# HEPiX Fall 2013: October 28th to November 1st, 2013

## *University of Michigan in Ann Arbor, MI, USA*

# Trip report


## *Helge Meinhard / CERN*


## General

For the first time, a North-american HEPiX workshop was organised by a university in the United States hosting a large WLCG Tier-2. The University of Michigan in Ann Arbor, (major) part of the ATLAS Great Lakes Tier-2 centre for WLCG, took the challenge with a small team lead by Shawn McKee and assisted by Benjeman Meekhof. A total of 116 participants registered – I am pretty sure that this is a record for a North-american meeting, even outperforming the very well attended 20th anniversary in October 2011 in Vancouver – and probably one of the highest attendances at HEPiX ever. Fermilab was a major contributor this time to the attendance with a total of 12 participants! Equally remarkable and a novelty for HEPiX: a total of 27 attendees came from eight North-american universities, mostly Tier-2 sites of WLCG, a community that the HEPiX board had long tried to address in vain.

Yet at the beginning of October, the U.S. Government shutdown caused a major doubt on whether the registrants from U.S. National labs would be able to attend. Other conferences such as CHEP, two weeks before HEPiX, had been significantly hit, and for HEPiX more than 20 attendees were at risk. Fortunately, a good two weeks before the workshop, the shutdown ended, and most labs could sort things out – in total only three registered participants were unable to travel, two of which gave their scheduled presentation remotely – a possibility that the team at UMich provided on short notice.

A total of 66 abstracts were submitted, again a very high level for a North-american meeting – that high that the initially allocated slots of 30 minutes (except for site reports) needed to be cut back to 25 minutes in order to accommodate all talks.

The meeting took place in a very nice auditorium just sufficiently large to fit the audience. Network services were working perfectly, and as a CERN participant with painful experiences in the past, I was delighted to just open the laptop, log into Eduroam and be set.

Social events included a copious welcome reception on Monday night at the university, and a dinner on Wednesday in the Henry Ford Museum some 50 km away from Ann Arbor. In the museum, participants strolled around and could admire the numerous cars on display (including the one in which John F. Kennedy was assassinated in 1963) as well as other transport vehicles including a huge steam locomotive and a number of aircrafts.

As mentioned in the wrap-up presentation, the workshop showed a few clear and interesting trends: identity federations, CEPH, large private clouds reaching full production status, HTCondor, Puppet, and log file analysis.

## *Monday 28 October*

## Welcome address (Homer Neal / U Mich)

Homer Neal was introduced by Sandy Philpott, North-american HEPiX co-chair. Homer acted as

interim president and vice-president for research of the university, is now distinguished professor and head of the physics department as well as vice-president and president-elect of the American Physical Society. Homer explained that the university was founded in 1817, and has been located in Ann Arbor since 1837. Today the university counts 51'000 students and 5'600 faculty members; it runs one of the largest health care complexes. The university has obtained very high ranking results consistently. It's total budget is around $5 B. Physics accomplishments include the invention of the bubble chamber, first ideas for racetrack cyclotrons, the discovery of the proton spin, major participation in the discovery of top and Higgs, and the discovery of the Xi_b and Omega_b hyperons. The university has a strong (about 30 persons) group in ATLAS, of which about 1/3 are based at CERN. Providing access to US students to CERN is an important part of the university's activities. The university has also been intimately involved in developing networking over the past three decades. Homer finished with a warm thanks to the organising committee led by Shawn McKee.

## Workshop logistics (Shawn McKee)

Shawn welcomed all participants, explained networking and the daily schedule. The social event will take place at the Henry Ford museum, buses will be available just after the session, and will return to different key places in town. A group photo will be taken on Tuesday just before lunch, with a fallback on Thursday if the weather is inappropriate on Tuesday. He reminded all speakers to upload their slides to Indico in advance.

## *Site reports*

## INFN-CNAF (Andrea Chierici / CNAF)

Andrea explained that the capacity of the general IP links and the LCGOPN links had been doubled; worker nodes are connected with higher throughput as well. Some 195 kHS06 have been installed, corresponding to 17 k job slots. Migrating the farm to SL6 has resulted in significantly increased performance with older worker nodes. The next tender for worker nodes will take into account TCO, in particular energy consumption. They overhauled the monitoring and accounting, replacing home-made Perl scripts by a solution based on the Graphite open-source solution. They have met some issues with Grid accounting as well as with the WnoDeS (worker nodes on demand) upgrade to SL6. They are investigating GridEngine as a potential replacement of LSF, and are considering Zabbix for monitoring. For CDF, they are involved in long-term data preservation, copying data over from FNAL. For the storage nodes, they will replace Quattor by Cobbler and Puppet. Grid middleware has been upgraded to EMI-3.

## NIKHEF (Paul Kuipers / NIKHEF)

Paul started by announcing personnel changes: Wim Heubers will be replaced by Ronald Starink as computer group leader; his responsibility for the Amsterdam Internet Exchange point will be taken over by Koen Keijer. At about the same time, some more persons from the computer group will retire. NIKHEF has successfully moved to single sign-on based on password authentication. They have moved to Windows 7 under a new, streamlined AD domain; they are considering Windows 8 for laptop computers. Some 100 farm nodes were added to the Linux cluster, running cloud services in a dedicated subnet. For the storage, Gluster is used, which they find flexible and easy to set up and manage, even though the fine-tuning requires more expert knowledge; there are occasional issues with client side stability, load balancing, and full file systems. Network and storage systems have been rejuvenated in the framework of a campus challenge (Juniper router, Hitachi storage server). Concerning networking, they are making smooth progress with IPv6. For mail, they had been using Horde as web interface for mail access, which they are replacing by SoGo because of its

additional support for calendaring, address books and Exchange. Finally, Paul mentioned CHEP that was a significant amount of work resulting in a very successful conference.

## Caltech (Dorian Kcira / Caltech)

Dorian presented the Caltech site report (I think it is the first one ever). They are a Tier-2 for CMS and were indeed the first LHC Tier-2 site, now running 3 PB of Hadoop storage and 2'000 job slots. A total of 60 users, of which 25 are very active, use the facilities for a wide range of physics analyses. They use xrootd, CVMFS, and PerfSONAR, and have replaced Rocks by Puppet. They regularly run high-throughput WAN tests that they report to the Supercomputing conference. For the additional small Tier-3 they are running, GridEngine is being replaced by HTCondor. About the same amount of data get sent as are received (about 300 TB over a year). The data centre was fully renovated, new chillers were installed, and all servers were moved into APC racks, giving the opportunity for a clean, orderly installation. They have just added Supermicro quad servers adding some 29 kHS06, and three 4U 60 drive storage units. Networking-wise they are involved in US LHCNET, LHCONE, Ultralight, PlaNetS, DYNES, OliMPS. They are a major contributor to ANSE (Advanced network services for LHC experiments) in close collaboration with ATLAS and CMS. Dorian finished presenting results from network tests between Salt Lake City and Victoria, BC.

## ATLAS Great Lakes T2 (Benjeman Meekhof, U Michigan)

Benjeman explained that the capacity includes 4616 job slots and 3.5 PB of storage distributed over the University of Michigan in Ann Arbor and Michigan State University in East Lansing. Most Tier-2 services are virtualised via VMware. For installation, they still use Rocks and have recently deployed Rocks 6 that supports SL6 installs. Additional post-build configuration is applied via CFEngine3. They now use Cobbler rather than a home-grown script to generate kickstart files, aiming for pushing as many operations as possible into the configuration stage. BIOS/firmware updates are applied in the post-install phase by RPMs. They took the decision to migrate everything to CFEngine3 after careful evaluation of Puppet//Foreman which they found to have interesting features, but in the end the existing CFEngine expertise and its simplicity made them decide for it. They moved to ZFS on Linux for the /vicepX storage of AFS because of compression, snapshots and data integrity. Apart from file system occupancy issues, their experience so far is very good. They are running dCache 2.2.17. They use iSCSI storage at UM and DAS at MSU, and are working on site resiliency details, addressing the transparent move of services. Networking links are being upgraded, for example UM now has a 100 G connection to Chicago. The inter-site link will grow to 2 x 40 G. They have fully deployed PerfSONAR with thee instances each at UM and MSU. The upgrade to SL6 is ongoing. He finished by mentioning a number of networking projects they are involved in, many of them in common with Caltech.

## BNL (Ofer Rind / BNL)

Ofer started by recalling what the role of RACF is, and then turned to business continuity with respect to the recent government shutdown, referring to it as a man-made business continuity issue. They had a contingency plan for running RACF in degraded mode with 18 FTE, targeting various levels of service. The processor farm was beefed up by 90 + 220 Sandy-Bridge hosts (the 220 hosts for RHIC are equipped with more disk space). The total number of job slots will soon be about 40'000 (15'000 for ATLAS, 26'000 PHENIX/STAR). The upgrade to SL6 is going on, with no major issues identified so far. They use HTCondor as batch scheduler. With the increasing core count, they have started to see issues with disk access on the worker nodes, and have hence investigated using Flashcache and Bcache, which were found not to be the best solution for their use cases, as described in a presentation given at CHEP. Concerning facilities, they have replaced ageing Liebert rack-top cooling units, and have deployed wireless 3D probes for temperature, humidity and pressure (see presentation at CHEP). The subnets of PHENIX and STAR have been separated

entirely, reducing the inter-switch traffic significantly; they are deploying 100 G to New York City. They are looking at Infiniband IPoIB fabrics as a potential alternative to 10 GE; although technically interesting, because of needs and funding priorities, this will probably not be deployed at scale before 2015. Concerning mass storage, they hold 36 PB on about 42'000 tape cassettes within HPSS that will be upgraded next month to 7.4.1p2; the NFS storage appears to hit the limit of what is currently possible with NFS, even though investigations are ongoing. The ATLAS files in dCache are being renamed according to the Rucio convention; a catastrophic Nexsan RAID controller failure resulted in the loss of more than 1 M files. A coordinated effort has begun to evaluate other options for ATLAS such as MapR, CEPH, GPFS. They are deploying Openstack Grizzly on a test basis. They use Puppet, Cobbler, and RHN Satellite for configuration, considering replacing Cobbler by Foreman.

