtual Tier 2 sites spread across the UK linked by Memoranda of Understanding to support particular VOs. They have noted issues with trying to get immature middleware to run in production mode on a large number of sites, a well-known problem within LCG. They have tried many tools for installation and configurations but simple tools are not powerful or flexible enough for the complexity of configuring grid nodes and complex ones are resisted by local site administrators. The preferred method in the UK is YAIM.

Another issue is release management – steady drip of bug fixes and updates or big bang? There are pluses and minuses for both. The configuration and support of VOs must be defined in a more standard method. John then covered some of particular topics where GridPP has taken the lead such as R-GMA<sup>16</sup> which is the foundation of a number of monitoring tools; a grid operations database (GOCDDB) holding information about sites across the grid; an accounting package called APEL and the LCG Monitor tool which shows the flow and execution rates of jobs across LCG.

Lessons learned include

- “good enough” is not good enough
- Exceptions are the norm
- Bad error reporting wastes a lot of time
- Very few people understand how the whole system works
- Communications between sites and users is an N times M problem

In summary, work during the next year before LHC starts operations must concentrate on making the middleware more mature, improve communications and build middleware and operational procedures to expect and take care of failures.

### **The EGEE Production Grid** – Ian Bird

Ian reviewed briefly the history of EGEE starting with the [Monarc](#) network modelling study in 1999 through EDG and EGEE 1 to today’s EGEE 2 service. He described in detail the infrastructure part of EGEE to which more half the project funds are dedicated. The goal of Ian’s team is to produce a

---

<sup>16</sup> Grid Monitoring Architecture

middleware distribution which can be widely deployed and he listed the steps in this procedure. He showed lots of colourful plots which showed that the service is currently approaching 200 sites, 32,000 CPUs across 42 countries worldwide. They are now running at around 50,000 jobs per day of which nearly 10,000 are non-HEP and he listed some of the other application areas and in which way they differed from HEP in their use of the grid. Typically there are around 20,000 jobs executing at any one time. Moving on, he presented the management and operational structure with the Regional Operations Centres, the portal for user support, etc. Monitoring must be pro-active but in the end problems are almost always passed to the local site managers for repair. There is a framework for site functional tests (SFT) leading to a site availability monitor (SAM) and he showed recent plots for September for the LCG service, illustrating how one could drill down to the machine level at an individual site to trace a problem. He described the interoperability with other grids such as OSG in the US, Nordugrid in the Scandinavian countries and so on and also the relationship to various related infrastructure projects. Finally he noted that we need to start preparing now for the next phase of EGEE and also how to establish a longer-lived grid support structure.

### **Virtual Machines in a Distributed Environment – Mauricio Tsugawa (Uni of Florida)**

He started by describing the [ACIS Lab](#)<sup>17</sup> at the university. He then gave a rather academic one hour tutorial on resource sharing issues and challenges and how the concept of virtual machines could be a solution to these issues. He showed how virtualisation technologies can be applied to grids creating a virtual space above physical space on the network. Then logical links can be partially or totally isolated by using multiple physical links, virtual routers, tunnelling and so on. This leads to what he described as a Grid building recipe and “[In-VIGO](#)”<sup>18</sup>, a virtualisation middleware package for grids. He presented the In-VIGO portal and described how it can be used to build a virtual workspace. He also showed how web services are used to link the applications to the infrastructure and are essential to intercommunications (see summary slide). In summary, he claimed that virtualisation technology allows the user’s application requirements to be satisfied independent of the constraints of physical resources.

### **Site Manageability Issues for LCG – Ian Bird**

LCG has started to measure site and service availability and the results are not good even considering only the Tier 0 and Tier 1 sites; it has been a problem for many months and it is not getting better. Apart from missing the set targets, this is a poor example for the Tier 2 sites. Noting again that SFT is gradually being replaced by SAM<sup>19</sup>, he explained SAM in a little more detail including the sensors providing the raw monitoring data. He noted that a test job is sent at intervals to the Tier 1 sites but that sometimes, instead of letting the batch system select a random worker node, the local site administrator selects where the test jobs will run and this means one cannot be sure that the results returned are representative of the service from that site. LCG proposes a site availability measurement based on the availability of a defined set of core services (CE, SE, BDII and VO-specific services). An algorithm would transform that into a metric representing the availability of a site. Also, Ian would like to see this data further used to trigger alarms when the metric drops below a threshold. In the future, he hopes more tests can be added, better documentation of the tests is required, false alarms must be identified and the metric itself needs to be reviewed.

