

# HEPscore Deployment Task Force status

D. Giordano (CERN/IT), Randy Sobie (UViC)

WLCG MB


02 Aug 2022

# Overview

- ❑ Task Force Objective:
  - Propose an **HEPscore22** based upon the mix of HEP workloads collected in 2021/22 (“Run 3” workloads)
- ❑ Randy and Domenico took over from Helge the HEPscore task force convenership starting from June 15 ([meeting minutes](#))
- ❑ Recent activities:
  - Concluded benchmarking campaign using the available HEP workloads provided by the Experiments. **Analysis ongoing.**
  - Organized a survey addressed to the Task Force members to gather the "individual" expectations about HEPscore benchmark ([results](#))
  - Organizing an hybrid workshop on 19-20/09/2022 ([indico](#))

# Measurement campaign


- ❑ 11 HEP Workloads provided by the Experiments' experts
  - Alice gen-sim, Atlas gen/simMT/recoMT, Belle2 gen-sim-reco, CMS gen-sim/digi/reco, Grav Waves, Juno gen-sim-reco, LHCb gen-sim
  - in addition: HS06, SPEC CPU 2017
- ❑ ~7 months-long measurement effort from ~10 sites
  - ~30 CPU models across ~10 WLCG sites
  - Large amount of data collected (12k measurements), analysis performed in parallel to the data collection




# Analysis focus

- ✓ Confirm findings on SPEC benchmarks
  - Correlation between 32 vs 64 bits, HS06 vs SPEC 2017, cpp vs intrate
  
- ✓ HEP workloads Vs SPEC benchmarks
  - Which HEP workload better scales vs SPEC
  
- ✓ HEP workloads matrix
  - NB: Every element value comes from the fit on the scatter plot of  $WL_a$  vs  $WL_b$  over all benchmarked servers
  
- 🔍 Evaluate combinations for HEPscore22
  - Hundreds of combinations are possible.
  - We are restricting the set to the most reasonable ones

Mean deviation from fit  
2022-07-08 16:18



Mean deviation from fit  
All CPU-arch 2022-07-08 16:43



# Survey

- ❑ 50% participation
- ❑ Main questions
  - Should the WLCG keep HEPspec06 or use SPEC2017 or HEPscore?
    - Ans: **90%** “HEPscore”
  - How HEPscore should be composed
    - Ans: **60%** “Weighted by the experiment utilisation of WLCG”
  - Optimal time duration to run HEPscore
    - Ans: **<6h.** 40%: 3h (as HS06); 50%: 6h
- ❑ Interesting to read the motivations in the report

# HEPscore Workshop

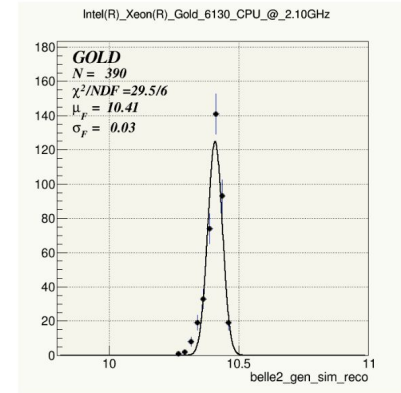
- ❑ 2 days workshop 19-20 of September
  - Registration is welcome to estimate the attendance on site
  - Agenda in backup slides
- ❑ Expected attendance not only from Task Force members and WG members
  - But also: Software experts from the experiments, GDB/HEPiX participants, etc
  - **Please advertise the meeting across the different WLCG channels**
    - **Already done via GDB**
- ❑ Some known clashes
  - C-RSG meeting involving the Computing Coordinators (only?) will be on the same days
  - CMS week