## NDGF (Mattias Wadenstein / Umea)

Mattias started by mentioning some personnel reassignments. Gerd Behrmann is now on-board as a full-time developer, as are Jon Kerr Nielsen and Magnus Jonsson. Central services are being moved from physical blade servers to virtual machines running on two fat hypervisors; only the name space database will remain on a physical machine. In terms of computing, there is less and less dedicated resources to fulfil the WLCG pledges; the majority comes from shares from larger HPC systems on which CVMFS, compatibility libraries etc. are installed. However, the WLCG community does not control which operating system runs on the nodes. The situation of storage with eight distributed dCache pools is similar. At NSC Linkoping the first quarter of a new 3000 m2 machine room has been put into production. At HPC2N Umea a new TSM server under Linux has been commissioned.

## DESY (Wolfgang Friebel / DESY Zeuthen)

Wolfgang reminded that the focus of DESY has shifted from particle physics to include photon science (Hamburg) and astrophysics (Zeuthen). For Grid work, 7'700 slots and 1'600 slots are provided in Hamburg and Zeuthen, respectively; the majority of resources in Hamburg are 4-socket AMD systems. All server nodes in Hamburg run on XEN VMs; there are no plans to virtualise worker nodes. Hamburg is planning an extension for an additional 50...60 racks for an additional power consumption of 1.5 MW. Concerning networking, DESY has now 2 x 10 GE into the public Internet, and 2 x 10 GE into LHCONE; for the interconnect between the two sites, a second 10 GE link will be added. Zeuthen is involved in Icecube, being an important Tier-1 centre (running a CVMFS stratum 1 for example). Dcache is holding 13 PB in total. The concept of a distributed NAF has been replaced by a single instance in Hamburg, with a total of 4'000 cores and related storage. At Hamburg, in view of the large number of VOs, CVMFS on all worker nodes had the issue of needing too large caches, hence they consider having few CVMFS client nodes export the file system via NFS to all worker nodes. The migration of their Exchange 2003 server to Zimbra is going on; initially the deployment will use Zimbra 8, an upgrade to 9 is foreseen for later. The networking tools presented at HEPiX Fall 2012 have evolved, some of them implementing new functionality. Wolfgang finished with a number of miscellaneous points including the purchase of a Bomgar appliance, the new additional role of Stefan Wiesand as project manager for AFS, and a student project to investigate log file analysis similar to what was presented by GSI in spring.

## RAL (Martin Bly / STFC-RAL)

Martin recalled the assets of the Tier-1: 10k cores, 8 PB disk, 10'000 slots in SL8500. In FY 2013/14, they are adding 7.0 PB of disk space and 46 kHS06. They have observed that the vendors have done a much better benchmarking job than previously. They have decommissioned systems procured in 2007, and are preparing for retiring the 2008 procurements. Concerning networking, RAL has been migrated to the new UK backbone (SuperJanet 6), with a dual 30 Gb/s active/passive

failover link. The Tier-1 link to the boundary has been established at 20 Gb/s. Following painful experience with Torque/MAUI, they have decided to deploy HTCondor and to migrate from CREAM to the ARC CE. All batch resources are on SL6, about 50% are migrated to HTCondor. FTS3 is fully up and running, with servers on VMs with a MySQL backend. They use Quattor and Aquilon for configuration, and are running important CVMFS services. Most services run on virtual machines, but there have been issues with Ganglia and BDII. Concerning storage, 64 million files are in Castor with 14 PB on tape and 8 PB on disk. They are looking at a disk-only storage complementing Castor, without any firm conclusions yet; however, CEPH testing continues as storage option for cloud infrastructure. The RAL AFS cell will be discontinued soon. Then Martin described their uprating of the essential power board capacity, which means that there was no UPS supply to protect HA services. A major test will be done on November 5th causing interruption of most services; batch will only be restarted when all other services are back in stable operation. Other users at RAL ("facilities") are using similar, if not identical, tools to the Tier-1, however in separate instances. They have experienced a number of UPS generator failures (failed to start, failed to assume load etc.), and have moved to a much more rigorous testing regime.

## U Wisconsin/Madison CMS T2 (Ajit Mobahatra / U Wisconsin)

Their first site report ever explained that the facility started out as a grid3 site with strong collaboration with the HTCondor team; it was subsequently selected as one of seven Tier-2 sites of CMS in the US. In addition, they support all OSG VOs. They are using three machine rooms with 16 racks for an overall power of 650 kW. Compute resources are running SL6 on a total of 7'400 cores, to which 1'000 cores will be added soon (Ivy Bridge). The storage was migrated from dCache to Hadoop 3 years ago, providing 3 PB across 350 nodes, to which 1 PB will be added soon; Hadoop is being upgraded to 2.0. The networking will soon deploy a 100G connection to Chicago; in addition there are dedicated links to FNAL, Purdue and Nebraska. PerfSONAR has been deployed for monitoring. Thanks to the good collaboration with the campus network team, good transfer rates have been achieved for all CMS data transfers. In terms of services, the list is rather classical – AFS, NFS, CVMFS, Frontier/Squid; HTCondor; Globus etc; Hadoop, BestMan, gridFTP, … The cluster is managed by Puppet (migrated from CFEngine this summer), Nagios, Ganglia as well as dedicated tools from OSG and CMS. They are working on a project making all CMS data available transparently anywhere, anytime; the underlying technologies include xrootd, CVMFS and Parrot as well as glideinWMS/HTCondor. The technology has been demonstrated already by using campus capacity and other sites for overflow CMS jobs. They also made experience with Amazon EC2, configuring nodes as Wisconsin worker nodes. About half the cost was for the data transferring data (the output of the jobs). About 5...15% of the jobs were lost due to spot instance termination.

## PDSF at NERSC (James Botts / LBNL)

As James was unable to travel, he gave his presentation remotely. James reminded people that NERSC is the primary computing facility for the US DoE Department of Science, featuring large Cray and IBM Dataplex systems as well as GPFS and HPSS data storage. PDSF takes advantage of the GPFS storage system; it has been in continuous operation since 1996, and now comprises 160 compute nodes, to which 48 nodes (Sandy Bridge, Infiniband FDR) will be added. In GPFS some 900 TB are being used for PDSF. The nodes are managed with xCAT; for the most part, servers and compute nodes are booted diskless. The workloads are of the serial high-throughput type; they use UNIVA GridEngine as batch system. Main users are ALICE, ATLAS and a number of non-LHC experiments. Scheduling uses a fair-share mechanism, projects buy into PDSF and get their share adjusted accordingly. ATLAS, ALICE and STAR account for 85% of the PDSF usage. In terms of user support, they have some 160 tickets per year. PDSF is special as compared with other clusters by its high number of dedicated service nodes; virtualisation is hence a very natural way to go, as is

unifying the resources into a common pool with a shared network, for which the migration to the NERSC global file system is being considered. He described the changes in detail that this implies, including the migration from CFEngine 2 to CFEngine 3. Finally, he discussed the NERSC move from Oakland to CRT at LBNL foreseen for 2015.

## *Facilities and business continuity*

## The CSC Kajaani datacentre (Ulf Tigerstedt, CSC Helsinki)

The presentation started 15 minutes late because of technical problems – an unreadable PowerPoint file, no PDF, echo on the mike etc. Ulf put the story into context: classical industries in Finland were on the way down, but emerging industries took over. Google entered the country with a plan to establish a large data centre in the north. CSC got funding for a new large machine, but was asked to run the machine outside the Helsinki region. An open call for tenders was won by Kajaani, a town in the middle of the country with a population of 38'000. Temperatures are between -15 degrees in February and +15 degrees in July. A suitable site was identified as well – a paper mill that had closed down shortly before. However the room was not adapted to computers, not least because of fire protection. Hence the idea of modular data centres came up, but talking to vendors, they indicated a number of problems, including the cold temperatures, the snow, and the use of water for the cooling. They finally decided to put the modular data centre into the paper warehouse, and settled on two different types, one for a water-cooled Cray machine, and the other one for a cluster with free air cooling. The cluster consists of dual 8-core Xeon machines (HP SL230); commissioning has been delayed significantly by problems with the Infiniband network. However, it is expected that ALICE and CMS can start using the cluster in 2014. In hindsight, it was not a wise decision to run a project of that size and complexity, trying to install clusters and supercomputers at the same time. In addition, remote management works for computers, but not for humans; and as everybody outsources, in a rural region, all express couriers will end up with the same man and the same van.

## Operating dedicated data centres – is it cost-effective? (Tony Wong / BNL)

Tony explained that the starting point was a DoE review of the RACF, during which they reported on successful tests of Amazon EC2 and Google Compute Engine, hence the question naturally arose how the dedicated BNL data centre compares with these commercial offerings. The RACF consists of 2'200 servers, 23'000 physical cores, 16 PB of worker-node based storage as well as dedicated storage on disks and in tape robots. Typical usage profiles include Monte Carlo, analysis, and interactive work, with different requirements on the CPU/IO ratios and response times. Tony then explained the EC2 offerings, estimating the compute power in terms of HS06, and reported on how they had conducted the measurements with EC2. With Google Compute Engine, things are similar, except that only Google-provided images can be used. In order to compare with costs of dedicated facilities, he considered direct (hardware, software) as well as indirect (staff, infrastructure) costs. He then explained the cost estimates of servers, networking, software (they only pay for Ksplice, Synapsense, Sensaphone – this results in a negligible amount), electrical costs, and space charges for the overall data centre. This results in estimates of about $350 per core per year (a little more for RHIC, a little less for ATLAS), corresponding to about $0.04 per hour.

A related question is resiliency of data storage, hence BNL have investigated what would happen if duplication of derived data at the Tier-1 would be a requirement. Not surprisingly, tapes are cheapest, but have the largest disadvantage in terms of latency.

Tony concluded that the dedicated data centre has a cost advantage over commercial offerings of about a factor 3.

## Safety in the data centre (Tony Wong / BNL)

Tony reminded the audience that BNL is a multi-disciplinary lab with projects at various stages, diverse workforce, and an ageing infrastructure. New facilities still depend on old infrastructure. Over the last few years there have been some isolated high-profile accidents such as an arc flash, a building explosion, and a 16 ft. fall from a scissor lift. The safety culture at BNL has been changing, giving more emphasis to reducing downtime by avoiding accidents.

