

# **WLCG HEP-SCORE Deployment Task Force**

Helge Meinhard / CERN-IT  
04 November 2020

# Reminder of History

- Pre-2009: WLCG used a benchmark based on SPEC CPU2000
- At EOL, WLCG started looking into SPEC CPU 2006
  - C++ applications of both SPEC CPU2006 int and SPEC CPU2006 fp matched applications well
- Defined benchmark in 2009 (?) as HEP-SPEC06
- At EOL, WLCG started looking into SPEC CPU 2017
  - Found to be a bad match



# HEP-SPEC06

- Benchmark: defining workload is not enough
  - Also need to describe the conditions of running
  - Chose conditions in 2009 that were as realistic as possible in view of WLCG CPU farms
    - gcc version and flags, 32bit app, as many concurrent processes as cores, ...
  - Scaling behaviour of real workload within about 10% of benchmark
  - Conditions have changed much since then
    - 64bit, new compilers/versions, flags, SMP processors, multi-threaded applications or pilots launching identical binaries multiple times, VMs and containers, ...
  - Still maintained the initial choices
    - Scaling behaviour matching real workload still surprisingly well (~ 20% or better)

# HEP-SPEC06 Usage

- From WLCG perspective, most importantly
  - Experiment requests and site pledges
  - Accounting of CPU usage
- Many sites also use it for procurements
- Initially designed as a tool for WLCG, found widespread use in other communities (not limited to HEP)

# HEP-SPEC06 Criticism

- Individual reports of scaling deviations of 40% and more
  - Some suggested replacements turned out to be worse for typical workload mixes
- Benchmark workload not typical of HEP applications
  - Would the reasonable scaling persist with non-x86 CPUs, for example? Well, we didn't have that issue (yet)...
- Running HEP-SPEC06 requires a software licence from SPEC
  - Strong desire to consider licence-free benchmarks as successor



# HEP-SPEC06: A Success Story

- Key reasons IMO:
  - Benchmark defined as one single number
  - Definition did not change during the HEP-SPEC06 lifetime



# HEP-SCORE

- Pretty much like in 2007/2008, benchmark experts got together and worked on a new benchmark
  - Without fixing the details of how to use it for a given purpose
- Landscape has changed completely, which implies new challenges and new opportunities

# Application Domains to Consider

- Compute facilities at WLCG sites
  - Still very much x86 dominated (mostly Intel, some AMD)
- Compute facilities used (quasi-)opportunistically
  - E.g. HPCs: Various processors (x86, POWER, ARM), various GPUs in various relations with CPUs
  - May see some of this soon at WLCG sites, too – even as part of the pledges
  - More may be coming, e.g. FPGAs
- We now have typical workloads for x86 CPUs, but not for non-x86 CPUs nor GPUs
  - Multi-step procedure?





# In Addition...

- HEP-SPEC06 already used outside WLCG
- Large communities want to profit from WLCG's structure and experience (e.g. Belle II, DUNE; SKA)
- Including astrophysics risks to become too broad, but doing something not useful to Belle II and/or DUNE would be stupid
  - That's why we proposed that they contribute to this TF

# Some clarifications

*There are some recurring questions/doubts that can be clarified here*

1. Does HEPscore favor a WL or an Exp. respect to another? Example the fastest one?
  - [Question received at the HSF-WLCG workshop]
  - **Answer: NO.** It's based on speed factors (see slide 4)
2. How to track the configurations?
  - **Answer:** All configurations are **versioned** (by hash string), and the versions are part of the **score report**
    - NB: The versions of the WLs are also part of this versioning. E.g.: two LHCb gen-sim versions  $\Rightarrow$  two score versions
3. Does HEPscore weights all WLs in the same manner?
  - **Answer:** It depends on **WLCG policy**
  - Technically both possibilities are in place: it can weight all WLs equally, or adopt tuned weights (see slide 4,5)
4. How often shall the HEPscore change **default** configuration?
  - **Answer:** It depends on **WLCG policy**.  
It could be never.  
Or only when a new production sw provides a major change in performance to be reproduced by the benchmark score.
5. How to assess pledges of old CPUs with new software?
  - **Answer: WLCG policy**, to be addressed in the **accounting infrastructure**. Scenarios: fixed score, or evolving score, etc
    - This aspect is true for any benchmark