HEPiX Autumn 2009

NERSC, 26-30 October 2009

CERN-IT-Note-2009-022

Contents

The Autumn 2009 HEPiX was hosted by LBNL NERSC at the UC Campus in Berkeley. The meeting was held in a large conference room high on the hills above the town with a spectacular view over the San Francisco bayshore. Attendance was around 70 persons from a variety of US, Canadian and European sites. As usual almost all of the overheads are available on the (Indico-based) web site - http://indico.cern.ch/conferenceTimeTable.py?confId=61917&showDate=all&showSession=all&detailLevel=contribution&viewMode=parallel.

The views expressed below are mine, as are the errors and omissions. I thank several of my CERN colleagues for corrections to the original draft.

Executive Summary

- Good attendance for a US meeting, good representation as usual from European sites.

- Storage by far the most popular subject in terms of submitted talks
- Interesting presentation of the newly-announced Magellan project (see under Virtualisation)
- Most intriguing presentation so far – Walter Schoen's description of a 3D cube being proposed for computing at the FAIR project at GSI. A steel frame 26m to a side and 26m high to house 1000 racks of computers.
- Another hot topic was virtualisation. There were a number of interesting talks and Tony's "vision for the future" presentation created lots of discussion, one conclusion of which is that HEPiX can be an important vehicle for sites to share experiences and perhaps gradually help us grow towards a common offering for the experiments. The idea of a dedicated working group was raised and the Board agreed and asked Tony Cass to run it. Initial target will be experience sharing.
- Another subject mentioned several times was research into using GPUs for HEP computing.
- The acronym of the week was ARRA - American Recovery and Reinvestment Act. Grants are awarded for new projects (not previously planned) and for spending now, and in the US. The NERSC/ANL Magellan project had received $32M, ESnet got $62M for an "advanced networking initiative" to build a prototype 100Gb network and JLab $5M for a lattice QCD project. How it would be nice to be faced with an unexpected grant for something new and asked to spend it quickly!
- Next meetings: April 19 to 23, 2010 in LIP, Lisbon; autumn in Cornell (choice of 3 weeks under consideration, complicated by CHEP timing). Would like to hold Spring 2011 meeting in Asia after the success of the Taiwan meeting but we need to identify a host site. Autumn 2011 likely to be TRIUMF.

Keynote Talk

The **Opening Talk**, on Hardware and Software in the Multi-core Era, was given by Kathy Yelick, the Director of NERSC. She covered some current projects at NERSC including one on cosmic microwave background, dealing more with data reduction than computing problem, having enormous 32TB data sets. She noted that data sets grow with Moores's Law. Another example is KamLAND, a neutrino analysis experiment; it is recording 200GB/day every day. Then ALICE expects 600TB of data in 1GB files, some 25% of US ALICE data collection. And the Palomar sky survey collects 300GB per night and has already discovered some 40 supernovae and finds a new "event" every 12 minutes! A common theme is that it's not just about providing computing or data storage, more about organisation to offer secure public access to that data, especially where some of the data sets cannot be copied to "home". Fast I/O is key and manipulation and analysis of data is a growing problem. Hence NERSC has built "science gateways", custom hardware and software to offer remote access to the data and/or computer services. They also have developed new models of computational access, including a "computational beamline", a scheme of reserving resources for anticipated needs.

NERSC has a Cray XT4 and they are currently assembling a Cray T4. Plus PDSF, more familiar to HEP users. The speaker then briefly introduced the Magellan project, a $32M project at NERSC and ALCF, 100 Teraflops per sec, with PB scale storage. The DoE has decided to explore clouds for science computing. Can a cloud serve's DoE's mid-range computing needs? What part of the workload can be served on a cloud, what features are needed? And how does a science computing cloud differ from a commercial cloud? [More details of this under Virtualisation.] Coming to multi—core, she explained why these are becoming more common but wondered for how long we can double concurrency in the algorithms and software every few years to maintain performance increases in the chips coming from the vendors. She suspects traditional sources of hardware performance improvement are flat-lining, exponential clock speed growth is slowing and hardware to extract parallelism has limited promise for the future.

She compared the power and energy needed for chips, from an IBM Power 5 chip to the processor typically used in mobile phones. The latter may run at less than one-third the speed of the former but it consumes 1/400th of the power so it is far cheaper to run – a farm of mobile phones? And the loss in compute power is less dramatic than may be expected since a lot of it is due to dropping many X86 instructions no longer actually used by modern codes. In general, power will dominate other concerns. The new trend is to double the number of cores every 2 years although clock speed will not change so much; cores will be very simple but with more data parallelism and more use of wide SIMDs, GPUs and accelerators. The programming model will remain MPI with the addition of "X" for multi-core but what is X? Could be OpenMP but too much serial thinking. Could be Berkeley's Unified Parallel C project or a similar language but these tend to have very lightweight communications. Could be CUDA (Compute Unified Device Architecture), a parallel computing architecture developed by NVIDIA, or the related OpenCL; developments on these are being targeted to multi-cores.

Her conclusions were that single processors are not going to get much faster and memory per core will drop. Parallel flops are "free" but memory is expensive. Adding parallelism will be important for all applications. Power and energy costs will dominate and we will need to work with more experts on software, algorithms and applications. Finally she noted in passing that she does not consider that virtualisation will be the answer.

Site Reports

CERN: Helge gave brief details about LHC plans as now known, progress on ITIL and the problems discovered in the Linux kernel around CPU timers which are only going to be fixed in Redhat 5. He listed recent procurements plus those in progress or planned. He mentioned the plans to introduce Microsoft's ILM2 for identity management and how this should offer better support for single sign on. Other items covered were the STEP challenge on LCG this summer, the new desktop support contract and the increasing interest in virtualisation.

