

HEPiX CERN

5th-9th May, 2008

Introduction

For the first time since 1992 (which was only the second HEPiX meeting ever), HEPiX met at CERN for the Spring 2008 meeting. Some 95 people registered and there was a record 65 submitted talks, completing filling the available time slots. The meeting was organized by a small team of IT staff led by Helge Meinhard. The format was as usual, the days split into blocks by topic, with the addition this meeting that each day was started by a plenary talk on some aspect of LHC or the computing for it. The overheads of the talks can be found on the [web](#) and the following are my personal notes so don't blame the speaker if you find an error.

Some highlights:

- Having the meeting at CERN gave many young and not so young IT staff the opportunity to present their own work and all of them gave good, informative presentations, some very good.
- Once again many talks on new data centres (RAL, BNL, St. Louis) or upgraded ones (NIKHEF, and SLAC if you include the addition of SUN Black Boxes). Another coincidence was the repeated tales of electrical problems, leading to fires and burnt equipment in three cases.
- If previously Lustre appeared an interesting file system, according to a large number of presentations this week, it is now the file system of choice across much of the HEP world. This was confirmed by the summary of serious testing performed by the HEPiX Storage Working Group for whom Lustre is now the recommendation for a cluster file system.
- Scientific Linux futures were discussed in detail and a number of decisions were taken regarding support lifetimes and future directions (reported fully below)
- There was an invited Google talk for which at least 50 people came in specially from IT and PH. The speaker gave a most interesting review of data storage at Google; the problems and issues were familiar but the numbers were a different scale to ours.
- This, the best-attended talk of the week, was followed by the most entertaining talk of the week, by Charles Curran on tapes. I recommend the overheads but they will lack the flair with which he delivered them. We should have recorded this talk.
- A growing number of Labs are adopting or adding new sciences and it may or may not be surprising that many are meeting again common problems and HEPiX is an excellent forum for sharing these experiences as it was, and is, for the pure HEP labs. This is not to say that we should consider expanding its scope (although two of the world's leading Genome Sequencing Centres were present for the second consecutive time), but rather than those labs who are now not solely HEP can still find good reasons for continuing to support and attend HEPiX.
- Organisational details: the room chosen, the Council Chamber, was the correct size (except for the Google talk); only two crashes of the projector, one of which was solved by a power cycle and the second by A/V support; all projections were direct from Indico except for one. Big

thanks to the support team (Helge, Yasemin, Fabienne, Kristina, Vlado and AGS). Attendance matched our target (93) although US labs are severely affected by DoE budget cuts.

- Next meetings – ASGC in Taiwan in October. Next year (probable): Nordic host in the Spring (May?), NERSC in October.

Plenary Talks

LHC Status Report (Jos Engelen)

Brief review of the physics of the LHC and the experiments. Final cool down has started and is expected to be finished by July. In parallel, hardware commissioning of the magnets is ongoing. Full cool down takes about 6 weeks plus about 3 weeks for stabilization. When the whole ring is cold, there will be about 130 tonnes of helium circulating. Jos described the status of commissioning in some detail for various sectors before turning to the testing of the various experiments. The current plan is to restrict the first run to 5 TEV beams. Finally, he turned to possible future scenarios – for example Super-LHC, with a luminosity upgrade by replacing the inner triplets. He also briefly listed some options for upgrading various parts of the accelerator chain.

WLCG Update (Ian Bird)

After a brief review of the structure of the project, Ian showed a plot of recent CPU usage which demonstrated the spread of CPU cycles between the Tier 0 centre (CERN, 11%), Tier 1 centres (35%) and Tier 2 sites (54%). In 2007, WLCG ran over 44M jobs and is currently supporting over 300K jobs per day, around the level expected during LHC running. He explained the reasons for the Readiness Challenges (see later plenary talk) and noted that during phase 1, the data transfer rate via the LHC OPN had reached target levels for realistic production conditions. Another concern had been CASTOR and again this proved to be capable of sustaining the rates to be expected for data taking. One area which remains to be addressed is resource ramp-up whereby the procurement process is usually longer than expected and installation and making the systems ready for stable production running is often hit by unexpected problems, often not under the control of the sites (late delivery, failing systems, etc). Reliability and availability of sites is steadily improving but remains a worry, especially the mass storage systems at Tier 1 sites, and this had led to a continuing effort in monitoring tools, in particular visualization to improve understanding and automation. Finally he noted that more operations load must be decentralized and middleware teams must concentrate on support rather than development and he showed how the various phases of the EGEE project should evolve into a permanent European grid infrastructure, EGI.

Openlab (Sverre Jarp)

Sverre explained the structure of openlab and introduced the partners and contributors listing which projects they collaborated on. Among those covered were

- the gridmap tool for grid monitoring developed along with EDS
- grid scheduling survey with HP
- virtualization benchmarks with various hypervisors
- many individual collaborations with Oracle in using databases in the physics data chain such the optimisation of the rules affecting Oracle streams throughput and radical improvements in database access speeds by the PVSS tool used heavily in Controls
- a project (called CINBAD) to study anomalies in network behaviour with HP Procurve
- investigation of the use of multi-core systems in HEP codes and the promotion of multi-threading in collaboration with Intel
- runs dedicated to submitting performance ratings to the list of the Top500 computers in the world

Finally Sverre described the openlab summer student programme.

WLCG Common Computing Challenge (Jamie Shiers)

The goal of CCRC'08 is to understand where WLCG stands with respect to the needs of the experiments and to identify and fix any problems that remain. One known bottleneck for example is at the level of support personnel. A first phase of the challenge was held in February where not all the 2008 resources were in place and a second phase in May, targeted to use all the committed 2008 resources. However phase 2 has started and in the event all resources are not yet available and it is now envisioned to do a third round in July. Specific metrics and targets have been preset and are measured on a daily basis during the challenges. Generally things have worked reasonably well although, one of the largest remaining issues, communication still needs to be improved. But even simply the fact of scheduling these phases has focused the effort into getting production level code in place, support structures readied, etc. If the successes noted in phase 1 are repeated in phase 2, where the bar is set much higher, then WLCG will be well prepared for first LHC running.

