




Science and
Technology
Facilities Council

HEPScore

...and why you should care!

Alastair Dewhurst



Introduction

- Since 2009 the official accounting for WLCG CPU resources has been HS06.
 - This is 32 bit and no longer an accurate representation of HEP workflows.
 - Since 2018 HS06 is no longer supported by the Spec organization.
- A [WLCG task force](#) was created to solve this.
 - The first step in performance evaluation is to select the right measures of performance”
 - “The types of applications of computers are so numerous that it is not possible to have a standard measure of performance [...] for all cases.

HEPscore definition







Ingredients:

- a set of reference workloads (**WLs**)
- a measure of performance per WL (m_i): work done in unit of time
- a reference server

The score **S** of a server (srv) is defined as the geometric mean of the speed factors $x_i(\text{srv}, \text{ref}) = m_i(\text{srv})/m_i(\text{ref})$ respect to the reference server (ref)

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

https://en.wikipedia.org/wiki/Weighted_geometric_mean

	WL ₁ 	WL ₂ 	WL _n 	Score $\left(\prod_{i=1}^n x_i \right)^{1/n}$	S(A,B)		
Ref. Srv 	$m_1(\text{ref})$	1 (by def)	$m_2(\text{ref})$	1 (by def)	$m_n(\text{ref})$	1 (by def)	
Srv A 	$m_1(\text{A})$	$x_1(\text{A}, \text{ref})$	$m_2(\text{A})$	$x_2(\text{A}, \text{ref})$	$m_n(\text{A})$	$x_n(\text{A}, \text{ref})$	$S(\text{A}, \text{ref})$
Srv B 	$m_1(\text{B})$	$x_1(\text{B}, \text{ref})$	$m_2(\text{B})$	$x_2(\text{B}, \text{ref})$	$m_n(\text{B})$	$x_n(\text{B}, \text{ref})$	$\frac{S(\text{A}, \text{ref})}{S(\text{B}, \text{ref})}$

Workloads

Run3 workloads for LHC experiments

alice_gen_sim_reco

atlas_gen_sherpa

atlas_sim_mt

atlas_reco_mt

belle2_gen_sim_reco

cms_reco

cms_digi

cms_gen_sim

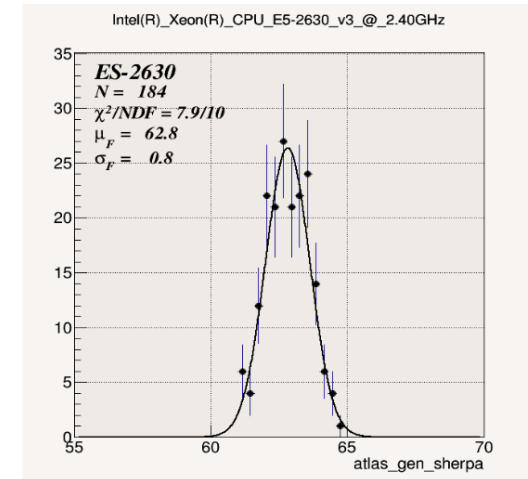
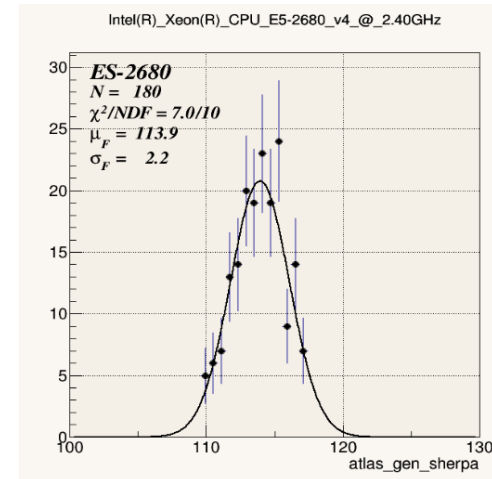
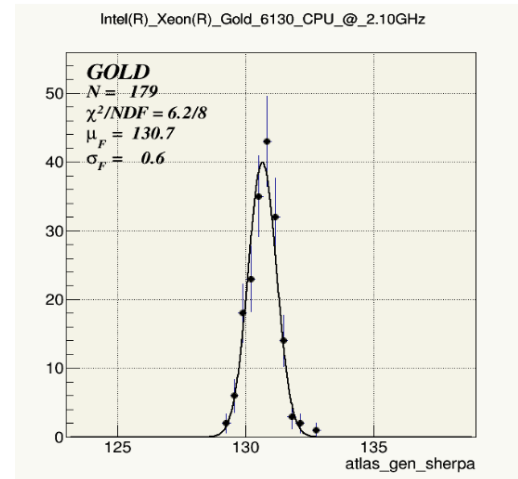
juno_gen_sim_reco