More recently, on 30 September 2013, a complete rack fell on its front face during installation. No one was injured; however as the servers were loaded from the front, servers could not be de-installed, hence the complete rack of 454 kg needed to be lifted entirely. Some power cables, at the time fortunately without power, were damaged, as was a CRAC. Analysis showed that like for many incidents, a sequence of unfortunate conditions was responsible. The rack needed to be raised up with specialised equipment. The front door was destroyed, one hot-plug disk was re-seated, but otherwise no damage to the servers could be found – fortunately as the cost of a fully configured rack is about $100k. Three power cables were replaced on 14 October; in total a delay was incurred of three weeks, not least as remediation must be auditable. Following the accident, they verified the training documentation and found it clearly inappropriate. Remediation procedures have been established – a third person must now assist, training is being modified, a work plan to be approved must be established in advance, and the rack wheels must be immobilised when nearby tiles are removed. They are also trying to understand best practices in the field. Finally, Tony was inviting other data centres to let him know what they do in order to prevent similar kinds of incidents.

## *Basic IT services*

## CERN-IT monitoring (Massimo Paladin / CERN)

Massimo gave an overview of the rejuvenation project of monitoring at CERN-IT, which is part of the larger Agile Infrastructure project. A new setup was needed in order to cope with changing and growing requirements (e.g. virtualisation), to allow for more synergy between teams requiring monitoring services, and to provide better analytics. In line with other AI areas, a tool-chain approach was used – break the problem space down into small blocks and select a component, preferably open-source, for each block. The basic architecture identifies producers, transport, and consumers, the latter comprising notification (alarms) and a repository facilitating analytics and visualisation. Blocks chosen include Flume, HDFS, ElasticSearch, and Kibana, each of which he described briefly including the experience with them in the context of the project. The infrastructure has been deployed at CERN, for each block a number of (virtual) dedicated servers has been set up. Massimo showed sample graphs, and gave high-profile examples where services have used the infrastructure for their own, dedicated purposes: CERN's internal Openstack-based cloud services, and the large-scale batch services.

## Building a Puppet infrastructure at DESY (Jan Engels / DESY)

Jan started by giving a short introduction into the functions and workings of Puppet, a by now very popular configuration management tool, and then explained that DESY had chosen it for its cross-platform support, large repositories of re-usable elements, the widespread use etc. They started in 2012 with two Puppet servers, one Git server, one Puppet CA, a Netapp filer to provide shared storage, and an F5 load balancer. Meanwhile they have added two database (PostgreSQL) servers and two servers for PuppetDB, Gitlab and Foreman. They chose Gitlab as Puppet development platform because of its large span of functions needed. Jan then described the workflow from development to production. For bootstrapping (host registration, kickstart file, network registration, Puppet node definition) they use a home-grown tool (WBOOM). Host grouping takes a hierarchical approach with few top-level hostgroups. They mirror external

repositories to DESY and use yum to interact with the copies. For node management they are considering Mcollective, but for the time being are still using home-grown tools. Secrets management is an area that requires further work, both the DESY tool and Hiera encryption backends are considered inadequate. Currently some 500 client nodes are managed by their infrastructure, more are constantly being added.

## Automatic server registration and burn-in framework (Afroditi Xafi / CERN)

Afroditi described the problems motivating the work done in the context of CERN's Agile Infrastructure project: Previously massive deliveries were hand-registered in the network database and in the system administration toolkit, the OS was installed, the burn-in test was run, and the results analysed mostly by human intervention. The process was error-prone, hard to scale, and very person-power intensive. In order to allow for more automatisation, CERN required suppliers to record the CERN order number and the serial number on a barcode sticker as well as in NVRAM of the machine. The PXE-booting node is now assigned a temporary IP address, with which it loads a live image analysing the machine and registering it definitely in the network database based on its MAC address. It then obtains the required certificates, and registers itself with all its components in the hardware database, after which the burn-in test suite is started automatically. Afroditi then explained some details of the burn-in tests that had been presented previously at HEPiX (spring 2013); they cover memory, CPU, disk endurance, and disk and CPU performance testing. She reported that a number of systematic issues was found due to this test. More than 1'000 machines were installed and stress-tested in 1.5 weeks, an operation that would previously have taken several months. For the future, they plan to add a P2P network test, better support for RAID cards, and to factorise out CERN-specific elements to make the development useful for other sites as well.

# *Tuesday 29 October 2013*

## *Site reports*

## GridKA (Andreas Petzold / KIT)

Andreas started stating that the migration to Univa GridEngine was successful, there are no major issues. The total farm capacity is 150 kHS06 with 10'000 job slots. Machines rolled in as replacement capacity are based on Sandy Bridge. The migration of worker nodes to SL6 is complete; benchmarks have shown a gain of 5% in terms of HS06. New Ivy Bridge systems have also been benchmarked, resulting in a linear scaling with the number of cores; however the power efficiency is superior to the Sandy Bridge systems. KIT are running six production dCache instances with a total of 9 PB capacity. Most instances were upgraded to 2.6 straight away in order to profit from SHA-2 and xrootd monitoring. The disk-tape separation has been achieved for CMS. Disk storage is based on DDN systems for a total capacity of 14 PB usable; they are evaluating new storage solutions including running complete dCache instances inside the controller, which required quite some fine-tuning to achieve an acceptable level of performance. A second SL8500 tape library has been added recently that is being equipped with LTO drives. They plan to migrate from TSM to HPSS in 2014. They currently run 7 10G network links, but are testing 100G by participating in tests for SC2013; a migration to 100G for LHCOPN and LHCONE is foreseen for 2014. For configuration management, they plan to replace CFEngine 2 by Puppet. Apart from their services to WLCG, KIT are involved in a number of activities in the Bundesland for scientists including Sync&Share (dropbox-like), file storage and block storage; for this purpose a federated identity system has been set up. PowerFolder has been selected for the dropbox-like service.

## Jefferson Lab (Sandy Philpott / JLAB)

Sandy reported that in 2012 JLAB was awarded computing hardware for US Lattice QCD; the money was split between a classical IB cluster (276 dual Sandy Bridge nodes with QDR IB) and accelerators. The cluster just missed to get into the TOP 500... They have experienced quite some issues between the Mellanox IB adapters and the QLogic IB switches. As accelerators, they chose a mixture of Nvidia Kepler and Intel MIC. With the 42 Kepler nodes, they did achieve position 364 of the TOP500 list; however some instabilities have been observed. For the physics analysis farm, they consider moving away from Torque/MAUI, possibly to Slurm. The Lustre storage is still on 1.8.8; backup services are being added now. ZFS is being seriously considered for the five Thor servers that are still in operation even though there were issues over IB. New LTO6 drives have been added to the TS3500 library. Finally, Sandy mentioned that there is a documented plan for data preservation, and described the status of the accelerator upgrade to 12 GeV; once in full operation in about two years, an estimated 15 PB per year will be produced.

## GridPP Tier 2s (Chris Brew / STFC-RAL)

Chris started by reminding the audience what GridPP is about, and how it is funded. There are 19 individual sites with 13 PB of disk, 33'000 batch slots providing 300'000 HS06, making the UK contribution the second most important one to WLCG Tier-2s. There is a very significant spread in size between the participating sites as well as in the chosen software tools. In winter/spring 2012, the UK government awarded 3 MGBP to GridPP for Digital Research Network infrastructure that was spent on increasing the flow of data to the jobs, tying in well with the upgrade of the Janet backbone. The Tier-2 have formed a storage group organised mostly as a self-help group, and are involved in the HEPiX IPv6 working group; the concept of a diskless Tier-2 is being tested. Chris finished by stating that many sites are moving to Puppet; a group has been established to that end.

## Fermilab (Keith Chadwick / FNAL)

Keith started by a flash overview of the lab, referring to a Chicago Tribune article on Scientific Linux, and then explained some structure and personnel changes. Work is going on on implementing ITIL service management; Keith showed graphs on incident management and change management as well as key performance indicators (KPIs). The facilities have been running stably. A number of network initiatives are going on such as a new 100G metropolitan area network and IPv6. The experience with the Dell managed services is rather good. Some 78 PB of data are stored on tape, of which 47 PB are active; they are migrating from LTO4 to T2. They added 6.5 PB of disk cache storage. Like JLAB, they are running Kepler and MIC accelerators. An initiative has been set up to have other frontier experiments profit from the experience gained with the energy frontier ones.

## CERN (Arne Wiebalck / CERN)

Arne started by referring to the Physics Nobel price, and giving a short overview of the OpenDays held at the end of September. The Wigner data centre in Budapest is now fully operational; work is going on about shut-down procedures for the computer centre. The Zenodo launch triggered a lot of interest; attractive features are being added to Vidyo. An Eduroam pilot service has been set up; the primary connection to Geant is now at 50 G in total. Oracle 12c has been set up for testing, and PostgreSQL has been added to the catalogue of the database-on-demand service. The “lost-write” bug has been identified as a combination of issues in Oracle and Netapp that have both been fixed. A public instance of EOS has been opened; Oracle T10000D drives are being deployed. They are investigating OwnCloud for a dropbox-like interface, and have successfully finished some initial testing of CEPH. The virtualisation layer is now using Openstack Grizzly, using the experimental cell feature. Unified messaging has advanced, including the integration of Lync and Skype. The

preparation of Windows 8 has advanced, more news will be reported in spring. The configuration services have been migrated to Puppet 3 and introduced a new work flow. Git and JIRA are increasingly popular, CVS has been stopped. The SL6-based batch and plus services are mostly using virtual machines. Multi-factor authentication is ready for use; obsolete firewall openings are being detected, as are mis-configured and infected devices. The integration of services into ServiceNOW is progressing well.

## IHEP site report (Jingyan Shi / IHEP)

Jingyan said that IHEP are a Tier-2 site for CMS and ATLAS providing 1'500 job slots and 400 TB of storage split between dCache and DPM. Worker nodes have been upgraded to SL6; disk arrays for Atlas have been replaced. DPM and Cream are running in SHA-2 compliant versions; dCache will be upgraded in due time. A new CA server has been deployed that supports SHA-2, next to fixing a number of other problems. The local farm at IHEP supports several Chinese HEP experiments and some biomed activities on 7'500 job slots managed by Torque/MAUI that gives rise to some concern, hence they are considering replacing it. These experiments also use 3 PB of disk space and 5 PB of tape capacity. For configuration management, Quattor that had been used for years is being replaced by Puppet (and Foreman). Most users have access to a 3PB Lustre instance to which 1 PB will be added soon (hardware burn-in is running now); they are monitoring access patterns in order to understand and tune the Lustre performance. They also run a small Gluster installation (186 TB) for cosmic-ray experiments with good experience concerning performance; however some change needed to be applied following user requests. The tape libraries are managed by a modified version of Castor I, serving experiment on-line data and backup storage; they have started systematic performance measurements for copies from tape to the Lustre capacity. They are worried by the expiry of the hardware warranty for two thirds of their hardware; they are maintaining memory modules, power supplies and disk drives themselves, which is person-power intensive. Jingyan finished with a status description of SDN, a national research network infrastructure that fully supports IPv6.