EGEE Status and Plans (Bob Jones)

EGEE should deliver infrastructure operation, middleware and user support for grids. Take-up of grids in general is perhaps being retarded by a feeling in other disciplines that grids are targeted at physics only or mainly. So EGEE has dedicated resources to seeking out, encouraging and helping non-HEP communities to invest in grid activities and Bob showed a long list of Virtual Organisations which are inside or somehow associated with EGEE. Phase III of EGEE has just started with the goal of expanding the infrastructure and preparing for the migration from a project to a sustainable utility model based on National Grid Initiatives. Bob then summarized the current status and work going on inside the EGI design study which should define a blueprint for this sustainable infrastructure. This should be published in the summer and discussed over the autumn.

LHC Networking (David Foster)

The OPN (Optical Private Network) for LCG was conceived in 2004, based on 10Gb lines as a simple end to end network, CERN-centric out to the Tier 1 sites. Since then, additional links have been created between the Tier 1 sites. The links are provided by a variety of suppliers, some institutional

(GEANT for example), others commercial, and others NRENs. David then described the structure of some of these suppliers and how CERN links to them. The variety of suppliers only adds to the complication of operating and monitoring the high level service; for example, simply going to different suppliers for duplicated links (for redundancy and backup) may not exclude that both suppliers are using the same physical cable. In some cases beyond our control, backup paths are just not available and this showed up in backup tests in April. SA2 is the work package in EGEE which covers operational aspects of grid networking in EGEE and they are working on a number of topics within this. These include the interaction between the different actors involved with LHCOPN (service providers, users, operators, managers); monitoring and alarms; ticketing services and problem responsibility (who owns a given problem, avoid finger-pointing loops); etc. He then listed and illustrated some of the monitoring tools in use. David ended by predicting a major expansion in the demand for networking services, explaining a number of reasons for this.

Site Reports

GridKa – FZK has recently merged with the Uni Karlsruhe and created a combined computing centre – SCC - and this operates GridKa for grid computing. During the first half of 2008, over 500 new systems were added with nearly 7000 KSI2K units. As is becoming common, the tender was based on a request for a given number of SPEC units as measured by SPECint codes and the adjudication included premiums for power, space and network ports needed. For their next procurement later this year, early testing shows that power requirements are going down significantly. Of course disc capacity has increased in parallel with CPU but this is covered in a later talk.

LAL – Recent efforts have been dominated by the installation of 250 8-core systems, mainly as grid worker nodes. This was complicated by harmonics in the power circuits and the work on fixing this is delaying necessary upgrade of the computer centre cooling system. Other work includes upgrades to their network, both internal and connection to other GRIF¹ sites by a 10 Gb/s optical private network (GRIFOPN). Main activity in GRIF is in scaling up to 2000+ concurrent jobs; there is now a unified view of storage whereby all CEs will have a uniform view when GRIFOPN is fully installed. The next step at LAL is to replace their Tru64 cluster which is the cornerstone of many internal services. This will happen slowly, creating dedicated systems per service with failover for reliability, but this is proving non-trivial. Another major task is to replace their backup scheme, moving off Networker on Tru64, in place since 15 years, to a scheme based on Linux. Again, commissioning was difficult, including poor vendor support. Their distributed Lustre tests (CARRIOCAS, see previous meeting) continue and a new project has been accepted by the EU under FP7 for integration of institutional grids with desktop grids such as BOINC.

¹ GRIF – a grid between various sites in and around Paris – see previous HEPiX reports.

TRIUMF – additional 560TB of disc for ATLAS, nearly tripling the total installed and this should double next year, as should tape capacity. CPU capacity is also rising steadily. With mail traffic exploding, and mostly spam of course, they are adding mail filters to perform spam rejection earlier in the chain and he described these “milters” (mail filters) in some detail – see overheads. TRIUMF is now supporting Vista and Office 2007 for Windows desktops and SL5 on the Linux side, both servers and desktops. Most of TRIUMF uses Xen for virtualization except for ATLAS who use VMware for historical reasons. After looking at a number of tools they are investing in Bedework for shared calendaring across Windows and Linux systems.

NDGF – 10 production sites across the Nordic countries running a variety of Linux flavours. They operate a distributed dCache facility based on a central service in Copenhagen. In the process of setting up connection to LHCOPN but currently blocked in the approval cycle. NDGF has applications beyond HEP: for example in Finland, their Cray XT4 is being upgraded in a project to archive Finnish TV and radio programmes and the site in Umea, Sweden is adding a 5376 node IBM blade system for general computing for Swedish researchers. For this, plus some 432 systems for grid work, they are building a new computer room. The following talk described one NDGF site in particular – HIP in Helsinki.

RAL – the speaker started by updating his previous report on building a new computer centre: the structure is now in place and although internal fitting-out has not yet started it is on track for completion in the autumn of 2008 with migration of existing equipment timed to coincide with LHC shutdown in 1Q09. A major new UK government concern is the encryption of data stored on “mobile devices”, e.g, laptops, and this has led to a significant change in behavior by laptop owners with respect to what they store locally on their laptops. The Tier 1 centre at RAL is now accessible for jobs only via grid submission. The migration from dCache to CASTOR continues and the former is due to stop by end May although this may be slightly extended. They have added 182 9TB disc servers, a total of 1638TB additional capacity. They had 3 backplane burnouts on their new acquisition, one of which provoked a call-out of the local fire brigade! The vendor responded well and all faulty pieces have been replaced. The 2007/2008 CPU upgrade involves the addition of ~3000 KSI2K units. Finally, he presented a post-mortem of the RAL Centre’s first power interruption in over 3 years in early February. The restart effort is estimated to have cost 200 FTE hours spread over many days until the full range of services were up and stable. One interesting lesson was that the building access system, computer-controlled of course, refused to let staff back into the building (!) and staff already inside had to open access via fire exits.

