HEPscore benchmark

D. Giordano (CERN)
on behalf of
HEPiX Benchmarking WG & WLCG HEPscore Deployment TF

HEPiX Autumn 2022 Workshop
01 Nov 2022

Two teams collaborate for this project

HEPiX Benchmarking WG

Roles

- Evaluation of benchmark alternatives
- Design and development of the **HEP Benchmarks project**
- Validation of the HEP workloads
- Analysis of benchmark measurements

Team of ~13 people

Active (again) since 2016

WLCG HEPscore deployment TF

Roles

- Recommend the HEPscore composition
- Strategy for HS06->HEPscore migration
- Coordinate the collection of new workloads
- Onboard WLCG sites for validation

Team of ~20 people

Started on Nov 4. 2020

HEP Benchmarks project

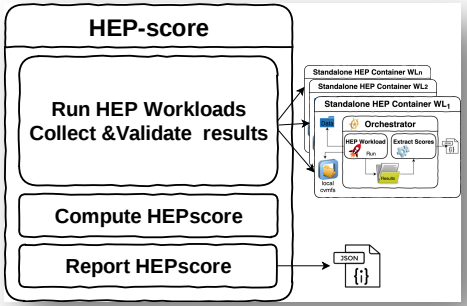

HEPscore has been proposed by the HEPiX Benchmarking WG as replacement of HS06

- Uses the workloads of the HEP experiments
- Combine them in a single benchmark score

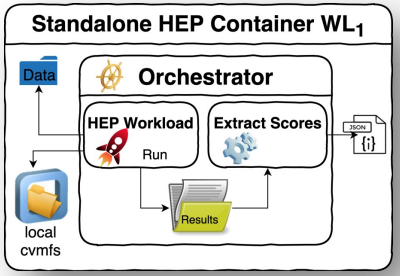

HEPscore relies on HEP Workloads

- Individual **reference** HEP applications

In addition, HEP Benchmark Suite

- Orchestrator of multiple benchmark (HEPscore, HS06, SPEC CPU2017)
- Central collection of benchmark results on OpenSearch

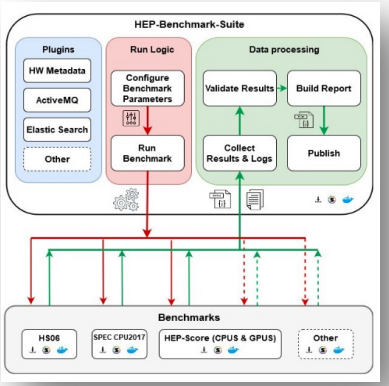

All released under GPLv3 license

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

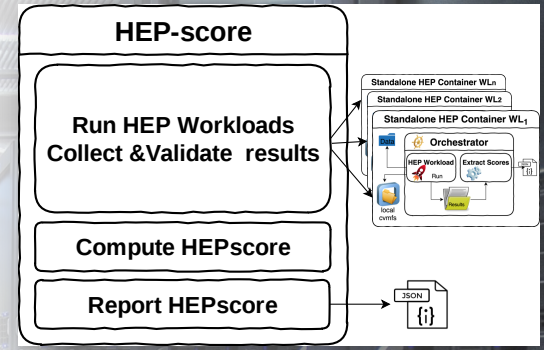

The score S of a server (**srv**) is defined as the **geometric mean** of the **speed factors** x_i(**srv**,ref) = m_i(**srv**)/m_i(**ref**) respect to the reference server (**ref**)

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

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

	WL_1		**WL_2**		**WL_n**		**Score** $\left(\prod_{i=1}^{n} x_i\right)^{\frac{1}{n}}$	**S(A,B)**
Ref. Srv	m_1(ref)	1 (by def)	m_2(ref)	1 (by def)	m_n(ref)	1 (by def)	1 (by def)	
Srv A	m_1(A)	x_1(A,ref)	m_2(A)	x_2(A,ref)	m_n(A)	x_n(A,ref)	S(A,ref)	$\frac{S(A,ref)}{S(B,ref)}$
Srv B	m_1(B)	x_1(B,ref)	m_2(B)	x_2(B,ref)	m_n(B)	x_n(B,ref)	S(B,ref)	

Brief history

WG activities - short overview (3)

Summer 2018

- Proposal of building a set of **HEP reference workloads** (*WLCG MB*)
  - Enable feature studies of the experiments' workloads
  - Build a HEP benchmark suite

Fall 2018

- Collect instructions from LHC experiments to run reference WLs
- **Prototype** the build of HEP reference benchmarks in containers
- Studies on **hardware performance counters** (using Trident)

☞ *HEP WLs have same characteristics*
and differ more respect to HS06 and SPEC CPU 2017 workloads

2019

- Start the **HEP Benchmarks project**

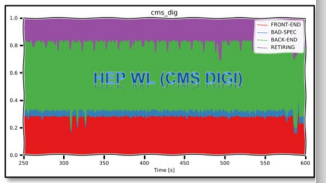

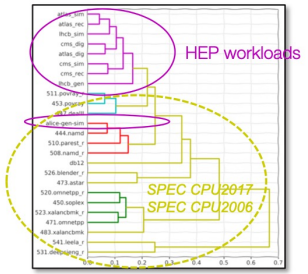

DENDROGRAM OF WL'S SIMILARITY

D. Giordano (CERN) HEPiX 15/10/2019 9

From a challenge to a routine

Collect, maintain, extend **workloads** from several HEP experiments

- Not affordable with ad-hoc recipes for each workload
  - More than **30 workloads from 7 experiments** prepared in the last years
- Experts from the Experiments focus on providing the workloads: software, data, result parser
- Experts on benchmarking focus on implementing a **unified** approach

Requirements

- Provide consistent CLI, report structure, error logging
- Reproducible results
- Zero burden from accessing remote data, databases, etc
- Not too large package distribution
- Portable
- Long term support

Experiment	WL repo	Description	Latest Container	Lastest Built Version	Latest Pipeline status
alice	digi-reco-core-run3	link	docker	v1.0	passed
alice	gen-sim	link	docker	v2.1	passed
alice	gen-sim-reco-run3	link	docker	ci-v0.6-aod	passed
atlas	digi-reco	link	docker	v2.1	passed
atlas	gen	link	docker	v2.1	passed
atlas	gen_sherpa	link	docker	v0.2	passed
atlas	gen_sherpa-ma	link	docker	v2.0	passed
atlas	kv	link	docker	v2.1	passed
atlas	reco_mt	link	docker	v2.1	passed
atlas	sim	link	docker	v2.1	passed
atlas	sim_mt	link	docker	v2.1	passed
atlas	sim_mt-ma	link	docker	v2.1	passed
belle2	gen-sim-reco	link	docker	v2.1	passed
belle2	gen-sim-reco-2	link	docker	ci-v0.4	passed
cms	digi-reco-run3	link	docker	v2.1	passed
cms	digi-reco-run3-ma	link	docker	ci-v0.7	passed
cms	gen-sim	link	docker	ci-v0.6	passed
cms	gen-sim-run3	link	docker	v2.1	passed
cms	gen-sim-run3-ma	link	docker	v2.1	passed
cms	hlt	link	docker	v1.0	passed
cms	hlt-run3	link	docker	v0.1	passed
cms	hlt-run3-ma	link	docker	v0.1	passed
cms	reco	link	docker	v2.1	passed
cms	reco-run3	link	docker	ci-v0.6	passed
cms	reco-run3-aarch64	link	docker	v0.1	passed
cms	reco-run3-ma	link	docker	v1.1	passed
dune	reco-fd	link	docker	v2.1	passed
hella	world-ma	link	docker	v2.1	passed
igwn	gen-sim-reco	link	docker	v2.1	passed
juno	gen-sim-reco	link	docker	v2.1	passed
lhcb	gen-sim	link	docker	v2.1	passed
lhcb	gen-sim-2021	link	docker	v2.1	passed
mg5amc	madgraph4gpu	link	docker	v2.1	passed
mg5amc	madgraph4gpu-2022	link	docker	v2.1	passed