## *Computing and batch systems*

## HS06 performance per watt and transition to SL6 (Michele Michelotto / INFN Padova)

Michele cited rumours (in spring) of significant discrepancies of HS06 across SL distributions on the same hardware. He has subsequently run dedicated tests on an AMD Opteron 6272 system, showing a gain of about 5% just changing the OS, and another 5% by using a more recent gcc compiler. When running the benchmark on a single core, differences were even as large as 30%. Tests on an Intel Sandy Bridge platform gave similar results. Michele used all results published on the HEPiX benchmarking Web page, exposing differences between about 4% and 10%. He then described measurements of a new dual 12-core Ivy Bridge system that reaches a HS06 performance of about 550 HS06 in 64-bit mode! Michele then described measurements performed on an ARM platform with a 1.7 GHz quad-core Cortex chip, mentioning that the CMS core software had been ported within just a few days. CMS had indications that the power efficiency is very much better (by factors!) than with the x86 platform. The measured HS06 performance is about 3...4 per core depending on how many cores are loaded. A detailed study of the power efficiency with a dual 8-core Sandy Bridge system shows that the more processes are run, the better is the power efficiency. Michele finished by announcing that he plans to look soon at the new Intel Atom server processor, and at more ARM-based models.

## Future of batch processing at CERN (Jerome Belleman / CERN)

Jerome started by describing the current batch service setup at CERN. It runs IBM LSF 7.0.6 on 4'000 nodes; the migration from SL5 to SL6 is underway, as is a move to virtual worker nodes. Up to 400'000 jobs are submitted per day, resulting in 65'000...70'000 concurrent jobs. For future expansion, they are looking at 30'000...50'000 nodes, dynamic behaviour, a dispatch rate of up to 100 Hz, and a query rate of up to 100 Hz, all of which points LSF may have problems with. They have worked with a consultant, who suggested some minor improvements, but did not identify any major shortfalls of their setup, hence the identified limitations of LSF are to be taken seriously. They hence started looking at alternatives, including Slurm 2.5.7, HTCondor 8.1.0, Son of Grid Engine 8.1.3, and (perhaps) LSF 8 or 9. Slurm was particularly interesting because of its similar look-and-feel and because of reports of very good scalability. A test framework has been set up that allows for simulating loads both in terms of submission and queries. Jerome then showed preliminary results of submission rates, where Slurm rated highly, but rather erratically, while Condor and SoGE were more stable at or above 50 Hz; probably the Condor setup was not fully optimised, further tests will be done. For scalability tests, they are considering piggy-backing on existing LSF batch worker nodes. Slurm has shown issues when starting it up in cloud mode (i.e. without pre-defining the list of worker nodes) with large numbers of workers; HTCondor is much more made for this use case. Jerome concluded that the general impression was rather negative on Slurm, while HTCondor has worked well except that updates to the configuration could be a bit hard. SoGE is hard to configure, as it requires a shared file system for the configuration; SoGE was “rough around the edges”. Additional features they need to consider are Grid support, support for AFS and Kerberos, accounting, host normalisation, fairshare scheduling, support for commercial applications and for IPv6. The next topics to be studied in detail are host scalability, query load and the feature set.

## HPC activities at CERN (Ioannis Agtzidis / CERN)

Ioannis explained that some 5% of the CERN applications do not fall into the massive / embarrassingly parallel category, but require special HPC support. These applications include physics (lattice QCD), engineering, and accelerator physics applications, a mix of commercial and community-developed applications. He then reminded the audience of Amdahl's law limiting the potential gain from parallelism. They have deployed a set of standard tools for performance analysis such as iostat, dstat, sar, netstat and the Intel counters next to mpiP, the data of which are analysed with Matlab and Excel. Preliminary results include an almost perfect scaling (of Ansys mechanical) with Amdahl's law; however not all applications behave in this ideal way. Ioannis listed the factors that determine the behaviour, and showed studies of all individual factors; for example, using 1 GE interconnects significantly limits the performance of Fluent with respect to 10GE. He then explained the performance differences of the QCD application between a low-latency (RDMA) 10GE interconnect, and an IB network, which are significant on the communication level, resulting in a 20% performance impact on the overall application. CERN has chosen to accept the 20% overhead and stay with Ethernet. Plans for the future include iWARP MPI scaling from 20 to 60 nodes, performance tuning, a more detailed analysis of MPI, and more tools for users for performance analysis.

## Condor at the RACF (William Strecker-Kellogg / BNL)

William explained that there are three HTCondor pools at RACF: One for ATLAS, one for STAR, one for PHENIX, with different characteristics for RHIC and ATLAS. They run four central managers each supporting 10...20 k cores with modest hardware requirements; there are six submit nodes for ATLAS and 20 and 10 interactive nodes for PHENIX and STAR, respectively. They have much appreciated a change allowing for splitting configuration files in Condor in a directory. He then explained the structure of the ATLAS setup in more detail, explaining how spillover between

different resources works. About two years ago they started to support multi-core slots, initially by statically partitioning machines, with a number of issues they have had to address together with the HTCondor team. Now they can fully utilise partitionable slots. He then discussed potential problems with multi-core jobs, for example mismatches of CPU-RAM ratios that cause wasting resources. They are addressing this issue by a new concept called defragmentation, the principle of which he explained. Finally he discussed in which direction HTCondor may evolve in order to accommodate requests for more flexibility.

## Batch system status at RAL Tier-1 (Andrew Lahiff / STFC-RAL)

Andrew explained that the batch system features 93 kHS06 on 9'312 slots that support all LHC VOs and in addition many non-LHC experiments. For many years they had been using Torque/MAUI, which however has given rise to many issues that he listed in detail, leading to spending significant effort just in order to keep the system running. In August 2012 they decided to look for an alternative, considering Torque 4 and MAUI, LSF, GridEngine, Slurm, and HTCondor. Criteria included the integration with the WLCG community and with the RAL environment, scalability, robustness, software support, procurement costs, maintenance costs, essential functionality and desirable functionality. Some products were quickly rejected: LSF, Univa GridEngine, Oracle GridEngine for being commercial; they also rejected the open-source Grid Engines because of the competitive situation with an uncertain prospect. They also rejected Torque 4 and MAUI, as there were still issues with job submission. That left Slurm and HTCondor in the game. Andrew described the tests run, and reported that HTCondor could easily run 10'000 jobs and deal with 200'000 pending jobs, while a number of issues were found with Slurm that despite serious efforts could not be resolved. Hence they settled on HTCondor. However, there is no official support for Cream with HTCondor, a problem they could easily overcome with a bit of effort. They also investigated using the ARC CE instead, which they found to have certain advantages such as easier configuration and setup. Testing with ATLAS was successful, hence they chose it over Cream. The only VO not able to deal with this is ALICE, however they can submit directly to HTCondor. By late May 2013 they had an almost production-quality service set up using resources beyond WLCG pledges. The testing was mostly done with ATLAS, CMS joined in a little later. This test went very well after having fixed some teething problems. The setup uses the high-availability central manager, hierarchical fair-share, partitionable jobs, concurrency limits etc. The migration is scheduled to be completed in November 2013. Their experience with HTCondor is very good so far, they did not have any major problems, and have enjoyed very good support; the expectations concerning performance have been fully met.

## GridEngine: One roadmap (Cameron Brunner / UNIVA)

Cameron started by giving an overview of the history of GridEngine that started in 1992. Univa has recently acquired all Grid Engine assets from Oracle, which has implications on the copyright, trademarks etc. Hence Univa is the home of GridEngine now, removing confusion among users. He then described the offerings Univa makes to GridEngine users and mentioned areas of product development: support for new hardware and operating systems, access to add-on products such as UniSight and Hadoop, security and maintenance patches etc. New products around GE include unicloud, a licence orchestrator, and a native Windows implementation. Cameron then described the highlights of release 8.0 as well as the actual 8.1; the latter includes high-availability features, performance improvements, and time-saving features. He then turned to case studies mentioning customers from oil and gas, chip design, and healthcare modelling. The next minor revision 8.1.7 is scheduled for December 2013; 8.2 is to be released in 1Q2014, featuring a native Windows implementation, cgroups support, a read-only qmaster thread for better query performance, DRMAA v2, and further performance optimisations.

In response to a question, Cameron indicated that there may be another release of a downlevel

version into open-source, but no decision has been made by the company yet.

## *Security and networking*

## Mobility at CERN (Sebastien Ceuterickx / CERN)

Sebastien started with an overview of the wireless evolution at CERN, comparing the situation in 2007 with today. 7'000 unique devices are seen today with 4'000 simultaneous users served by 1'400 access points. Deploying such large a structure is very difficult due to the variety of clients as well as environmental constraints. Sebastien then mentioned some of the client issues: hidden nodes, co-channel interference, power mismatch, hardware failure, driver failure etc. In addition, running one slow station in a cell may slow down all devices in that cell. Wifi is no longer an extension of the cabled network, but a basic service that is the only communication means for an increasing number of devices. It is hence important to optimise the RF spectrum, for example by avoiding the 2.4 GHz band wherever possible. However this limits the size of the cells; in addition, a large deployment requires careful planning and tuning. To that end CERN's networking group have deployed efficient monitoring and analysis systems, of which Sebastien showed examples. BYOD (bring your own device) brings additional challenges – some devices are limited to a single spatial stream in order to save power, but this limits the achievable throughput. Sebastien then discussed special use cases – conference rooms and underground caverns – and how CERN's networking group deals with these challenges. Finally, Sebastien covered Eduroam and its deployment at CERN; for CERN users only certificate authentication is supported.

