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

CHEP 2023, Norfolk, May 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

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

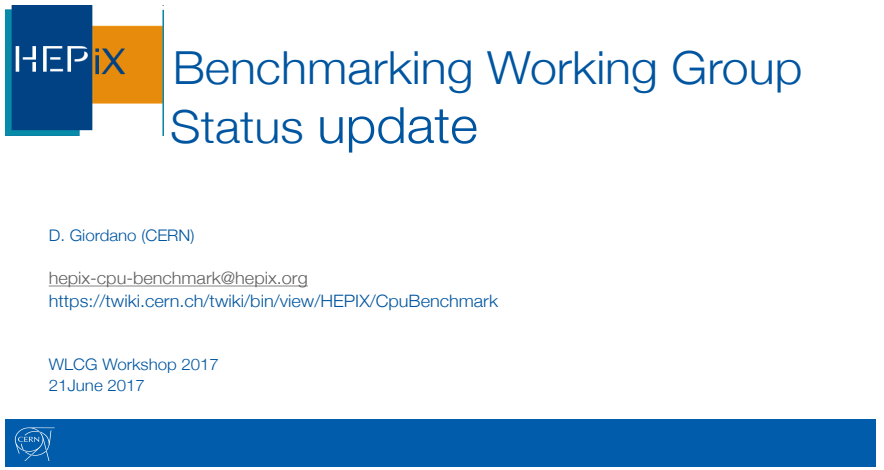

Evolution of experiment software
New programming approaches (multithreading and vectorization)