CNAF – currently expanding their infrastructure to accommodate additional CPU, disc and tape capacity needed as an LCG Tier 1 site. On their farm, they are using Quattor for installation and Lemon for monitoring. He showed a graph showing that available computing power at CNAF had remained more or less constant for some years at 2000-3000 KSI2K units but was about to be tripled with their next acquisition. They remain skeptical about some aspects of CASTOR and are evaluating alternatives including a combination of IBM’s GPFS and TSM and CNAF’s StoRM and he offered some reasons for this choice (see also later talk on this).

GSI – their computer centre has reached maximum capacity and there are no funds for a new building so new racks are being installed in random empty spaces around the lab but these are often not properly cooled or powered and this cannot continue for long. More successful is the introduction into production of their Lustre cluster, already used by ALICE. He then reported on a very serious incident which burnt a backplane of a file server, eventually blamed on leak currents. But despite the vendor replacing 39 file servers suspected of having the same problem, there was a similar incident recently and more replacements are scheduled.

NIKHEF – There is a new name for the lab – FOM Nikhef - reflecting more emphasis on sub-atomic physics. Changes to their Windows environment, introduction of desktop search and WSUS had a negative effect on performance initially; and Vista use is growing. On the Linux side, SL5 is appearing (as opposed to SLC3 and SLC4 previously). For their grid service they are undertaking a tender operation supposed to be valid for 4 years. They have installed Indico locally and found the transition easy although work has been necessary with CERN to improve its stability.

PSI – PSI is a Tier 3 site for LCG covering CMS groups at ETHZ, Uni Zurich and PSI; they have 64 Intel processor cores this year and 75TB of SUN ZFS disc space and these should triple next year. They are transitioning from 1U dual-core Opteron systems to quad-core Xeon blade servers and are very happy with the change. They use Scientific Linux, currently moving from SL4 to SL5 and in parallel moving away from cfengine to configure systems, adopting puppet instead. On their accelerator beamlines, they use GPFS and get good performance from this – 300-400 MBps sustained aggregated throughput per beamline. For backup, they use a disc cache in front of LT04 devices. This allows the drives to be streamed at full speed from the disc cache and fast restores but they have concerns about the reliance on having good monitoring and requiring in-depth knowledge of the backup engine (EMC Networker).

BNL – installed a new SUN/STK silo for HPSS and migrated their existing HPSS file base to the new one with a certain amount of pain, largely caused by the old filebase containing a very large number of small files. They have moved away from Panasas (see previous HEPiX reports) for their NFS cluster to BlueArc which is thought to offer better performance. BNL offer a dCache service to their ATLAS users and the speaker described some issues with that and some optimizations performed. They are actively evaluating a number of advanced developments from different vendors to replace their existing SUN X4500 Thumper systems on which their dCache is currently based. The BNL network shows signs of strain and is being re-architected to improve performance.

SLAC – local developments are dominated by enforced layoffs due to budget cuts of which computing was hit by a 20% cut and the other cuts were shown on a slide and he showed the new structure of the computing department. Nevertheless, their acquisitions have continued steadily, older systems being relegated into their second SUN Black Box whose commissioning is due to complete this week. They have had generally good behavior with the SUN Black Boxes except for a recent incident when a chiller failed. Last meeting they reported problems with harmonics in their power supply. After tracing the source they installed a harmonics suppressor but the first one burned out and a second, larger

one, is now being acquired. They have finally switched off Kerberos 4 on May 1st and now only use Kerberos 5. As would be expected, the final K-4 users were the hardest to migrate and required considerable effort. Until now, SLAC has used 2 ticketing systems, Remedy and RT, but they have now they will converge on RT since Remedy support was one of the areas affected by the budget cuts.

Operation Systems

Virtualisation on Windows (Juraj Sucik)

CERN's IT/IS group runs around 45 services on 17 physical servers offering a variety of flavours of Windows and Linux. Uses include production, development and test. There is a self-service scheme to create a new virtual service which takes about 10 minutes. Work has started on moving the service from Virtual Server 2005 to the new Hyper-V in Windows Server 2008 and Juraj described this in some detail, in particular new features for improved service such as quick migration for planned downtime, failover for unplanned downtime and volume shadowing for data protection.

Vista at CERN (Rafal Otto)

Rafal updated the talk he gave at the previous HEPiX on the use of Vista at CERN. As before, Vista is still only proposed for the more recent PCs and XP is still the default. Nevertheless, the number of Vista systems in use has doubled since the St. Louis meeting last November to around 300, some 5% of Windows installations at CERN. Work is ongoing to prepare the roll-out of Vista SP1 which was released in February but it is a major update, taking 2 hours and requiring 3 reboots. And some peripheral drivers are not compatible with SP1. But it is desirable to update since it fixes over 300 bugs and significantly improves performance, especially that of IE and of file copy. Rafal showed a number of benchmark results comparing Vista and XP; some showed Vista performing better, others showed XP to be the winner. Their preferred update service would be via WSUS (much smaller download) but SP1 is not yet available on WSUS, expected soon. For now they don't support 64 bits O/S but this is coming for some applications which require it, for example CAD, again based on Vista. Bearing in mind the declared end of support dates for hot fixes for XP (April 2009, security fixes until April 2014), sooner or later they will be forced to make Vista the default on new PCs. Finally, Rafal updated his story about the inevitability of moving to Vista based on not staying on old technology too long. As in St. Louis, there followed a long question period, showing that not all his audience were convinced and that migration will not be as fast as the support group might desire.

