HEPScore
A new CPU benchmark for WLCG computing

Randall Sobie, University of Victoria

on behalf of the

HEPiX Benchmarking Working Group

Domenico Giordano and Michele Michelotto (co-chairs)

Luca Atzori, Jean-Michel Barbet, Gonzalo Menéndez Borge, Christopher Hollowell, Ladislav Ondris, Andrea Sciaba, Emanuele Simili,
Randall Sobie, David Southwick Tristan Sullivan, Natalia Szczepanek, Andrea Valassi

WLCG HEPscore Deployment Task Force

Domenico Giordano and Randall Sobie (co-chairs)

Julia Andreeva, Giuseppe Andronico, Gonzalo Merino Arevalo, Tommaso Boccali, Gonzalo Menendez Borge, Concezio Bozzi, Simone Campana, Ian Collier
Alessandro Di Girolamo, Michel Jouvin, Walter Lampl, James Robert Letts, Zach Marshall, Helge Meinhard, Andrew Malone Melo, Bernd Panzer-Steindel
Stefano Piano, Danilo Piparo, Fazhi Qi, Matthias Jochen Schnepf, Oxana Smirnova, Jeff Templon, Andrea Valassi, Joshua Lee Willis, Tony Wong

HEPiX Workshop, Taipei, March 2023

Introduction

CPU benchmarks are an important part of the WLCG infrastructure

Experiment requests and site pledges
Accounting of CPU usage
Procurement of new hardware

The current WLCG benchmark, HEPSpec06 (2009), has several drawbacks

Not representative of HEP workloads (HEP workloads are more performant on newer hardware)
HEPSpec06 is the 32bit version
SPEC stopped supporting the underlying SPEC-CPU 2006 benchmark (2018)

WLCG needs a benchmark for other processors (ARM and GPUs)

We have HEP workloads for ARM from a number of experiments
Workloads with GPUs are just emerging

The current HEPSpec06 benchmark

HS06 is a suite of 7 C++ applications
Subset of SPEC CPU® 2006 benchmark
SPEC's industry-standardized

CPU-intensive benchmark suite
However, none of the applications are an event-based detector simulation or reconstruction

Correlated with HEP Workloads in 2009

Execution time of approximately 3 hours

Bmk	Int vs Float	Description
444.namd	CF	92224 atom simulation of apolipoprotein A-I
447.dealII	CF	Numerical Solution of Partial Differential Equations using the Adaptive Finite Element Method
450.soplex	CF	Solves a linear program using the Simplex algorithm
453.povray	CF	A ray-tracer. Ray-tracing is a rendering technique that calculates an image of a scene by simulating the way rays of light travel in the real world
471.omnetpp	CINT	Discrete event simulation of a large Ethernet network.
473.astar	CINT	Derived from a portable 2D path-finding library that is used in game's AI
483.xalancbmk	CINT	XSLT processor for transforming XML documents into HTML, text, or other XML document types

J. Phys.: Conf. Ser. 219 (2010) 052009
CHEP-09

Experiment workloads as a CPU benchmark

First proposal of HEP Benchmark with containerized HEP applications
WLCG Workshop Manchester 2017

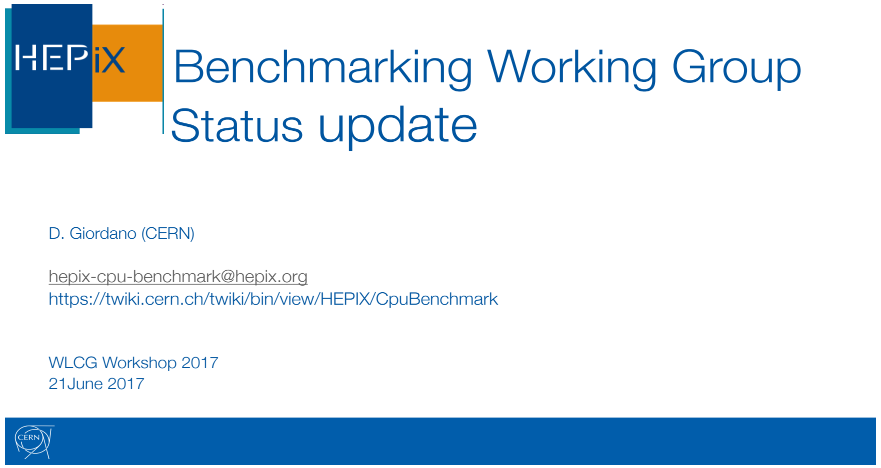

Experiment workloads
Complex systems with hundreds of algorithms
Event based
Event generation, digitization, simulation, reconstruction
Analysis applications not considered for benchmark

Evolution of experiment software
New programming approaches (multithreading and vectorization)

Need for a benchmark that adapts to the emerging technologies
HPCs, GPUs, ARM

https://indico.cern.ch/event/609911/contributions/2620190/attachments/1480455/2295576/WLCG_Workshop_2017_benchmarking_giordano.pdf

HEP Benchmarks Project (HEPiX WG)

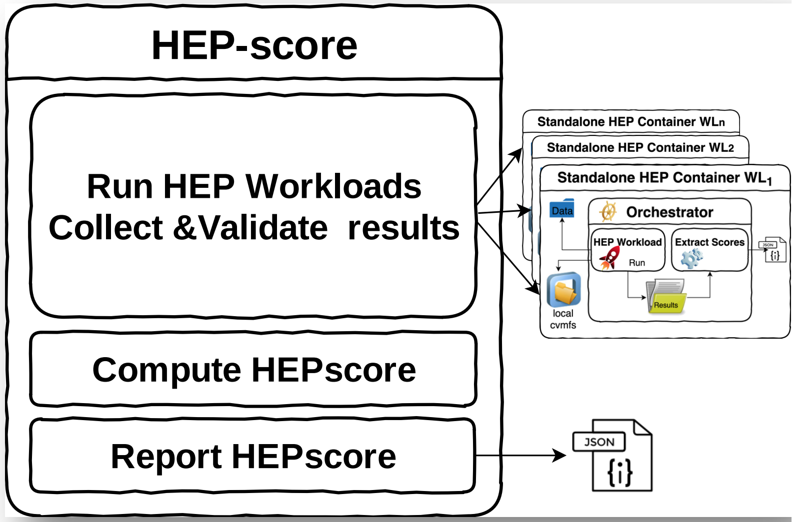

Goal is to develop a single benchmark, called *HEPscore*

Based on experiment workloads

HEPiX WG developed the *HEPscore Suite* infrastructure

Orchestrator of multiple benchmark (HEPscore, HS06, SPEC CPU2017)