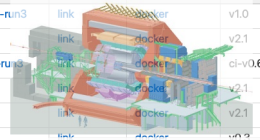
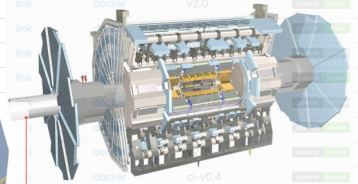
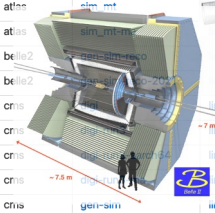
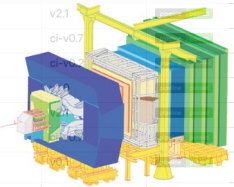
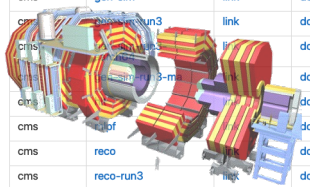
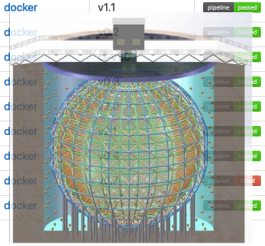

Workloads' repository

Realized an effective infrastructure to **build and distribute** the HEP workloads

Containers are built for multiple architectures: x86, aarch64, (potentially) Power, GPUs

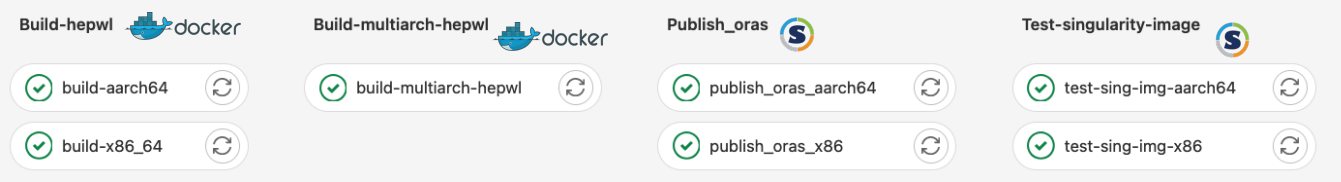

- GitLab **CI/CD** for fully **automated** continuous integration
- GitLab **Registry** for container distribution (Docker & Singularity/Apptainer)

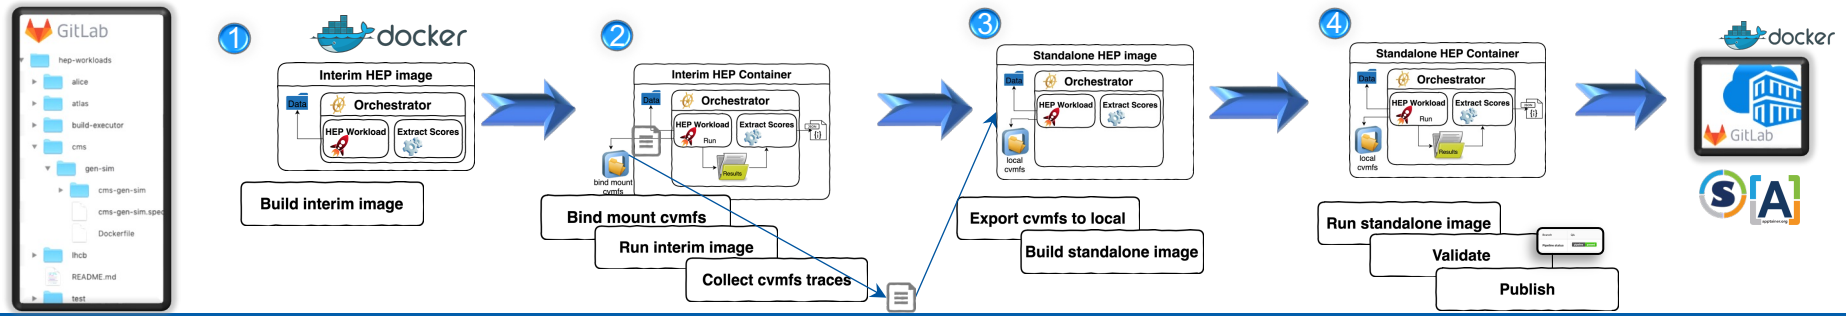

Brief history

WG activities - short overview (3)

Summer 2018

- Proposal of building
  - Enable feature st
  - Build a HEP ben

Fall 2018

- Collect instructions fr
- **Prototype** the build
- Studies on **hardware**
  - ☞ *HEP WLs have s*
  and differ more r

2019

- Start the **HEP Benc**

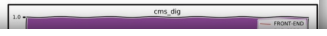

D. Giordano (CERN)

Extensive validation process of WL

- ❑ Validating reproducibility, robustness, run duration, disk space
- ❑ Continuously running in a number of virtual & physical machines
- ❑ Evaluated a different number of events per WL to shorten the runtime

Workload	ATLAS gen	ATLAS sim	ATLAS digi-reco	CMS gen-sim	CMS digi	CMS reco	LHCb gen-sim
Robustness	✅	✅	✅	✅	✅	✅	✅
Reproducibility	0.8%	2%	0.6%	1.5%	1%	1%	1%
Memory	✅	✅	✅	✅	✅	✅	✅
Image size (unpacked)	1.65 GB	6.0 GB	6GB	5.4 GB	11 GB	8.4 GB	2.6 GB
Readiness	✅	✅	✅	✅	✅	✅	✅

✅ okay
❌ blocker

CPU Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz (32 cores, SMT ON)

WL	# threads or proces. (default)	# Evts/thread (default)	Duration of a single WL run on ref machine [hh:mm]	Wdir size (per running copy)
Atlas gen	1 (SP)	200	~12	50MB
Atlas sim	4 (MP)	10	~1:32	100 MB
CMS gen-sim	4 (MT)	20	~0:15	70 MB
CMS digi	4 (MT)	50	~0:09	400 MB
CMS reco	4 (MT)	50	~0:15	100 MB
LHCb gen-sim	1 (SP)	5	~0:40	15 MB
Total			**~3:30**	

