

# Agenda for the GSI Computing Users Meeting on 2012-12-03

## Current Action Items

Assigned to	Due date	Description
LinuxGroup	Sat, 1 Dec 2012	<b>Prepare interactive Wheezy test nodes</b> CH upgraded his desktop and it works a test node for general testing will be 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 (go to action)
LinuxGroup	Sat, 1 Dec 2012	<b>Investigate the implementation of GridEngine group admins via sudo -u group_member [qdel/qmod/...]</b> HPC will check if the implementation has been done already. This is needed by ALICE and HADES. HPC will test up to next meeting. There are problems involved connected to the group membership of the experiment members. In principle it can be possible this way that HADES eliminates ALICE jobs and vice versa (in such a case "Rufbereitschaft" should not be called). One mainly responsible person should be identified for each experiment and communicated to BN. Eventually the group KP1 for ALICE members can be removed completely. (go to action)
ThomasRoth, HorstGoeringer	Sat, 1 Dec 2012	<b>Status of Hera/gStore connection</b> HG tested the setup with 200-300 MB/s via Lnet router. Opening to public will be done at the beginning of September after the current beam time HADES should test this (JM) tests have not been done yet. The feature is available and has been tested by Horst Goeringer. More bandwidth (more data mover) will be available after the HADES beam time (October 2 up to the coming Monday). (go to action)
LinuxGroup	Sat, 1 Dec 2012	<b>representatives from GSI and Frankfurt should sit together to identify how the new /hera mount at Frankfurt can be used most efficiently.</b> 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. HADES and ALICE jobs suffer from too many open file handles. Since this is a problem for Lustre a solution has to be found (HADES, ALICE, IT) the problem did not appear anymore because ALICE did not run simulation jobs. JM mentions that the scheduler is unstable JT wants to test once more. (go to action)
LinuxGroup	Sat, 1 Dec 2012	<b>Squeeze desktops have a resolution problem with 16:9 monitors</b> 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. (go to action)
KilianSchwarz	Sat, 1 Dec 2012	<b>lxalitransfer1 - 3 should be moved from /lustre to /hera</b> a discussions started if the boxes should be moved to different hardware and to a different location. This should be clarified within the IT. - lxalitransfer1 will be replaced by lxb563. Ticket has been created - lxalitransfer2 and 3 have /hera mounted, but \$HOME for ali-data is still missing. Do we need a different location for those boxes ? (go to action)

KilianSchwarz	Sat, 1 Dec 2012	<b>transfers via 10 Gb link should be activated.</b> – ALICE::GSI::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 – as a first test case the data set AOD115 (150 TB) will be copied (go to action)
KilianSchwarz, ChristopherHuhn	Sat, 1 Dec 2012	<b>The LSDMA topics need to be discussed with the FAIR experiments.</b> Therefore a set of concrete questions should be formulated beforehand and circulated The LSDMA gap analysis and application document should be sent to this list (KS) information about LSDMA has been sent to the cluster–mailing list. (go to action)
KilianSchwarz, ThomasRoth	Sat, 1 Dec 2012	<b>Overview of computing resources</b> 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. (go to action)
CarstenPreuss	Sat, 1 Dec 2012	<b>Blackhole detection on Gridengine</b> 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.  The idea is that the script should discover black wholes and if one is discovered the node should be disabled and a ticket should be created automatically.  Problems :  a) how to discover black holes? –general, node/user/job dependent problem? –short jobs  b) how to deal with disabled nodes afterwards? –just reboot or more actions needed?  c) what if a large fraction of the cluster is disabled? –automated call for the Call on Duty?  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 informations so that a better view helps to decide whether a node is a black hole or not. (go to action)
KilianSchwarz	Sat, 1 Dec 2012	<b>Squeeze upgrade of Grid storage elements</b> 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) (go to action)
LinuxGroup	Sat, 1 Dec 2012	<b>Sketch lustre–hera migration time line</b> estimate of exact time lines won't be possible before Decembre. (go to action)

## Pending Als

Assigned to	Due date	Description
LinuxGroup	Fri, 14 Dec	<b>Plan the LSF decommissioning and hardware recycling</b> Shutdown of Lenny farm scheduled for the end of 2012 (14th of December 2012) SGE test cluster and old

	2012	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. (go to action)
sc@gsi.de	Wed, 2 Jan 2013	<p><b>MPI jobs never scheduled outside the default queue</b> a patch for GridEngine will be developed work is in progress (Anar Manafov (AM))</p> <p>Decision was made to have a look on SonOfGridEngine if the problem still exists there, before starting to develop a patch for GE.</p> <p>Also SLURM is now in production (good experiences reported by Victor Penso), so the question raises if it makes sense to invest more time in GE (SoGE)? (go to action)</p>
LinuxGroup	Fri, 1 Feb 2013	<p><b>Hardware upgrade for CVMFS server(s)</b> 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?! (go to action)</p>

— ChristopherHuhn – 03 Dec 2012