## Follow-up on network projects (Sebastian Ceuterickx / CERN)

Sebastien gave an update of ongoing network projects at CERN on behalf of colleagues in the networking group. The data centre has finished its 2-year migration to Brocade routers, which allows for 100 G links and more scalability; the routers also support MPLS and virtual routing. He then described the topology of networking, noting that the bandwidth of the major networks at CERN has been increased significantly. CERN's strategy is to skip the 40 G technology and move to 100 G immediately wherever needed. Concerning the top-of-rack, there is still some uncertainty about the medium for 10 G – can copper be used? Sebastien then described the two links CERN runs between Geneva and Budapest for its Tier-0 extension. At Wigner, the same network structure as at CERN has been implemented. The Wigner centre offers some possibilities for implementing business continuity; in addition, in the same spirit, efforts are going on to add a second network hub at CERN. He then described the status of IPv6 at CERN, stating that the network is basically ready for dual stack and dual routing. In order to support massive numbers of VMs, the choice is between IPv6 only and a dual-stack configuration with public IPv6 and private IPv4 addresses; probably the latter option will be picked. Finally he described Tetra, a secure digital radio network, and the status of its deployment at CERN.

## Deploying PerfSONAR-PS in WLCG: an overview (Shawn McKee / UMich)

Shawn explained the motivation for the PerfSONAR deployment: WLCG critically depends on the network, but network problems can be very hard to diagnose. PerfSONAR is intended to find and isolate “network” problems and alert timely, characterise network use, and provide a source of network metrics. PerfSONAR-PS is a specific implementation that is designed to be standalone and federated. Its deployment started in US-ATLAS. Shawn showed sample dashboards demonstrating the value of the tool; negotiations are going on with OSG about hosting the modular dashboard service. The service is now configured via a mesh configuration – instances refer to centrally maintained configurations rather than relying on a local configuration file. For WLCG, all sites are

supposed to deploy PerfSONAR-PS. Shawn described the setup that uses the concept of regions; the configuration has been chosen carefully in order to be scalable. He then described how the metrics can be used to identify and debug network problems, and presented example success stories from US-ATLAS. There are a number of high-level PerfSONAR-PS issues that are being addressed, profiting from close links with the PerfSONAR-PS developers. Shawn finished with an outlook of what direction network monitoring could take in future.

## *Wednesday 30 October 2013*

## *Security and networking*

## The HEPiX IPv6 working group (David Kelsey / STFC-RAL)

David recalled the situation of the IPv4 addresses which get exhausted. In July 2012, Google and other content providers enabled IPv6 and kept it switched on; in the Netherlands, 2.3% of Google traffic is already IPv6, probably without end users noticing. In Switzerland, the number is even as high as 11.6%. The CERN infrastructure for IPv6 is well advanced, but many smaller labs don't even seem to have started. In 2011 the working group was launched in order to consider how IPv6 should be deployed in HEP, do a readiness and gap analysis, and test applications. Since the last HEPiX meeting, CERN has stated that they run out of IPv4 addresses; new testbed sites have joined; there is more engagement from the LHC experiments, and a lot of testing is going on. David then described the testbed now comprising twelve sites, and presented results of GridFTP mesh transfers (which had a 87% success rate without any special attention). The next step was to test end-to-end transfers with IPv6-only DPM installations with PhEDEx which went fine; on the way a number of performance and tuning issues were solved. At Imperial College, dual-stack services were configured and left running when no problems were found. The group is currently running an “asset” survey, the results of which David presented. There is still a lot of work to be done... David then mentioned a number of use cases identified by the WLCG IPv6 task force, with which the HEPiX working group is collaborating closely. David finished by presenting an overview of the plans for the next year or two, focussing on establishing dual-stack services, and discussing the time lines. In summary, David noted that while good progress has been made, the group is still effort-limited, hence more contributions would be highly welcome.

## Security update (Romain Wartel / CERN)

The slides of this presentation were not available on Indico (a sanitised version has been uploaded after the presentation), and Vidyo was switched off. Romain explained that his talk will describe the changes since the last presentation, which emphasised already that a site-centric security model is bound to fail. The only way is traceability and international collaboration and trust. Latest trends include distributing malware (e.g. blackhole) via advertisement banners of well-trusted sites, Android malware that is designed to intercept the SMS sent by the bank as part of the two-factor authentication for online banking, very sophisticated government-class attacks such as Stuxnet and the Syrian Electronic Army (the British government even started hiring cyber-security experts for a cyber-army reserve), and national security agencies sniffing everywhere (even if this is not a surprise per se, the scale and scope is). Romain then turned to the impact on HEP labs. Protection is difficult, as all traffic goes through routers, the firmware of which the HEP labs do not control. In addition, fighting sophisticated attacks requires expertise that goes far beyond the average site and system administrator, hence it is important to be in touch with relevant experts both at a policy and an operational level. Romain gave an example of an attack to CERN during October 2013. Romain urged sites with local expertise to invest in their expertise and training, and those without expertise to establish strong relations with relevant bodies that can help. Romain finished by discussing a security issue with BMC/IPMI software. Most vendors are affected (Supermicro, Dell, HP).

## Identity management in future scientific collaborations (Bob Cowles / U Indiana)

Bob started by clarifying the scope of identity management, and putting the effort into its historical context, which the project strives to take to the next level. In order to do so, a trust relationship needs to be described between resource providers and collaboratories. As the first step, trust needs to be defined, which already is not easy... there is a broad spectrum, from which the project picked a a suitable definition. Then interviews were held with stakeholders in order to clarify the conditions under which they would trust others. Bob described the spectrum between early and late user identification by resource providers, reflecting different levels of trust into the collaboratories. This lead to a detailed understanding of the factors affecting the design of identity management. They considered different use cases with rather specific requirements: HPC, cloud (Bob reported on the results of a survey among XSEDE members). For HEP, the question arises primarily for the LHC experiments and for Belle-2; beyond that there are no concrete plans yet. As an illustration, Bob described the computing model of the Belle-2 experiment. Areas of future growth include astrophysics, biomedical, chemistry, earth sciences, other physics, all of which differ from HEP. Bob finished by giving an outlook of the future work of the project.

## Federated identity management for HEP (David Kelsey / STFC-RAL)

David Kelsey, similar to the preceding speaker, introduced the subject by recalling the history and explaining how the idea of an identity provider came about, and then referred to a keynote talk by Kostas Glinos (European Commission) at CHEP that mentioned an integrated European e-infrastructure. He then talked about FIM4R (federated identity management for research), an initiative emerging from EIROforum and including a very large range of sciences, remaining open to any other users who wish to join. For WLCG, there are practically two interests: Web-based and CLI-based applications. Currently work is focused on the Web use case. While interesting for the experiments, the issue does not seem to have high priority. Federated identities will necessarily assign a much more important operational role to the identity providers than they had before, not least for security reasons. Finally David covered the IOTA initiative within EUGridPMA. The next steps are a discussion in IGTF next week, discussions with others about Horizon 2020, the potential creation of an interest group in RDA, and further discussions in HEP.

## Evolution of OSG grid authentication model (Kevin Hill / FNAL)

Kevin explained that OSG has historically used X.509 certificates (proxies) for authentication. DoEGrids are shutting down their CA, hence OSG had to plan setting up its own CA. He described in some detail how this CA has been set up, and that it is fully integrated with the existing OIM system and the GOC ticketing system. He then described the approval process. SHA-2 support is ensured for all OSG-provided software, other software may need testing; the current recommendation is to start issuing SHA-2 certificates as of 01 December 2013, but OSG pushes for 15 January 2014 in order to avoid changes during the holidays. Then Kevin mentioned CILogon basic certificates as an alternative source of X.509 certificates for users (even though not approved by IGTF). In addition, OSG are considering certificate-free job submission, which needs to address user identification, access control, and blocking of unwanted access. These can be addressed if job submission gets moved from end-user systems to VO-managed portals. In a similar spirit, OSG Connect is a project providing a web portal for job submission. It uses CILogon federated authentication, and flocks jobs to existing OSG VO front-ends. In future, they will continue with Digicert CA signed certificates, using CILogon CA signed certificates for CIcommon member sites.

## Technical security tips and techniques (Romain Wartel / CERN)

“Be a lot safer and prepared in six steps” - under this title Romain gave practical tips what to do in

order to improve security, again without slides in Indico, and without remote transmission. Romain explained that there will always be compromised accounts, hence root escalation must be made as difficult as possible. That's why it is important to keep up to date with security patches. That doesn't protect against 0-day or private exploits... to limit the damage, disable LKM loading at run time, and run rpmverify (the latter will detect most user-space root kits). Once somebody obtained root access – how did they get there? A typical attack vector is now via stolen credentials (SSH) from infected windows desktops. Traceability is the most important asset. To that end, use remote syslog, and safely archive your logs for a long time (beware of legal requirements!), making sure that data is accessible only to authorised staff. Also, keep a record of user actions – shell history, including commands and options. Accounting information is insufficient. In addition, keep a detailed record of network traffic. (CERN has developed a kernel module to this end; auditd is a good alternative.) In addition, rely on relevant experts, stand-alone incidents are history! Then Romain covered the CERN IDS; whenever there is a login from an unknown location, the user gets mailed. Initially this lead to discovering a substantial number of compromised accounts. Finally he gave an update of the Ebury attack: it is ssh malware found on many systems (rpmverify will detect it). The tool sniffs incoming and outgoing connections, and grabs passwords as well as ssh keys that get sent, in a rather convoluted way, to a collector. The objective is obviously to build a large botnet for profit purposes. The follow-up is extremely difficult due to the potential scale problem, and as it requires full collaboration of the users concerned.

## *Basic IT services*

## Using control system tools for operations and debugging (Gabriele Carcassi / UMich)

Gabriele first introduced himself – he was at BNL from 2002 to 2012, having been involved with the synchrotron light source since 2008. In 2012 he joined the University of Michigan, where he worked on the monitoring of the IT/grid operations. He made a first attempt by taking out data from Ganglia, feeding them into rrd and making summary plots and correlations. He found that this approach was not flexible enough – the queries must be known in advance. His second attempt was based on control system studio – an integrated system for control systems operations, which provides for accessing real-time data as well as creating and deploying user interfaces. In order to develop the latter, no programming is required. He used this system to create a prototype obtaining data from Ganglia as well as HTCondor. This prototype has successfully been demonstrated.

## Logstash and ElasticSearch deployment at GSI (Matteo Dessalvi / GSI)

Matteo explained that this is an update of a presentation he gave at HEPiX in spring 2013. For monitoring, GSI uses Nagios/Icinga, Netdisco, collectd, SNMP/MRTG/Torrus. The weak point used to be log analysis; they have started to use logstash, a tool written in JRuby (distributed as a single .jar file) that provides hooks for inputs, filters, and outputs. Use cases at GSI include access logs from ssh servers, apache access logs, GridEngine accounting, network device logs, Postfix MTA logs, … Matteo gave a demo highlighting the flexibility and ease-of-use of the tool, walking the audience through some complete records. He then explained how they are using logstash for analysing ZFS logs. Concerning visualisation, they have moved from Kibana 2 to Kibana 3 that is completely written in HTML and Javascript (earlier versions were written in Ruby).They have introduced Graphite and statsd. Finally, Matteo explained that ElasticSearch is not meant to be a long-term data repository tool; in addition, security is still a significant issue, and the Java virtual machine needs careful tuning.