Central collection of benchmark results

Presented at previous HEPiX meetings and published in CSBS

***Workloads* provided by many experiments**

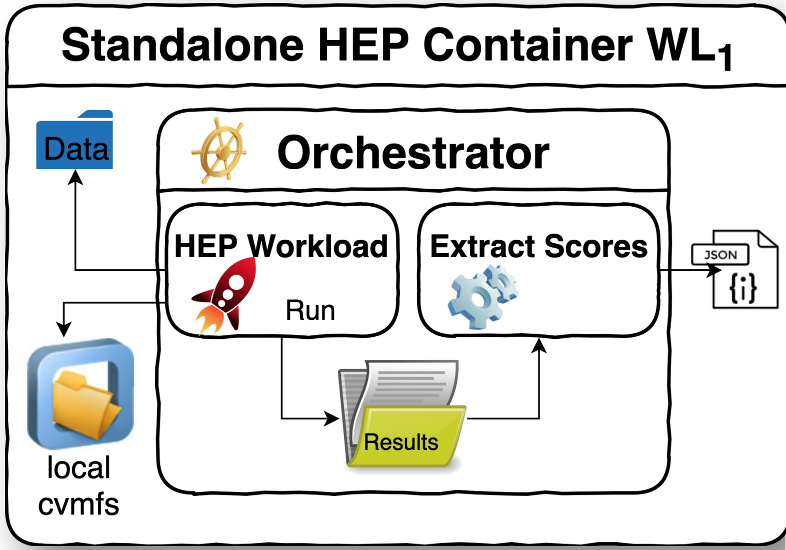

Experiment Workloads

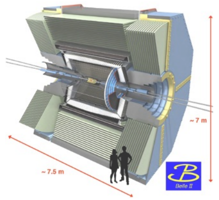

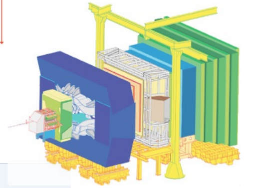

Requirements

Provide consistent CLI, report structure, error logging
Reproducible results
No access to remote data, databases, ..
Portable with a modestly-sized package distribution
Runs in a reasonable period (tuneable with the number of events)
Long term support

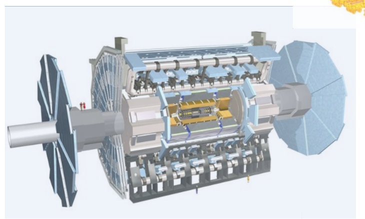

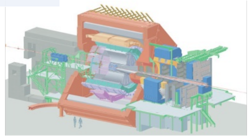

Workloads provided by 7 experiments

ALICE, ATLAS, CMS, LHCb
Belle II, Gravity Wave (LIGO/VIRGO), JUNO

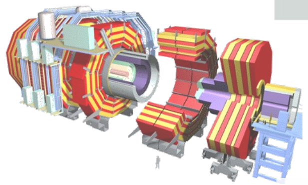

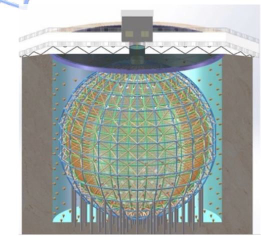

Typically, event generation and digitization, MC simulation and reconstruction
Often using very complex and busy event topologies

Initial set of workloads provided in 2021-2022
Updated for the latest software and ARM-compatibility in late-2022 and 2023

HEP Benchmark Container

Standalone containers for each workload

Encapsulates all and only the dependencies needed to run the benchmarks

Build for x86 and ARM architectures

Components of an HEP workload

SW repository (in general distributed via CVMFS)

Input data (event and conditions data)

Orchestrator script per workload

Set the environment

Runs the application

Parse the output

GitLab Registry for container distribution

Docker & Singularity/Apptainer

Documentation and instructions provided

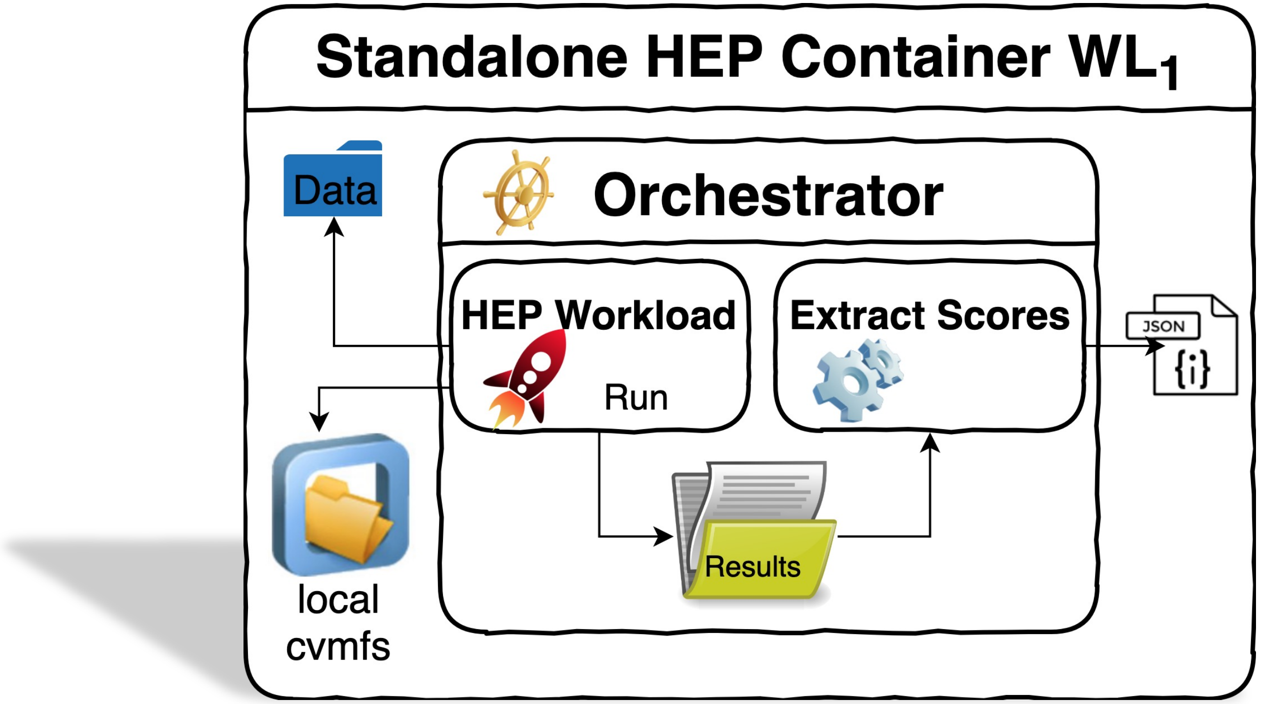

Initial set of workloads

alice_gen_sim_reco

atlas_gen_sherpa
atlas_sim_mt
atlas_reco_mt

belle2_gen_sim_reco

cms_reco
cms_digi
cms_gen_sim

juno_gen_sim_reco

igwn_pe (Gravity Wave)

lhcb_gen_sim

Each workload has been run and validated on a set of CERN servers
Reliable/reproducible to < 1%