Jefferson Lab: celebrating 25th anniversary plus their 12GB upgrade electron beam approval and a $5M grant for their Lattice QCD project which will increase their computing capacity from 3 TFlops to 20 TF. The new cluster will include significant storage and they will probably run Lustre. They have been unable to spend any funds on their Auger analysis cluster since 2006 but purchasing is now restarting. Their physics users remain lab-centric as opposed to grid oriented and the computing team is trying to persuade them to benchmark their codes to better tune the cluster. They will switch from Redhat/Fedora to the CentOS flavour. They managed to add to their mass storage system with a new IBM TS3500 tape library and they have replaced their StorageTek Powderhorn silos. They use Sun Fire X4500 (Thumpers) for cache today but may change with the QCD cluster experience being gained. The QCD acquisition has also meant a lot of physical moves to guarantee access of the new systems to sufficient power. They have some 80 production VM servers with 30-50 more planned. For Cybersecurity, they have an "off year" (external audit every 2 years) and are performing a self-assessment.
Q: why CentOS not SL? A: the decision maker does not like the packaging of SL; but does now accept CentOS packaging!

GSI: concentrating on the FAIR project on which there is some progress. The computing needs of the largest experiment, PANDA, is comparable to LHC computing so they have a long way to go to be ready for first beam in 2016. The first problem is the lack of physical space so they are proposing a €9M 3D cube for stacking 1000 19 inch racks in a minimal footprint, built round a steel skeleton 26m on a side and 26m high. Expected to host 40,000 nodes. No chance for air cooling, only water cooling. A pilot mini-cube is under construction). Within the cube, the

racks will be stored such that there are gaps between racks for human access. but the current working temperature within the cube is 30° which does not make for a comfortable environment for the staff working in there, e.g. to exchange a disc.

IN2P3: have now decided to migrate away from their home-grown BQS and currently studying alternatives. Also migration to SL5 is underway. Able to run some 70K jobs per day on their main Anastasie cluster. A second 544 core cluster runs MPI and PVM jobs. And they have installed a dedicated 20 worker node analysis farm based on PROOF but this is still under test; ATLAS is the main user but CMS are studying it. HPSS is alive and well and version 6.2 has now been installed. Other mass storage systems in use include dCache, AFS, SRB, GPFS (although this is stopping and may be replaced by Lustre) and a stand-alone xrootd service interfaced to HPSS (under evaluation for ALICE but already in use for some non-HEP experiments). There have been some enhancements to the computer building including a doubling of power to 3MW and an additional chiller; the trend is towards water-cooled racks. A new building has been proposed to house a new 800 sq.m computer room.
Q: why change batch systems? A: oversight committee decided that the lab could no longer justify a dedicated development team for BQS.

IRFU Saclay: various IRFU sub-sites being incorporated into the GRIF grid. Windows 7 being installed and all new Windows systems will get this in 64bit mode (Vista is forbidden).

RAL: new building now fully operational since April and new Tier 1 procurements were shipped in along with elements of the existing Tier 1 over the summer. Since then, there were a couple of air conditioning failures and a moisture source was uncovered (above the tape robots!) and this is being repaired. Planning changes resulted in a shortage of chilled water so this had to be solved with some urgency. Traditional increase in CPU and disk capacity and a second SL8500 silo was added and there is more to come. One storage delivery from last year is still failing acceptance tests, despite ongoing support from the supplier. Several controller crashes on CASTOR and other storage and investigations are ongoing, possibly an electrical interference. Bought a Force10 C300 switch for policy-base routing only to discover the current version does not support this; coming in a future upgrade. Decided to adopt Quattor after tests, replacing their hand-crafted PXE/Kickstart scheme for system installation. After the changes over the summer, they have now frozen all service updates ahead of LHC data taking.

PDSF: Minor upgrades to PDSF over the summer with the addition of 90 quad-core Dells and 130TB of SATA storage, a mixture of Fibrechannel and Serial-attached SCSI connected. Rolling out SL5 and maintaining use of GPFS on all file systems. Investigating Nehalem- and Istanbul-based systems, no conclusion yet but more effort appears to be going into the former. Also investigating the use of Fuse for ATLAS users mainly but unsure how to allow general use because many requests imply privileged operations. Moving to xCAT (originally from IBM) for node installation, moving away from Kickstart; should allow image-based node installations and management.

NIKHEF: Windows file backup was done via NFS to TSM and Windows ACLs were lost until a recent update to TSM. Also they required to split the volumes into active and archive to be able to complete one backup cycle before the next was due to start. Problems deploying Symantic AntiVirus version 11. No work yet on Windows 7. Still running a little SLC3 but most experiments are moving to SLC5. Added a Dell 16 blade server for virtualisation work, using XEN as the base for now. Their new data centre, built inside a room previously used as a library, is now operational with 47 racks, 400KW of power. Teething troubles included a leaking (in fact open) valve which caused a major flood and many alarms; and a UPS circuit failure. Migrations of the 15 existing Tier 1 and 5 Nikhef racks took some 2 weeks in total. Added 176 quad-core HP blade servers and more disc servers.

LIP: site report from the host of the next HEPiX meeting in Lisbon. Involved with a variety of fields from HEP to astro-physics to medical physics. Small computer support team of 8 people. Main site is in Lisbon with a small site in Coimbra, both attached to a national grid. Aggregated across the 2 sites, some 600 cores and 350TB of disc space compromise their Tier 2 site, based on SGE for batch and Lustre file services. XEN is used to provide redundancy on core network and grid services. Other resources sit on the national grid. They consider Lustre as a concern – always afraid to upgrade it. Another concern is whether or not to stripe the recently-announced ATLASHOTDISK. Despite being a small team, they list an impressive and ambitious range of activities and initiatives in which they are involved.