---

<sup>17</sup> Advanced Computing and Information Systems Lab at the University of Florida

<sup>18</sup> Virtual Information Grid Organisations

<sup>19</sup> See Ian’s earlier talk for explanations of SFT and SAM

From an experiment view, the [ARDA](#) project investigating distributed analysis for LHC experiments has tried to gather information about the efficiency of job submission/execution. Apart from anything else, they show that many of the error messages are not particularly useful and often hard to understand or analyse. From this work, ARDA has produced dashboards by LHC experiment of site efficiency and they show there is no consistency (neither regularly good nor bad efficiency) of any given site from day to day. It is proposed to produce a simplified set of tests to be run on all worker nodes at a site measuring the local environment of the node and its capability to run real jobs. Such tests should run without creating a noticeable overhead but giving a more complete picture of each site's availability – a procedure he described as job wrapping monitoring.

What is needed, Ian claimed, is NOT more tools, there are plenty of these already and he listed many. Some of the smaller sites are only recently discovering what it means to operate a production computer centre and they should be assisted by better documentation and more training. There is also a need for better security management. Sites need more feedback on problems found by SFT/SAM and the job wrapper monitoring described above.

Ian closed by asking if HEPiX could help define how to measure and indeed improve site availability and performance. What follows are some of the responses from the audience to this proposal:

- John Gordon – yes, HEPiX can and should help but it cannot wait until the next HEPiX meeting (April 2007) but there should perhaps be a report then.
- Ian – any study or discussion need not necessarily take place at a HEPiX meeting.
- There was a discussion about how to initiate any study – a volunteer is needed. All agreed that it could not be limited to one grid community which is why HEPiX appears a more appropriate forum but a driver is almost certainly required initially.
- Lisa Giachetti recalled that she had seen the beginnings of a Grid Cookbook but Ian replied that it had not gone forward.
- John noted that false results meant that some sites did not yet trust the tests and Ian agreed
- Chuck Boenheim noted that an earlier recognition that error reporting should have been better architected, or even architected at all. Again Ian agreed
- Michel Jouvin said that an LCG workshop already planned for January could be a marker for a first discussion.
- Alan invited the sites represented in the audience, especially the non-LCG sites, who are interested in this activity, to propose to Ian by the end of October the name of a participant in the study.

### **FermiGrid – Keith Chadwick**

He listed the components of the grid. He described how the site gateway accepts and routes jobs. He also described the interface to OSG and how this offers opportunistic use of FNAL resources. They use [GUMS](#)<sup>20</sup> to map recognised OSG and some LCG VOs and they developed a Site AuthoriZation (SAZ) tool to map job authorisation according to the username, VO, role and certification information. He listed various site policies applied for joining a VO and for submitting jobs (see slides for details). He listed the various grid services, their status and how service redundancy was provided. Monitoring shows a maximum of around 6000 jobs per day.

### **Open Science Grid (OSG) Progress and Vision – Keith Chadwick**

---

<sup>20</sup> Grid User Management System

OSG is currently operating 70 sites, 24,000 CPUs. A new phase started last month with funding in place for 5 years with 34.3 FTE funded in the first year. There is list of non-HEP application areas including biophysics, nanotechnology, molecular dynamics and so on. He listed the components of the middleware stack and noted that they plan for 2 major releases per year with minor updates between and he showed what is expected to be in the next release. He also presented major area of current developments, such as work on VDT components. In relation to the naming convention of the releases, John Gordon suggested that calling the next release version 0.6.0 rather than 1.x may be read as a lack of confidence in the software but Keith replied that it may be more that the developers are too realistic about its state.

### **WLCG and EGEE Grid Security – David Kelsey**

He first reminded the audience how users register once (with renewal) with a Virtual Organisation (VO) and each VO registers once (with renewal) with each grid. Sites provide resources to grids and decide which VO to support. The model is based on authentication (who you are), global authorisation (what rights you possess) and local authorisation (what you can access at a site). The International Grid Trust Federation, founded a year ago, has established 3 regional Policy Management Authorities (PMAs) to accredit identity providers, own and maintain authentication profiles, coordinate the X509 namespace and distribute the highest level trusts. CAs are members of PMS. Authentication profiles include the classic PKI service and there is also a new short-lived certificate service (SLCS) where the certificates are generated by local site authentication services such as Kerberos; their short life means there is no need for revocation lists for these. He then described the VO Membership Service (VOMS), how global and local authorisation service works and how a user joins a VO.