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**

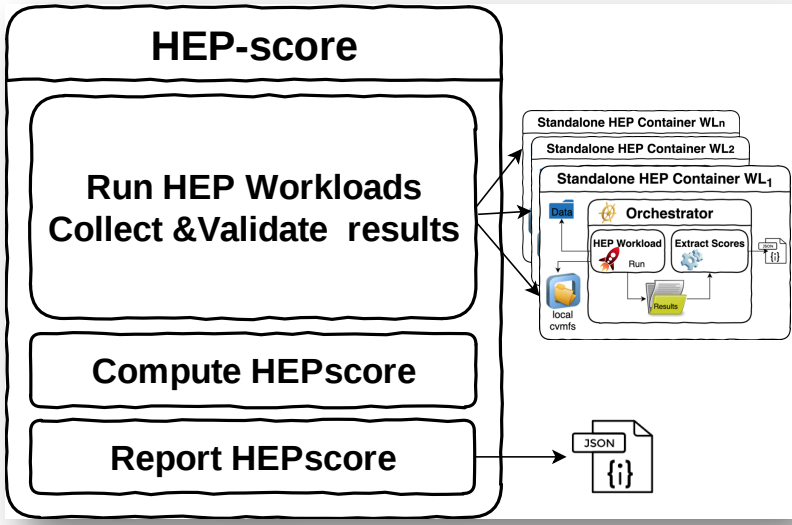

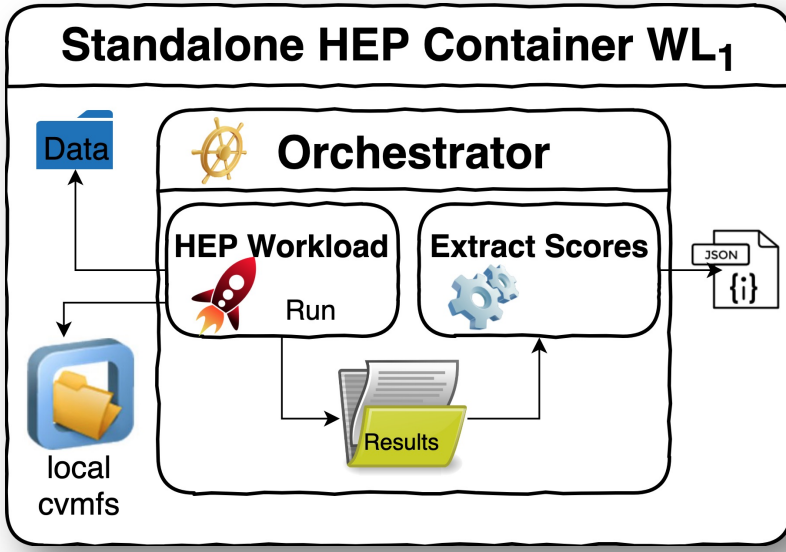

Experiment Workloads

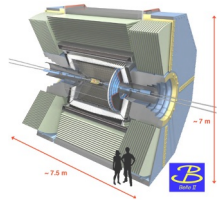
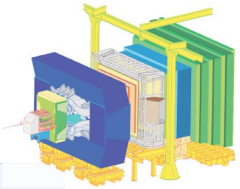
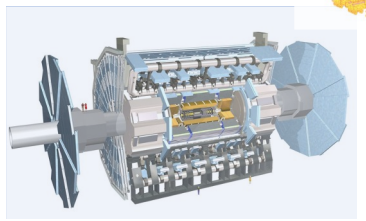
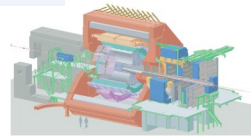
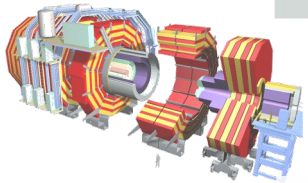
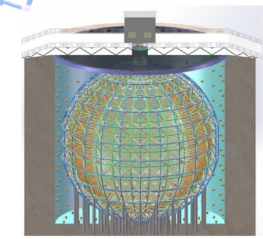