INFN Tier 1: major upgrade in Bologna this year, tripling the available rack space as well as greatly expanding their chiller and UPS capacity. New farm tender just adjudicated for 118 twin Nehalem systems. More systems as well as more disc and tape are planned and they expect to achieve their WLCG 2010 pledged resource level by May 2010. A new accounting scheme allows them to separate tier 1 work from that of tier 2. Investigating GPU computing based on an NVIDIA chip; target apps include bio-informatics and physics. Factors of 10 improvement have already been seen and these tests will continue and extend to other areas. Still find CASTOR cumbersome to manage (too many db interventions needed and perceived lack of management tools) and CMS have migrated to StoRM for D0T1 data. Adding YAMSS (Yet Another Mass Storage System) – a scripting and configuration layer to interface to GPFS and TSM. It can work stand-alone or be driven by StoRM.

TRIUMF: just celebrated its 40th anniversary by adding 3 new member universities and confirming the funding of the SRF e-linac project. The Canadian Westgrid has been expanded with more servers and Lustre disc space. The TRIUMF Computing group was externally-reviewed and was well assessed, gaining authorisation for additional staff while being encouraged to make more use of commercial vendors for commodity computing and to explore alternatives for in-house management information services. The General Scientific Computing service was advised to reduce the scope of their activities and to charge the experiments for DAQ development and support. [Despite this former recommendation, another part of the group has developed their own new RPM packaging scheme.] They added a SUN storage solution based on ZFS in the spring and they have upgraded their ATLAS Tier 1 systems. They believe the ATLAS centre will need a new computer room by 2012 and a new building is in the next 5 year proposed plan although not yet approved. A piece of HSM code has been developed to re-order and optimise tape requests. Finally photos were shown of a flash flood caused by construction outside the site but luckily there was no damage to equipment inside the buildings.

INFN: lost 100 staff positions and squeezing funds. They are creating the IGI within the EGI framework. They have a pilot project on an Authorisation-Authentication Infrastructure between various INFN sites with goal of SSO across all of INFN by the end of 2010. The IPv6 working group continues its studies; they have an IPv6 addressing plan and they are ready to move to IPv6 “when necessary”. On Batch, 73% of batch slots in INFN are LSF and 24% PBS. Most sites use GPFS for their distributed file system, one site uses Lustre. For remote access to data, dCache is most often found on Tier 2 sites. The GARR upgrade planned for 2009 is delayed by budget cuts and will not now be complete until September 2010, at which time all Tier 2 sites will have dedicated 10Gb links to the Tier 1.

SLAC: LCLS (Linac Coherent Light Source) had a spectacular startup with 94% efficiency and the Fermi Gamma Ray Space Telescope is in orbit collecting data as expected, 15B events so far, 225M of serious interest. New CIO, Donald Lemma and Randy Melen is now his Deputy. The computing services have been renamed to a Division. New HPSS servers and migrating to a new version. Moving from RHEL 3 and most of batch RHEL 4 nodes to RHEL 5. Starting now to move to Exchange 2007. Business IT is migrating to a mixture of new servers including Solaris, RHEL 5 and

Windows XP. Each of the physics experiments is busy upgrading their farms and KIPAC is heavily into GPU use. SLAC is participating in the SciDB collaboration, an open source DB for scientific research. Seeking ARRA funds.

NDGF: ARC has started to use cache pre-staging so they can once again perform cache-aware job brokering. They have mostly moved the caching service from Dell servers to HP blade servers. The HPC2N machine room update is complete, including "district cooling" with heat recycling on a city level; power costs 0.08Euro per kWh, cooling costs 0.025 Euros. Making growing use of flash drives, for example for the root file system; so far incident free. It is not fast but good enough for /root and it frees disc slots.

Storage

OpenAFS: Jeffery Altman, an OpenAFS gatekeeper, updated us on OpenAFS. With a fellow gatekeeper, he has created a company (YFS, Your File System) to develop OpenAFS via dedicated funding initiatives, notably US Government grants. Goals for YFS are to boost server scalability, better networking, including IPv6 and TCP) and network traffic optimisation. In the process of this work, OpenAFS received some significant spin off and this shows up in improved AFS performance; other planned improvements can be found on the OpenAFS web pages, for example use of X509 certificates for easing the use of AFS on grids. In conclusion he begged the audience to trust him despite the need to create a for-profit software company and he listed desirable features whose development is still unfunded.
[These new features and improvements were presented in Rome last month and Rainer Toebbicke has confirmed CERN's interest in some of them. CASPUR is coordinating a joint European contribution (including contributions from CASPUR, DESY, GSI, IN2P3 and RAL) to this development in which it is proposed that CERN's contribution is €25K. We are working to make this concrete.

Experience with Storm and Lustre on an ATLAS Tier 2 site (QMUL): these are based on benchmarks using Iozone for file system performance and the ATLAS Hammercloud for the performance of user analysis jobs. The detailed results are shown in graphs in his talk on the web. Among other things they should Lustre relatively poor when handling large file sizes. The Hammercloud tests show different results at different sites in the UK, Glasgow being the best in terms of CPU time to wall clock time[1]. Comparing WMS and Panda jobs shows the WMS ones are more efficient, because of no checksum? He has ideas on how usage of the file protocol can be improved and had passed these to the developers. They decided not to use Lustre striping (fear of the effect of losing a disc). Storm should benefit from Lustre performance but only if the file protocol is used and he wondered if and when lcg-utils would do that. His conclusion is that the Storm/Lustre scheme gives good performance as long as the file protocol is used properly by the tools and that there is some tuning still to do.

Lustre at GSI: as previously exposed at HEPiX, Lustre is heavily used in GSI for online storage. In tests with ALICE data, they reached 50 Gb/s aggregate network throughput using 2000 worker cores on an 120Gb/s network. Next year they expect to double the number of cores so they expect to put strain on their network. In a year of operation, disc failures are running at about 1%, all fully recovered and not seen by the users. But one new file server failed and this was seen, with a 1 hour downtime for server replacement and disc recovery. Three RAID controller failures, one caused a 40 minute break and one damaged a few files. On the software side, Lustre seems much more stable than when he last reported. In a practical test, he was able to expand the cluster from 0.7 to 1PB without the users noticing any service interruption. There are an impressive number of audit reports available but a risk of loss of performance if too many users request reports in parallel; so they use a package of auditing tools (the Robinhood

[1] Had to include that quote.

tools) produced by the CEA. On the downside, Lustre is a complex system and is vulnerable to network communication problems. Plus he reported that quotas are not yet working in V1.6 and he says the situation is not changed in V1.8. In summary however GSI is quite pleased with Lustre and is more and more relying on it with expansion plans in terms of capacity and network bandwidth.
Q: GSI has plans for a 2.5PB file system but where is the limit? A: the US military have clusters up to 10PB at least it is believed, although on high quality configurations. Limits can also be pushed using larger capacity discs, 2TB for example but then rebuild times go up of course.