## Migration from ELFms to Agile Infrastructure (Vitor Guveia / CERN)

Vitor started explaining that the configuration management services cover the tools, processes and infrastructure for service managers and system administrators. The previous system, Quattor, administers 8'000 machines managed by 260 persons. The system is however not scalable and dynamic enough to cope with a remote data centre and an increasing usage of virtualisation; in addition maintaining home-grown tools is very expensive, as there is no large active community outside CERN. The Agile Infrastructure (AI) project addresses these issues, revamping server provisioning, configuration, and monitoring. The configuration services are based on open-source tools: Puppet, Foreman and Git. For the migration from Quattor to Puppet, it was decided not to provide any automatic translation, but to implement the configuration in the new system from scratch. The new system was developed with active participation of major users according to the Agile guideline “release early, release often”. Regular meetings and training sessions have been held. One of the results was that the process needed to be reviewed, and the workflow for configuration changes was changed significantly. Documentation has been a significant issue, keeping it up to date was a real challenge given the rapid pace of changes. The system is hence maturing as it is being used. The users have in general been appreciative of this approach; the use of off-the-shelf tools has proved the right choice. Currently 4'000 hosts are being managed with the new tool suite.

## Managing a large heterogeneous environment with Puppet (Edward Simmonds, Tyler Parsons / FNAL)

The speakers represent the Fermilab Experiments Facilities (FEF) Department faced with Linux server and workstation management and Scientific Linux development. It was the first unit at Fermilab to deploy Puppet; their installation went into production already in 2010. Since then, a number of other departments have deployed Puppet as well. Puppet is a very actively developed toolkit with frequent releases (and changes). FEF now run 2'127 Puppet agents; the configuration uses version control. They use 27 custom facts, 134 modules with 1602 classes and 136 definitions. The details of node configurations in Puppet were presented. Tyler then emphasised the importance of some of the Puppet features such as determining if a particular resource is defined (beware of pitfalls – the outcome of an inconsistent usage may be unpredictable...), various ways of managing files (there are different ways to achieve the same goal), and the pros and cons of using Augias, a configuration editing tool. Finally he covered the area of package installation, where they have written their own custom package provider.

## Puppet at USCMS-T1 (Tim Skirvin / FNAL)

Tim started by stating that Puppet is not the only choice for configuration management, but a pretty good one. Both Puppet and Ruby are moving targets; PuppetLabs have the usual startup choice between revenue and market share that they try to address by enterprise additions that are not useful for Fermilab. At Fermilab, there are lots of teams with separate infrastructures and many commonalities. The USCMS-T1 features three major classes of hosts: 800 workers, 300 storage hosts, 100 server hosts. Rocks is used for system installation, most systems run SLF5. With the upgrade to SLF6, they will move to Cobbler. They had trainers come on site, and then started on the 'cattle' style nodes. Over time they have developed a number of 'best practices' (see links in slides). They do not put scripts / code nor secret data into the main code reponsitory. A Puppet user group has been created at Fermilab, maintaining e-mail lists and holding regular meetings. He then described how they use git to implement their workflow, and reported on very positive experience with Hiera that, among other things, manages the secret information. They use Puppet roles and profiles, again with good experience. Their changes are fed back into Puppet Forge so that the community can use them. Locally they have written modules for HTCondor, EOS, dCache, CVMFS, and Puppet itself. He then described Cobbler and how they use it. Upcoming challenges

include more and better internal documentation, scaling to 1'500 nodes, retire the Rocks server, tracking Ruby and Puppet, and finding the right balance between Puppet and Cobbler.

## Unified communication and IP telephony (Fernando Moreno Pascual / CERN)

Fernando started by describing the problem: collaboration tools, messaging systems, PABX systems, GSM services started out all separately. Microsoft Lync is a potential remedy to this that has been chosen at CERN. He described the main features: indicating availabilities, providing instant messaging, voice mail on Exchange, desktop and presentation sharing, connection from any internet connection as if the user was in his office. They have chosen Lync because it is compliant with Windows and Mac and supports mobility. There was competence in CERN-IT, and because of no roaming charges, it was cost-effective. Currently up to 8'500 users and up to 400 simultaneous calls can be handled. He then described how Lync had been deployed in the existing CERN environment. Lync has allowed to replace the “classical” IP phones deployed at CERN by those providing unified communication features such as access to the CERN address book, information about the contacts, management of incoming calls etc. With an additional USB connection to a Windows PC or Mac, the phone synchronises. Voice mail is saved in the Exchange inbox (as an MP3 attachment) and can hence be consulted whenever the user has access to e-mail. User feedback by the 130 pilot users has been very positive; in particular the collaborative features such as whiteboard and desktop sharing have been much appreciated. Federation with other institutes allow for entirely free calls; Lync 2013 will provide total integration with Skype and an XMPP gateway as well as VoIP on portable phones.

# *Thursday 31 October (Halloween!)*

# *Storage and file systems*

## Hard disk drive – reliability overview (Amit Chattopadyay / Western Digital Corp.)

Amit started by stating what the current issues with HDD field reliability are: HDD complexity leads to a multitude of different failure modes, the head-disk interface means macroscopic objects at molecular dimensions, the reliability needs to be re-modelled. The core problem in HDD reliability is to avoid an unmanageable catastrophy midway or beyond through a product's warranty, and to allow for a predictive analysis of failures. The large increase in areal density seen over decades was largely due to reducing the head-disk spacing that is now at 1...2 nm and hence molecular dimensions. Now some 70% of all HD failures are due to low-clearance operation. The small clearance also requires its standard deviation to be controlled to the 0.35 nm level – the diameter of a carbon atom! Modelling HD reliability is a highly complex task with new challenges for every generation of drives. Amit then explained that reliability was often expressed as a function of time (MTTF for example), but that this is the wrong parametrisation, as failure depends a lot on usage patterns. Temperature plays a very important role – the failure rate doubles for each 15 deg increase in temperature. Hence a model was used over many years that accounts for power-on hours, duty cycle, and temperature. However, duty cycle is not well defined – even a claimed 100% usage does not mean that data are read and written with maximum speed; workload in terms of sectors read or written are a much better parameter, it is directly proportional to the data volume transferred. Typical workloads range from 55 TB/year (desktop) to 800 TB/year (Enterprise XE drives). There is an explanation for the dependence on workload – inactive heads fly at 10 nm over the platter, only when the head becomes active thermal expansion reduces the distance to 1...2 nm. Amit showed results of studies that prove that “Mean Petabyte to failure” leads to predictions that are

largely independent from the workload. Amit then turned to the question how HDD failure can be predicted, and showed that their drives are equipped with sensors that allow for such predictions. These predictions, also known as “virtual failures”, are an important means to select and qualify materials and design changes. Collecting these metrics in the field is hence very interesting; they could complement or even replace SMART counters that have proven not to be very useful for HDD failure prediction. While they are ready with the data collection part, they still need to work on optimising the health monitor algorithms.

## HEPiX bit preservation working group (German Cancio / CERN)

German explained what the background of the group is, referring to Jamie Shiers' DPHEP presentation at the Bologna meeting. The proposed mandate is in short collecting and sharing knowledge on bit preservation, provide advice to DPHEP, and provide recommendations for sustained data archival. German then reported on the survey that had been conducted with 20 large sites in terms of total archived volumes, number of VOs supported, age of archive data, the expected retention period (only two sites have a data expiration policy, only four sites confirm that budget is secured for the complete preservation lifetime), existence of a formal agreement (SLA or QoS agreement), fraction of data not accessed over the last 12 months, tools used for archived data (most sites do not have any dedicated tool for archiving, but use the standard MSS solution), availability (online or nearline), accessibility via standard POSIX or HEP protocols, observed data losses (due to a variety of reasons), archive protection and auditing (checksums, file-level replication for important data, media scans), procedures in case of data loss, archive migration (some sites require active user confirmation to migrate, otherwise data expire; two sites will not be able to migrate because of lack of funding). German finished by summarising the survey results and giving an outlook – the group intends to extend the survey to other communities, collect, document and share best practices and tackle automation of data recovery for distributed archives. Volunteers are highly welcome to contribute.

## CMS: Storage life and times (Lisa Giacchetti / FNAL)

Lisa started by listing the requirements: The Tier-1 holds about 11 PB on disk and 24 PB on tape, the pledges for 2015 are 12 PB and 32 PB, respectively, accessible via xrootd, srm, phedex. In addition, there is the USCMS LPC (LHC Physics Centre), similar to a Tier-3. The challenges they have been facing are heavy random access (scale up does not work, scale out does but is expensive), combining small (home) areas and larger data storage areas on the same storage instance led to performance issues. They are running an HDS HNas (formerly Bluearc) comprising 300 TB of disk; it holds home areas, per-user data areas with quota, per-physics group data areas and scratch space without quota. The system had a history of poor performance under load. Eventually the filer was mounted read-only on the worker nodes. Then Lisa described the dCache system currently holding 13.5 PB of data on disk, using an Enstore tape backend. They also run a Lustre instance with about 150 TB of disk space, which has created a significant management load, and have established an EOS testbed of about 600 TB accessible via fuse mounts and xrootd. They then started a formal storage evaluation looking at the same time at home/data areas and a project to separate dCache disk from tape, aiming to reduce the number of systems, provide better performance and accessibility, and reduce costs. The goal was to end up with three areas: a dCache-like area for Tier-1 production, a POSIX-compliant online area for LPC analysis, and a very reliable area for home areas. Products considered were dCache 2.2.7, Hitachi HNas, Lustre 1.8.6, EOS 0.2.29, Isilon, Hadoop 2.0, Nexsan 5000, Overland SnapScale, Netapp, GPFS, many of which have been tested in the Fermilab environment. She described how the tests were done for the three respective use cases and their results. EOS performed very well; dCache is not entirely POSIX compliant. They decided to split out the requirements for home areas from the data areas. Their decision was for dCache 2.2 for nearline, EOS for online, and CCDs (core computing division's) HNas for the

home area. She then described the status of implementing these decisions. Finally, Lisa stated her surprise about the numerous vendors to offer a Linux NFS solution without adequate Linux knowledge, and the small list of vendors that were willing to provide test equipment.