Requirements

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 2023

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

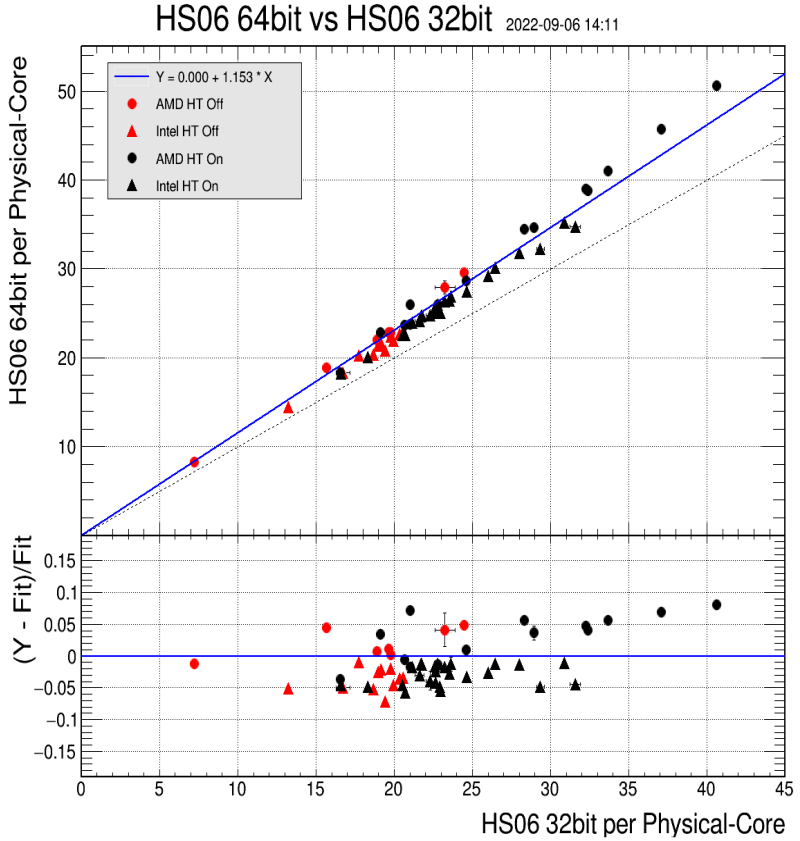

Plots

Normalized to the number of physical cores

Red	HT Off
Black	HT On
Circle	AMD
Triangle	Intel

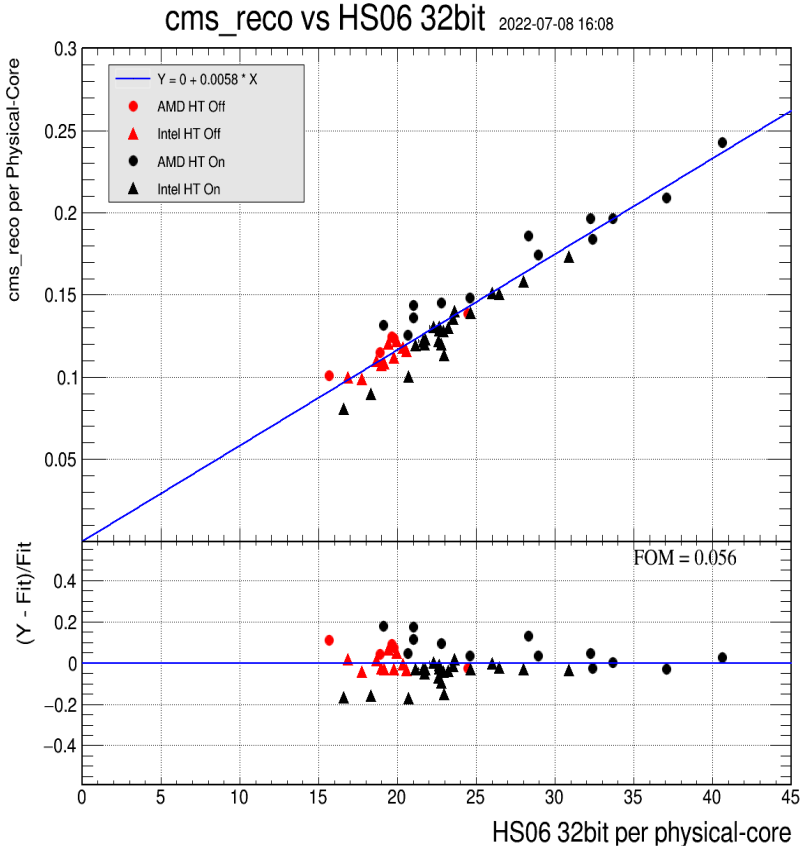

Selection of Workloads for HEPscore

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)