Hadoop File system at FZK: Hadoop is a distributed file system designed to run on commodity clusters and making use of their local hard drives. This was a brief report on tests being performed at FZK under the aegis of the HEPiX working group, comparing Hadoop to Lustre. Lustre consistently performed slightly better except in the maximum load tests (8 jobs on an 8 core node) and the details can be seen in the slides. But if other considerations are taken into account, such as complexity and ease of deployment and administration, the choice becomes more open. Of course Hadoop would need a lot of effort to be considered for production Tier 1 storage but perhaps is more suited to Tier 2 sites, as evidenced by the growing number of OSG tier 2 sites adopting it in the US. FZK will continue studying it in future releases.

R & D at CERN: Helge started by asking why CERN is looking at iSCSI. Possible replacement of expensive FibreChannel is one answer; possible replacement for small disc servers is another; use in cases where absolute performance is not the most important criterion is a third. The storage target has to be either a storage appliance able to respond to block-level requests, some of these are available on the market; or storage in a box, adding a software stack (such as tgtd in Redhat 5 although this does not scale well because it is single threaded) to provide the block-level interface; CERN uses the IET stack although it does not run on the RHEL 5 kernel yet. He listed some configurations used in these tests. His graphs show that the IET solution compares very well with an F/C SAN. Some iSCSI/tgtd configurations are already running as "Storage in a Box" servers in production for a couple of grid services and other configurations are being investigated for other services. Properly configured, storage appliances can show remarkable resilience to, for example, random network failures. Helge compared storage appliances to home-made solutions.
Moving on to the Lustre evaluation currently underway, he explained CERN's interest in it (possibility to reduce the variety of file systems supported) and what we are looking at in detail (installation, backup, strong authentication and fault tolerance). What we have found is that

- strong authentication is still missing and apparently slipping
- strong client/server coupling on recovery
- missing support for life cycle management
- moving targets compared to the published roadmap.

HPSS: there followed 2 talks on HPSS, migration to a new version at IN2P3 and experience at LBNL, but I won't comment on them given our lack of interest in this. The slides are on the web for those who care.

Optimising Tape Access: this is important because the LHC experiments use a tape backend with a dCache front end and they want around 400MB/sec read speed and reprocessing runs in particular will expect to access thousand files at a time. This talk concentrated on HPSS tape storage (used at IN2P3). HPSS treats requests FIFO and dCache restores files from HPSS blindly so there is no chance to group requests per mounted tape, resulting in chaotic tape access. One solution is manual pre-staging but this is obviously heavy. BNL built an automatic tool which analyses the HPSS metadata and orders requests per file and IN2P3 have tailored this to their site, for example making it database-driven and adding metrics for monitoring and accounting. The tool has been developed as a stand-alone tool and although the work described was a dCache interface, it could be interfaced to another HSM.

PROOF on NFS 4.1/pNFS: first results from trials at DESY, a spin-off of the initial tests of the HEPiX working group. Why look at NFS again? It seems that NFS 4 has interesting new features such as strong authentication and NFS 4.1

should offer a parallel NFS. The pNFS component separates out the metadata. Given the similar architecture, the speaker claims, of pNFS and dCache, he feels NFS 4.1 is worth another look. The only major protocol difference is in file close and this was fixed by adding some code to the NFS MDS (metadata server) to cut the client - server connection after file close. They believe this is safe but it was questioned in the Q&A (some small risk data is cached on the wrong side of the cut). PROOF was chosen for a proof of concept demonstration and they are rather pleased with the results, no code changes were needed, performance is slightly lacking but no serious tuning was attempted, and the setup was straightforward. He stated that this can be considered another front end to dCache, in particular for legacy applications. Further tests will be done against xrootd and AFS.

Virtualisation

Batch Virtualisation at CERN: Ulrich presented the current work in this area in FIO. On the question of why do this, it offers customisation of images for specific tasks, a better framework for the pilot jobs desired by some experiments and it allows to hide changes in operating system, for example allowing to run old code for a longer time. He listed some operational and infrastructure requirements. They have chosen to concentrate, at least initially, on XEN because of its stability but have recently added support for KVM. They use "golden nodes" to serve as templates for the creation of images and worker nodes are derived from these. The worker nodes are drained after 24 hours and shutdown when idle. Stage 1, proof of concept, has been completed supporting SLC4 and SLC5 images with 1 CPU and 2GB or RAM. Support of more specialised images will follow. They are now looking at commercial (ISF from the Platform VMO platform) and open source (OpenNebula) solutions for image placing and he listed pros and cons of both solutions. Only minor changes were needed to adapt their model to the Cream CE for use on WLCG. CERN is now setting up the lxcloud cluster with 10 powerful servers as a host for this virtual service but the choices for the VM Kiosk (image repository and distribution) are still under discussion. He closed with some possible ideas for future phases (use of images outside of CERN, experiment-specific images, cloud-like operation) and he listed some challenges which would need to be addressed to implement these goals.
Q: why run down after 24 hours? A: VM jobs not managed and there may be Quattor scalability issues in adding these. Also security, avoid gradual build-up of jobs in a bad state.
Q: performance of VM jobs compared to jobs on bare metal? A: not measured yet. [But Fermilab reported no performance degradation under XEN and only 2% under KVM.]