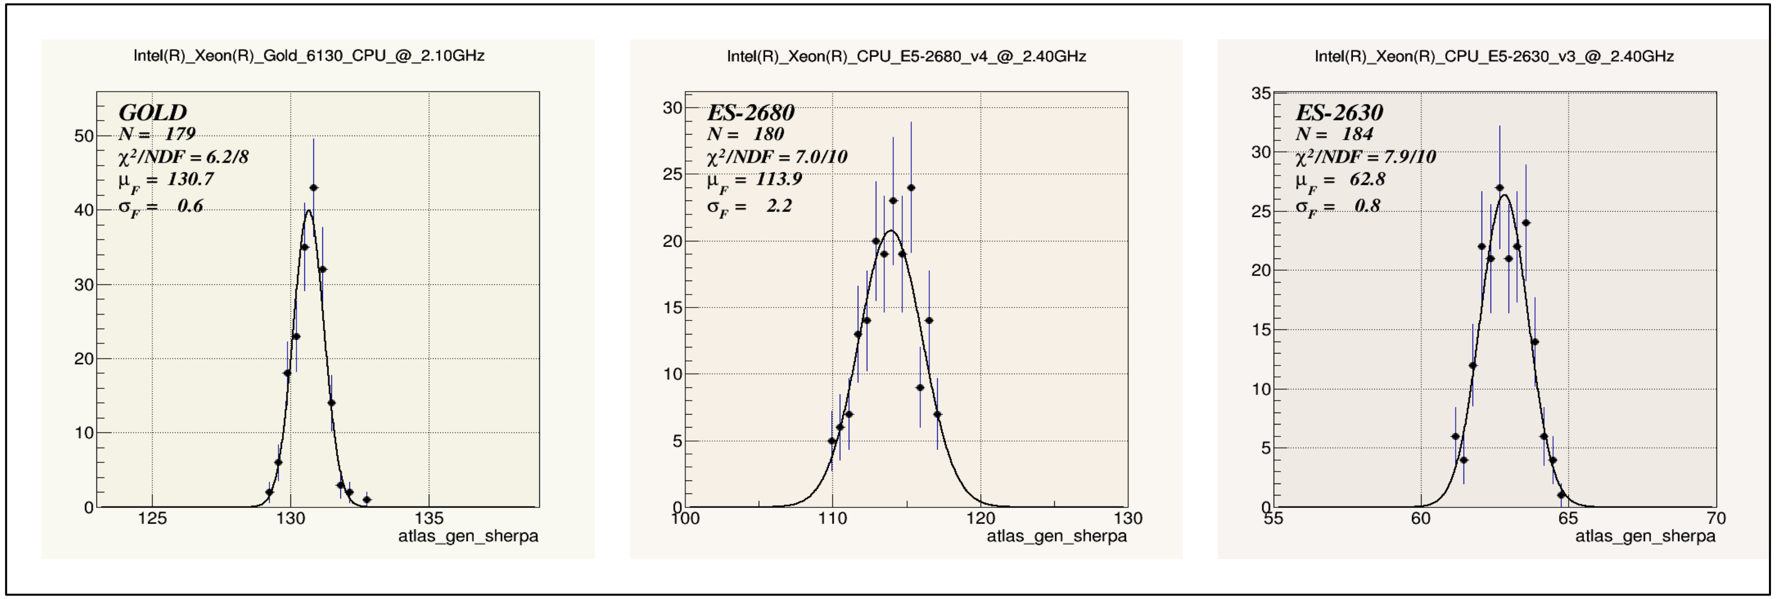

Measurement campaign 2021-2022

Accumulated a large set of measurements in 2021-2022

HEP-SPEC06, SPEC2017 and HEP Workloads

Approximately 70 different “systems” (CPU, cores, site, hyper-threading) around the world

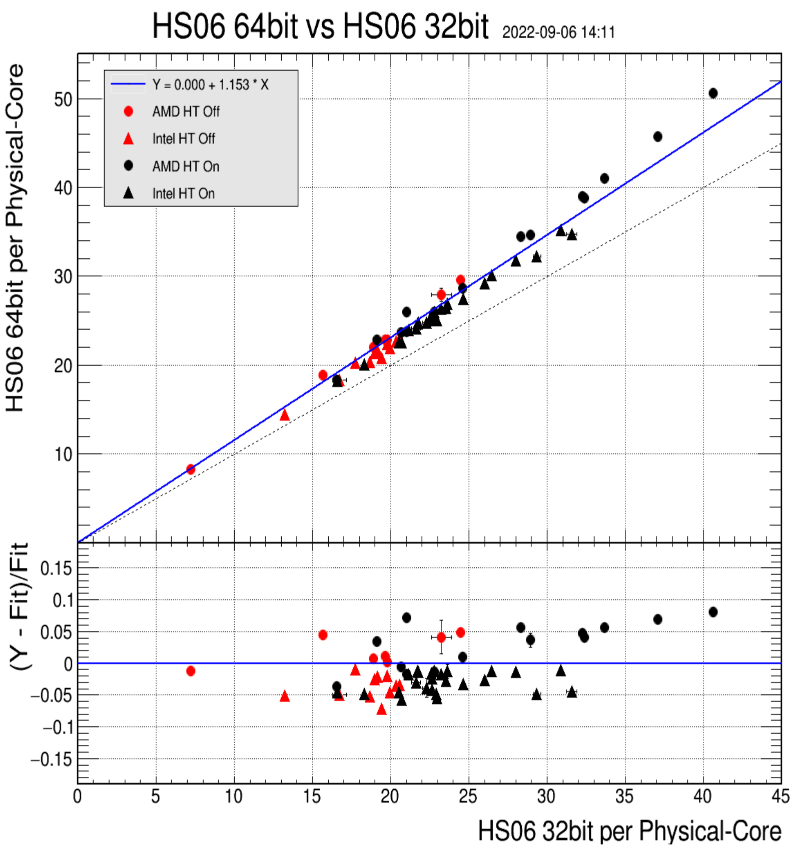

Plots
Normalized to the number of physical cores

Red	HT Off
Black	HT On
Circle	AMD
Triangle	Intel

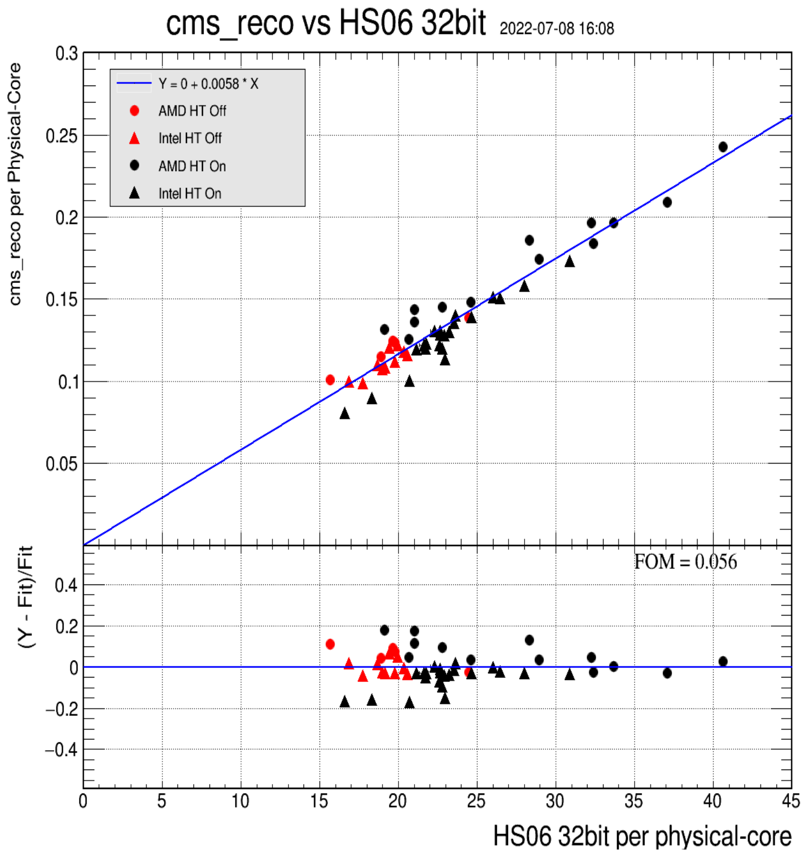

Selection of Workloads for HEPscore

HEPscore workshop (19–20 Sept 2022)

Representatives of the experiments, sites, WLCG MB

Feedback on the composition of HEPscore and strategy for adoption of the new benchmark

Consensus of the TF and Workshop participants:

Reflect the relative CPU usage of the experiments and application

Run in a timely manner 3-6 hours

Valid for one or more LHC beam period

HEPscore candidate proposed at the Workshop

7 workloads proposed (2-CMS, 2-ATLAS, LHCb, ALICE, BelleII)

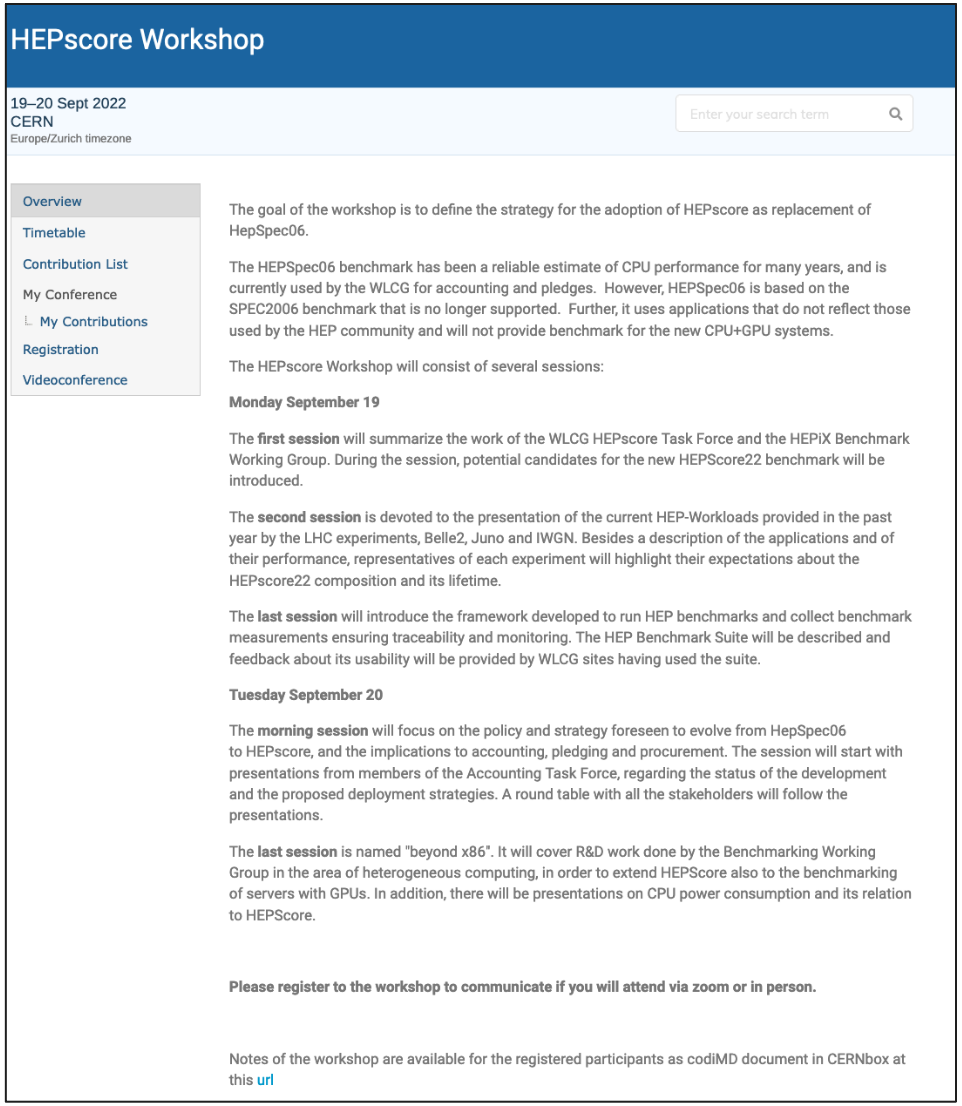
HEPscore Workshop

19–20 Sept 2022
CERN
Europe/Zurich timezone

Enter your search term

Overview
Timetable
Contribution List
My Conference
My Contributions
Registration
Videoconference

The goal of the workshop is to define the strategy for the adoption of HEPscore as replacement of HepSpec06.

The HEPSpec06 benchmark has been a reliable estimate of CPU performance for many years, and is currently used by the WLCG for accounting and pledges. However, HEPSpec06 is based on the SPEC2006 benchmark that is no longer supported. Further, it uses applications that do not reflect those used by the HEP community and will not provide benchmark for the new CPU+GPU systems.

The HEPscore Workshop will consist of several sessions:

Monday September 19

The **first session** will summarize the work of the WLCG HEPscore Task Force and the HEPiX Benchmark Working Group. During the session, potential candidates for the new HEPScore22 benchmark will be introduced.

The **second session** is devoted to the presentation of the current HEP-Workloads provided in the past year by the LHC experiments, Belle2, Juno and IWGN. Besides a description of the applications and of their performance, representatives of each experiment will highlight their expectations about the HEPscore22 composition and its lifetime.

The **last session** will introduce the framework developed to run HEP benchmarks and collect benchmark measurements ensuring traceability and monitoring. The HEP Benchmark Suite will be described and feedback about its usability will be provided by WLCG sites having used the suite.

Tuesday September 20

The **morning session** will focus on the policy and strategy foreseen to evolve from HepSpec06 to HEPscore, and the implications to accounting, pledging and procurement. The session will start with presentations from members of the Accounting Task Force, regarding the status of the development and the proposed deployment strategies. A round table with all the stakeholders will follow the presentations.

The **last session** is named "beyond x86". It will cover R&D work done by the Benchmarking Working Group in the area of heterogeneous computing, in order to extend HEPScore also to the benchmarking of servers with GPUs. In addition, there will be presentations on CPU power consumption and its relation to HEPScore.

Please register to the workshop to communicate if you will attend via zoom or in person.

Notes of the workshop are available for the registered participants as codiMD document in CERNbox at this url

Lancaster WLCG Workshop (Nov 2022)

Presentation of the HEPscore candidate

Key outcomes of the Workshop
And approved by the WLCG Management Board in December 2022

Existing equipment at the sites does not need to be re-benchmarked
Sites with heterogenous x86 resources will continue to calculate a site-average that is posted to the WLCG Accounting System

New hardware can be benchmarked with either HEPSPEC06 or HEPscore in 2023

HEPscore will be normalized to HEPSPEC06 on a reference machine at CERN

Aim for a HEPscore than be used for both x86 and ARM processors but default to an x86-HEPscore

https://indico.cern.ch/event/1162261/contributions/5092745/attachments/2543843/4380269/Sobie-WLCG-HEPScore.pdf

Update of workloads Jan-Feb 2023

Workloads for ARM processors and refresh to latest releases
Multi-architecture (“ma”)
Many are multi-threaded (“mt”)
All LHC workloads for Run3

Significant revisions:

lhcb_sim_run3_ma
identification of a CPU-consuming Geant4 geometry location function – simple fix reduced simulation time by x2

atlas_gen_sherpa_ma
Significant speed up of generator

alice_digi_reco_core_run3_ma
Pb-Pb event reconstruction

x86 and ARM workloads

Emergence of the importance of power consumption and environmental impact

Presentation at ACAT 2022 Conference (Oct 2022)

Showed that power consumption of an ARM processor was 45% lower than x86 and processing time was shorter for the atlas_sim workload

https://indico.cern.ch/event/1106990/contributions/4991256/attachments/2534801/4362468/PoW_ACAT2022.pdf

The experiments provided updated workloads for both x86 and ARM processors

Concerted effort by all experiments to have an x86/ARM compatible workload – completed Mar 1 2023

Currently in the validation phase (March 2023)