Lifecycle Management of Desktop Applications on Windows (Sebastien Dellabella)

CERN has some 6000 installed Windows PCs with a wide variety of installed applications from over 100 available on demand, all with their own patches and update cycles. The use of these apps is monitored by the home-built CMF (Computer Management Framework) tool. CMF provides an inventory of what is installed and its use. Based on group policies, CMF is used to deploy applications and relevant updates. He ended with several case studies where a combination of CMF and group

policies allowed them to control how and when particular upgrades were first made available for test and subsequently recommended to or even forced on users.

Scientific Linux Update (Troy Dawson)

Downloads of version 5.x continue to rise but that of 3.x has started to fall. Looking at downloads by country outside US, Germany has jumped into first place with around 60% of all downloads, more than all other countries combined by a factor of 2! SL 5.1 was released in January for i86 and X64 and work has already started on the 5.2 betas to prepare for the eventual release of the source from Redhat. Current work includes introducing more automation to downloading and package building. Work is also going in automating the fastbugs scheme for pushing out non-security patches on a weekly basis. Future releases include security fixes for 3.0.9 and the ports of 4.7 when it comes from Redhat, 5.2 as noted above and eventually 6.0 when released.

After the talk, Troy introduced an interactive session about what to do next in SL.

- For how long should they support SL 4? Would early 2010 be too little (5 years after initial release, as promises Redhat for their releases)? DESY requested at least 2011. FNAL have set an ultimate limit of 2012 because that is when Redhat drops their support of RHEL4. Another speaker noted that the answer depends on when grid middleware could support SL5 (or SL6). Tony Cass quoted that, at least for worker nodes, SL5 middleware support was planned to be available this autumn. After further discussion, Troy then proposed October 2010, the current end-date also of support for SL3. Jan Ivan worried that Troy was proposing to support too many systems in parallel but Troy noted that support of SL3 was very low work and he expected the same to be true of SL4 soon.
- Add XFS into main SL? Generally yes it seems but previous attempts hit a corruption problem, is that fixed? Troy believes it is. Jan suggests waiting until a new rumoured update of XFS is finished. Also, add it only into SL5. It would be an option and not interfere for example with sites preferring AFS.
- Should SL6 be based on [CentOS](#) instead of building everything from scratch. It would avoid having to rebuild rpms which never change anyway? CentOS people appear not to object. No objections from the audience so it looks like they will try this but this is not a final decision. Inevitably, someone, Alf Wachsmann in this case, asked why in that case do SL6 at all, why not simply use CentOS? Troy noted that even if the base was CentOS, FNAL would make the changes they currently make to produce SL. This question was left hanging for offline discussion.
- Firefox, Thunderbird and openoffice are currently unstable on SL but Troy intends to wait for Redhat's expected ports of the next versions of these on SL5 and hopefully SL4, due end 2008.

Data Centres

St. Louis Genome Centre Update (Gary Stiehr)

The new centre was under active construction during the last HEPiX meeting which was held in St. Louis and the first phase, already expanded from the original plan, is now mostly complete and installation of equipment has started. Two more phases are planned. Standards have been defined for all racks with necessary consideration for fire, safety and seismic regulations as well as power and cooling requirements. Seismic risks also affected floor layout. Racks receive 3 phase 208V power and they are currently debating if they need to worry about balancing between the phases and if anything these days really needs 120V power. He closed with a list of perhaps banal questions (what to clean and how frequently) but nevertheless details which cannot be overlooked. His summary was that you need to start preparations early, that dedicated effort is needed and that you need to be continually measuring and checking progress.

Remote Administration (Stefan Haller, GSI)

The problem – how to perform admin tasks at a distance? Simple tools are available (ssh on Linux, and Remote Desktop on Windows) but these are often not sufficient, no BIOS access for example. IPMI (Intelligent Platform Management Interface) service modules satisfy this last requirement at the expense of extra hardware, namely a Baseboard Management Controller (BMC). IPMI was designed by a collaboration of PC vendors specifically to managing and monitor PCs remotely and is now at version 2.0. He then went through a number of service modules, their functions and options, all illustrated with screen shots. Access to IPMI is possible via the web or by command line and the latter is useful for scripting or for use across low-speed modem lines. GSI created a dedicated sub-domain for their IPMI network. In summary, IPMI is a cheap (few hundred euros per service module) and easy solution for remote management, independent of O/S with comfortable and powerful front ends. In the discussion, Helge noted that partly for security reasons and partly for simplification, CERN switches off IPMI access via the web and uses only encrypted serial line interfaces

NIKHEF (and SARA) Data Centre (Wim Heubers)

Whereas the St. Louis centre has to take account of possible seismic events, the Amsterdam Science Park is 4 metres below sea level so the computers themselves are on the first floor to be above sea level. Unfortunately, all the power systems remain at risk of “the great flood” as do the offices of the staff! Co-located, NIKHEF and SARA share some things (building, LCG OPN, etc.) but not others (power, sys admin, operations). This can only complicate administration. NIKHEF also provides infrastructure to the Amsterdam Internet Exchange which demands zero downtime. In their plans to expand the data centre, the constraints were to use the existing building and not disturb the Internet Exchange. To assist in this, they contracted commercial technical experts. During this preparatory phase they had an incident in power distribution and they decided to invest in temporary measures to guarantee power and extra cooling during the coming work. Wim then explained the plans which include a second UPS, more cooling on the roof, converting their library (on the second floor) into a new computer room. Extra monitoring has been added to the power circuits. An open question is whether to upgrade or not their current fire protection which today is only based on smoke detection. All

things considered, from idea to realization is going to take 2 years, the person(s) in charge must get involved with many infrastructure matters (cf previous talk). His final suggestion was – “if you can start all over (like RAL and St. Louis who are building new centres), do it”.

