



# IHEPCCC/HEPiX benchmarking WG

Helge Meinhard / CERN-IT  
Grid Deployment Board  
09 January 2008



# History (1)

- Autumn 2006: IHEPCCC chair contacts HEPiX conveners about help on two technical topics
  - File systems
  - (CPU) Benchmarking
- HEPiX meeting at Jefferson Lab (Oct 2006) reacted favourably
- Participation sorted out
  - Chairs: Maslennikov (CASPUR): File systems, HM (CERN): Benchmarking
- File system WG started immediately... not so the Benchmarking WG
  - Mea culpa... and my apologies to everyone

## History (2)

- Ad-hoc discussion at HEPiX at DESY (April 2007)
- HEPiX in St Louis (November 2007)
  - Benchmarking track
    - INFN move to SI2006 (Michele Michelotto)
    - Update on GridKA benchmarking (Manfred Alef)
    - Topics around CERN benchmarking (HM)
    - AMD and Intel performance for lattice QCD codes (Don Holmgren)
  - 2 more formal face-to-face meetings of people interested in benchmarking, formal start of WG
- Two phone conferences since (30-Nov, 13-Dec), next one scheduled for this afternoon

# Motivation

- Used so far (at least in LCG): int\_base part of SPEC CPU 2000
  - Inconsistent usage across institutions
    - Measurements under real conditions vs. looking up values
  - Experiment code shown not to scale well with SI2k
    - In particular on multicore processors
    - SI2k memory footprint too small
  - SPEC CPU 2000 phased out in 2007
    - No more support by SPEC
    - No new entries in table of results maintained by SPEC
- New benchmark required
  - Consistent description of experiment requirements, lab commitments and existing resources
  - Tool for procurements
    - Buy performance rather than boxes
- “Natural” candidate: SPEC CPU 2006 (int\_base or int\_rate part)



# WG composition

- Currently on mailing list:
  - Ian Gable (U Victoria)
  - Alex Iribarren (CERN)
  - Helge Meinhard (CERN)
  - Michael Barnes (Jefferson Lab)
  - Sandy Philpott (Jefferson Lab)
  - Manfred Aef (GridKA)
  - Michele Michelotto (INFN Padova)
  - Martin Bly (RAL)
  - Peter Wegner (DESY)
  - Alberto Aimar (CERN)
  - Sverre Jarp, Andreas Hirstius, Andrzej Nowak (all CERN)
- Actively looking for participation from FNAL and SLAC
- More expressions of interest to collaborate very welcome



# WG status – Agreements

- Focus initially on benchmarking of processing power for worker nodes
- Representative sample of machines needed; centres who can spare a machine (at least temporarily) will announce this to the list
- Environment to be fixed
- Standard set of benchmarks to be run
- Experiments to be invited (i.e. pushed) to run their code
  - Check how well experiments' code scales with industry-standard benchmarks

# Machines

- A number of sites announced that they could make machines available, at least temporarily
  - Aim for interactive logins for experiments
    - At CERN, environment much like lxplus/lxbatch
  - Keep track of number, type, speed of CPUs, chipset, number, config, CAS latency and clock speed of memory modules
    - Potentially BIOS settings (and even temperatures) to be added ☹
- At CERN: Ixbench cluster with currently 7 machines
  - Nocona 2.8 GHz, Irwindale 2.8 GHz, Opteron 2 x 2.2 GHz, Woodcrest 2.66 GHz, Woodcrest 3 GHz, Opteron Rev. F 2 x 2.6 GHz, Clovertown 2.33 GHz
  - Harpertown 2.33 GHz to be added soon
  - Some more machines (similar to one type or another) added temporarily for consistency checks
  - Standard benchmarks run
  - Perfmon to be installed

# Environment

- SL 4 x86\_64, 32-bit applications
  - Some cross-checks with SL 5
- Gcc 3.4.x (system compiler)
  - Cross-checks with gcc 4; will need gcc 4 for SPEC FP 2006
- Compilation options by LHC architects' forum (-O2 -fPIC -pthread -m32)
  - Check whether other communities (CDF/D0, BaBar, ...) prefer different options
- Multiple independent concurrent runs
  - Multi-threaded benchmarks (SPECrate) checked as well – although normalisation is the same, values too low by ~5% on fast machines



# Standard benchmarks

- **Compulsory:**
  - SPECint\_base 2000 (concurrent, transition only)
  - SPECint\_base 2006 (concurrent)
  - SPECfp\_base 2006 (concurrent, gcc 4)
- **Optional, cross-checks only:**
  - SPECint\_rate 2006
  - All above as 64-bit applications
  - SPECint compiled with gcc 4
  - All above under SL5
- **Notes:**
  - “concurrent” – as many instances as there are cores run in parallel, results added over all cores
  - Applications are 32 bit unless noted otherwise



# Experiment participation

- Important in order to check how well real-life applications scale with industry-standard benchmarks
- Discussed in LCB Management Board
- December 2007: Alberto Aimar contacted reps, warned them that benchmarking will come up, and asked for people to contact
- Replies received from ATLAS, ALICE, LHCb
- Detailed information about CERN cluster to be distributed this week



# Conclusions

- WG has taken long to get started, but is now up and running
- Momentum is there, interest and (some) person power as well
- Needs a bit of focussing in order to ensure results are comparable
- Needs active participation from experiments
- Interim report expected at HEPiX at CERN May 2008
- Some interest in looking at power consumption as well