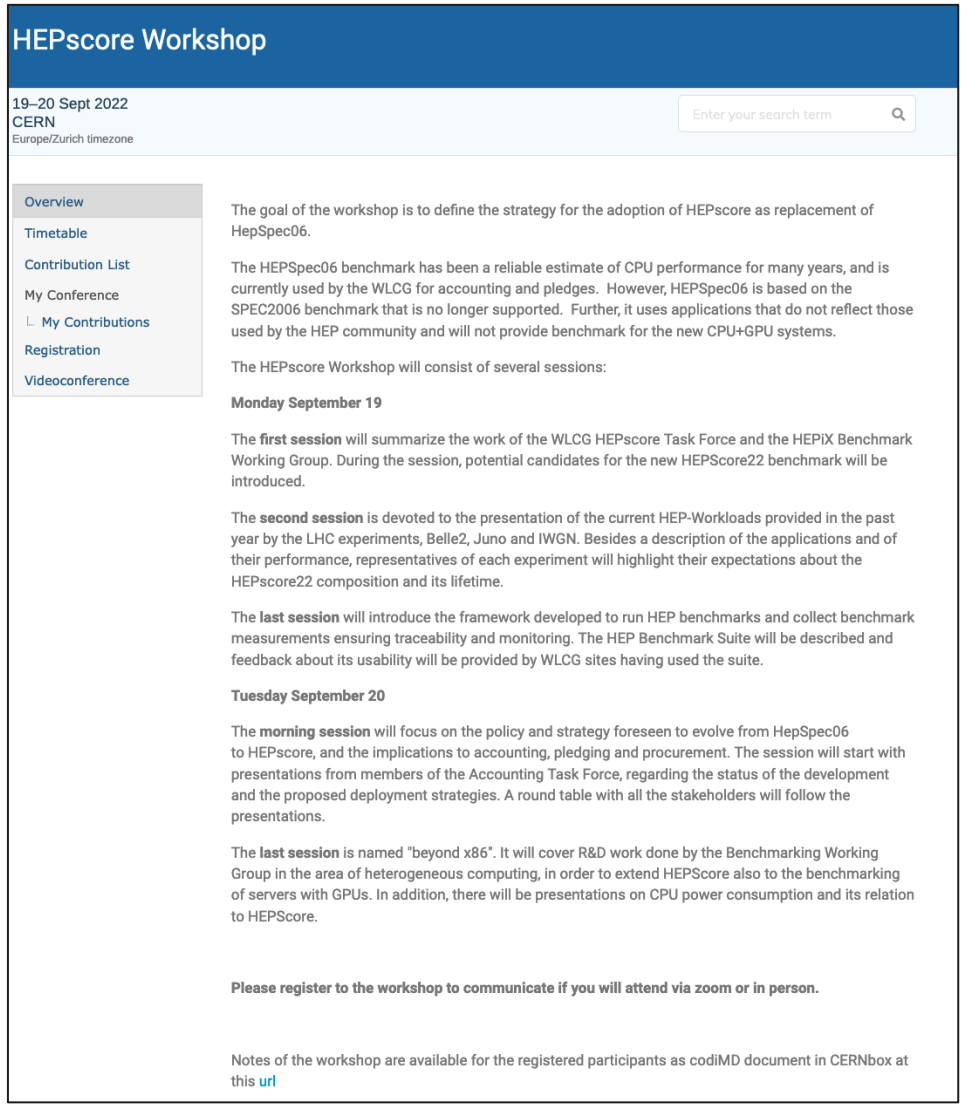

Lancaster WLCG Workshop (Nov 2022)

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

Emergence of the importance of power consumption and environmental impact
Experiments developed workloads that can run on x86 and ARM processors

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

HEPscore23 vs HEPSpec06

HEPscore23 is the production workload (x86/ARM)
Validated with measurements from a variety of servers

SMT	CPU Model	Year	Site	HS06	HS06 std	HS23	HS23 std	HS23/HS06
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 7551P 32-Core Processor	2017	Nikhef	611.73	1.03	637.44	8.57	1.04
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	360.95	0.88	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 E5-2650 v4 @ 2.20GHz	2016	Nikhef	522.73	0.81	473.50	1.35	0.91
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	332.11	1.57	222.50	1.18	0.67
On	Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz	2012	GridKa	344.99	9.17	238.93	1.15	0.69
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 5118 CPU @ 2.30GHz	2017	GridKa-Tier3	551.68	1.17	549.20	0.81	1.00
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	676.38	2.40	0.94
On	Intel(R) Xeon(R) Silver 4216 CPU @ 2.10GHz	2019	CERN	769.76	1.32	714.71	3.28	0.93