D. Giordano (CERN) HEPiX Autumn 2020 Online workshop 13/10/2020 6

Brief history

WG activities - short overview (3)

Summer 2018

- Proposal of building
  - Enable feature st
  - Build a HEP ben

Fall 2018

- Collect instructions fr
- **Prototype** the build
- Studies on **hardware**
  - ☞ *HEP WLs have s*
    and differ more r

2019

- Start the **HEP Bench**

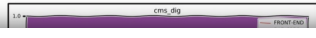

D. Giordano (CERN)

Extensive validation process of WL

- ❑ Validating r
- ❑ Continuous
- ❑ Evaluated a runtime

D. Gior

Demo: results in the central DB @ CERN

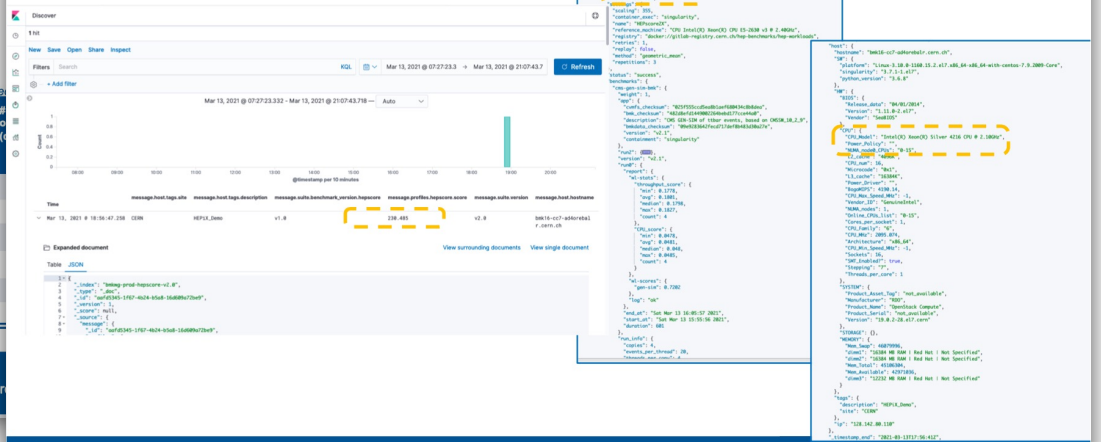

D. Giordano (CERN) HEPiX Spring 2021 Online workshop 16/03/2021 17

Brief history

WG activities - short overview (3)

Summer 2018

- Proposal of building
  - Enable feature st
  - Build a HEP ben

Fall 2018

- Collect instructions fr
- **Prototype** the build
- Studies on **hardware**
  - ☞ *HEP WLs have s*
    and differ more r

2019

- Start the **HEP Benc**

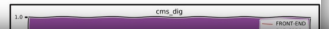

D. Giordano (CERN)

Extensive validation process of WL

- ❑ Validating r
- ❑ Continuous
- ❑ Evaluated a runtime

CPU Intel(R) Xe

WL	
Atlas gen	
Atlas sim	
CMS gen-sim	
CMS digi	
CMS reco	
LHCb gen-sim	
Total	

D. Gior

Demo: results in the central DB @ CERN

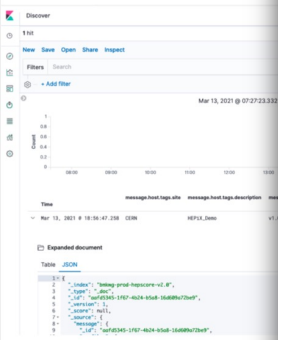

D. Giordano (CERN)

$HEPscore_{\beta}$ vs $HS06_{64bits}$ measurements (II)

- ❑ In a few occasions **hundreds** of servers with same HW have been benchmarked and compared
- ❑ E.g.: measurements done on 180 identical servers from a CERN delivery
  - One measurement per server
  - $HS06_{64bits}$ also scaled to the reference server
  - $HEPscore_{\beta}$ resolution is narrower than $HS06_{64bits}$

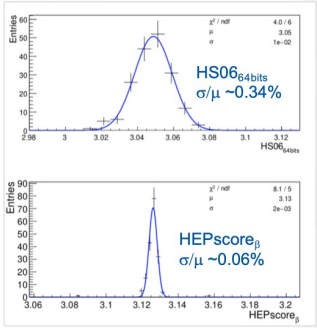

D. Giordano (CERN) HEPiX Autumn 2021 Online workshop 26/10/2021 12

Brief history

WG activities - short overview (3)

Summer 2018

- Proposal of building
  - Enable feature st
  - Build a HEP ben

Fall 2018

- Collect instructions fr
- **Prototype** the build
- Studies on **hardware**
  - ☞ *HEP WLs have s*
  and differ more r

2019

- Start the **HEP Benc**

D. Giordano (CERN)

Extensive validation process of WL

- Validating r
- Continuous
- Evaluated a runtime

CPU Intel(R) Xe

WL
Atlas gen
Atlas sim
CMS gen-sim
CMS digi
CMS reco
LHCb gen-sim
Total

D. Gior

Demo: results in the central DB @ CERN

D. Giordano (CERN)

HEPscore$_{\beta}$ vs HS06$_{64bits}$ measurements (II)

- In a few occasion HW have been b
- E.g.: measureme servers from a CE
  - One measureme
  - HS06$_{64bits}$ also s the reference ser
  - HEPscore$_{\beta}$ resol than HS06$_{64bits}$

D. Giordano (CERN)

Action 1: Extend list of HEP workloads

- The WG to build the standalone containers
  - In close contact with the experiments' experts
- 10 workloads for x86 to enter in the matrix
  - 8 ready; 4 are Single Process or Single Thread
- In addition, 2 prototype workloads for GPU (Madgraph generator and CMS HLT-like)
  - Demonstrate the HEPscore usability on other arch. (longer term goal!)

Legend: Ok / In progress / Not started yet

	WL	Responsible	OS	Platform	WL developed in a git fork (if relevant)	Merged in HEP-Workloads repo	Built	Validated	Reference score	Ready for the "metrix"
2	Alice Gen-Sim-Reco	S. Piano	cc7	x86						
3	Atlas gen sherpa	W. Lamp	cc7	x86						
4	Atlas simMT	W. Lamp	cc7	x86						
5	LHCb gen-sim 2021	A. Valassi	cc7	x86						
6	CMS gen-sim Run3	A. Sciabà	cc7	x86/arm			x86/arm			
7	CMS Digi Run3	A. Sciabà	cc7	x86/arm			x86/arm			
8	CMS Reco Run3	A. Sciabà	cc7	x86/arm			x86/arm			
9	CMS HLT-like	A. Sciabà	cc7	x86 & GPU			x86			
10	Belle2	R. Sobie	cc7	x86						
11	Dune	A. Mc Nab	cc7	x86	https://gitlab.cern.			On hold for lack of time from Dune experts		
12	Juno	X. Yan	cc7	x86	https://gitlab.cern.					
13	IGWN	J. Willis	cc7	x86	https://git.ligo.org					
14	Madgraph	A. Valassi	cc7	x86 / GPU						