# Workshop Goal

- ❑ Review the workloads provided by the Experiments and their foreseen software evolution
  - **Main question:**
    - Are the current workloads stable?**
    - Do Experiments expect to radically change them in the near future?**
  - **Consequence: Set the feasibility of an HEPscore22**
- ❑ Identify a final HEPscore candidate benchmark to be presented to the WLCG community
- ❑ Learn from the Accounting WG how the migration will be organized operationally
- ❑ In addition
  - Discuss benchmarking of systems with accelerators (e.g. CPU+GPU)
  - Measurement of power consumption

# Summary

- ❑ The measurement campaign has been long
  - Slowed by the readiness of some workloads and unexpected issues on other workloads
- ❑ The experiments' workloads show excellent event/throughput resolution
  - A “Perf”-based analysis is ongoing to check the hotspots in each workload
- ❑ Current focus:
  - Identification of the HEPscore22 candidate and Workshop organization







# Monday 19

## Session legend

HEP Benchmark Suite   HEP Workloads   HEPscore

10:00	<b>Welcome and purpose of the workshop</b> <i>6/R-012 - conference room, CERN</i>	10:00 - 10:05
	<b>Overview o the work done by the Task Force and the Working Group</b> <i>6/R-012 - conference room, CERN</i>	10:05 - 10:35
	<b>Analysis of the Measurement Campaign</b> <i>6/R-012 - conference room, CERN</i>	10:35 - 11:15
11:00	<b>HEPscore candidates</b> <i>6/R-012 - conference room, CERN</i>	11:15 - 11:45
	<b>The Benchmark DB (demo how to use it)</b> <i>6/R-012 - conference room, CERN</i>	11:45 - 12:05
12:00	<b>Lunch Break</b>	

14:00	<b>Profiling of HEP Workloads</b> <i>6/R-012 - conference room, CERN</i>	14:00 - 14:30
	<b>Alice Workload</b> <i>6/R-012 - conference room, CERN</i>	14:30 - 14:45
	<b>Atlas worloads</b> <i>6/R-012 - conference room, CERN</i>	14:45 - 15:00
15:00	<b>CMS worloads</b> <i>6/R-012 - conference room, CERN</i>	15:00 - 15:15
	<b>IGWN workload</b> <i>6/R-012 - conference room, CERN</i>	15:15 - 15:30
	<b>Break: Coffee break</b> <i>6/R-012 - conference room, CERN</i>	15:30 - 15:50
16:00	<b>Juno workload</b> <i>6/R-012 - conference room, CERN</i>	15:50 - 16:05
	<b>LHCb workload</b> <i>6/R-012 - conference room, CERN</i>	16:05 - 16:20
	<b>The experience of some data centres in running the suite</b> <i>6/R-012 - conference room, CERN</i>	16:20 - 17:00



# Tuesday 20

Extension of HEPscore beyond x86  Selection of HEPscore

09:00	<b>Integration of HEPscore in the accounting system: Multiple talks</b> <i>6/R-012 - conference room, CERN</i> 09:00 - 09:45
10:00	<b>Golden candidates for HEPscore</b> <i>6/R-012 - conference room, CERN</i> 09:45 - 10:15
	<b>Madgraph on GPU</b> <i>6/R-012 - conference room, CERN</i> 10:25 - 10:45
	<b>CMS ParticleFlow with ML</b> <i>6/R-012 - conference room, CERN</i> 10:45 - 11:05
11:00	<b>Coffee break</b> <i>6/R-012 - conference room, CERN</i> 11:05 - 11:25
	<b>CMS HLT-like</b> <i>6/R-012 - conference room, CERN</i> 11:25 - 11:40
	<b>Experience running workloads on CPU+GPU at HPC</b> <i>6/R-012 - conference room, CERN</i> 11:40 - 12:00
12:00	<b>Run HEPscore "light" on grid</b> <i>6/R-012 - conference room, CERN</i> 12:00 - 12:20
	<b>Lunch Break</b>

14:00	<i>6/R-012 - conference room, CERN</i> 12:20 - 13:50
	<b>Measurements of Power Consumption</b> <i>6/R-012 - conference room, CERN</i> 13:50 - 14:50