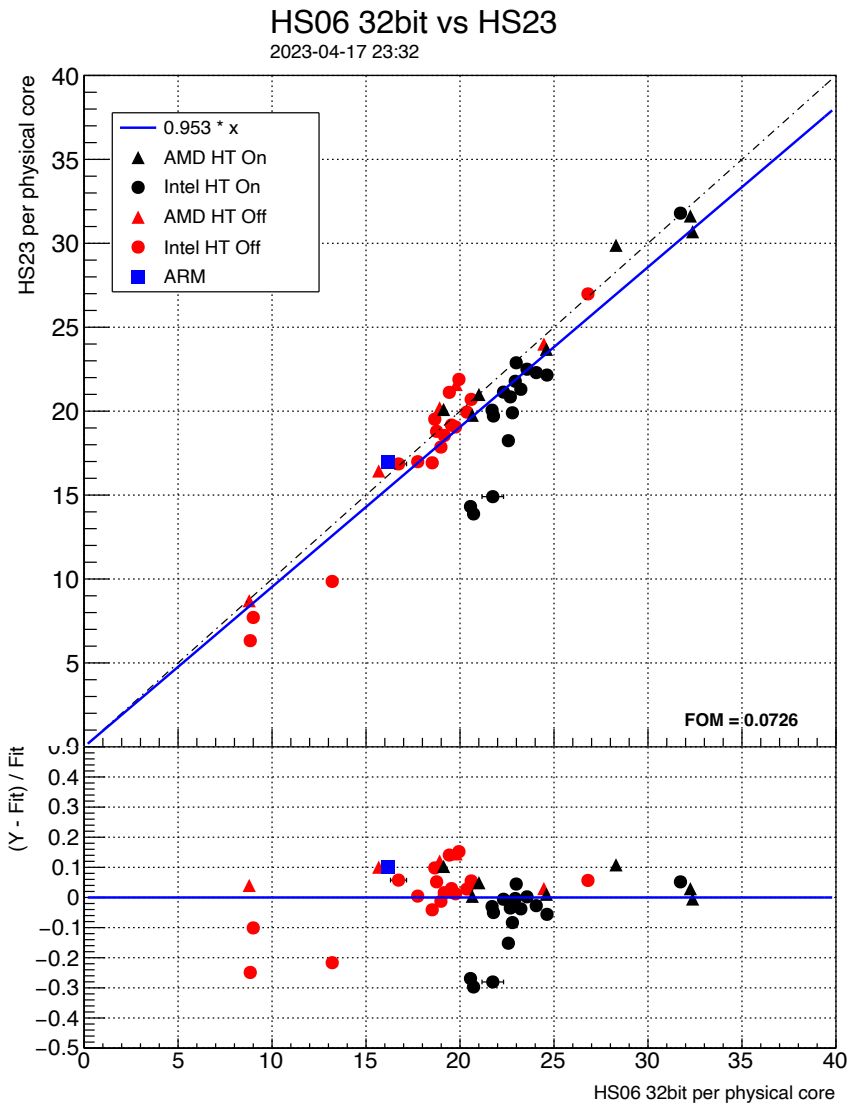

Solid-blue line is a fit to the results (constrained to the origin)
Dashed-line has unity slope – normalized to reference machine

