



## Benchmarking Working Group

---

# Beyond HS06: Toward a New HEP CPU Benchmark

Manfred Aef (KIT), Michele Michelotto (INFN-PD)



# Benchmarking Working Group

---

## History – Before HS06:

- Benchmark scores had been looked up at [www.spec.org](http://www.spec.org) (SPECint95, SPECint2000)
  - Performance scores supplied by hardware vendors
  - SPEC run rules don't dictate terms of optimization, vendor may choose compiler and flags (and high-performance libraries)
  - Hard to compare performance of different machines
    - ◆ Do we want to measure the performance of:
      - system software
      - compiler, system libraries, compiler flags
      - hardware
      - ... ?



# Benchmarking Working Group

## History – Before HS06:

- Benchmark scores had been looked up at [www.spec.org](http://www.spec.org) (SPECint95, SPECint2000)
  - Performance scores supplied by hardware vendors
  - SPEC run rules don't dictate terms of optimization, vendor may choose compiler and flags (and high-performance libraries)
  - Hard to compare performance of different machines
    - ◆ Do we want to measure the performance of:
      - ~~system software~~ (→ default: Linux distro)
      - ~~compiler, system libraries, compiler flags~~ (→ specified)
      - ✓ hardware
      - ≡ !



# Benchmarking Working Group

---

## Interim Solution:

- Providing rules to run SPEC CPU2000 (SPECint2000):
  - Operating system: Linux (distro+release in use at site)
  - Compiler: GCC (default compiler)
  - Flags: `-O2 -pthread -fPIC`
  - 1 benchmark copy per core of the system under test
    - ◆ Analogous to the batch scheduling



# Benchmarking Working Group

---

## Advantage of Strict Benchmarking Rules:

- Easy comparison of the performance of machines
  - Not depending on installed software



# Benchmarking Working Group

---

## HS06 benchmark details:

- Based on industry standard benchmark suite (SPEC CPU2006)
  - Inexpensive for academic use
    - SPEC CPU2006: non-profit pricing=200\$ (site license: no fee per host, and no annual renewal)
  - ♦ Vendors are familiar with it
- A subset of 7 benchmarks is used for HS06 scoring (= all C++ packages coming with SPEC CPU2006)
  - ♦ 3 integer, 4 floating point
  - ♦ Matches the percentage of floating point operations of HEP job mix



# Benchmarking Working Group

## HS06 Benchmark Run Rules:

- Operating system:  
Linux (same distro and release running at your site)
- Compiler:  
Default compiler coming with the operating system (GCC)
- Optimization flags are mandatory:

`-O2 -pthread -fPIC`

**Defined by LCG  
Architects Forum**

`-m32`

**Proposed by  
Benchmarking WG**

- Running 1 benchmark copy per job slot  
of the system under test



# Benchmarking Working Group

---

## HS06 Development:

- HEPiX Benchmarking Working Group
  - Selection of benchmark tools and configs
  - Proof of scaling with SPECint2000 scores
- WLCG Benchmark Task Force
  - Proof whether the benchmark scores scale with typical mix of HEP applications



# Benchmarking Working Group

---

## Next generation benchmark (working title: "HS14"):

- Since the hardware has become more and more powerful (faster chips, increased number of cores, huge cache size) the benchmark should be redesigned from time to time to bring it up to date
  - Increase in single-core performance much less than 8 years ago
  - No inflation of cache sizes like at that time
  - But should move to 64-bit applications



# Benchmarking Working Group

---

## Next generation benchmark (working title: "HS14"):

### → Benchmark candidates:

- SPEC is working on the next generation of CPU benchmark
  - ◆ Will probably become available by end of 2014
  - ◆ Good candidate for HS14 benchmark!
- Free alternatives like Geant4 will be considered as well
  - ◆ Open source based benchmarks can be distributed from public software repositories
- Open for suggestions for alternative candidates



# Benchmarking Working Group

---

## HS14:

→ Fields of application:

- Metric of installed capacities, accounting
  - ◆ Should scale with (mix of) real HEP applications
    - ✓ HS06: Proven by WLCG Benchmarking Task Force
    - ✓ HS14: Identify suitable applications to check
- Performance of assigned job slot



# Benchmarking Working Group

---

## HS14:

### → Requirements (cont.):

- Easy to use
  - ◆ Prefer either standard benchmark which is widely used like replacement of SPEC CPU2006, or other benchmarking framework which must be easy to run
  - ◆ Free, or at least inexpensive for academic use
  - ◆ Vendors must be familiar with it (no knowledge of HEP software)
  - ◆ Acceptable execution time of benchmark runs



# Benchmarking Working Group

---

## Wishes:

- Users want to know performance of assigned job slot e.g. to schedule payloads of pilot jobs
  - At least a rough guess
  - WLCG Machine / Job Features Task Force proposes to use \$MACHINEFEATURES/hs06 file
    - ◆ No problem on bare metal, but possibly unavailable in virtual environments / clouds
  - Full HS14 benchmark run will most probably take several hours, can we design a lightweight tool for quick performance estimate?
    - ◆ Possibly less exact than HS14 but "good enough"



# Benchmarking Working Group

---

## First steps – looking for volunteers:

- Manfred Alef, Michele Michelotto (chairmen)
- VO representatives:
  - ◆ Alice: Peter Hristov, Costin Grigoras
  - ◆ Atlas: Alessandro Di Salvo, Franco Brasolin
  - ◆ CMS: Gabriele Benelli
  - ◆ LHCb: Ben Couturier
- Others:
  - ◆ Martin Bly
  - ◆ Ian Gable
  - ◆ Fotis Georgatos
  - ◆ Andrew McNab
  - ◆ Romain Wartel



# Benchmarking Working Group

---

## Time scale:

- Collection of representative HEP applications to compare with benchmark candidates (from now)
- First investigations of benchmark candidates (e.g. next SPEC CPU benchmark suite) 2014/2015
- Recommendation of final benchmark candidate(s) by mid 2015
- Validation of benchmark by LHC VOs till end of 2015
- Approval by WLCG MB: H1 2016
- Creation and support of wiki of benchmark scores
  - <https://w3.hepox.org/benchmarks>



# Benchmarking Working Group

---

## Contact:

- [manfred.alef@kit.edu](mailto:manfred.alef@kit.edu)
- [michele.michelotto@pd.infn.it](mailto:michele.michelotto@pd.infn.it)

## Mailing list:

- [hepixon-cpu-benchmark@cern.ch](mailto:hepixon-cpu-benchmark@cern.ch)