Running the workloads on a wide range of servers: x86 (Intel/AMD), ARM, hyperthreading on and off

Comparing old and new workloads, checking impact of each workload, other cross checks

HEPscore23 workloads

Follow measurement strategy used in HEPSPEC06 benchmark
Each workload is run 3-times and the geometric mean is calculated
Workload output is typically events/second

Server score
The "score" of a server is geometric mean of the WL-scores

Reference machine
Normalize to the score from a reference machine at CERN (Intel Gold-6326)

$$\bar{x} = \left(\prod_{i=1}^{n} x_i^{w_i} \right)^{1/\sum_{i=1}^{n} w_i}$$

Time to generate HEPscore23 benchmark on reference machine
Intel(R) Xeon(R) Gold 6326 CPU @ 2.90GHz
"Time" for the second of three runs of the workload

Workload	Threads	Events/thr	Time
alice_digi_reco_run3_ma	4	10	930 s
atlas_gen_sherpa_ma	4	200	350 s
atlas_reco_mt_ma	4	10	910 s
belle2_gen_sim_reco_ma	1	50	320 s
cms_gen_sim_run3_ma	4	10	500 s
cms_reco_run_ma	4	50	740 s
lhcb_sim_run3_ma	1	10	620 s

Time to measure HEPscore23 approximately 3.5 hours

HEPscore23 validation – March 2023

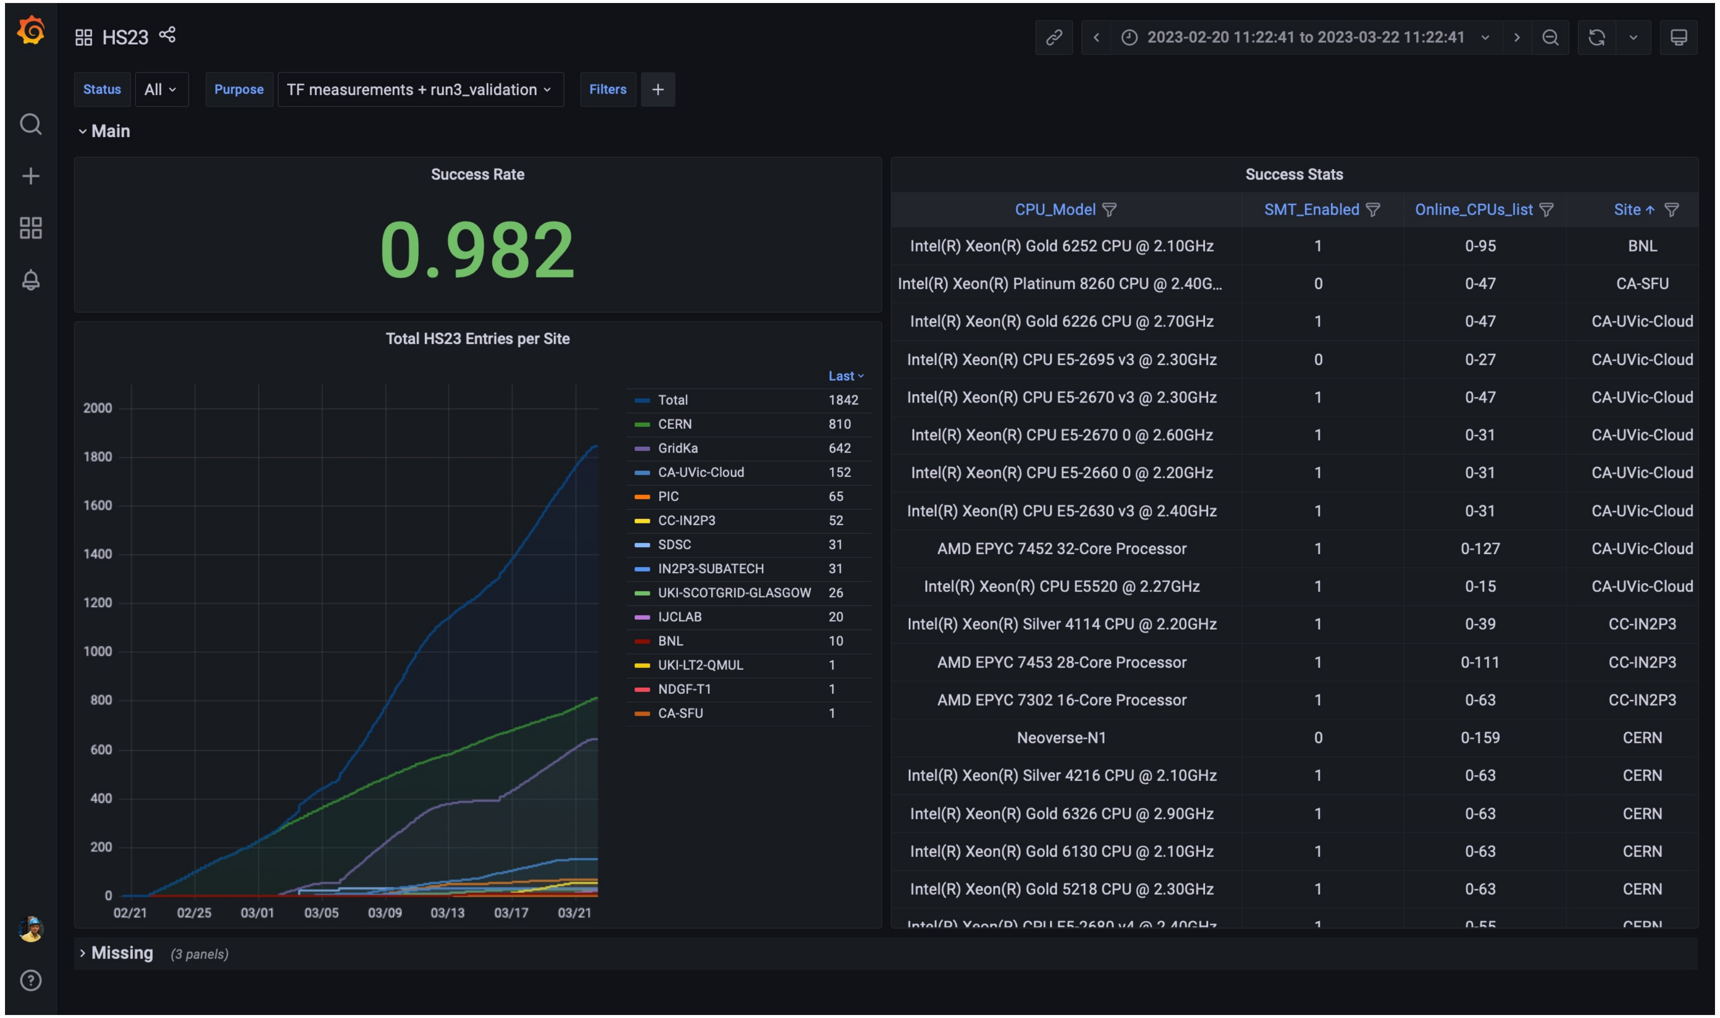

Running the workloads on a wide range of servers at sites around the world

Intel, AMD and ARM

Confirm reliability and stability on CERN Testbed

We have a set of servers at CERN for initial testing of workloads and HEPscore23

“Reference machine”: Intel® Xeon® Gold 6326 CPU @ 2.90 GHz (HT=On)