The CERN AFS Console (Arne Wiebalck)

AFS at CERN today serves over 20,000 users, 250 million files using 15TB, spread across 35,000 volumes on 35 servers. Most servers are SUNs. There are 1-2 billion accesses per day. In comparison, CASTOR stores 95 million files in 12TB. The AFS support staff is estimated at 2.75 FTE. The AFS Console is a web-based tool used by this team to monitor the service but there is also limited access for users. Data collection is via the Lemon² infrastructure and it makes not only current information accessible but also historical data and it can also generate alarms. He described the two Lemon sensors which had been developed for AFS, to monitor partitions and volumes respectively. These sensors use standard AFS calls to gather the requested information and a small benchmark routine acts as a performance probe. Arne then detailed some cases where the tool was successfully used to recognize and/or debug a service problem. Other uses include service level status and volume placement. On the downside, it is not fast in data collection and data access can be heavy. The tool is currently CERN-specific but a code revision is underway to remove CERN dependencies.

New BNL Data Centre (Tony Chan)

Growth is being limited by power, space and cooling issues and this would affect plans for ATLAS ramp up. In the short term they are making serious efforts to husband these resources plus renovating unused space in the building while a new medium-sized centre is built (2008-2009). A new computer centre of the planned size (600 sq.m) should normally take 2.5 years but they aim to complete it in less than 18 months and for under \$5M, largely to avoid having to go to Congress to ask their agreement which would considerably delay the work. In the longer term, plans are under discussion for a single large centre to be ready by around 2015.

A Service-Based SLA Model (Tony Chan)

Facility operations is obviously manpower-intensive. Whether a component of the service gets primary shift cover only or 24x7 support depends on its classification in a service level table. However, as they grew, some systems came to support multiple services. So BNL have implemented a system dependency matrix and the SLA was revised accordingly. At the same time, they worked on better integration of the various public domain and BNL-written monitoring tools. A lot of work in particular was done using NAGIOS as well as cacti for network monitoring and Condor. All this was integrated with RT via an Alarm Monitoring Layer (AML). AML not only creates an RT ticket when triggered, it monitors the processing of that ticket and escalates it if it is not responded to in time. AML is driven by configuration files for the triggering and escalation rules.

CluMan (Sebastian Lopienski)

CluMan is a design for a cluster monitoring tool. Today, there are a variety of tools, many CERN-developed for cluster mgmt, including CDB, SMS, Lemon and SLS. Given the size, number and

² A CERN-developed monitoring suite, described several times in previous HEPiX meetings

complexity of the clusters in the CERN Computer Centre, keeping to the same standard is not at all simple and inconsistencies creep in. Often reconfigurations or hardware interventions are multi-step operations requiring mastery of different tools and he gave an example of adding a disc server to the CASTOR service. For these reasons as well as scaling purposes, automation is vital and this implies scripting of actions, parallel and/or chained triggering of actions and monitoring their progress and error handling. The ideas behind CluMan are based on the usage of the nodes (production, test, etc), node flags to represent these states (inconsistent; up or down, etc) and the properties of the nodes such as load, free space, etc. Actions would typically be reconfiguration scripts launched from a web application. The technologies which are planned to be used include the Google web toolkit, Apache iBATIS to map SQL queries to Java objects and YAML for data serialization.

Problem Tracking at CERN (Eric Grancher)

CERN/IT is currently reviewing its use of Remedy. A single instance supports several independent workflows and the current effort is intended to migrate from Solaris to Linux but this has hit severe problems. Eric presented a list of requirements for a problem tracking tool and a number of tools which could satisfy some or all (?) of these requirements. For each of these, Eric asked which sites used which tool. He then asked which sites used SLA to trigger alarms (none other than CERN), which integrated a helpdesk and incident tracking (none other than CERN) and so on. In summary, it appeared that CERN had invested more effort in their ticket system than most or perhaps all other sites. However it was suggested that even for sites present running the various systems, the person able to answer Eric's questions may not be present; perhaps the questions should be sent offline.

Storage Working Group Reports

File System Benchmarks (Andrei Maslennikov)

The WG was originally suggested by IHEPCCC in Oct 2006; it consists of 26 site representatives; they have prepared, issued and analysed a questionnaire and, on the basis of the results, which he presented graphically and which can be consulted on a so-called "HEPiX Storage Technology" [web site](#), they decided to investigate several of the architectures in use in various labs. A test cluster was established at CERN equivalent to a typical Tier 2 site. Coordinators were named for each file system to be tested and the environment was carefully defined with boundary conditions such as to make comparisons as meaningful as possible. Full details of these are given in the overheads. For simple parallel writes, most of the file systems gave similar results except AFS (10% lower) and GPFS (20% lower). A sequential read test gave a matrix of results but all file systems seemed to hit the same limit, except possibly Xrootd whose limit was notably lower. On the other hand, a pseudo-random read test showed much more variation; here AFS did rather badly (as might be expected) while Lustre did particularly well. dCache, Xrootd, GPFS and DPM had broadly similar results but always less than Lustre except for GPFS for a small number of jobs. In summary, the working group recommends serious consideration of Lustre as the file system of choice for a shared data store in large compute

clusters. The group will now formalize and publish the final results on its web site and they would like to perform one last production test if time and available resources permit.

Options to Improve LHC Data Management (Dirk Duellmann)

Dirk presented some options being looked at to optimize the data management of LHC data. The CCRC results show that the current situation is fairly stable and now, before the experiments start the more chaotic physics analysis stage, is a good moment to look again at future evolution and ideas for performance optimization. Also new technologies are on the horizon, such as Lustre (see previous talk) and the risk of a paradigm shift in media during the lifetime of the LHC. In the short to medium term, there is notable room for improvement in the areas of file size and tape formats. File clustering on tape could improve write performance, at the cost of some inefficiency however for reads. The data management team is proposing joint work with the CASTOR team to develop a scaling model for validation of production data. Other work is going on with the central database team to optimize Oracle database use. Another area which needs study is the environment for analysis which will forcibly be disc-based and a first disc pool model has been prototyped.