igwn_pe (Gravity Wave)

lhcb_gen_sim

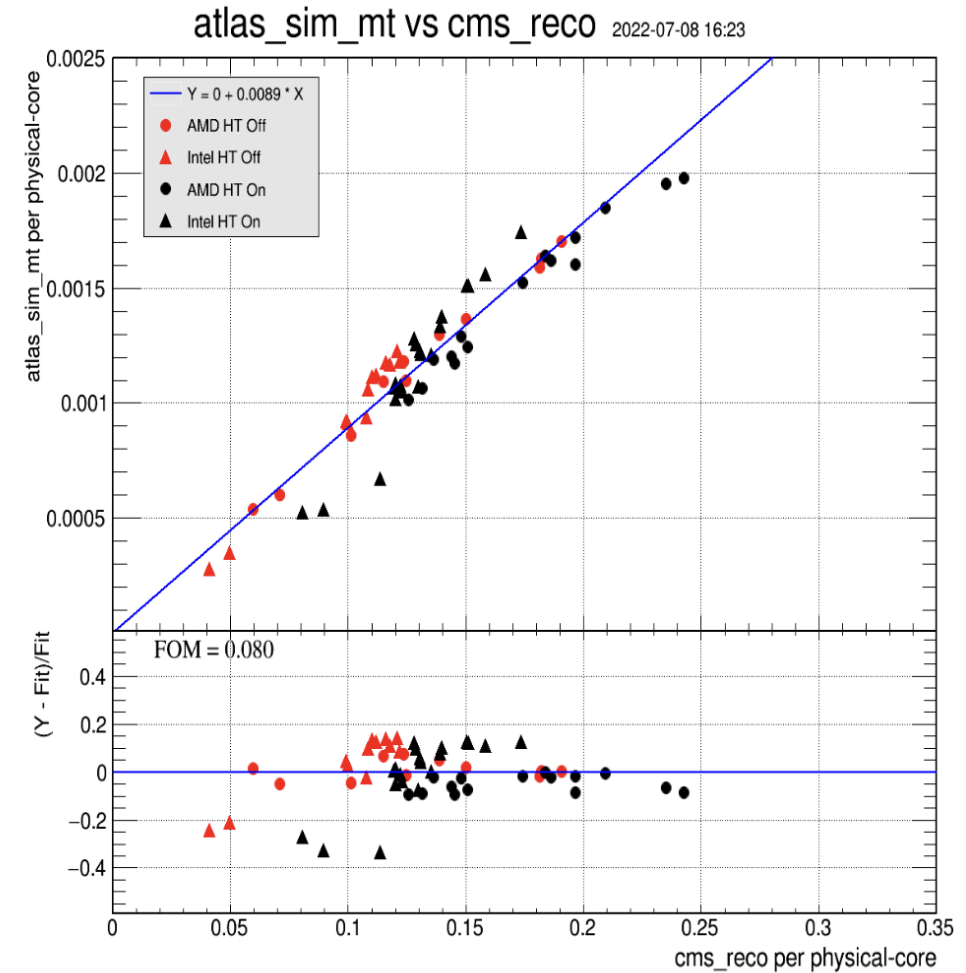
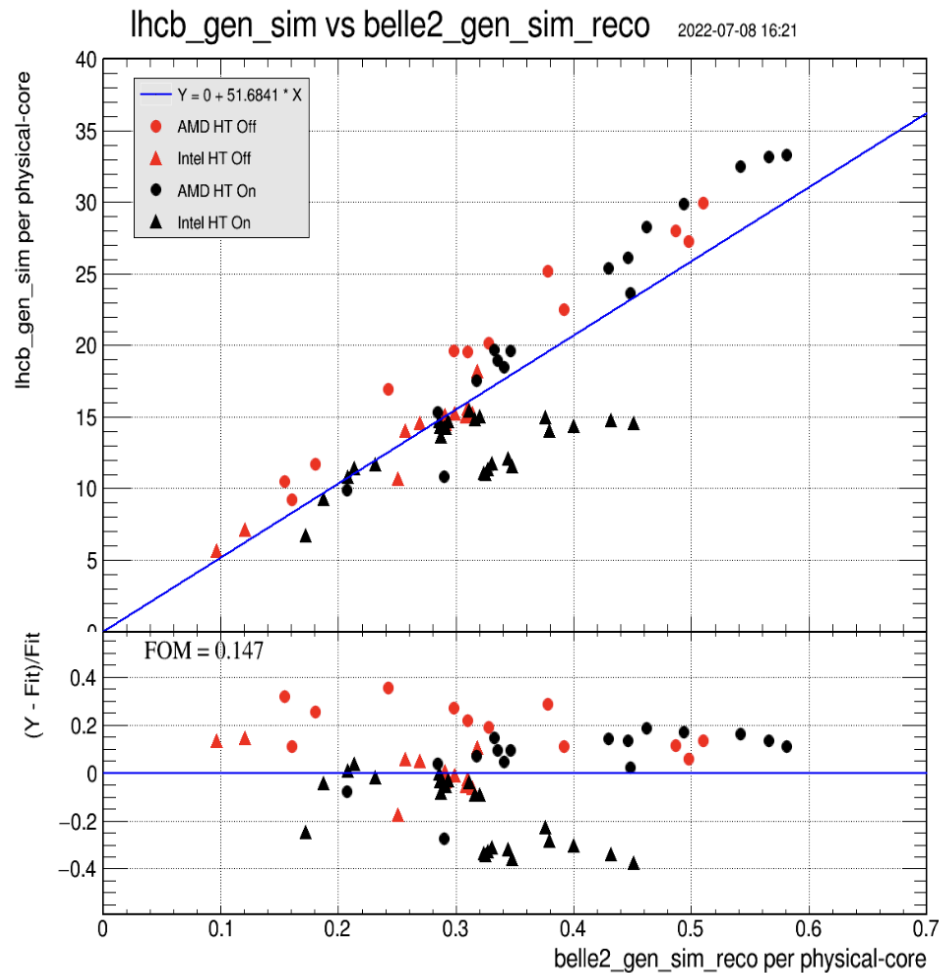
Each workload has been run and validated on a set of CERN servers