HEPiX | M.Michelotto (INFN) | HEPiX Spring 2022 | 26/04/2022 | 15 | INFN

Progress in the last 6 months

- Finalized the analysis of the 2022 large measurement campaign
- Agreed on the composition of HEPscore
- Converging on a procedure to migrate from HS06 to HEPscore
- Progress on x86/aarch64 HEP workloads
- **HEPscore workshop** https://indico.cern.ch/event/1170924/

HEPscore workshop

2 days of workshop (19–20 Sept 2022)

- Good representation of the different bodies involved:

  Experiments, Sites, WLCG Collaboration

- Registered participants: 51

Valuable feedback from our WLCG community on

- Proposed composition of HEPscore
- Usability of the HEP Benchmark Suite
- Strategy for the adoption of HEPscore as WLCG benchmark

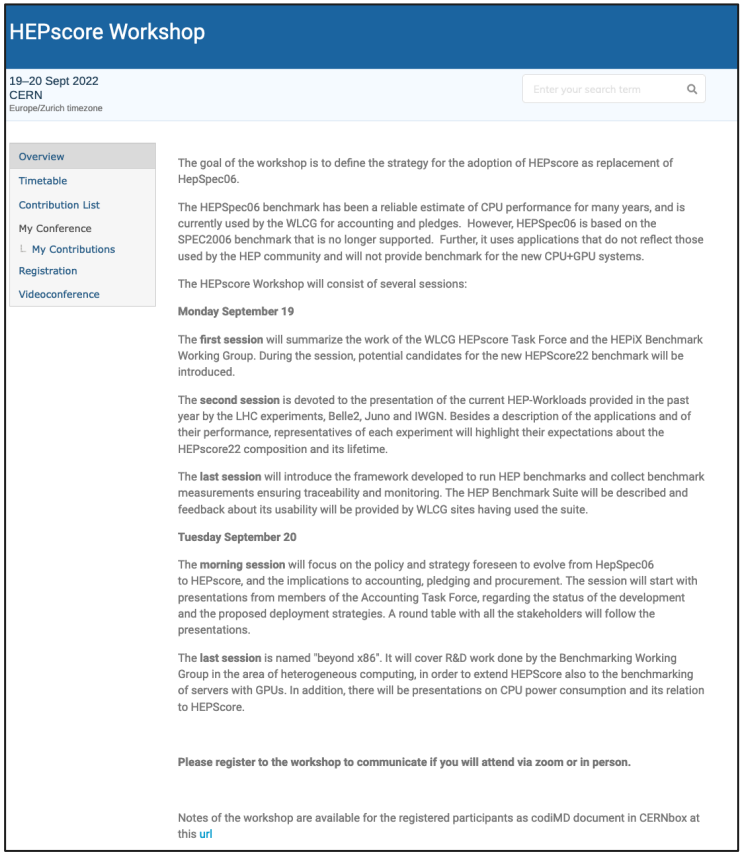

https://indico.cern.ch/event/1170924/

Agenda

- 5 sessions to address interconnected topics:
  - HEPscore TF results
  - HEP Workloads
  - HEP Suite
  - HEPscore Deployment
  - Beyond x86
- 30 contributions
- Allocated time for discussions

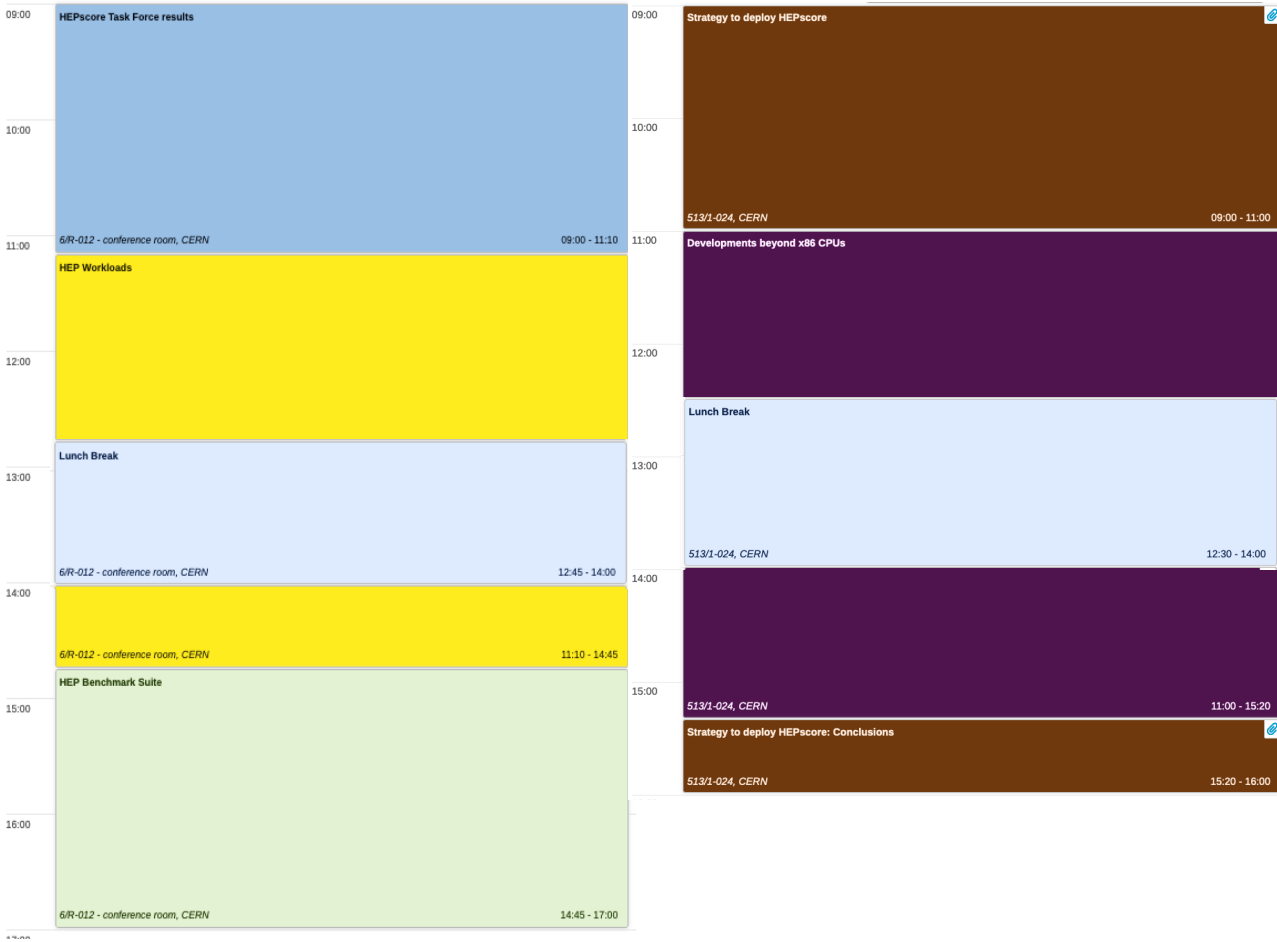

HEP Workloads 2021/22

SW from LHC experiments (Run 3) + Belle2, Juno, IGWN

- Alice (gen-sim), Atlas (gen, sim_mt, reco_mt), CMS (gen-sim, digi, reco), LHCb (sim)
- Belle2 (gen-sim-reco), Juno (gen-sim-reco), IGWN

Repeated measurements

- On ~40 different CPU models from ~15 WLCG sites
- Thanks to all site admins that made this possible!

