

GSI-Computing Users Meeting
date: February 4 2013

Participants:

SC	Kilian Schwarz (KS)
SC	Mohammad Al-Turany
SC	Peter Malzacher
SC	Carsten Preuss (CP)
HPC	Thomas Roth (TR)
HPC	Thomas Stibor (TS)
HPC	Bastian Neuburger (BN)
HPC	Jan Trautmann
HPC	Victor Penso
ALICE	Benjamin Doenigus (DR)
ALICE	Markus Fasel
Theorie	Thomas Neff
PANDA	Klaus Goetzen
HADES	Jochen Markert (JM)

TOP 1: Status Update HPC

- new hardware did not arrive yet and there is no defined delivery date yet.
- TS did the second Lustre patch
- the hardware for Saudi Arabia is still built in

Action Items (AI):

- Prepare interactive Wheezy test nodes (HPC)
 - CH upgraded his desktop and it works
 - a node for general testing (lxbuild03) has been created
 - CH has done successful testing.
 - requests should go to BN
 - requests came from FU and JT. BN is working on them.
 - So far nothing has been done, also because Lustre modules for Wheezy are not yet existing.
 - It should be possible, though, to create Wheezy test nodes without Lustre or with Lustre being mounted via NFS
 - Distinct CVMFS servers may be required for Wheezy
 - Wheezy desktops should be made available ASAP
 - Also other Linux flavors should be investigated
 - WS promised an answer up to this meeting but he was absent
 - if someone wants to build software an e-mail should be sent to BN
 - TR plans to provide interactive Wheezy nodes. They will have no Lustre mount, though
- representatives from GSI and Frankfurt should sit together to identify how the new /hera mount at Frankfurt can be used most efficiently. Connected topics are e.g. CVMFS mirror for software (IT, ALICE)
 - a CVMFS mirror from GSI to Frankfurt might be problematic due to different operating systems involved.

- According to HPC the system is currently not in a usable state. The problem is that /hera is currently mounted only from a small number of clients. Big data throughput is not possible yet.
- PM suggests that one should sit together with representatives from the Frankfurt side. The timeline until production state should be defined. Currently GSI has a bandwidth of 80 Gb to Loewe CSC.
- KS contacted Philipp Luettig at Frankfurt University. Result: several scientists in Frankfurt are interested in the Loewe - /hera connection. They want to use it within a time scale of 2 months. Probably first a test run with not more than 5 users.
- Philipp Luettig should contact WS directly. WS still needs to do latency tests. Tests on small scale should be feasible.
 - So far there came no request from Frankfurt
 - HPC is currently not able to login in Loewe CSC so further testing is not possible.
- Squeeze desktops have a resolution problem with 16:9 monitors. A solution has to be found (HPC)
 - Stefan Haller (SH) did not know that he was supposed to fix the problem.
 - CH started to look for a solution. When using the correct parameters booting with proper graphic resolution is possible.
 - CH and SH will drop by and solve the monitor resolution related problems of ALICE and HADES.
 - HADES still has problems with that
 - the new Kernel does not support Lustre. This should be solved until end of 2013.
 - drivers have been updated. HADES is happy with the current solution, so is ALICE.
 - This AI will be closed.
- Printing from some of the Squeeze desktops does not work. (February 4, 2013, HPC)
 - This is a problem for ALICE and HADES, e.g. P444 does not work.
- AI: Blackhole detection on Gridengine (CP)
 - a solution needs to be found for problematic worker nodes which are not automatically excluded from production and thus lead to serious problems.
 - As long as it can not be distinguished if problems come from jobs or from the machines a fully automatised solution is difficult.
 - The idea is that the script should discover black holes and if one is discovered the node should be disabled and a ticket should be created automatically.
 - Problems:
 - how to discover black holes (general: node/user/job, dependent problem: short jobs)
 - how to deal with disabled nodes afterwards (just reboot or more action needed ?)
 - what if a large fraction of the cluster is disabled ? (automated call for the Call on Duty?)
 - current status: the script is in semi automatic test operation. A cron job sends e-mail to CP in case of problems. The script runs on the SGE server.
 - the current tool is just a plain Python script which shows some statistical numbers
Current focus is on a web site with multiple graphs showing the collected information so that a better view helps to decide whether a node is a black hole or not.