Virtualisation on Fermigrid: considerable experience and now run over 150 virtualised servers across their different high-availability services as well as some non-HA ones. XEN is their primary platform but they are looking at KVM (mainly because Redhat is promoting it). The goal of this is to achieve five 9s availability (99.999%). They nearly got this in the first half of 2008 (99.9969% - 10 mins of downtime in 7 months) but more recently some major hardware failures have forced this down to 99.<something>% for different services. They have made a proposal for funding for using virtualisation for clouds.

Magellan Cloud Computing Project at NERSC: Funded under ARRA (American Recovery and Reinvestment Act), the goal is to determine the role of public and private cloud computing in mid-range workloads. What is a cloud? Something opaque, with on demand capacity, it is easy to acquire, offers pay-as-you-go, has low up-front costs, high availability, etc. Looking at 500-way parallel jobs with performance around 5 TFlops. Areas of investigation include

- ease of use requirements – how to build a science-friendly cloud?
- different delivery models - what is the most appropriate business model – fund the cloud or fund the application and let the users chose the cloud(s)

- computational efficiency of the applications - some science apps already run on clouds but how does this compare to current DoE workloads. They will experiment by initially applying their own workload to clouds (or is the other way round?) and broaden this out to other sites
- cost efficiency analysis (TCO) - how cost effective to use a commercial cloud or should the DoE build their own cloud and how do both of these compare to today's grids. How to measure success or cost effectiveness
- applicability of cloud computing APIs to science apps - how to integrate cloud computing APIs with data storage facilities
- the practicality of multi-site clouds – hence the reason to involve NERSC and ANL.

They will explore

- hardware as a service (HaaS): access to batch, interactive and virtual clusters with on-demand division of the workload.
- data as a service (DaaS): shared storage with enough bandwidth, is this practical? Can an application use this transparently?
- software as a service (SaaS)

First purchase is a 61Tflop cluster of 720 computer nodes (5,760 cores) with 6GB memory and 1TB discs, based on the IBM iDataplex chassis, linked by Infiniband. They will also look at flash storage with 10TB deployed in the NGF cluster and later 16TB into a scalable unit in the main cluster as well as 2TB into HPSS to accelerate access to HPSS metadata. The timescale is to have a preliminary report to the DoE next summer and a final report 9-12 months later. Expect to involve 5-6 FTEs.

Hyper-V at CERN: Juraj Sucik presented the virtualisation work in IS group. Hyper-V supports 4 cores per VM, up to 32GB of memory, snapshots, failover and has a scriptable interface and it comes with an advanced VM manager. VMs created by Hyper-V using the interface developed by IS behave like normal physical servers – registered in LANDB, backed up by TSM, etc. Enhancements in Hyper-V 2.0 include live migration between hosts in a cluster (useful for maintenance or load balancing), network optimisation, improved Linux support and cluster shared volumes. Live migration is particularly interesting, no dropped network connections, no perceived loss of service. He compared the target audience and features of the batch virtual service described the previous day by Ulrich and the Hyper-V service (see his slides). Hyper-V is heavily used for VOBox servers (222 instances today) as well as a wide variety of infrastructure servers (web, licence, mail, etc). Future plans include integration with LEAF, Quattor, etc, better integration with SLC and eventually to offer a virtual desktop.

On-Demand Virtualisation and Grid/Cloud Integration: an INFN talk, this started by explaining the why and how of virtualising a farm cluster. From this arose the on-demand virtualisation service although images are not created by users directly but only by the resource provider from the user's declared requirements. This mechanism has been extended to standard grid jobs, allowing to offer cloud-like services. Worker Node on Demand (WNoD) is in production at the INFN Tier 1, with several hundred VM servers in use as the talk was given. SLC5 is only offered via WNoD and CMS are already heavy users with the other LHC experiments expected to follow soon. They are writing an API to interface to Amazon EC2 and to the Open Cloud Computing Interface (OCCI). As far as the chosen VM, they started with XEN but moved to KVM for the WNoD implementation. They are investigating some security aspects, how to bill for cloud usage and support for multi-core jobs.

A Vision for Virtualisation in WLCG: Tony asked how the experiments can be allowed to choose their preferred environment at the same time as allowing the sites to control and be able to trace resource usage. We need to use

success stories to build trust and he proposed a 5 step approach, of which perhaps the last two could be considered kite-flying today.

1. Let users select their preferred virtual images created at the sites, similar to the WNoD scheme just presented. Users need a way to express their required resources.
2. Distribution of virtual images between sites; the image must be minimalist to reassure local sites and be able to hook into local monitoring and batch scheduling.
3. Distributed virtual images must include experimental software environment. This could remove the need for today's pilot jobs. Is the current CernVM project (joint IT/PH collaboration) a suitable model? How scalable is that model?
4. Distributed virtual images to include a client to connect directly to experiment pilot job framework such as PanDa or Dirac. This should allow experiments to use their own central scheduling scheme rather than send jobs to a site to be scheduled by that site.
5. Replace experiment pilot job frameworks by commercial or public domain schedulers. For example there could be a virtual LSF cluster for ATLAS, a virtual SGE cluster for CMS, and so on.

An interesting discussion ensued but the HEPiX audience consists almost entirely of system administrators rather than experiment representatives so the broad agreement could have been predicted. Sharing experience at HEPiX can only encourage sites to work towards being able to offer this.

Benchmarking

HEP-SPEC06 on latest processors: An INFN presentation of the now-standard HEP benchmarks on the latest generation of processors. Tests on Nehalem are complicated by the configuration choice – 32/64 bit, hyper-threading on/off, turbo mode on/off (some cores speed up if other cores inactive) and memory sizes and he showed these individual comparisons. The results are too detailed to cover here but the reader is referred to the slides on the web. The graphs also show some data collected on the latest Opteron 2427 and Istanbul chips. The speaker plans more tests with real experiment code and to check the effects of the chip settings on power consumption.