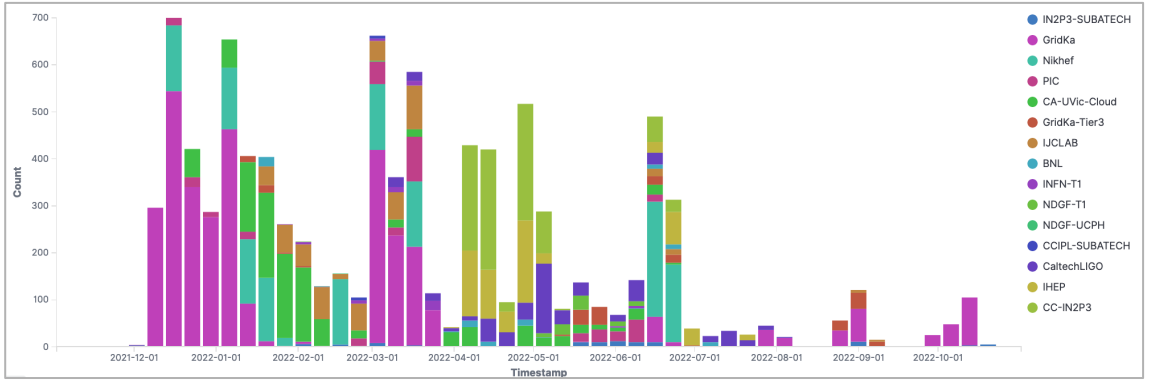

2022 measurement campaign

Measured

- Robustness against failures
- Resolution (σ/μ typically < 1%)
- Performance

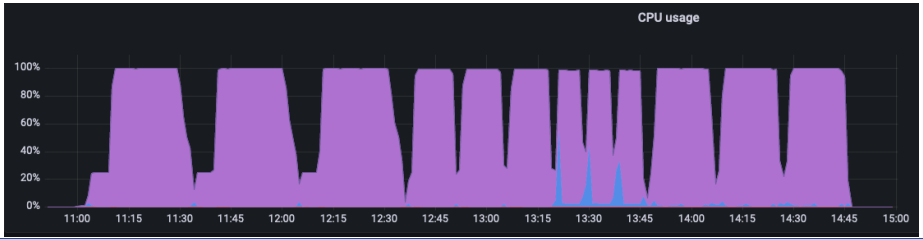

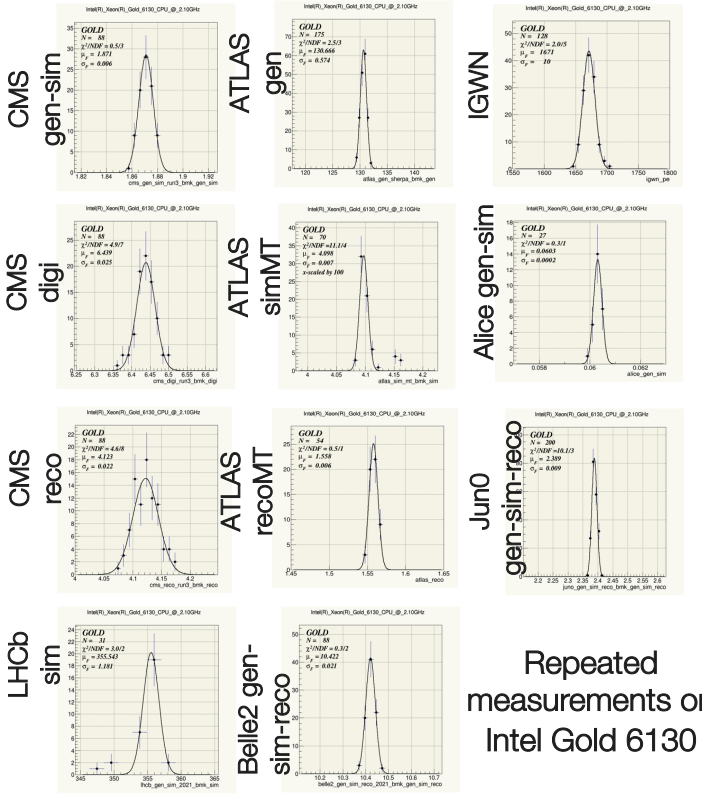

Analysis results at a glance

- (Re)confirmed HEPSpec06 (HS) and SPEC2017 (SP) findings, i.e. high correlation

- HEP Workloads vs HS/SP
  - Measure scaling and mean deviation from fit

- HEP Workload Vs Workload matrix
  - Indication of the correlation between the workloads