A lot of work has been done by the Joint Security Policy Group ([JSPG](#)), which now includes representatives from OSG, into producing common Acceptable Use Policies (AUPs) and many of these are probably now due for updating. The JSPG is very pleased to have produced a common policy document between OSG and EGEE on what conditions or rules a user must accept to join either grid community and he listed other recently-agreed AUPs, others being worked on and yet others which must be worked on soon, a number of which should if possible be shared by other grids.

The newly-formed Operational Security Coordination Team (OSCT) led by Ian Neilson from CERN is planning service challenges on the security of the grid, for instance by creating an event and monitoring how it is dealt with. They will perform risk assessments with the aim of incrementally improving the vulnerability of the grid. The high level e-Infrastructure Reflection Group (eIRG), appointed by European Union ministers, has a long-term roadmap which includes plans for various federations to consider authentication and authorisation matters and there is work underway already in EGEE 2 in this area. The Open Grid Forum (OGF, the successor to GGF) ([LINK](#)) has created a Grid Interoperability Now (GIN) project and EGEE runs a VO for this already.

## **Friday**

### **UK Tier 2 Data Storage – Chris Brew (RAL)**

In order to prepare for smooth(er) production running when LHC starts, it was decided to work with UK Tier 2 sites to establish and evaluate end-to-end connections between Tier 2 sites<sup>21</sup> with each Tier

---

<sup>21</sup> Although it is not currently planned to have direct T2 to T2 connections in LHC production mode, it was felt that involving the UK Tier 1 site at this time was not possible because of Tier 0 to Tier 1 testing going on in parallel.

2 site being assigned to one of three Tier 2 sites designated as reference sites and effectively behaving as a surrogate Tier 1 site. Apart from testing the infrastructure, advantage could be taken by having UK data storage experts pass on some of their knowledge to Tier 2 system administrators. He described the FTS file transfer software used (from the LCG middleware stack) and the procedure used. Among the lessons learned were

- One of the reference site previously showing good performance could not sustain a reasonable performance level
- Sites were allocated fixed time slots and there was not enough time allocated for all the tests if there was a problem such as a delayed start
- Network variability
- Test procedures not fully explained or understood
- Problems with e-mail communication
- Suspicion that Ganglia mis-reports bandwidth rates

Despite all the above, some positive results were achieved on 14 out of the planned 19 sites, the most important of which was the community spirit of the site administrators, a feeling of ownership by local site administrators and confirmation that FTS is much more of a stable service than in previous tests.. So far, 60TB has been transferred at a maximum rate of 330Mb/sec. Work is continuing on sites with problems and with gradually building up the network bandwidth efficiency. [As someone with no previous knowledge of these tests, I came away with a somewhat negative impression but I was assured by both other RAL and CERN colleagues that the presentation under-sold the positive results achieved. Perhaps this is a lesson to present the successes of such a test before presenting the (negative) lessons learned.]

#### **IHEPCCC – Randy Sobie (IHEPCCC Chair)**

Via an audio link from Victoria, Randy reported on a recent discussion in IHEPCCC about the possibility of joint working groups with HEPiX in areas of mutual interest, the first two of which could be benchmarking (already a hot topic in recent HEPiX mailings plus a recurring theme of one or more presentations at all recent HEPiX meetings including this one) and global file systems (also mentioned by a number of sites earlier this week). IHEPCCC proposes the so-called RTAG<sup>22</sup> model which was used successfully by LCG; (most) RTAGs were short-lived study groups generating brief reports. Randy explained the background to these two topics and what could be of interest to study. The overheads contained a possible template for the working groups, make-up, mission, timescale.

Randy noted that the initial problem may be finding chairs for each group but maybe the role could be split. I suggested that IHEPCCC themselves could propose a chairman for each and Randy agreed to investigate this with his committee. Although it was agreed that the benchmarking working group should initially discuss if SPEC2006 was representative and sufficient for a HEP benchmark suite, Helge suggested that ongoing review of this area should continue and this is not consistent with the RTAG format; Randy agreed but we should at least start with the RTAG format to provide short-term results. IHEPCCC's proposed timeframe was for reports next January but HEPiX would prefer to aim at reports by the next HEPiX meeting in April. Randy will now talk to his committee about participation and chairmanship and he has been added to the HEPiX Board mailing list for further discussion around the establishment of at least a core team for each group before opening the discussion to the full HEPiX mailing list. Randy Melen of SLAC noted that some HEP labs have or are branching into other fields and these should be included in such discussions even if the needs may not always be the same. Randy (Sobie) agreed and welcomes such participation.

---

<sup>22</sup> Requirements Technical Assessment Group

Alan Silverman  
18<sup>th</sup> October 2006