Plots of HEPscore on Intel (Reference Machine), AMD and Neoverse (ARM) servers

Scatter of results <0.5% for all systems

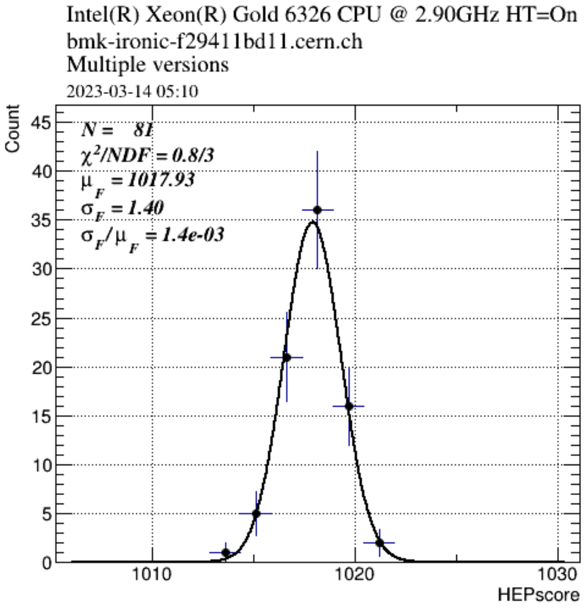

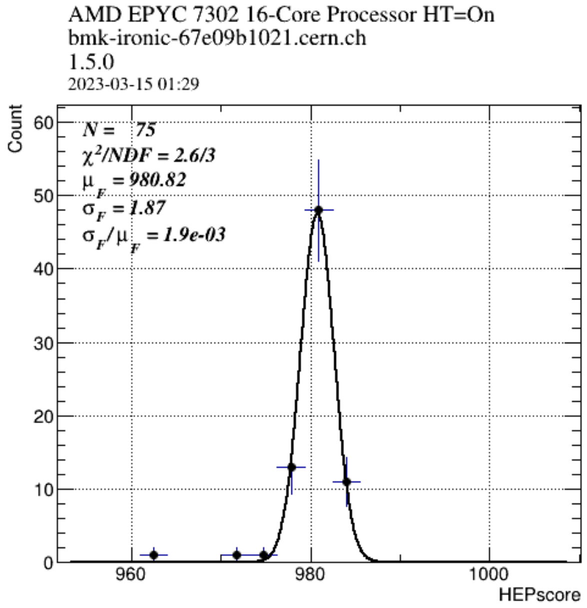

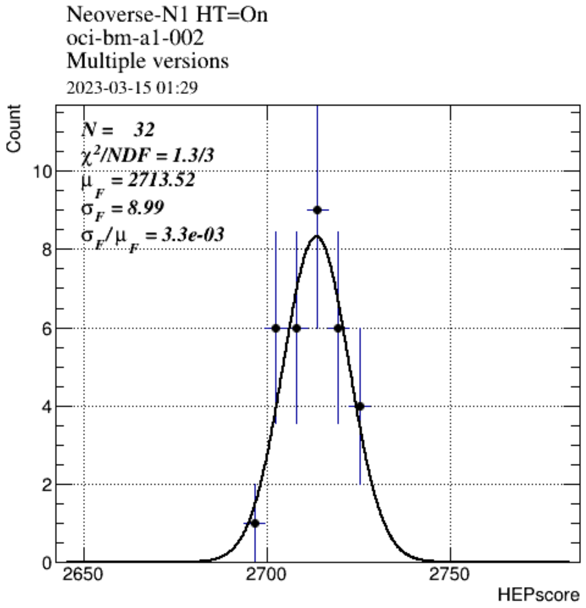

HEPscore23 vs HEPSpec06

Current results with measurements from a variety of servers

Solid-blue line is a fit to the results (constrained to the origin)
Dashed-line has unity slope – normalized to reference machin

SMT	CPU Model	Year	Site	HS06	HS06 std	HS23	HS23 std	HS23/HS06
Off	AMD EPYC 7302 16-Core Processor	2019	CERN	782.63	0.89	767.80	1.53	0.98
Off	Intel Core Processor (Haswell, no TSX, IBRS)	2013	NDGF-T1	403.79	10.22	404.37	1.31	1.00
Off	Intel(R) Xeon(R) CPU E5520 @ 2.27GHz	2009	IN2P3-SUBATECH	105.86	0.49	79.36	0.88	0.75
Off	Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz	2014	CERN	295.73	1.42	270.69	0.41	0.92
Off	Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz	2016	CERN	425.92	1.82	408.55	0.93	0.96
Off	Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz	2012	GridKa	282.88	21.80	200.81	2.19	0.71
Off	Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz	2016	CERN	547.38	1.63	537.21	0.99	0.98
Off	Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz	2019	CERN	632.38	1.13	611.10	0.70	0.97
Off	Intel(R) Xeon(R) Gold 6130 CPU @ 2.10GHz	2017	CERN	600.20	1.41	601.35	1.27	1.00
Off	Intel(R) Xeon(R) Gold 6326 CPU @ 2.90GHz	2021	CERN	857.85	1.91	864.16	1.51	1.01
Off	Intel(R) Xeon(R) Silver 4216 CPU @ 2.10GHz	2019	CERN	613.28	1.27	592.04	1.83	0.97
Off	Neoverse-N1	2020	CERN	2587.89	2.91	2713.75	7.24	1.05
On	AMD EPYC 7302 16-Core Processor	2019	CC-IN2P3	1031.43	2.73	1012.19	0.85	0.98
On	AMD EPYC 7302 16-Core Processor	2019	CERN	1036.27	2.26	981.15	4.18	0.95
On	AMD EPYC 7452 32-Core Processor	2019	PIC	1573.55	6.83	1516.77	4.71	0.96
On	AMD EPYC 7453 28-Core Processor	2021	CC-IN2P3	1584.59	4.24	1675.02	9.63	1.06
On	AMD EPYC 7702 64-Core Processor	2019	IJCLAB	2686.15	5.82	2690.69	11.95	1.00
On	AMD EPYC 7702 64-Core Processor	2019	GridKa	2643.00	6.35	2539.81	9.19	0.96
On	AMD EPYC 7742 64-Core Processor	2019	GridKa	2917.07	25.68	2944.49	10.54	1.01
On	Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz	2014	CERN	364.59	0.69	319.12	1.97	0.88
On	Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz	2014	GridKa	361.43	0.81	292.34	2.53	0.81
On	Intel(R) Xeon(R) CPU E5-2640 v3 @ 2.60GHz	2014	PIC	371.89	1.21	340.78	2.42	0.92
On	Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz	2016	CERN	521.14	0.97	482.16	2.14	0.93
On	Intel(R) Xeon(R) CPU E5-2660 0 @ 2.20GHz	2012	CA-UVic-Cloud	328.64	1.28	229.08	0.23	0.70
On	Intel(R) Xeon(R) CPU E5-2665 0 @ 2.40GHz	2012	GridKa	332.09	1.59	222.50	1.18	0.67
On	Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz	2012	GridKa	350.57	7.90	238.93	1.15	0.68
On	Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz	2016	CERN	659.85	1.97	631.28	3.82	0.96
On	Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz	2016	PIC	661.91	1.93	621.69	2.86	0.94
On	Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz	2019	CERN	788.18	1.49	708.95	1.01	0.90
On	Intel(R) Xeon(R) Gold 6130 CPU @ 2.10GHz	2017	CERN	734.61	2.15	691.03	8.04	0.94
On	Intel(R) Xeon(R) Gold 6326 CPU @ 2.90GHz	2021	CERN	1015.30	2.58	1018.11	2.31	1.00
On	Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GHz	2017	CC-IN2P3	453.66	1.35	416.88	0.49	0.92
On	Intel(R) Xeon(R) Silver 4216 CPU @ 2.10GHz	2019	IJCLAB	716.83	8.91	675.52	3.24	0.94
On	Intel(R) Xeon(R) Silver 4216 CPU @ 2.10GHz	2019	CERN	769.76	1.32	714.71	3.28	0.93