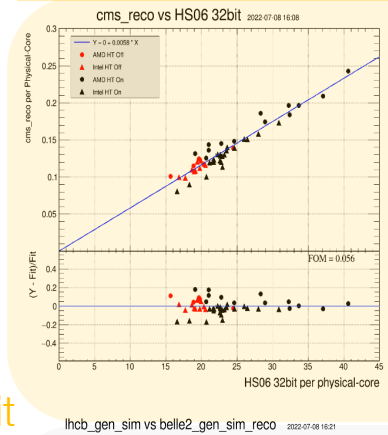

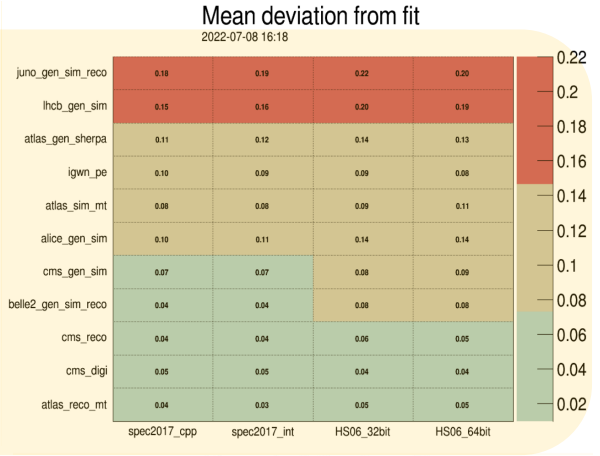

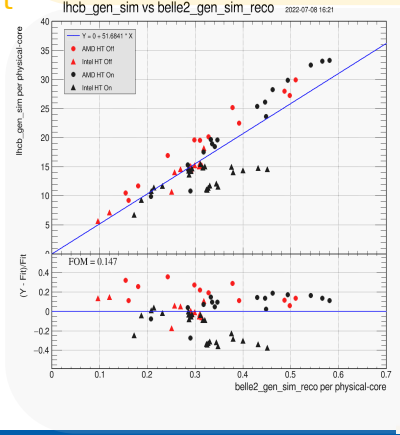

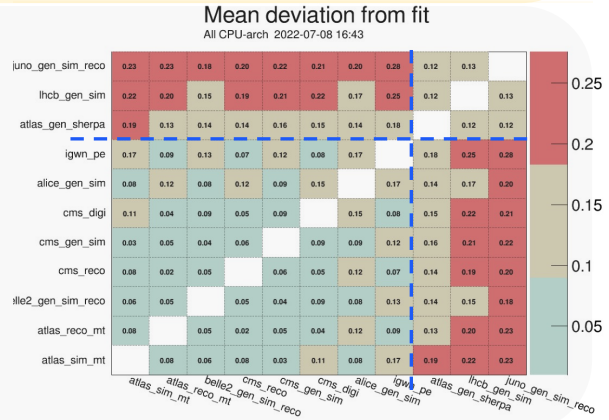

HEPscore candidates

Several combinations investigated:

- All workloads/Exclude the long-running ones/Use a subset/No weights/Weights from Grid fraction of jobs
- Little difference (**few %**) between the candidates

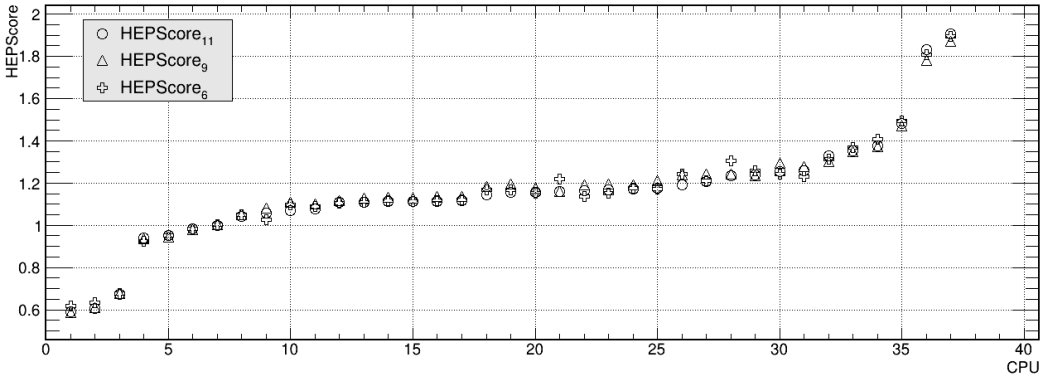

Preference for

- A small workload set for shorter runtime
- The simplest approach of unweighted WLs
- HEPscore composition of **7 workloads**:
  - **Alice (digi_reco),**
    Atlas (gen_sherpa, reco),
    CMS (gen_sim, reco),
    LHCb (sim), Belle2 (gen_sim_reco)
  - Important to include the Alice Reco workload: reconstruction of Pb-Pb events. Workload being prepared

Desiderata: **at least some workloads should run on ARM**

Benchmark DB

- OpenSearch instance @ CERN
  - https://es-hep-bmk.cern.ch/
  - Access restricted: contact me for read access
- Benchmark reports
  - Modular JSON structure
  - Suitable for non-relational DB