Benchmarking at CERN: Ulrich presented similar tests done at CERN. As well as testing performance improvement from multi-threading alone and switching turbo mode on alone, he measured the effect of having both turned on and found that together they did not quite add up to the separate effects of each alone. Scaling tests showed that performance rises linearly to 4 processes but then falls off beyond that but over-committing the number of processes per core can have some advantage, even with multi-threading off. Comparisons of tests run on a set of 4 systems appeared to show some right-side bias (systems on the right consistently performed slightly better) and this is not understood – statistics, bad cabling, poorly tuned system? Finally Ulrich reported some power measurements he had performed. Using only one power supply instead of two seems to be advantageous for power saving but he saw variations of up to 4% between apparent and effective power measurements. Also, he saw notably different power factors between the different units of a given model.

AMD Roadmap: An invited talk from the Opteron product manager. He started by describing today's Istanbul 6 core processor in some detail. New socket platform next year – Maranello and will be populated by Magny-Cours processors (no longer using the same Sao Paolo). Magny-Cours will be offered in 8 and 12 core models. Believes some Linux kernels do not take sufficient advantage of AMD's NUMA features. Also looking at harnessing the multi-teraflop computational power of GPUs, for example adding programmability. They are betting on openCL as the

programming language for GPUs. They are also working on a new GPU targeted at the HPC market where the GPU would drive one or more displays. [A lot of what he said can be found on blogs.amd.com.]

Benchmarking at LIP: tests performed in recent procurement exercises. Described this exercise in great detail. Only memorable snippet was a disagreement between the speaker and Helge on the mathematics being used in the HEP-SPEC06 benchmark. They were seen going into a corner at the following break.

Monitoring

WLCG Open Source Solutions: presented by Wojciech Lapka. The SAM framework has proved its usefulness but better tools are requested by site administrators and the future European grid organisation pushes us towards a more devolved solution. Hence the move to include popular open source tools. Nagios was chosen and some 50 sites have already installed the egee variation of it. An integral part of the monitoring service is the messaging service, and we need something able to scale to the LCG grid, for example a service based on an SOA architecture. Fuse, based on Apache ActiveMQ, was chosen and deployment has begun. Performance tests performed by openlab and the EGEE monitoring team are encouraging. He then showed how SAM has evolved to use Nagios and the messaging service.

Nagios at NERSC: each service (PDSF, network, etc) has a dedicated Nagios node. Security is paramount so the nodes are firewalled, have (almost) no local accounts and offer https access only although outside access to data is allowed by ssh tunnels. They have developed local Nagios plug-ins, for example Cray cabinet temperature checks. They also built graphical displays for Nagios results.

Monitoring tapes and media at IN2P3: they have 3 Sun/STK SL8500 silos with a total of 90 drives, 30,000 slots and 25,000 tapes. And a fourth, larger silo is coming with another 40 drives. They need a system to identify failures, I/O errors and poor performance. They tested 2 commercial products for a month each – RVA from Crossroads and StorSentry from Hi-Stor. Criteria included raising alarms based on sensed data, make recommendations based on the alarm; data should be stored in a database; it should not be disruptive. The selection was for StorSentry (but he did not explain in detail why) and some screen shots were presented. The tool sends daily reports and recommendations by mail. It includes tape pool management tools with thresholds for maximum mounts and error counts. Sun accepts its recommendations and performs maintenance according to its recommendations. Hi-Stor gives good support although the user interface could be improved.
Q: Tony and I asked similar questions: cost basis and comparison details. To the first he said licence cost is related to drive and media usage. On comparison he was obviously reluctant to say anything in public but he did say the RVA product was notably simpler, offering alarms but no recommendations for example.

Quattor at RAL: current variety of system installation tools in use does not scale and there were issues of consistency between installations. A review earlier this year narrowed down to a choice between Quattor or puppet. Quattor had been rejected by the RAL Tier 1 in the past but as a newcomer to the RAL team the speaker felt able to revisit the question and he chose it because (a) it was built for exactly the job in hand; (b) its automation; and (c) support for gLite. However, he reported a steep learning curve and found it complicated to set up. Implementation started over the summer with the arrival of the latest acquisitions. Some teething problems getting used to it but very appreciative of the Quattor support community, in particular Michel Jouvin. Summary: high entry costs but real benefits.

Nagios at IN2P3: Nagios was felt to satisfy their needs although it did need quite some amount of configurations. Rest of talk was mostly a description of Nagios and how it is installed and configured at Lyon.

Lustre Monitoring at NERSC: set of plug-ins to open source product Cerebro, a spin-off of Ganglia. Data received by modules inside Cerebro and stored in a MySQL database. Measures the usual parameters. And he has produced scripts to do the usual analysis.

Security and Networking

Cyber Security Update: presented by Sebastian Lopienski presented the HEPiX-traditional scary list of vulnerabilities, attack vectors, malware, etc. in products ranging from Adobe, Windows products, MacOS, browsers, Oracle, etc. The main attack vector remains web surfing. Attackers are getting more sophisticated, using hot news to attract hits on to sites under their control or to come within range of drive-by downloads. E-mail attachments with hidden exploits are also "popular". Always beware of scareware – typically what looks like a window encouraging the user to buy a product to clean up the PC. On the other hand, happily, the conficker worm appears not to have lived up to our fears for it, at least so far. Not yet much targeting of mobile devices but Linux is not exempt from attacks. CERN has a campaign of user awareness and security training, including his own course for coders. And there has been some encouragements for web masters to review what material is visible from outside CERN.

Grid Security Update: another HEPiX tradition is this talk from Dave Kelsey. No attack directly aimed at grid middleware so far but "normal" security incidents are looked at seriously and sites are strongly encouraged to keep their systems properly patched, with some sites eventually suspended for not doing so. There has been significant progress in dealings with NRENs as well as between different grids. 170 possible vulnerabilities have been submitted to the Grid Security Vulnerability Group (although few in the past 12 months) of which many have been closed already. On the policy side, a new policy framework is being developed for EGI. Some NGIs may have their own policies but the framework should be such that individual policies should not be incompatible with the framework. Turning to the International Grid Trust Federation, he described some of the work being undertaken in Europe on making identity vetting and certification operations easier, cheaper and more scalable.