BNL Storage (Robert Petkus)

BlueArc is used exclusively for their NFS needs. They have 222TB on fibre channel-connected discs. Upgrade plans are not yet decided and the eventual choice, between more fibre channel or SATA discs may be decided by available funds. The supplier has indicated interest in implementing dCache access to the system. In the non-NFS storage area, in 2008 they intend to add 2PB of storage for ATLAS and RHIC and another 4PB in 2009. Limited power, space and cooling forces them to select high density solutions and SUN X4500s are the favourite for consideration at this time but other vendors are expected to offer similar solutions very soon. On their PROOF/Xrootd service they have tested using solid state discs to improve performance. This proved to be very interesting and they expect more use in the future as larger (and cheaper) SSDs become available.

CASTOR Status and Plans (Sebastien Ponce)

Since the last CASTOR report at HEPiX (2 years ago in Rome), the monitoring and I/O scheduling has been rewritten and there has been a full implementation of disc-only pools. More checking was added for file checksums at the disk level to detect any corruption earlier. An SRM2 interface to CASTOR was built at RAL. There are 5 instances of CASTOR on the CERN Tier 0, one for each experiment and one public. These total 486 disc servers, 2493 TB and 8.1M files. The full CASTOR namespace at CERN has 95M files, with 18 PB of tape storage capacity accessed via 125 tape drives. In addition, three Tier 1 sites run a CASTOR service. In the short term, few changes are planned, mostly in the areas of security, better admin tools and some performance improvements. In the medium and longer term, the plans were already described by Dirk earlier.

New CNAF HSM Solution (Vladimir Sapunenko)

As noted during the site report, this will be based on IBM products – GPFS and TSM - and CNAF's StoRM. This solution is already in production for files on disc but not on tape (D1T0) and now in test (during CRCC'08 for LHCb!) for files on disc and also on tape (D1T1). For files only on tape (D0T1) they

will stay, for now at least, on CASTOR. The work has been done in collaboration with IBM US and IBM Germany. CNAF's role was to interface GPFS to TSM, until now independent IBM products. First T1D1 tests were sufficiently encouraging to permit them to start the May LCG readiness challenge with this solution. For DOT1, more work is needed and an update to TSM is due from IBM and promised for the next major release. In the meantime, this is where StoRM will be used as an interim solution. Were changes required to some of the IBM tools? Answer: yes, "improvements". Are these going to be maintained by IBM in future versions? Answer: yes, as options on future releases.

Large Datasets at Google (Sascha Brawer)

Offline jobs at Google deal with large datasets (PBs) with bulk read/writes (MB) and where short outages are acceptable. Online apps deal with smaller datasets (TBs) and outages are immediately visible. Google uses "truckloads" of low cost systems, storage and networking and they care about performance per dollar but not per machine. They have found that even so-called reliable hardware fails at their scale. They have multiple data centres around the world, all interconnected. Their software must tolerate failure and apps must not depend on the host system. They have an in-house rack design; they run Linux 2.6 plus in-house software. He showed the failure stats for a typical cluster's first year of operation; all the problems were recognizable in the HEP world but the numbers were off-scale (e.g. thousands of disc failures). Their file system is based on some 200 (Redhat) GFS clusters with the master servers dealing in 64MB chunks with the chunkservers. The largest clusters may have 5000+ servers and be serving 10,000+ clients with 5+PB of files. He then described the Oracle database layout and transaction scheme, for example only single-row transactions to simplify operations, use of "bloom" filters to avoid seeks for non-existing data. Obviously replication of their databases is of primary importance. Table usage again shows impressive numbers: a single cell may manage up to 6000TB of data spread across 3000+ nodes supporting 500,000+ operations per second round the clock, with sometimes higher peaks! He then showed how pieces of data are broken down into (key, value) pairs and how these pairs are merged with similar data to produce a single pair for a given term (see overheads). Their current scheduling works well within a single cluster but there is no cluster-interworking and this is an area they are looking at – global name spaces; also more automation, lower end-user latency.

FZK Storage (Silke Fastenberg)

1700 nodes; >4PB of online (disc) storage on 5000 discs on 110 storage controllers and other 1500 drives about to be added; 3PB nearline (tape) storage; file system GPFS, SRM is dCache and archival is TSM. They use water-cooled closed racks. Servers are connected directly to storage and cluster file systems connect pairs of servers. The GPFS filesystem is visible on both nodes and dCache data management allows fast failover by starting a dCache instance on the non-failing node. The speaker showed a series of slides listing throughputs for various I/O interconnects.

The Unbearable Slowness of Tapes (Charles Curran)

Charles showed the planned evolution of tape equipment over the coming years. There are two plans, one based on expected new drives and a fall-back plan based on LTO drives although media costs would be much higher. Charles feels (strongly) that tape performance and throughput could be

dramatically improved if we considered more carefully the characteristics of tapes and the consequences of these when reading or writing files, especially small files, even despite some tricks played by some vendors to speed up small file operations. Charles plotted in detail the time taken to write LHC 1GB files onto a tape on both a SUN and an IBM drive. He showed that the actual write time was typically around only 50% of the time taken to complete the operation from tape selection to replacing it in its slot. First of all, he suggested dropping the one label per file scheme. Next, better drive selection, especially on SUN robots, could also save considerable time. With changes such as these, he believes we could improve useful time to 75-88% of elapsed time for writes. The same exercise for reading multiple files is even worse because of the file searches for multiple files, even on the same tape; here the useful cycle drops to 2-25% for both SUN and IBM. If we could arrange reads to correspond to the order of reading the tape, plus delay tape mounting until there are a certain of pending requests for that tape, we could hope to improve the useful cycle to 40%. Another problem is what happens when the media change and we are required to repack existing data. At today's volumes and speeds and dedicating half of today's installed drives, this would need at least 120-150 days but in practice it has taken far longer, usually over a year! This will be impossible in parallel with LHC running. Charles is proposing using a block by block copy rather than by programme and handling the typical 1% of errors individually, perhaps trying another drive or finding a copy on the Grid.