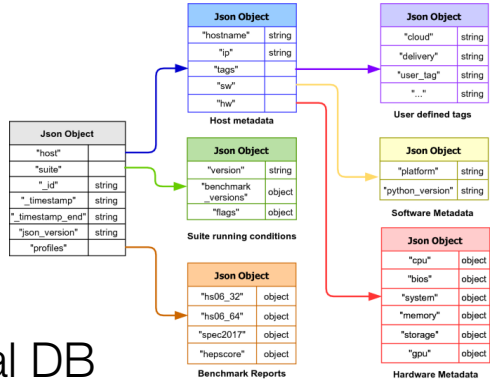

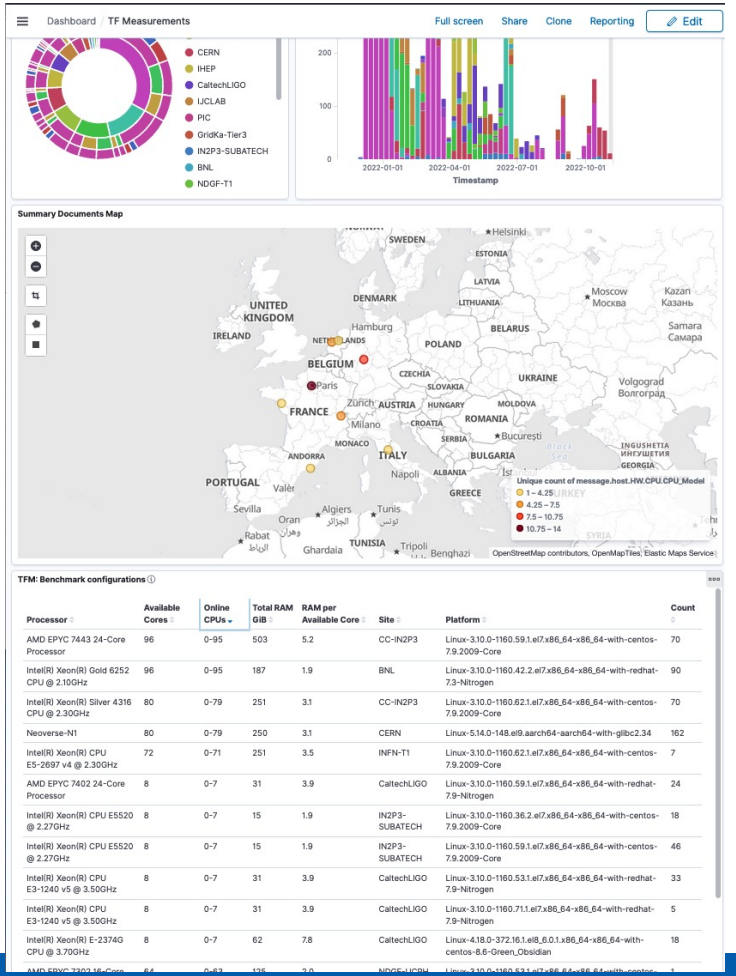

Processor	Available Cores	Online CPUs	Total RAM GiB	RAM per Available Core	Site	Platform	Count
AMD EPYC 7443 24-Core Processor	96	0-95	503	5.2	CC-IN2P3	Linux-3.10.0-1160.59.1.el7.x86_64-x86_64-with-centos-7.9.2009-Core	70
Intel(R) Xeon(R) Gold 6252 CPU @ 2.10GHz	96	0-95	187	1.9	BNL	Linux-3.10.0-1160.42.2.el7.x86_64-x86_64-with-redhat-7.3-Nitrogen	90
Intel(R) Xeon(R) Silver 4316 CPU @ 2.30GHz	80	0-79	251	3.1	CC-IN2P3	Linux-3.10.0-1160.62.1.el7.x86_64-x86_64-with-centos-7.9.2009-Core	70
Neoverse-N1	80	0-79	250	3.1	CERN	Linux-5.14.0-148.el9.aarch64-aarch64-with-glibc2.34	162
Intel(R) Xeon(R) CPU E5-2697 v4 @ 2.30GHz	72	0-71	251	3.5	INFN-T1	Linux-3.10.0-1160.62.1.el7.x86_64-x86_64-with-centos-7.9.2009-Core	7
AMD EPYC 7402 24-Core Processor	8	0-7	31	3.9	CaltechLIGO	Linux-3.10.0-1160.59.1.el7.x86_64-x86_64-with-redhat-7.9-Nitrogen	24
Intel(R) Xeon(R) CPU E5520 @ 2.27GHz	8	0-7	15	1.9	IN2P3-SUBATECH	Linux-3.10.0-1160.36.2.el7.x86_64-x86_64-with-centos-7.9.2009-Core	18
Intel(R) Xeon(R) CPU E5520 @ 2.27GHz	8	0-7	15	1.9	IN2P3-SUBATECH	Linux-3.10.0-1160.59.1.el7.x86_64-x86_64-with-centos-7.9.2009-Core	46
Intel(R) Xeon(R) CPU E3-1240 v5 @ 3.50GHz	8	0-7	31	3.9	CaltechLIGO	Linux-3.10.0-1160.53.1.el7.x86_64-x86_64-with-redhat-7.9-Nitrogen	33
Intel(R) Xeon(R) CPU E3-1240 v5 @ 3.50GHz	8	0-7	31	3.9	CaltechLIGO	Linux-3.10.0-1160.71.1.el7.x86_64-x86_64-with-redhat-7.9-Nitrogen	5
Intel(R) Xeon(R) E-2374G CPU @ 3.70GHz	8	0-7	62	7.8	CaltechLIGO	Linux-4.18.0-372.16.1.el8_6.0.1.x86_64-x86_64-with-centos-8.6-Green_Obsidian	18

HEP BMK Suite in Procurement and Enrollment

During the tendering and evaluation phase:
- Bidders usage to provide HS06 score
- Sample evaluation

Lifecycle of HW at CERN - Ironic enrollment process:
- Benchmarking as Ironic clean step
- Where data goes: ElasticSearch, inventory

Conclusion
- Future work
- Wishlist
- Feedback

Luca Atzori

Conclusions

- We are running smoothly the benchmark components during the measurement campaign.
- We are enthusiastic ! with the data population and aggregation mechanism.
- If we have in the future Jupyter notebook direct access to data will be very helpful for custom graphs and better inspection of the data.
- Sanity check of the wallclock, job efficiency and system I/O are important to detect corner cases in hardware configuration, runtime environment condition or benchmark nature (e.g. cpu bound or I/O bound).

Emmanouil Vamvakopoulos

6

Conclusion / Suggestions

- HEPScore has changed for the better what we decide to purchase.
  - Compared to HS06 it's easier to run, more flexible and more insightful.
  - The 32bit requirement of HS06 was skewing results.
- HEPScore will allow us to compare diverse types of hardware.
  - E.g. x86 vs ARM vs GPU
- HEPScore is still "just" a benchmark and doesn't account for operational reality.
  - E.g. RAL has purchased more memory per core than VOs specify for 5+ years.
- We don't want to over optimize for one fixed benchmark. Suggestion:
  - Release a new HEPScore annually (in April) that sites can use to guide procurement if they so choose.
  - Pick a Golden HEPScore for official use per X years / LHC Run.

Alastair Dewhurst

Alastair Dewhurst, 19th September 2022

Task Force Measurements

- **Much easier**
  - **A little trial and error to get started**
  - **Supplied scripts were simple to customise and run**
  - **Easily repeatable for the different scripts**
  - **Work flow:**
    - **run all repetitions for a script with publish false**
    - **then upload afterwards**
- **Ran the HS06 benchmarks – straight forward**
  - **Create the HS06 tarball**
  - **Copy to benchmark machine**
  - **Run benchmark script**

Andrew Pickford

19/09/22 A. Pickford 4

Envisage the WLCG sites running benchmarks via the Suite with the data (and its metadata) written to OpenSearch at CERN

Metrics gathering and analysis

Mattieu Puel

Publishing to the message bus
- Extensive data structure reported (host configuration, results…).
- An API key might be handier to manage compared to SSL certs (generation proc

ELK
- Published data is useful to sites : processor powers, memory population…
  → that's a net improvement compared to former Hepix website publication.
- Some new smart graphics recently (TF Measurements, processors distribution…) ! Thank you for that.
- Ability to save searches, vizualizations and dashboards would be a plus.

Concluding remarks

Carles Acosta Silva

- We run the benchmarks in PIC WNs: Intel Xeon E5v3 and v4 architectures and more modern AMD EPYC 7452 machines
- We consider a great improvement to run the benchmarks using containers and collecting the data in a DB to be visualized in the kibana dashboard

Power consumption measurements

- Several similar initiatives ongoing at multiple sites
  - Same motivation: measure “events/watt”
- Opportunities for collaboration
  - Which tools to use, pros/cons, caveats
  - Use the HEPiX Benchmarking WG as exchange forum