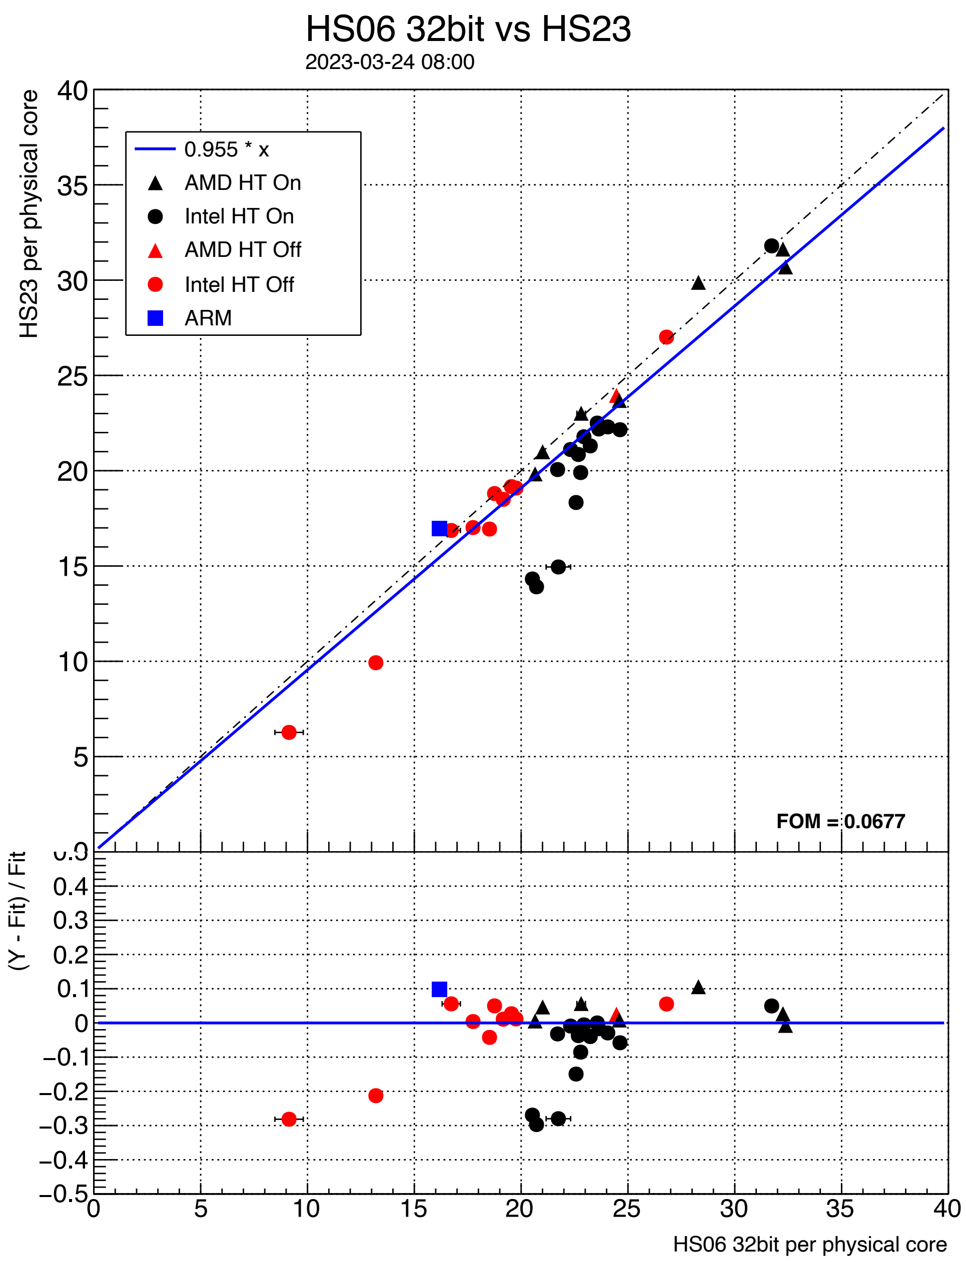

HEPscore23 by Year of Server

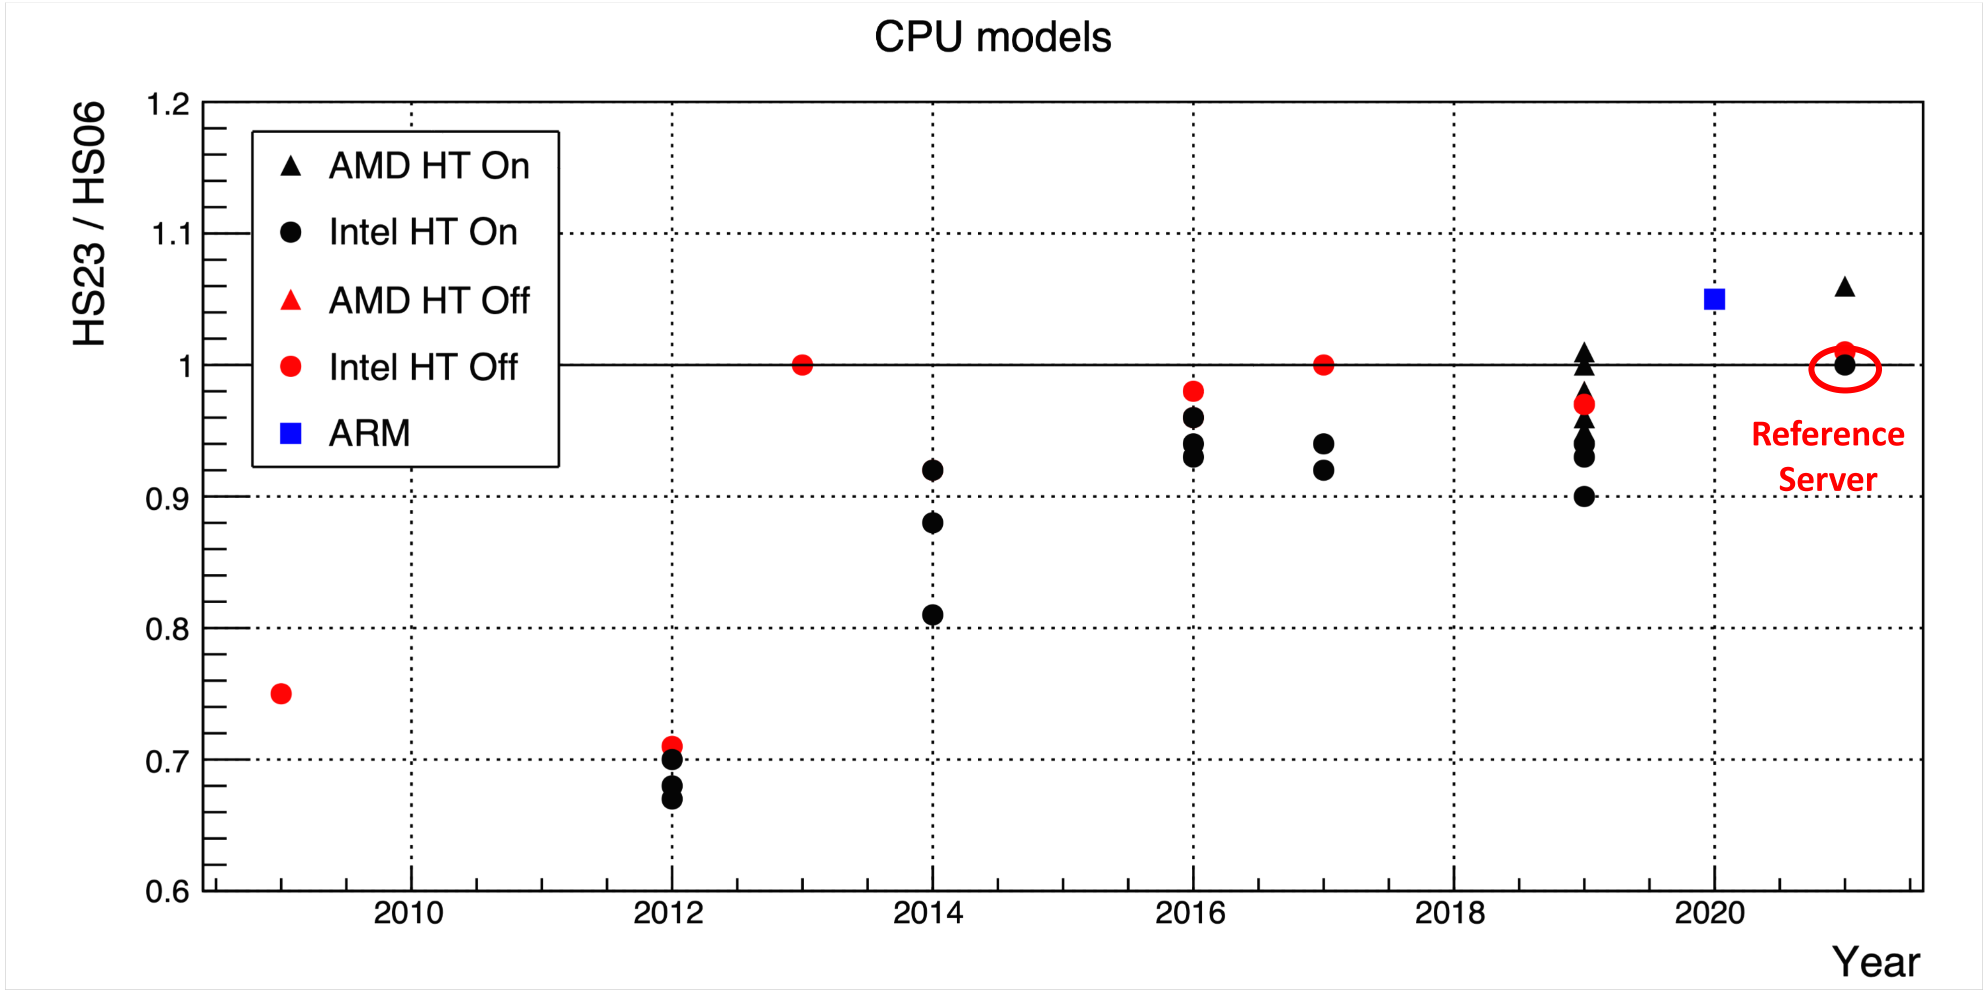

HEPscore23 is more performant relative to HEPSPEC06 on newer architectures

“Reference machine”: Intel® Xeon® Gold 6326 CPU @ 2.90 GHz (HT=On)

Transition plan

Plan presented at the WLCG Lancaster Workshop
"HEPScore - transition plan to the new benchmark (Accounting Group)" ***Julia Andreeva (CERN)***
https://indico.cern.ch/event/1162261/contributions/5117866/attachments/2544039/4380614/NewBenchmarkWS.pdf

Timescales are driven by the WLCG cycle for pledges (scrutiny group)
Pledges for FY2025 are made in Oct 2023

In 2023, sites are encouraged to run HEPscore23 for new hardware
(prior to the pledge deadline)

HEPscore23 will be normalized to HEPSPEC06 on the reference machine to simplify the calculation
(and to allow for smooth transition of tables and plots)

Sites will publish their information to the Accounting Group (instructions to be provided in April 2023)

Updates to HEPScore (from the WLCG MB meeting Dec 2022)
If " .. an experiment believes that an update to HEPScore is needed, they should submit the request to the MB, which will then task the relevant experts with understanding costs and benefits."

References

WLCG Benchmark Task Force (Nov 2020)

H. Meinhard CERN/IT (Initial Chair)
D. Giordano CERN/IT and R.Sobie Victoria (Co-Chairs July 2022)

HEPiX Benchmark Working Group

D. Giordano CERN/IT

Experts in software, accounting

Site administrators

Computing and Software for Big Science (2021) 5:28
https://doi.org/10.1007/s41781-021-00074-y

ORIGINAL ARTICLE

Check for updates

HEPiX Benchmarking Solution for WLCG Computing Resources

Domenico Giordano[1] · Manfred Alef[2] · Luca Atzori[1] · Jean-Michel Barbet[3] · Olga Datskova[1] · Maria Girone[1] · Christopher Hollowell[5] · Martina Javurkova[4] · Riccardo Maganza[1] · Miguel F. Medeiros[1] · Michele Michelotto[6] · Lorenzo Rinaldi[9] · Andrea Sciabà[1] · Randall J. Sobie[7] · David Southwick[1,8] · Tristan Sullivan[7] · Andrea Valassi[1]