- lxalitransfer1 – 3 should be moved from /lustre to /hera (KS)
 - a discussions started if the boxes should be moved to different hardware and to a different location. This should be clarified within the IT.
 - Lxalitransfer2, lxalitransfer3, lxalitransfer4 are configured transfer boxes with /hera mounted. On all 3 machines the cronjobs checking for new transfers to be executed are running. Proxies are automatically renewed.
 - Lxalitransfer1 will stay as transfer box on top of /lustre but should be upgraded to new OS.
 - This AI will be closed.
- transfers via 10 Gb link should be activated. (KS)
 - ALICE::GSE::SE2 on top of /hera, connected to the 10 Gb link, has been activated.
 - Disk status is “uncertified” so only data explicitly transferred to this SE will go there.
 - Soft links to AliEn LFNs have been enabled and also tested by local ALICE users.
 - As a first test case the data set AOD115 (150 TB) will be copied.
 - The transferred AODs, stored via softlink to a speaking name, equivalent to the LFN in AliEn, have been tested and everything works fine.
 - Transfer ongoing with an average speed of 17 MB/s
 - up to now unclear why the transfer rate is so slow
 - 60% of the traffic seem to go via HEPPI net
 - a MonaLisa monitoring client has been installed on lxalise1.
 - Data are incoming now. Mainly the inter site bandwidth tests should be of interest.
 - Current data seem to imply that the bandwidth from lxalise1 to other sites is not larger than 1 Gb
 - The results need to be analysed in detail.
 - The accounts via which the transfers are done are NIS accounts, but NIS was not running on lxalise1 nor is the corresponding network connection patched.
 - TR is configuring a switch to local accounts with the local accounts having the same Ids as the global NIS accounts so that writing on Lustre will be possible.
- The LSDMA topics need to be discussed with the FAIR experiments. Therefore a set of concrete questions should be formulated beforehand and circulated (KS, CH)
 - HPC will participate in WP1 (Federated Identity Management), WP2 (Federated Data Access) and WP5 (IO profiling).
 - KIT wants to participate in FAIR online trigger algorithms and pattern recognition.
- AI: within the technical Computing meeting shall be presented on a regular basis how many computing resources each experiment consumed within the last 4 weeks and also how much space is left on the Lustre storage (KS, CH).
 - Lustre fill stats and batch farm usage stats shall be prepared every month in cadence with the meeting and put on a web page.
 - Monitoring efforts of different groups and individuals shall be consolidated.
 - JT will compile a wish list.
 - This wishlist has not been communicated so far.
 - Access to the monitoring site from outside GSI shall be possible but with authentication being activated, e.g. via GSI Web login.

Closed and PENDING Action Items:

- Finden einer Lösung für problematische Maschinen, die nicht automatisch aus der Produktion ausgeschlossen werden und somit zu erheblichen Problemen führen können (vom 12.9.11: IT)
 - So lange nicht klar unterschieden werden kann, ob die Probleme von den Jobs oder von der Maschine kommen, ist eine vollautomatische Lösung schwierig.
 - Da ALICE hierunter nicht akut leidet wird das Problem auf PENDING verschoben.
- Um die Storage-Elemente von PANDA und CBM, die derzeit unter Etch laufen, auf Lenny oder Squeeze umziehen zu können, muss hierzu vorher geklärt werden, was von Grid-Seite aus vorbereitet werden muss (vom 16.01.12: KS)
- CVMFSServer for the test Cluster (B) runs on old hardware. Money for proper service deployment is missing. WS should investigate in the next IT coordination meeting who in principle is responsible for buying hardware for server in the Minicube (CVMFS server, Build server). VP mentions that hardware should be bought by the experiments. JM asks VP to specify what hardware would be needed. VP thinks that a server including storage for 3000 Euro in total would be sufficient. VP will make a concrete suggestion concerning the necessary hardware. WS will bring the topic to the IT coordination meeting (05.03.12: VP, WS)
 - together with the last hardware order a server for CVMFS has been bought
 - to some extent the ordered hardware will be at GSI in 1-2 months.
- Plan the LSF decommissioning and hardware recycling (HPC)
 - Shutdown of Lenny farm scheduled for the end of 2012 (14th of December 2012)
 - SGE test cluster and old LSF/SGE cluster shall be switched off. A list of computers to be switched off will be created (SC) and sent to the experiments for approval (KS). Affected will be Lenny and Etch boxes, e.g. some lxi machines.
 - work is in progress.
 - Lxi001.gsi.de and lx-pool.gsi.de should point to squeeze nodes.
 - Shared home dir still required for PandaGrid.
 - Small Gridengine convenience cluster may be helpful for small user groups.
 - Centrally mounted /u may be sufficient
 - /d file system can be phased out, too.
 - LSF has been switched off
 - SGE on the old cluster will run as long as the hardware stays alive.
 - This AI will be closed.
- MPI jobs never scheduled outside the default queue
 - a patch for GridEngine will be developed (SC)
 - work is in progress (Anar Manafov (AM))
 - decision was made to have a look at SonOfGridEngine if the problem still exists there, before starting to develop a patch for GE.
 - Also SLURM is now in production (good experiences reported by VP), so the question raises if it makes sense to invest more time in GE (SoGE)

TOP3: projects and shares for batch farm.

- The decay time of the fair share algorithm will be reduced to 1 week
- ALICE will now be able to account only on the ALICE share. All other users will account on “default”.

TOP4: Plans for next weeks

- HPC:
 - eventually the new hardware can be brought into operation
- HADES:
 - no larger activity planned.
- ALICE:
 - usual operation
- THEORY:
 - a few 1000 jobs are planned.
 - Problems with MPI jobs have been encountered. CP will look into this.
- PANDA:
 - nothing planned.

TOP5: AOB

- HADES requests user monitoring. JM would like to be able to see consumed resources, exit codes and GE failing codes of individual jobs.
 - In this context it has been shortly discussed if one should teach users, e.g. via a Wiki documentation, how to catch soft limit signals of their jobs.
- Last week /hera was slow (JM). BN recommends to send information or ticket immediately when the problem arises.

TOP5: next meeting

- the next meeting will take place at the usual time (Monday, March 4, 2013, from 2 pm to 3 pm) in the KBW building. BD will reserve the rooms in the KBW building for the whole year.