Reliable/reproducible to < 1%



Some of the workloads were found to have technical issues
(some are still being resolved)

Comparison of results



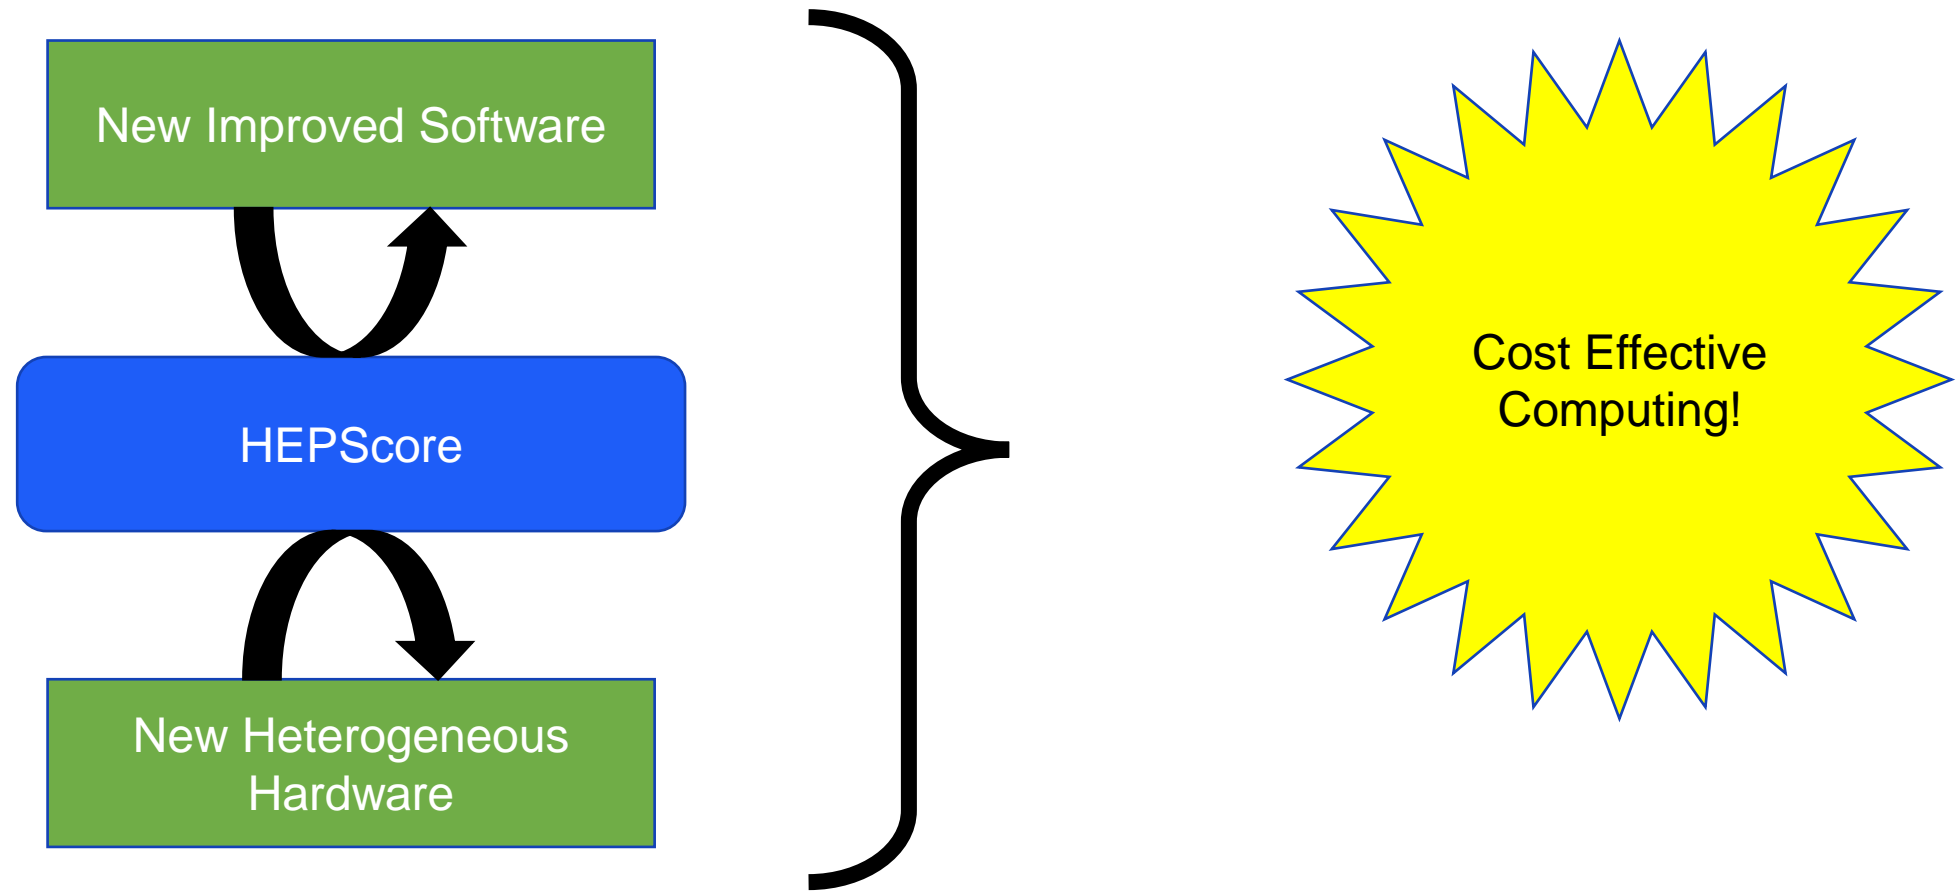
Migration

- From April 1st 2023, the official benchmark will be HEPScore 23 which will be referred to as HS23.
- To ease the transition HS23 has been normalized to equal HS06 for a reference CPU:
 - “Intel CPU Gold 6326 CPU @ 2.90GHz - 64 cores SMT ON”
 - New machines (e.g. most of our AMD), will score slightly better if re-benchmarked as HS23.
- All existing HS06 benchmarks will be just be assumed to be the HS23 score.
 - We do not need to make any urgent changes.
 - All new procurements will use HS23.
 - Over time we should re-benchmark a few WN in each older generation and update the accounting if significantly different.

APEL changes

- At the January GDB Adrian Coveney gave a talk: “[Accounting for HEPscore](#)”.
- APEL client
 - Support for retrieving benchmark type from BDII
 - Extend local benchmark setting to support HEPscore
- Messaging
 - New message format for normalised records to support specifying the type of benchmark used
 - Add support for HEPscore to other message formats (job and summary)
- APEL Repository (server)
 - Update schemas to record benchmark type for normalised records
 - Out of scope before April: including benchmark type info in data pushed to Portal

Summary





Science and
Technology
Facilities Council

Questions?