## OpenAFS and IPv6: Follow-up on the Bologna discussions (Arne Wiebalck / CERN)

Arne reminded the audience what the discussion at the spring 2013 HEPiX meeting was about, mentioning the various options such as rely on dual-stack, fund the IPv6 implementation, implement it ourselves, or look for some AFS alternative supporting IPv6. He then described the implications of a dual-stack implementation with non-routed IPv4 addresses and measurements of the real load as far as CERN is concerned. Indeed, the access from CERN to other AFS cells is at a very low level, and does probably not represent a real use case or even show-stopper. Concerning the option of funding the implementation, no other customer appears to be asking for IPv6 support; the implementation is now expected to take 4...6 months (compared to 18 months as estimated by the Bologna meeting). The two companies, YFS and Sine Nomine, would take very different approaches, as the YFS code base is closed source, hence an IPv6 implementation would require a licence subscription. Sine nomine would implement the functionality in the OpenAFS open source code, with uncertain prospects for getting the code into the core OpenAFS base, with a one-off payment. The survey Peter and he have run was answered by 14 sites; there was no clear view on whether the lack of IPv6 support is a problem. There is only very limited willingness to contribute financially to such development. To summarise, Arne said that a commercial implementation is feasible, but the community interest is limited, and a dual-stack approach appears to be an acceptable way out.

## OpenAFS status 2013 – Roller Coaster of Thrills (Derek Brashear / The OpenAFS project)

Derek explained that the last 12 months have been exciting for OpenAFS: 1.6.0 was released on 1st September 2011, the “rx ping” issue was found at CERN and DESY; there was also some corruption issues. 1.6.1, released in April 2012, fixed these problems that he described in a little more detail. OpenAFS 1.6.2 brought a number of further improvements. Concerning Windows, XP will be de-supported, Windows 8 will be supported (only on x86); NTFS symlink compatibility has been added. And then... a security vulnerability was found that is not really an OpenAFS one, but rather DES, but that is the only protocol supported by AFS. The effect was that through a brute-force attack, a general service key could be obtained by attackers. OpenAFS 1.6.5 brought two mitigations to rxkad since rxgk has not been completed yet: rxkad-k5 (backward-compatible) and rxkad-kdf (more secure, but harder to deploy). The former has the issue that the KDC must still support DES keys, but the service key is no longer vulnerable. Rxgk, which will resolve these problems properly, is still being worked on. Beware that whenever you run a server of level less than 1.6.5, your installation is vulnerable! Then Derek described the changes in 1.6.6 and further upcoming work, mostly around the MacOS version. Finally he described the status of the OpenAFS foundation; it has finally been created, the effort to obtain non-profit status is on-going. There is an agreement with IBM for the use of the name AFS. The business plan still needs to be defined.

## YourFileSystem (Jeffrey Altman / YourFileSystem Inc.)

Jeff started with a reminder of the discussion at HEPiX Fall 2007 in Saint Louis, where the OpenAFS gatekeepers presented a roadmap, and a community wish list was established, with some discussion about possible funding of HEP-wanted features. Six years later, YFS 1.0 is there... implementing dual protocol stack servers and clients, namely AFS for compatibility with OpenAFS and IBM, and YFS for enhanced functionality. The improvements they focused on include security,

networking, scalability, ease-of-use, and code quality. Jeff then described some of the enhancements in detail, and discussed interoperability with OpenAFS and the YFS road map.

## Small file aggregation in Enstore (Alexander Moibenko / FNAL)

Alexander gave a short introduction about Enstore and then mentioned the problem to be addressed: small files that are inefficient because of tape marks. Some tape drives offer solutions, but there is no common approach across all tape drive in use. The definition what is considered a small file depends; currently they consider files below 2 GB as being small. As no standard solution exists, they have worked on an aggregation solution: putting multiple small files into larger containers, with adjustable definitions of small and large. Of course there is no general solution, because only the users (e.g. an experiment) know about the read patters of their files. Alexander then described the policy entry used in their system. For the implementation they required the file caching functionality to be in Enstore transparent for users not affected by the aggregation. He then described the chosen solution, which implies making the whole container available for reading in a single go. They plan to implement transparent aggregation during media migration of files.

## Status of dCache (Patrick Fuhrmann / DESY)

Patrick started by recalling the most important facts about dCache. The project structure has slightly changed; there are contributions by Fermilab, Copenhagen (NDGF), and (as of recently) the Hochschule fur Technik und Wirtschaft of Berlin. There are links to OSG and USCMS, NordUnet, EGI/EMI, LSDMA. He then mentioned a number of new deployments and explained that the software continues to be distributed within the EGI UMD. He described the support for NFS 4.1/pNFS; deployment requires SL6, which is available now. They have tested the NFS 4.1 setup with a cluster of 60 worker nodes, on which CMS jobs were forced to use NFS; the results are being analysed, but in general things look good. A use case giving problems to NFS has been identified and is being addressed. The team is involved in xrootd federations as well as in federations built on HTTP/WebDAV. CMS are separating their disk and tape layers. Additional requirements have come from the photon science community. They attack the small file problem at the dCache level as well; this development in progress is currently available only at DESY, but not in the dCache distribution. Patrick described that mutable files cannot be aggregated nor replicated. They are working on extending gPlazma to support web-based authentication, but are meeting some legal constraints. They are also implementing CDMI for data transfers and for meta data storage, a feature the photon science community has asked for. Ongoing evaluations include running dCache pools on DDN storage boxes, and providing a dCache front-end to CEPH. Then Patrick turned to the scientific storage cloud, which is about a dcache.org service rather than the software. The motivation is to get students involved, in return they get “unlimited” storage space and a degree. This scientific cloud storage differs from other cloud storage by a number of key elements such as Kerberos authentification etc. The service is installed and in use, using registration through the DESY infrastructure. Sharing will be addressed next, as will more authentication methods, more protocols, user-determined data retention etc.

## Building an organic storage service at CERN with CEPH (Arne Wiebalck / CERN)

Arne explained that the talk is an update of what he presented at the spring 2013 HEPiX meeting, gave a brief overview of the CEPH architecture, and then recalled that the first test had been run on a small cluster of retired storage servers; setting up the system was simple and achieved quickly, all initial tests were passed. Hence a larger-scale prototype was set up with 12 x 4 storage servers for a total of 3 PB. The deployment is fully puppetised, using a community CEPH module with fully automated machine commissioning as well as Mcollective for bulk operations on the servers. The

cluster is configured as 11 data pools with three replicas each. Arne presented benchmarks that showed excellent performance. Currently the instance is used for Openstack Glance images; Openstack Cinder volumes will come soon, AFS and NFS, DPM, OwnCloud, and Zenodo are all testing. The Openstack-CEPH testing went well except for the missing cell support for Cinder in Grizzly. Issues are currently significant write latency (50 ms) – in-memory OSDs achieve 1 ms -, and the limit of processes due to the number of OSDs. To summarise, the killer application for CEPH would be to build upon it a general-purpose network file system – cephfs has been released and has been available for quite some time, but has not been regarded as production quality yet by its developers who nevertheless strongly encourage testing, which CERN will start doing soon.

## Experience with CEPH at the US ATLAS midwest Tier 2 (Lincoln Bryant / U Chicago)

Lincoln started with an overview of the facility split over three universities. They were interested in CEPH because of the full POSIX implementation of its file system, contrary to HDFS. The installation was easy, and the performance was according to expectations in view of the hardware used. They then bought some dedicated hardware and set up for a CEPH benchmark. He showed results of the internal RADOS testing, demonstrating good scaling with the number of threads on reading; writing was limited by SSDs (one SSD is recommended per 5...6 spinning drives), the removal of which allowed for good scaling once more. For the further tests, XFS was used. The file size did not appear to matter much in terms of transfer speed. Lincoln then described projects that are being considered for CEPH: FAX, stash, Openstack... A number of issues has been found with CEPHfs: the metadata server is currently not scalable, there is no quota support, and using a bleeding-edge kernel is required to mount CEPHfs. They may convert a large existing HDFS pool to CEPH, and may wish to consider adding a BeStMan gateway. In summary, their experience is very positive, hence they will certainly go ahead with it.

## *End-user IT services and operating systems*

## Building packages at CERN – from sources to RPMs (Thomas Oulevey / CERN)

Thomas explained that there are two main use cases, updates to Scientific Linux as well as providing flexibility to the Agile Infrastructure team. Concerning Scientific Linux CERN, an old script-based system had been used together with rpmbuild, requiring a number of manual steps for the QA. He then described the Agile Infrastructure use case; both call for automatisation. The solution consists of three elements: Mock (creates chroots and buildids packages therein), Koji (software building), Mash (query Koji, create repositories for specific targets). Thomas showed the Web interface as well as the CLI, and presented statistics of the usage of the system so far. Then he described the detailed RPM creation workflow; everything that can possibly be automatised has been. He then showed how repositories are created and RPMs get frozen. User feed back has been rather positive, even though the learning curve is a little steep. There are some limitations, including the single namespace, that Thomas described in detail. One of the challenges is to provide support, in particular as many users don't appear to read the documentation; in addition, existing RPMs may not have any sources. Thomas listed a number of nice “side effects” - he met a large number of packagers, observed a better quality of RPMs, traceability, benefits of centralised repositories, and more compiled add-ons. For the future, he intends a better policy integration with Active Directory, and enable SSO; SLC7, repository management tools, full puppetisation etc. are in the pipeline.

