

# GSI-Computing-Users-Meeting 02/07/2012

## Participants:

- Mohammad Al-Turany, SC
- Benjamin Dönigus, Alice
- Tetyana Galatyuk, CBM
- Paul Görger, Bephy
- Christopher Huhn, HPC
- Peter Malzacher, SC
- Jochen Markert, Hades
- Thomas Neff, Theorie
- Bastian Neuburger, HPC
- Victor Penso, HPC
- Carsten Preuss, SC
- Thomas Roth, HPC
- Jochen Thaefer, Alice
- Jan Trautmann, HPC
- Florian Uhlig, SC/CBM

## Current issues

- Network connection Prometheus
  - ◆ general purpose networking limited to 1 Gb/s atm
    - ◇ may be blown up to at most 10 Gb/s without new (network) hardware
    - ◇ would decrease the LNET throughput in turn
  - ◆ 2 separate use cases:
    - ◇ 1) interactive nodes with access to resources outside minicube, e.g. to the Internet, /lustre, gStore, ...
      - These nodes could be located outside minicube, only /hera mount as connection into minicube and ability to submit to Prometheus necessary
      - **Action item:** migrate submit nodes (pro.hpc.gsi.de) from TH to a sub-net with internet access and mount /hera via LNET routers (40Gb/s), add similar submit nodes: using older BG2 batch nodes
    - ◇ 2) Grid jobs: direct internet access from Prometheus worker nodes
      - not part of original plan for minicube
      - Absolutely undesirable for HPC
      - Would require further network setup (€€) and routing (NAT)
      - Working setup in Mainz w/o internet access
      - Perhaps access via proxies/vo-boxes is sufficient
        - **Action item:** find out more details about the setup in Mainz (SC)
- /lustre to /hera migration
  - ◆ both file-systems are full atm.
  - ◆ Alice will not delete significant amounts of data before Quark Matter (starts 12.8.)
  - ◆ **Action item:** sketch lustre-hera migration time line (TR)
  - ◆ **Action item:** create up-to-date fill stats for the Lustre file-systems (TR)
- Debian Wheezy
  - ◆ The next Debian version *Wheezy* has now been *frozen* ie. no major package upgrades will happen any more but only bugs will be fixed before the release
  - ◆ Everybody welcomes to start preparing and testing the Wheezy migration soon
  - ◆ **Action item:** Prepare interactive Wheezy test nodes (HPC)

## Action items

- Plan the LSF decommissioning and hardware recycling (HPC)
  - ◆ Hades still needs Etch nodes for software compilation for 3 PhD students that are near completion of their thesis'
    - ◇ compiled binaries also run on Lenny
  - ◆ Etch compute nodes can be phased out immediately
  - ◆ Shutdown of Lenny farm scheduled for the end of 2012
- Check if version control can serve as a backup replacement for Alice software development (JT)
  - ◆ No progress
- Investigate the implementation of GridEngine group admins via `sudo -u group_member [ qdel | qmod | ... ]` (HPC)
  - ◆ No progress
- Increase/remove job limits on Prometheus for hadesdst account (HPC)
  - ◆ managed by Jochen and Tetyana exclusively
- Status of Hera/gStore connection? (Horst Goeringer, Thomas Roth)
  - ◆ According to Horst Goeringer writing from gStore via LNET routers to Hera is implemented on test servers
  - ◆ Schedule for production deployment unknown
  - ◆ Hades will test this setup
- Investigate the possibilities for a backup of the CVMFS servers (HPC)
- Environment variables that indicate to which cluster a worker node belongs to (HPC)
  - ◆ Implemented as `$CLUSTERNAME`
- CVMFS on Squeeze desktops (06.02.12: JT, IT)
  - ◆ Work in progress
- Farm-wide availability of Modules (06.02.12: JT, IT)
  - ◆ already in use by Alice
- Funding for new hardware, e.g. CVMFS server (05.03.12: VP, WS, CH)
  - ◆ call for tenders in preparation
- MPI jobs never scheduled outside the default queue (SC)
  - ◆ SC will develop a patch for GridEngine

## Long-term and postponed tasks:

- Cleanup of jobs that are not scheduled during a given interval (HPC)
- Upgrade of PANDA and CBM Storage Elements from Etch to Squeeze (16.01.12: KS)

— ChristopherHuhn – 05 Jul 2012