Received: 30 August 2021 / Accepted: 8 November 2021 / Published online: 14 December 2021

Abstract
The HEPiX Benchmarking Working Group has developed a framework to benchmark the performance of a computational server using the software applications of the High Energy Physics (HEP) community. This framework consists of two main components, named HEP-Workloads and HEPscore. HEP-Workloads is a collection of standalone production applications provided by a number of HEP experiments. HEPscore is designed to run HEP-Workloads and provide an overall measurement that is representative of the computing power of a system. HEPscore is able to measure the performance of systems with different processor architectures and accelerators. The framework is completed by the HEP Benchmark Suite that simplifies the process of executing HEPscore and other benchmarks such as HEP-SPEC06, SPEC CPU 2017, and DB12. This paper describes the motivation, the design choices, and the results achieved by the HEPiX Benchmarking Working group. A perspective on future plans is also presented.

Keywords CPU benchmark · GPU benchmark · High throughput computing · WLCG · LHC computing · HEP experiments · High-Energy Physics · Heterogeneous computing

Presentations:

2021 CHEP (journal paper)

2022 HEPiX
2022 ACAT (proceedings pending)
2022 Benchmark Workshop (Indico)
2022 WLCG Workshop (Indico)
2022 WLCG GDB and MB

2023 HEPiX (this meeting)
2023 CHEP in Norfolk (May 2023)

Summary

On track to deliver a new CPU benchmark (HEPscore23) in April 2023
For x86 and ARM systems

Sites are expected to run HEPscore23 on newly procured hardware
Existing hardware need not be re-benchmarked
Instructions and documentation provided in April 2023

HEPscore23 will be normalized to HEPSPEC06 on the reference machine
Simplify the transition and the calculation of pledges and accounting statistics
At some point (TBD), HEPscore23 will become the default WLCG benchmark

Ongoing efforts to develop a benchmark for GPUs
Prototype workloads are under study (no time estimate for a new benchmark)

Other studies on power consumption and fast benchmarks