Setting up a Lustre Cluster (Stephan Wiesand)

After considering doing this for some time, DESY had delayed until some recent improvements to the product and its installation procedures before finally attempting it. The speaker described the main features of Lustre and then went through the steps and options for installing and tuning an instance.

High Availability on NAS (Nilo Segura)

The environment used to be Oracle running on SUN/Solaris with a mixture of JBOD or RAID disc. Recently we started to migrate to Linux servers and Oracle RAC (Real Application Cluster) and NAS as storage option. Why NAS? Easier file sharing; NAS features give you snapshots, failover, dynamic resizing, etc., all by default; and Ethernet technology - the NAS interconnect) is easier to manage (and cheaper) than fibre (which is used in SAN for example). Bonding, also known as trunking, allows us to aggregate multiple network interfaces which offers load balancing and redundancy and therefore improves reliability. The chosen configuration has 13 data discs plus a spare disc per shelf, aggregated into a single logical disc. We use RAID 6 (double parity). Oracle uses direct I/O over NFS, avoiding external cache since Oracle does its own caching. This arrangement has allowed several cases of replacing failed discs and other hardware interventions without service interruptions.

Lustre Cluster at GSI (Walter Schon)

Lustre metadata is stored in MDS which makes Lustre a good solution for large files and a bad one for small files. Data is stored in an OSS (object storage server). Other Lustre features include full POSIX compatibility, efficient use of bandwidth, scalable, quotas, etc. Various configurations offer varying degrees of failover. GSI spend considerable resources safeguarding the MDS but concentrate less on the OSS part. Lustre can drive a 1GbE line at 114 MB/s and, network permitting, performance scales linearly when you add clients. In both tests and real incidents, Lustre has proved to be rather resilient

and although there are some, mostly minor, problems; these are usually sorted quickly by the development team and GSI are convinced enough that they plan to migrate data from their existing NFS service to Lustre server by server.

Benchmarking

Perfmon (Andreas Hirstius)

Perfmon is a powerful open source monitoring tool for obtaining very detailed numbers from processors. Pfmon is the user interface, on which a lot of work has been applied to make it appropriate for use at HEP sites. In collaboration with the HEPiX benchmarking working group, some tests have been applied to Intel's latest Woodcrest processor. These tests were applied to SPEC2000 and SPEC2006 comparing 32 and 64 bit modes. For more details on the tests, the environment and the results, the reader is referred to the overheads. In answer to a question, Andreas noted that testing against SPEC tests is only part of the work. Similar tests have been applied to some experimental code to identify hot spots and bottlenecks.

Benchmarking Working Group Report (Helge Meinhard)

Was created some 18 months ago, to cover both file systems and CPUs. The former is reported above. Real work on CPUs only started in November 2007. Most sites until now had used SPEC2000 as the benchmark base but it was felt something new was needed; SPEC2006 (integer) looked like a good candidate. The working group contained representatives from mostly European labs as well as the 4 large LHC experiments. A set of benchmarks was defined, as was the working environment to run on worker nodes (the first set of processors to be tested) – see slides. In addition, each experiment was encouraged to run their codes for comparison. Helge presented some preliminary results which were expanded upon in the following talks. The working group may have started late but, having built up momentum, is now operational and should continue work, using standard benchmarks which indeed look adequate. One interesting suggestion is to do power consumption benchmarking also.

CERN Benchmarking Cluster (Alex Iribarren)

Alex then described in detail the benchmark systems used in CERN for the tests just described. They were selected from the range of systems currently installed in the Computer Centre. He noted that one cannot simply compare published SPEC scores because these are generated using the Intel compiler while HEP typically uses gcc; also the published data use an aggressive set of optimization flags.

ATLAS

ATLAS have created some benchmark tests which are made available to all members of the collaboration via a web portal. They run under the ATLAS ATHENA framework and the results also are accessible via the web to ATLAS power users via their personal certificates.

CMS

CMS wished to compare the SPEC tests to CMS codes within the CMS framework. Differences were seen in behaviour between AMD and Intel. The CMS apps scaled well with multi-core systems and work is continuing to make the tests more widely available.

ALICE and LHCb

Both presented preliminary results comparing their codes to SPEC for various chips.

NERSC (Iwona Sakrejda)

The benchmarks were performed on nodes extracted from the batch cluster and drained of jobs; this meant limited control of network and storage. User groups (including ALICE, ATLAS and STAR) were asked to submit characteristic loads. The results had to be normalized but appear to show that the PDSF workload runs well on multi-core systems as long as there is adequate memory and quad core in particular performs better than expected, probably due a re-designed memory caching architecture.

Networking

IPv6 (Mattias Wadenstein)

The perennial question: when will IPv4 address space become truly saturated? Mattias predicted it will happen “in a few years”. [Funny, I seem to recall hearing that some 10 years ago. Maybe this time ...] The UMEA site of NDGF has done some work on this with the target to roll out a service later this year. They found that, mostly, all installed O/S supported IPv6 as did most mainstream software. Only obsolete software, predating IPv6, failed. DNS did require some tweaking (see slide). Also any logging script which parses IP addresses – a warning to grid software developers. One show stopper is that NFS on Linux (at least on Ubuntu) only supports IPv4 today so both IP stacks must be run. In the questions, only two sites admitted to looking into IPv6 at any level.

Advanced Monitoring for the ATLAS TDAQ³ Network (Matei Ciobotura)

ATLAS's TDAQ is built on Ethernet switches and routers and must sustain rates up to 150Gbps. Given the size of ATLAS, the TDAQ consists of some 2500 computers installed in 90 racks with all the associated interconnections which could be expected. All equipment must be registered and stored into an inventory and a special web application has been developed for this and further tools compare the live system with what is listed in the inventory. Configuration changes and firmware updates are performed via scripts. For monitoring, Spectrum is used but it has some limitations (few diagnostic features for example) and some local software was developed to fill in the gaps. Spectrum produces nice traffic plots but there are 5000 ports and 15000 plots so a network browser was developed which can view, analyse, search and aggregate plots. Another lack in Spectrum is visualization and this has been added, both in 2D and 3D.

