

# GridKa site report **(Extract)**

Manfred Alef, Andreas Heiss, Jos van Wezel

Steinbuch Centre for Computing



Forschungszentrum Karlsruhe  
in der Helmholtz-Gemeinschaft



Universität Karlsruhe (TH)  
Research University · founded 1825



# Hardware Status on "Milestone 1<sup>st</sup> April 2008"

## ■ New CPUs:

- + 4438 kSI2k (338 boxes):

2x Intel Xeon E5345 (2.33 GHz Clovertown)

TYAN-Tempest-i5000PX-S5380 / Supermicro X7DB8

- + 2355 kSI2k (144 boxes):

2x Intel Xeon E5430 (2.66 GHz Harpertown) — D-Grid

IBM x3550

- 8x2 GB RAM (DDR2 FB-DIMM): 2 GB RAM per core

- 2x 250 GB local disk (system + local "homedirectories" for LCG jobs)

## ■ Old CPUs (retired):

- – 348 boxes: 2x Intel Xeon 2.66/3.06 GHz (Nocona)  
(0.5 GB RAM per core)

## ■ Sum total: **9520 kSI2k**

# Hardware Status on "Milestone 1<sup>st</sup> April 2008"

- Some remarks on the CPU procurement procedure at GridKa:
  - Call for tenders is based on the total compute power which is required
    - No fixed number of machines
    - Vendor has to deliver as many nodes as are required to achieve the total amount of compute power
  - Based on SPEC CPU2000 (CINT2000)
    - To be executed in the GridKa environment:
      - OS: Scientific Linux 4.x
      - Compiler: gcc-3.4.x
      - Fixed set of optimization flags
    - Simulation of the batch system:
      - Run as many SPEC runs in parallel as there are cores available
  - Adjudication: "The same procedure as every year"
    - Purchase price
    - + Estimated cost of electric energy for operation and cooling for 3-4 years.  
(4 € per  $W_{\max}$ )
    - + Estimated cost of space and racks (300 EUR per 19" unit)
    - + Network ports, software licenses ... (200 EUR per system)

- Some remarks on the CPU procurement procedure at GridKa:
  - Experiences:
    - Generally this procedure works quite well.
    - However ... (see the next slide)
    - Check tenders very carefully ...
    - ... Experiences with next procurement (milestone Oct 2008):
      - Benchmark report (SPEC and power) based on system with 4x4 GB FB-DIMM
      - Tender offered 8x2 GB FB-DIMM
      - Take into account the power consumption of FB-DIMM memory: 10 ... 15 W per module!
      - Makes a difference of about 50W per box!

- Acceptance tests at GridKa:
  - Running SPEC CPU2000 benchmark as an acceptance test:
    - The performance of all WNs, not only samples, has been checked with benchmarks.
    - Bad experiences with one batch of nodes:
      - Expected performance of boxes:  
11230 SPECint\_base2000 (prototype test machine)
      - Offer promises 11869 SPECint\_base2000
      - First benchmark results: 10560 SPECint\_base2000
      - BIOS upgrade + modification of BIOS settings by vendor:  
11042 SPECint\_base2000, that's an improvement of about 4.5 %!
      - Vendor delivered more machines than purchased for free

# Hardware Status on "Milestone 1<sup>st</sup> April 2008"

- Acceptance tests at GridKa:
  - Reproducibility of results:
    - Results of benchmark runs  
(after the BIOS modifications mentioned above):  
11002 ... 11092 SPECint\_base2000 (0.8 % bandwidth)
    - However, we found poor performance results on some boxes caused by hardware problems, e.g.:
      - OS found only 8 of 16 GB RAM
      - OS found only 4 (or 7) of 8 CPU cores
      - Wrong CPU type in /proc/cpuinfo, e.g.,  
"Intel(R) Xeon(R) CPU Q2407 (or Z0383) @ 1.33GHz"

# Hardware Status on "Milestone 1<sup>st</sup> April 2008"

- Acceptance tests at GridKa:
  - Electric power consumption:
    - Power consumption of one batch of machines higher than expected
    - Vendor replaced PSUs by 80plus<sup>1</sup> ones
    - Reduced power consumption around 10%
  
- Call for tenders for October milestone finished.
  - Machines will have very low power consumption compared to recent batch!

1) 80plus PSU: >80% efficiency when load >50%