## Update on Scientific Linux (Connie Sieh, Patrick Riehecky, Bonnie King / FNAL)

Patrick started by explaining that updateinfo.xml has been enabled for all SL6 releases. A successful collaboration with KISTI has been established. The ftp service was affected by an issue with the underlying NFS server. OpenAFS 1.4 packages now include the back-ported Kerberos security patch described in the OpenAFS talk during this workshop. Then he described what the constraints are of MySQL in SLC5. Software collections are now available for SL6 (not SL5), including new software and updated compilers and other programs. Also available is devtoolset; the new Eclipse is however quite tricky because of many levels of dependencies. SL has joined the UEFI Forum in preparation for supporting UEFI secure boot in SL7; there will be no support for RHEL6 (and hence SL6). Then Patrick recalled the handling of errata, security and other. Current work focusses on the SL7 build system. The scientificlinux.org site is beginning to show its age; there is lots of content for SL3, not much on SL5 and/or SL6, and the site runs on aging hardware. Patrick presented a possible design of a replacement service. Concerning SL6, RHEL6.5 beta was released on 08 October 2013; they plan to release SL6.5 as soon as the final release of RHEL is out. This minor version is expected to feature better PTP (precision time protocol) support, a new openssl, native hooks for GlusterFS and support for up to 1 TB RAM. Then Patrick turned to SL7/RHEL7; a beta version of the latter is expected later this year or early in 2014. It is expected to look somewhat like Fedora 19 and hence come with systemd, NetworkManager, Gnome 3 in classic mode, XFS as preferred file system over ext4, and possibly the new installer (Anaconda). Fedora 19 is still suffering from a number of important bugs that Patrick listed. Then he mentioned a number of conversation points: SL7 point releases, software collections, ZFS packages in SL6. He reminded the community about upstream community projects that may be of interest – RDO, Ovirt, GlusterFS. He then reminded everybody that fixes (as well as bug reports) must be submitted upstream.

# *Friday 01 November 2013*

# *Grids, clouds and virtualisation*

## CERN cloud status (Thomas Oulevey / CERN)

Thomas started with a short introduction into Openstack and then explained that the CERN deployment is currently based on Grizzly (Havanna has just been released), taking advantage from the RDO packaging. Community Puppet modules from Stackforge are being used. Currently the facility consists of some 700 hypervisors, to which about 100 are added per week. Every registered user is entitled to a personal project with 10 cores. For the support, standard CERN mechanisms are being used. He then explained the Grizzly components in use at CERN, covering in detail and showing examples of Horizon (dashboard), Nova cells exploited at CERN and Budapest, Glance, Ceilometer. They tend to run Windows virtual machines under Windows hypervisors and Hyper-V, even though Hyper-V in Openstack is still missing some desired features. Live migration and block storage (Cinder) will be introduced with the next release, Havana. Thomas showed fio test results for accessing block storage (CEPH) from Openstack VMs, and then explained how the service is monitored, using CERN-standard building blocks and infrastructure described in an earlier presentation at this workshop. The challenges ahead include an in-place migration from Grizzly to Havana, external authentication via Kerberos and/or X.509, Keystone domains (which will allow for delegating management responsibilities), and a federated cloud. The latter is being investigated together with Rackspace in the framework of a CERN Openlab project.

## IN2P3-CC cloud computing status (Mattieu Puel / CEA-IRFU)

Mattieu explained that the target use cases are test and development machines that can be procured in a self-service way, infrastructure services, computing services, and public cloud services. With the latter, they are addressing companies in the private sector, and achieve significantly better pricing than commercial providers such as Rackspace, Amazon or DSI. Like CERN's, the cloud is based on Openstack; it currently comprises some 1'000 cores. They are using Keystone, Nova, Glance, Horizon and Cinder, and are evaluating Swift. Dedicated tests on Neutron and Ceilometer are to follow. Mattieu gave details of their deployment for the various use cases, and then listed a number of features they have less appreciated: the Nova VLAN mode comes with serious limitations, updates are non-trivial, the documentation is sparse and not always up to date, and in general the system is rather complex and requires expert knowledge. On the positive side, they found the system reliable and full of features and configuration options; it is highly available and scalable; patches are easy to prepare and move upstream. In addition, they have liked being part of a huge dynamic community. Their plans for the future include (even on the short term!) establishing and testing a federation with other French clouds that are Openstack, Stratuslab or OpenNebula-based; in the longer run they, like CERN, want to virtualise everything that can potentially be virtualised.

## FermiCloud update (Gerard Bernabeu / FNAL)

Gerard started with an overview of FermiCloud that started in 2009, building on positive experience with FermiGrid and using some common elements, and then explained the phases of the project; they are currently in phase 4. It is running on 23 nodes with 16 cores each and features a private network bridged into FCC and GCC, a segmented IB network for 10 and 13 nodes, respectively (the VMs access IB via SRIOV, and has access to site storage services such as SAN, Bluarc, and dCache/Enstore. Typical use cases include VMs on the public network, a private cluster with a gateway VM to the public network, and VMs for storage (Gerard explicitly mentioned MySQL, PostgreSQL, Lustre/Hadoop). They have had Grid jobs bursting into FermiCloud transparently. Scalability is currently limited by public IPv4 addresses and the available RAM on the hypervisors. They use OpenNebula as orchestrator, which emerged as the winner in a competitive evaluation at the beginning of the project; they are currently preparing the migration to release 4.0. For fault tolerance, they have split the installation between two buildings, and taken steps to ensure maximum availability even if one building goes down. Then he described the collaboration with KISTI on integration and commissioning of a prototype federated cloud for scientific workflows. They have developed a mechanism to suspend VMs that are idle or under a working time agreement in favour of pending high-priority round the clock work; the mechanism consists of a probe running on the VM, and a number of modifications to OpenNebula. They will soon add six more hypervisors as well as PaaS and SaaS services. In the medium term, they will review their storage solutions.

## PanDA beyond ATLAS: Workload management for data intensive science (Jaroslava Schovankova / BNL)

Jaroslava presented PanDA (Production and distributed analysis system) that has been developed for ATLAS, but is now meeting more general interest as a workload manager for data-intensive sciences. Its first production release dates back to 2005 in US-ATLAS; in 2008 it was adopted as system for the whole ATLAS collaboration. During 2012 and 2013, AMS, CMS, LSST and Alice have used or tested the system to various extents. In ATLAS, it has very successfully managed the challenges of Run 1. Since September 2012 they have DoE funding for turning it into a general tool for data-intensive sciences not limited to HEP and not confined to a Grid infrastructure; the network is being integrated as a resource. The project is scheduled to take three years, and is split up into four work packages: factorise the core, extend the scope, leverage intelligent networks, improve

usability and monitoring. She then discussed details of the PanDA server and monitor. The pilot version has already been refactored in the context of the common analysis framework project; experiment-specifics are loaded via plugins. In a common project with Google, they have added support for Google Compute Engine, which has been used for production ATLAS jobs already.

## Experience with dynamically provisioned worker nodes at RAL Tier-1 (Andrew Lahiff / RAL)

Andrew explained that the cloud of the scientific computing department is a prototype intended to gain practical experience; it is using Stratuslab which in turn is based on OpenNebula. They are using iSCSI storage for the images, which can be instantiated as VMs in less than 30 seconds. Being a prototype, all hardware is rather old and out-of-warranty (18 TB storage, 100 hypervisors). The cloud is available to their staff on a self-service basis; even though there are 30 active users, it is not fully used, while the batch system is fully occupied, with many jobs in waiting state. Hence they investigated whether virtual batch workers can be instantiated in the cloud. This comprises provisioning of virtual machines as required, adding new worker nodes to the batch system, and deciding when to remove virtual worker nodes. Their recent migration to HTCondor has significantly helped, as adding and removing workers is much simpler than in most other batch systems. They use images provided by their Tier-1 admins; newly created worker nodes undergo the same health check as physical nodes. Provisioning worker nodes uses the HTCondor power management features, instantiating and suspending/shutting down virtual worker nodes as needed; the latter is controlled by the worker node itself. Andrew then reported on performance measurements: Both HS06 and a real CMS application run a few percent slower. Accessing network and storage shows significantly larger differences between physical and virtual worker nodes, even though the results are preliminary, and no tuning has yet been done. They ran a test with 11'000 real jobs representing all LHC and other VOs. Finally Andrew discussed monitoring, comparing the situation of physical and virtual workers; Nagios, their system of choice, does not easily support dynamic resources.

## CVMFS – beyond LHC computing (Ian Collier / STFC-RAL)

Ian started with a short introduction into CVMFS, highlighting its advantages over other software distribution mechanisms. At RAL they were one of the first to start deployment in 2010; in 2011 they set up Stratum-1 replica for LHC VOs (the Stratum-0 is at CERN); since 2012 they provide Stratum-0 services to a number of user groups, and are now leading an EGI task force to establish a network of Stratum-0 and Stratum-1 servers across EGI sites. He then described the WLCG deployment – clients connect to one of the Stratum-1s and can fail over to another. Their Stratum-0 services for enmr.eu, a large life science community, supports over 570 registered users; software is distributed to more than 25 sites; adding more sites is really easy, straightforward and fast. As an alternative to the classical method of releasing software into the Stratum 0, they have developed a Web application for uploads, using authentication via X.509 certificates.

## OpenShift on your own cloud (Troy Dawson / Red Hat)

Troy reminded the audience that OpenShift is Red Hat's PaaS (Platform-as-a-Service) for applications in the cloud. It is now available in three flavours: online where Red Hat provides everything; origin, where you build from source and put on your infrastructure; enterprise, where Red Hat provides the software and the user takes care of the infrastructure. Origin and enterprise run on top of major commercial clouds, Openstack, VMWare, RHEV, bare metal. He then introduced the key terms broker (host managing state, DNS and authentication), node, app, gear (allocation of fixed resources), and cartridge (framework to build applications). SELinux policies are used to securely subdivide the node instances; gears live in these subdivisions. Cartridges

include Java, PHP, Python, Ruby; MySQL, PostgreSQL; custom cartridges are supported. Troy then described the possibilities of scaling. The communication with external clients goes via a REST API; the broker communicates with the nodes through messaging services; Mcollective is being used. He finished by describing the complete workflow of an HTTPS request.

### *Miscellaneous*

## Workshop wrap-up (Helge Meinhard / CERN)

Helge started with some statistics that are partly reflected in the introduction to this report. 115 participants is definitely a record for a north-American meeting; participation both in terms of attendees and institutes was well balanced between north-America and Europe. A total of 65 contributions took 26 hours to present, leading to a rich and very interesting program. Among the trends identified from the presentations, Helge mentioned identity federations, CEPH, private clouds reaching full production status, HTCondor, Puppet, and log file analysis. He then covered the working groups: the benchmarking group is being revived, as a new release of SPECcpu will be released in October 2014, hence work with the experiments needs to start soon. The Web site needs attention; help by DESY is much appreciated. Sandy Philpott announced that a new North-American co-chair needs to be elected in autumn 2014. The spring 2014 meeting was confirmed to be held on May 19 – 23 in Annecy, France; for autumn 2014, the University of Nebraska is seriously interested. Proposals are always welcome. Helge finished with thanks to all participants, speakers, track convenors and chairs, sponsors, and in particular the local organising committee.