³ Trigger and Data Acquisition network

Security in the Grid Environment (Romain Wartel)

Security has a cost, is it worth it, do we do too much? But the same is true of lack of security and we could ask the same questions. Computer security is all about managing risks. Why do hackers hack? One reason could be to sell time on many nodes in an underground economy. Examples were shown using botnets, for example to send spam mails by the billions. Other reasons include denial of service attacks for political or other motives. Another example of the underground economy is the market in software kits available for \$500-1000 which allow to build one's own botnet or other malware. Web spamming tools are also widely available, for \$50 or an advanced version for \$450. Keylogger service? Of course – VirtualBreeze. Are grids especially vulnerable and why? So far not much but they have features which could be attractive to hackers – high availability, high throughput, large numbers of interconnected hosts and the transparent inter-access between sites. All these emphasise the need for strong security policies, appropriate access control, traceability and a defined incident response procedure agreed and followed by all the sites on the grid.

Cyber Security Update (Lionel Cons)

Lionel stepped in at the last minute to give this talk. He went through a frighteningly long list of vulnerabilities in the web, Windows, Linux, Mac, Controls and others. Web servers and browsers have long been a target but attacks on media players, for example Quicktime, are significantly increasing and web designers must be very careful in their coding because in Web 2.0, Javascript with AJAX is an easy target since code now runs in the browser. CERN has published some [hints](#). Turning to Windows, as Microsoft has increased security, hackers attack applications and user behaviour. And the myth that Macs are safe has been exploded. Under the banner of controls, articles have appeared recently discussing possible vulnerabilities of devices ranging from heart pacemakers to Boeing 787s (the attack vector is that the in-flight internet service is connected to the control of the aircraft, exposing the possibility of creating a “super flight simulator”). In summary, the cyber world is not becoming safer as attacks become more targeted. Coding must be made more secure and users, especially those on self-protected systems, must be constantly aware of the dangers.

Web Application Security (Romain Wartel)

Romain then returned to cover in more detail hints on coding web applications. Details should be obtained from the overheads but he covered cross-site scripting (XSS), SQL injection, file or code injection and forgery of cross-site requests. XSS alone was responsible in 2006 for two-thirds of common vulnerabilities. More recently, XSS is used during forgery attacks between sites (CSRF) and these are especially difficult to detect and protect against since CSRF exploits a design flaw implemented in many browsers to accommodate AJAX.

Miscellaneous

INSPIRE (Tibor Simko)

This is an attempt to merge various HEP information services such as SPIRES, CDS and others (including Google) into a single supported service. A collaboration was established in 2007 between

the largest labs in Europe and the US to investigate producing a single system based on CERN's Invenio system which would satisfy the needs of HEP researchers worldwide as currently offered by the existing systems. The current work is to reproduce the various existing user functionalities with a first scheduled release planned in a few weeks time for alpha testers. The next phase will add catalogue functions and next year they will develop added value functions such as user accounts as an option, collaborative features and others. Eventually they would like to expand to include researchers in other scientific fields and commercial publishers.

Offline Computing for LCLS⁴ at SLAC (Alf Wachsmann)

There will be two experimental halls and estimates for what will be needed for offline computing from "first light", scheduled mid 2009, are still surprisingly vague. The users have not the history of HEP researchers in modeling and estimating their needs well in advance. So the offline team requires to be flexible and not build in any known limits to extension. They have taken the best estimates from the physicists and prepared plans according to these, see slides for details. It is expected that data taking will ramp up slowly but there are no funds to build a major analysis and processing framework due to constrained resources, nor for significant increase in computing power beyond the initial investment. In the discussion which followed this, I was intrigued to note that at least two other sites present had experienced similar vagueness from other laser-based experimenters in their labs and they had suggestions or encouragements for the speaker. For years HEPiX sites have shared similar problems and issues and have benefited from experiences elsewhere; this appears to be still true as some sites switch to or adopt other sciences.

National Analysis Facility at DESY (Stephan Wiesand)

One of the work packages in a terascale alliance founded last year by a number of German universities and institutes. It should provide additional Tier 2 resources at both DESY sites, with extra emphasis on storage, for LHC and ILC physics. Requirements were collected from ATLAS and CMS for both batch and interactive work. Where initial plans called for separate grid parts, it has now been decided to merge the grid parts of the NAF into the existing (WLCG) grid and use VOs to dedicate shares for NAF. They have taken this opportunity to build a new facility from scratch with its own (sub-)domain and services and to introduce new concepts and techniques in system installation and configuration. AFS is used for home directories, Lustre for fast bulk data and SUN Grid Engine (SGE) is the batch system installed. A first round of acquisitions has taken place, installing some 6 HP blade servers for batch and interactive work as well as smallish AFS and Lustre services. Deployment is well advanced and the first beta users are testing it.

Managing Millions of User Accounts (GD Hoffmann)

The "millions" comes from multiplying the number of users registered at CERN (31,000) by the number of desktops (7600). A survey showed that several large labs use standard procedures such as LDAP or NIS to maintain common accounts between Windows and Linux. CERN had been looking at REGIS but LDAP offers some rather attractive features such as O/S-independence, scalable, many tools

⁴ Linac Coherent Light Source, currently being built and installed in the existing SLAC accelerator beamline.

available. Combining this with the existing Active Directory avoids data duplication and procedures have been written to use standard tools and protocols to configure Linux systems from the Active Directory user database. Care has to be taken for example not to overwrite particular local accounts such as root. On the Windows side, they make use of the “Services for UNIX” option. Putting these together, they are able to offer a single account between Linux and Windows, effectively single sign-on.

Alan Silverman
12th May 2008