Ratio of HEPscore23/HEPSpec06 by Year of Server

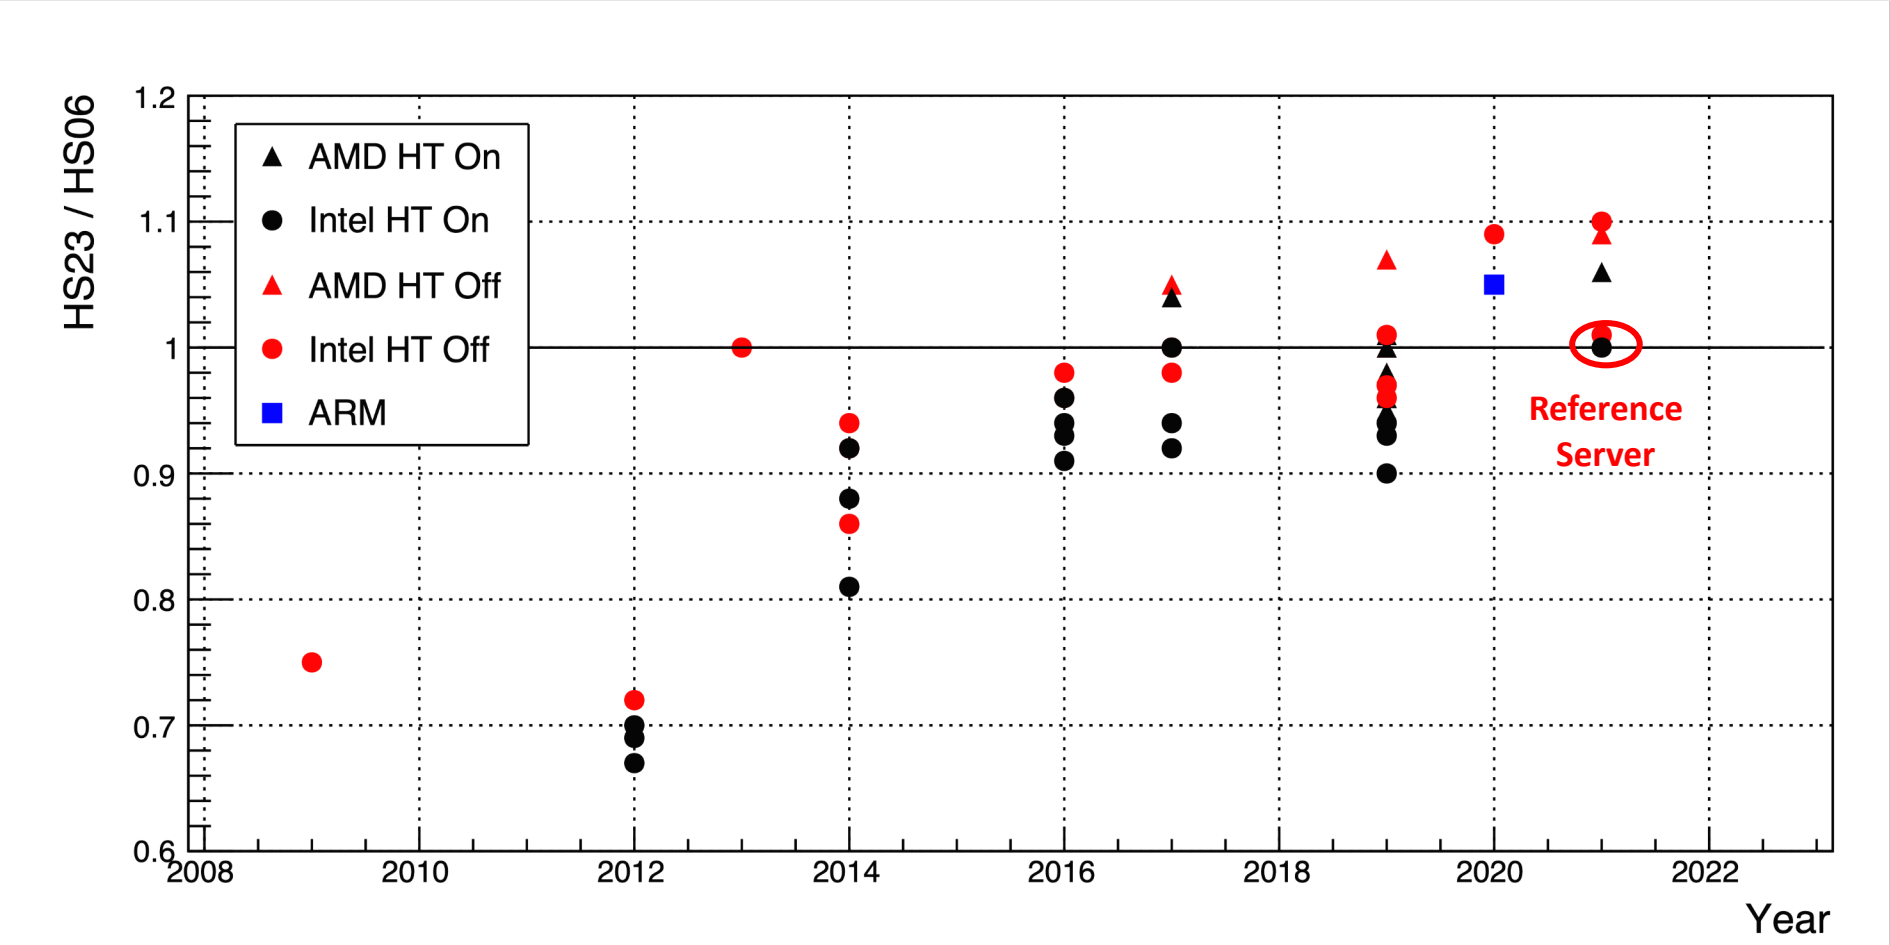

HEPscore23 is more performant relative to HEPSPEC06 on newer architectures

"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)

Preliminary requests for resources 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

Summary

New CPU benchmark (HEPscore23) is ready for use
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

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 and M. Michelotto CERN/IT

Experts in software, accounting

Site administrators

Note that instructions will soon appear in the official page of the HEPiX benchmark
https://w3.hepix.org/benchmarking.html

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)