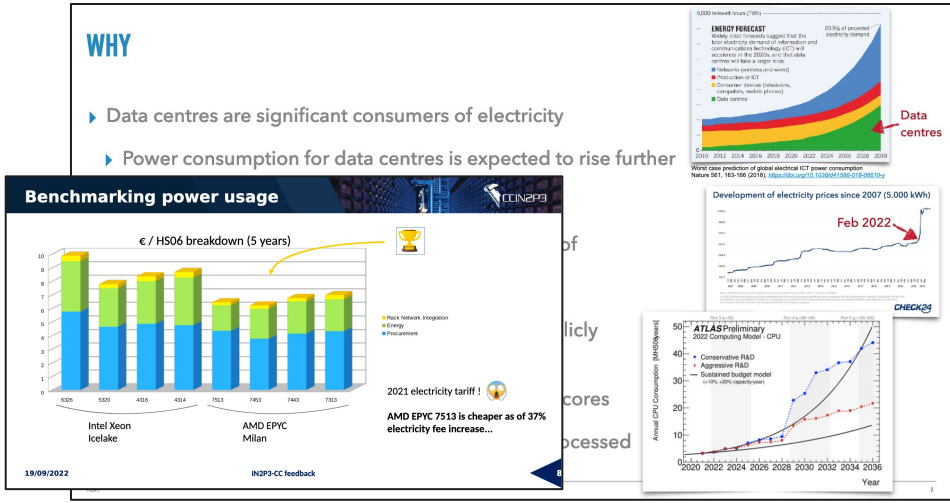

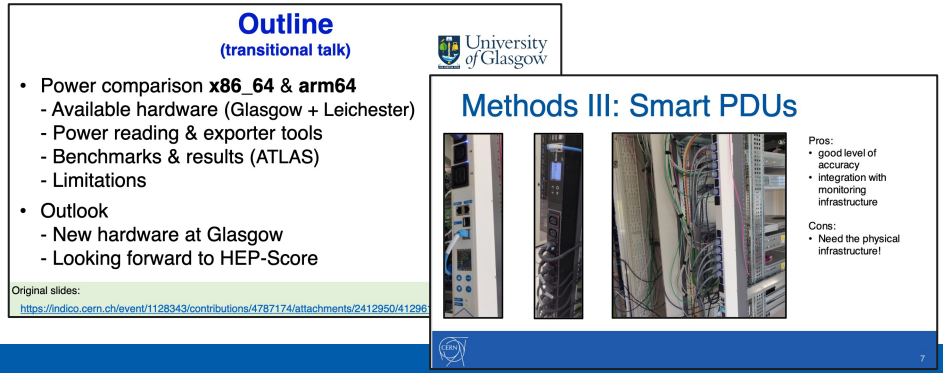

HEP Workloads containers for x86 and aarch64

- Increasing interest to build/run experiment's software on aarch64
- Having the HEP-workloads capable to run on x86/aarch64 enables studies like the one presented at ACAT by E. Simili

HEP-Score Containers

HEP-Score containers can run on **Singularity** (or **Docker**, which we do not use):

(x86) *Singularity-CE 3.9.9-1.el7* (previous version *3.8.7-1.el7* disappeared from EPEL and replaced with *AppTainer 1.1.0-1.el7*)

(arm) *singularity version 3.8.4-1.el7* (still available in EPEL)

Example execution of a containerised HEP-Score job:

```
$ mkdir –p /tmp/test/results
$ chmod a+rw /tmp/test/
$ singularity run -B /tmp/test:/results oras://registry.cern.ch/hep-workloads/cms-gen-sim-run3-bmk:latest
```

We used 5 HEP-Score containers from the container_registry (prev. slide):

gitlab-registry.cern.ch/hep-benchmarks/hep-workloads-sif/atlas-sim_mt-ma-bmk:v2.0
gitlab-registry.cern.ch/hep-benchmarks/hep-workloads-sif/atlas-gen_sherpa-ma-bmk:ci-v1.0 } ATLAS (2x)
gitlab-registry.cern.ch/hep-benchmarks/hep-workloads-sif/cms-reco-run3-ma-bmk:v1.1
gitlab-registry.cern.ch/hep-benchmarks/hep-workloads-sif/cms-digi-run3-ma-bmk:v1.0
gitlab-registry.cern.ch/hep-benchmarks/hep-workloads-sif/cms-gen-sim-run3-ma-bmk:v1.0 } CMS (3x)

Note: the HEP workloads are designed to scale with the number of available threads, therefore power consumption cannot be directly compared among machines with a different number of cores (/threads), as the machine with more threads would have done more work …

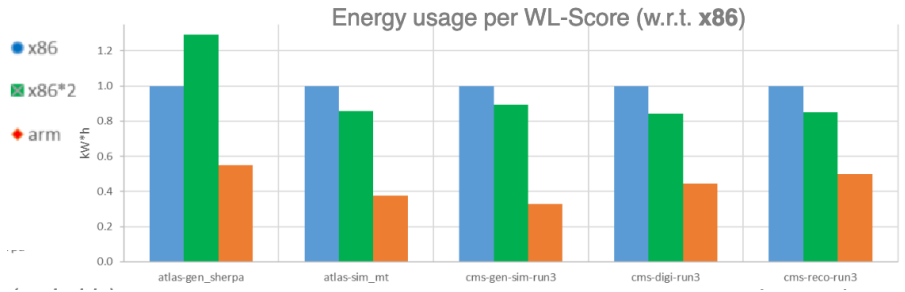

3 scenarios for HEPscore

1. **Conservative**: use the workloads agreed during the workshop
   - HEPscore just runs on x86. None of the composing benchmark run on aarch64
   - BUT: new workloads are already available (see case-2)
2. **Reasonable**: HEPscore just runs on x86, some composing workloads run also on aarch64
   - It's already the case, with the updated workload versions for CMS and Atlas
3. **Desirable**: HEPscore runs on x86 and aarch64
   - The composing workloads are all built for x86 and aarch64
     - Achievable if Alice, LHCb and Belle2 workloads are also built for aarch64

Discussion ongoing to evaluate the feasibility and timescale of case-3

Accounting transition HS06->HEPscore

Involves the Accounting Task Force

Strategy being designed around some principles/constraints

- Accounting/pledges transparency/continuity
- Lower impact on sites
- Progressive adoption
- Time constraint: draft resource requests made by Sept 2023 for year 2025

More details at the WLCG workshop (on Wed Nov 09, 9:00-10:00)

- Your feedback is welcome!

Extend HEPscore to heterogeneous resources

In the future WLCG resources will include GPUs

- This is already true for the online farms
- HEP experiments have/are re-writing their offline applications to use also **GPUs**

HEP Benchmark project:

growing support for heterogeneous workloads

- Madgraph4gpu
- CMS HLT-like
- ML/AI train AI model (e.g. MLPF)

Prototypes of analysis workloads are also available

All this is still too premature to be included in a production HEPscore

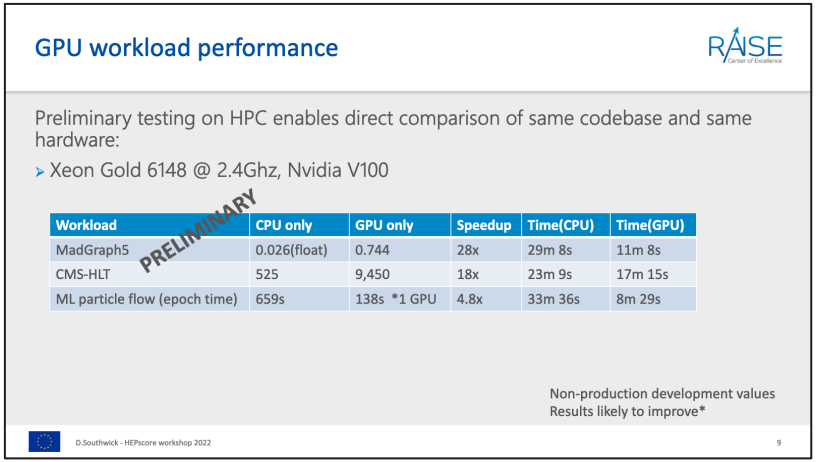

Conclusions

The replacement of HS06 with HEPscore for CPUs will very likely happen in 2023

- The HEPscore TF shall submit to the WLCG MB, by the end of 2022,
  a recommendation about the strategy for the adoption of HEPscore

There are still few open questions, in the agenda of the WLCG workshop

- Can we adopt rapidely workloads that run on x86 and aarch64?
- Converge on the procedure for the accounting transition

Input from this workshop is greatly appreciated!

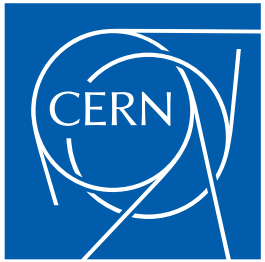
CERN