Web Application Security: another talk by Sebastian, this one concentrating on web applications. Web applications are widely and publically available and often easy targets for hackers who then can pose a variety of threats ranging from web defacement to data loss to unauthorised intrusion behind a firewall. At least 10% of CERN's 8800 centrally-hosted web sites include web applications. Of the various ways to avoid such threats, CERN has decided to concentrate on vulnerability scanning. There are a large number of commercial and open source tools to assist in such scanning and he presented the results of testing a number of them. In the end we chose a combination of open source tools, Wapiti and W3AF.

WLCG Security: next up was Maarten Litmaath. First area of concerns are proxies, for example the MyProxy server is currently a single point of failure and we find that different grid services treat proxies differently and he gave some examples. But perhaps the forthcoming ARGUS, now in certification, will resolve a number of these discrepancies. The next area is pilot jobs: here the hope is that glexec will permit the pilot job to run the payload under the user's own account but we need to get practical experience of this to be sure. Virtual machines have reduced opportunities for interference between jobs but VMs still share some common software and/or services. The various data storage services have their individual data security models. Vulnerabilities are also present in the information system. On a

brighter note, he believes the accounting is pretty secure but it is a rare clearing in a vast forest (see his slides for the analogy) and a lot of work remains to reach the path out of the forest.

ESnet, networking for Science: presentation by Joe Burrescia, ESnet general manager. ESnet is built for large scientific traffic flows. Lots of redundancy, runs mostly at 5 9s reliability to most sites. He showed the ESnet topology, the Metro area rings linking to the main HEP sites and also the links to other US and international nets. Growth is a factor of 10 in the past 10 years, crossing the 5PB/month level in May this year. He described OSCARS, a tool to perform advanced secure circuit reservation. In order to set up end-to-end circuits across multiple domains without violating security or allocation management policy of any of the domains, the process of setting up end-to-end circuits is handled by the Inter-Domain Control Protocol (IDCP) running between domains. There are currently 26 long-term production virtual circuits, of which HEP is a major user. He ended by describing how they plan to use a $62M ARRA grant to build a prototype 100Gb network as well as a network testbed facility for science and industry.

Network Performance Tuning: presentation from Brian Tierney, once a collaborator in the DataGrid project, now working in ESnet. To get good TCP performance, the TCP window has to be full. So TCP buffer size must be tuned, it is usually far too small. Over the past 5 years, socket buffer autotuning features in Linux and other O/S (including Vista and MacOS) have notably improved matters. Having set max buffer size, the next step to improve performance is to employ (a few, not too many) parallel streams but you also need some TCP congestion control. He then described perfSONAR – tuning for high performance networking. Typically getting maximum performance is a multi-domain problem, local testing will not find all problems. He listed a number of technical glitches which are hard to uncover but which affect performance. perfSONAR is an open web services framework to collect, manage ad share network measurements and it is already widely deployed (130 locations to date) including at major HEP sites (e.g. US ATLAS already and CMS are starting to use it). As well as measurement points and tools, there are specific troubleshooting tools and a topology service. He showed some real-life examples of performance improvement resulting from the use of perfSONAR.

Computing Centres

Intel HPC environment for silicon design: presentation by an Intel senior principal engineer. His team are under constant pressure to cut development costs. Their design systems consist of 64,000 servers running SuSE Linux mostly blade servers. Their newer data centres are typically 6000 sq.feet split into modules each capable of providing 500W per sq.foot. for a total of 3MW of power. They install some 200 racks or cabinets per module with 15-22KW per rack (enough for 48-64 blades). A few of the most recent data centres offer 30KW per rack (84 blades). The workload is made up during tape-in of perhaps 100,000 simulation jobs per chip design each week and during tape-out of 16-23,000 jobs per complex silicon layer. The simulation jobs are CPU and physical-memory bound while the second set are constrained by CPU, networking and storage. They broke down the HPC stack into components such as hardware, network, computer and infrastructure servers, batch scheme, clustering, systems management and analysed how to optimise each for their growing workload. The results, despite needing far greater power in the design cycle, they achieved dramatically improved performance. Of course part of this is benefiting from the newest chips which they are designing – e.g. using the power and features (e.g. hyper-threading) of the latest Nehalem - to design the next version. Silicon Valley power costs some 6.5 cents per KW which helps to keep their data centre costs low (8% of total computing costs). In passing he noted they rely on Panasas storage system. In New Mexico and Oregon when the weather is cold enough, they simply open the doors of their computing centres and reduce their cooling power needs.

Data Centre Expansion at BNL: needed for US ATLAS and RHIC. Currently housing 165 racks of equipment with an average of 650KW (60% of UPS capacity) and peaks of 790KW.Cooling capacity is around 1MW. Discussions on expansion started in 2007 because they could see limits in both space and power capacity in the near future. Short

term it was agreed to renovate an existing 2000 Sq.feet building while considering a new 6600 sq.feet building for the long term. The renovation stage took some $600K and 6 months. The new building cost $5M and 18 months and is just now being commissioned. The new building will allow to double capacity and there is a long term plan (2018) for a very large new 25,000 sq.feet facility. They are still unable to fund a second UPS for the new building, not a diesel generator to support the UPS. They believe the advances in technology will allow them to survive until a new centre can be approved, funded and be built by 2018. But they will need to add more cooling (rack top cooling units and rear door heat exchangers)

Update from CERN: I presented the latest details of our plans for hosting some systems in a commercial service in downtown Geneva and the plans to acquire and install on the Prevessin site a number of containers to house future server purchases.

