

# **Overview of the WLCG HEP-SCORE Deployment Task Force Activity**

Helge Meinhard / CERN-RCS  
WLCG Benchmarking Workshop  
19 September 2022

# Why a CPU Benchmark?

- From WLCG perspective, most importantly
  - Experiment requests and site pledges
  - Accounting of CPU usage
- Many sites also use them for procurements



# Why a New CPU Benchmark?

- Current benchmark, HEP-SPEC06 introduced in January 2009, has several drawbacks
  - No longer modelling performance scaling of HEP applications sufficiently well (order of 10%) in all cases
    - Running conditions good in 2009, now totally obsolete
  - SPEC stopped supporting underlying SPEC-CPU 2006 in 2018



# How to Benchmark?

- Benchmarking for HEP **is not** about what can be obtained from a given CPU model optimising all possible running conditions
- It **is** about how a complete machine configuration performs under realistic batch service and experiment conditions at HEP sites
  - Depends not only on CPU model, but also on memory size and configuration, disk configuration, ...
  - Depends on experiment environments, compilers, flags, optimisations, ...



# Future HEP Benchmarking (1)

- Starting in 2018, benchmark experts got together and worked on a new benchmark: HEPiX Benchmarking Working Group co-chaired by Manfred Alef (KIT), Domenico Giordano (CERN) and Michele Michelotto (INFN Padua)
  - Several reports to GDB, HEPiX, WLCG MB
    - Domenico's report to HEPiX 26-Oct-2021:  
<https://indico.cern.ch/event/1078853/contributions/4576275/>
    - Domenico's presentation at GDB 09-Mar-2022:  
<https://indico.cern.ch/event/1096028/>
    - Large parts of Michele's report to HEPiX 26-Apr-2022:  
<https://indico.cern.ch/event/1123214/contributions/4821948/>
  - Without fixing the details of how to use it for a given purpose



# Future HEP Benchmarking (2)

- Result of the HEPiX benchmarking working group:
  - “HEP Benchmark Suite”: Framework for running containerised benchmarking workloads
    - Automatises benchmark runs
    - Ensures structured delivery and storage of results
  - “HEP Workloads”: Collection of (mostly) HEP workloads for which it is desirable to obtain performance information
    - Rather dynamic – add improved workloads, new compilers/flags/OS, ...
    - Sometimes called “the matrix” or “the basket”
  - “HEP Score”: Single number based on a defined, stable combination of defined, stable reference workloads; sample implementation (“HEP-SCORE\_beta”) using various LHC experiment workloads (from Run 2)
    - Good consistency with HEP-SPEC06



# HEP-SCORE Deployment Task Force

- WLCG Management board discussed and decided to launch a task force
- Started in November 2020, bi-weekly meetings since then
- Membership:
  - Experts on pledge etc. process, procurements, accounting
  - Experiment experts
    - Four LHC experiments
    - Belle 2, LIGO/Advanced VIRGO(/KAGRA), JUNO/BES III, (DUNE) etc.
  - Site experts
  - Some MB members



# Task Force Members

Name	Function	Name	Function
Tommaso Boccali	CMS	(Jeff Templon	Nikhef)
Simone Campana	WLCG	Andrea Valassi	LHCb
Domenico Giordano	Co-chair; Benchmarking WG	Ian Collier	STFC-RAL; APEL team
Michel Jouvin	Tier-2s	Gonzalo Merino	PIC
Walter Lampl	ATLAS	Fazhi Qi	JUNO, BES etc.
(Andrew McNab	DUNE)	Oxana Smirnova	NDGF
Helge Meinhard	WLCG and previous chair	Tony Wong	US Tier-1s
Bernd Panzer	CERN	Josh Willis	LIGO/Adv. VIRGO
Stefano Piano	ALICE	Matthias Schnepf	KIT and Benchmarking WG
Randy Sobie	Co-chair; Belle 2	Alastair Dewhurst	STFC-RAL
Andrew Melo	Vanderbilt U; US CMS T2	Miltiadis Alexis	WLCG accounting
Gonzalo Menendez Borge	WLCG benchmark infrastructure	Yan Xiaofei	JUNO, BES etc.



19-Sep-2022

Helge Meinhard (at) CERN.ch – Overview of task force activity



# Topics to Cover

- Compute facilities at WLCG sites
  - Still very much x86 dominated (still majority Intel, increasingly AMD)
- Compute facilities used (quasi-)opportunistically
  - Various processors (x86, POWER, ARM), e.g. at HPC sites
  - Various GPUs in various relations with CPUs, e.g. at HPC sites
  - May see some of this soon at WLCG sites, too – even as part of the pledges
  - More may be coming, e.g. FPGAs



# Strategy

- Start with CPU benchmarking on x86-based systems; look at other CPUs and/or GPUs later
  - Aim: single benchmark with a stable definition for at least (a typical CPU server lifetime cycle | a complete LHC machine cycle)
- Framework by benchmarking WG is very attractive
  - Use it to record behaviour of (wide) selection of workloads (not limited to HEP) across machines
- Study behaviour of large range of workloads over large range of benchmarking platforms
  - Aim: Select a sample of workloads (possibly with weights)
    - Sufficiently representative of real workloads on installations we (WLCG or HEP) use
    - Sufficiently small to be practical as a benchmark
    - Sufficiently precise and reproducible
- Propose a transition scenario from HEP-SPEC 06 to HEP-SCORE 22 for WLCG



# Status – Viewed from 36'000 Feet

- Workloads:
  - LHC experiments: event generation, detector simulation and digitisation, reconstruction; (analysis)
  - Other (mostly HEP) experiments: Belle II, JUNO, GW experiments, (DUNE)
  - Non-HEP workloads: HS06, SPEC-CPU 2017, DB12
  - All used in major benchmarking campaigns
- Benchmarking platforms: plethora from BNL, CERN, FNAL, IHEP, IJCLab, KIT, LIGO, NDGF-T1, Nikhef, PIC, RAL, ...
  - Long list of server configurations
  - Lots of results on individual workloads available
- Completed cross-check of available benchmark server configurations with server configurations used by ATLAS and CMS workloads
- Matrix of workloads and platforms filled - now ready to move on to defining HEP-SCORE 22 – aim of this workshop
- Exercise very valuable beyond HEP-SCORE 22
- Infrastructure: Upgraded wherever needed, fully tested, functional, used for collecting a large matrix of results



# (My Personal) Perspectives

Slide from April 2022  
... still valid

- Realistic scenario:
  - End June: Matrix of workloads and server configurations completely filled
  - Soon after: Analysis of individual results completed
  - End September: Workloads (and, if necessary, weights) for HEP-SCORE 2x selected
    - Present it at HEPiX in autumn
  - End November: Migration scenario proposed to and discussed with WLCG MB
  - January 2023: Start accounting both using HS06 and HEP-SCORE 2x
    - Other migration details to be sorted out until end November



# My Personal Conclusions

- Defining a HEP-SPEC 06 successor is a complex, ambitious task
  - Has taken many months rather than weeks, still work to do – in particular on migration plan
  - ... and then we will have to tackle non-x86 CPUs and GPUs
- Discussions and work in a constructive and collaborative spirit
- Steady progress leading to ... HEP-SCORE 22 (HS22)
  
- Thanks to
  - all working group and task force members
  - all experiment and site representatives supporting this work
  - Randy and Domenico for having agreed to take over from me
  - all of you for a very nice personal experience with benchmarking from 2007 on!