ITIL

ITIL in KIT/FZK: Their work is based on ITIL V2 for historical reasons. The talk concentrated on the Service Desk and on the Incident and Configuration Management processes. The CM should produce a solid base for many other ITIL processes and he showed what information needs to be stored into the CM database. This db may need to be interfaced to other dbs, such as inventory dbs or monitoring dbs. An "incident" refers to an event which leads to a degradation of service. IM has to take into consideration the "service class level" of the different target services depending on their defined priority, hours of guaranteed service and incident response time. The IM is linked to the Service Desk, operating only in prime time at KIT, to receive its alarms, to assess and categorise alarms. The speaker presented various roles defined in ITIL and what they represent, such as Local System Experts; or the On-Call Engineers. There is a Local Monitoring System (Nagios in this case) which informs OCEs of problems via SMS (during prime hours) or mail outside these times – a workflow was shown. For both CM and IM he showed how internal and external views differ and some tools being used. In summary he noted that it had been hard to convince all his colleagues but worthwhile pursuing. He is convinced the benefits are apparent. The Service Desk in particular becomes more efficient. It has taken more than 3-4 years to get to the current level.

FNAL Software Acceptance Process: in fact this work pre-dates the ITIL interest at FNAL. The framework was developed in early 2008 in collaboration with the various software providers. Can be found at http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=2684. Some attitudes required to be changed, for example in the method and timing of the provision of documentation; and root access is in general not permitted. There are separate independent systems for software developers. Installation and testing, on dedicated integration services, and eventual installation on the production servers is performed by Fermigrid personnel. In October 2008, Fermilab announced a 5 phase plan to move to ITIL with the goal of gaining ISO 20,000 certification for the Computing Division and the scheme just describes meshes well with this, only minor modifications being required.

ITIL in IT: Tony Cass presented the current status of our ITIL efforts. He listed the high level goals as improving the coherence of communications with clients as well the coherence of service management. We are following a pragmatic approach, no big-bang, and aiming to build commonality across our services. He listed the issues currently being addressed such as service organisation and granularity; our customers and users; requests versus incidents; change management. He showed an extract of our Service Catalogue and listed the next steps – the new service meeting with our users, the new Service Desk implementation and a review of our change management process.

Desktop Support

Windows 7 at CERN: Juraj described IS group's experience since some months. Has been extensively tested in IS and with some engineering applications. He listed some issues encountered, mostly now fixed in the production release last week. With the experience gained with Windows Vista, the transition to Windows 7 should be relatively smooth. Hardware requirements are the same as for Vista but could even be less in memory if the AERO user interface is not used. IS expects a certain impatience to move to Windows 64 to take advantage of its 64 bit support. Finally he presented the current proposed roadmap for support for the different Windows flavours.

Linux Desktop Management at GSI: based on a scheme to boot via the network, not a new concept but a still-valid one. The desktops are by no means thin but nevertheless this network boot is used for some 400 desktops. Debian is the Linux of choice. They use PXE for the initial boot and a read-only O/S image is shared by the clients. If the root part of the file system is read-only, some writeable space must be provided elsewhere and this is achieved by including tmpfs and unionfs – the latter a stackable copy-on-write file system – both linked to NFS. Their Debian installation scheme has an option to create the client image. Configuration management is via cfengine. The advantages are good security (read-only kernel, system libraries and utilities), tight control of client desktops and fast updates and patches. On the other hand, the scheme is completely network-dependent and there have been issues with the various file systems and how they interwork.

Desktop Support Survey: I presented a short survey of different strategies used in the various HEP sites, what support is offered for visitor laptops, if security updates are forced, if and how Macs are supported.

Miscellaneous

Scientific Linux Update: Troy Dawson, one of the principal authors of SL presented this update. SL has currently 29 public mirrors and they know of a total of 178 mirroring sites. Number of known nodes is stabilising at slightly under 40,000. List of changes applied to the latest SL 4.8 release. Currently building 5.4, due for release next week. Turning to infrastructure, they plan to move soon (in about a month) to new faster servers and moving the web site from https to pure http has resulted in a notable performance improvement. Working on creating a Bugzilla, hosted offsite for free by a company, Infiscale, which uses it and wants to put some value back into the project. Support for SL3 will stop on 10 October 2010, at around 10:10 in the morning. Guestimate for 4.9 is Feb 2010, for 5.5 is May dependent on Redhat release of 5.5 2 months earlier. He believes the Redhat beta of 6 should be released before the end of this year and first official release around May/June 2010. He and Connie (Sieh, the other principal author) estimate it could take up to 6 months, perhaps less, to produce SL6.
Discussion:

- last Redhat 4 update is 4Q09 so they propose to call SL4.9 the SL4 legacy version, dropping support for previous SL4 versions in Oct 2010 and supporting SL4.9 until Feb 2012.
- Work with CentOS for SL6? Less clear than last time this was raised and some people even object to the idea.
- Should Lustre be included in the release? Little support.
- Include the Globus toolkit and utilities? No support, existing Globus distribution scheme is perfectly tuned for this.
- Open a wiki, offer to host by Infiscale? Mixed response. Would need a neutral moderator.

WLCG Technical Forum: Maarten presented the ideas behind the WLCG Technical Forum, including discussions on possible improvements to different WLCG services, defining longer-term needs, considering common solutions. It

needs to represent all the stakeholders, bringing in experts as needed, but not taking decisions which should be passed upwards (?) to the GDB and MB. The most interesting topics will be investigated by small working groups who should produce short reports. Extensive use should be made of wiki and mail lists.

Software Version Control at CERN: Giacomo Tenaglia presented the latest status of the migration from CVS to SVN. SVN has by now attracted some 110 projects and runs today at around 33,000 commits per month. On the other hand, only about one third of the 100 original CVS projects have migrated to SVN. So the ending of CVS goes backwards into at least 2011. They feel unable to force the migration so "may have to run CVS forever". There are some problems of version compatibility because the SVN server is version 1.5 but SLC5 has only version 1.4 clients; they work but miss some 1.5 features.

Alan Silverman
6 November 2009