Plans for HEPscore

D. Giordano (CERN/IT) Randall Sobie (UVic)

WLCG MB

20 Dec 2022

Consultations

Progress reported in several forums

- HEPscore Deployment Task force (2 meetings/month since 2020)
- HEPscore Workshop (19-20 Sept 2022)
- HEPiX Workshop (31 Oct - 03 Nov 2022)
- WLCG Workshop in Lancaster (7-11 Nov 2022)
- GDB (14 Dec 2022)

Large feedback received from our WLCG community

Facts

- HEPscore will contain the following workloads
  - ATLAS (gen, reco) , CMS (gen-sim, reco), LHCb (sim), Belle2 (gen-sim-reco)
  - ALICE (digi-reco) needs to pass the tests of stability and reproducibility
    - see backup slide for details
- HEPscore will be normalized on the reference server to the same HS06 value of the reference server (a.k.a 1:1 normalization)
- The proposal is that HEPscore will replace HS06 as of 1st April 2023
  - The HEPscore configuration will be named HEPscore23 (HS23)

Workloads Status

List only workloads part of HS23

- [1] Alice x86 is under validation
- [2] LHCb recent requests: update x86 version
- [3] Alice, Belle2 and LHCb are w.i.p. for ARM

Exp	WL	x86_64	aarch64
ALICE	digi-reco / reco	🔨 [1]	🔨 [3]
ATLAS	gen_sherpa	✅	✅
	reco_mt	✅	✅
Belle2	gen-sim-reco	✅	🔨 [3]
CMS	gen-sim	✅	✅
	reco	✅	✅
LHCb	sim	🔨 [2]	🔨 [3]

The resolution of the validated workloads is O(‰)

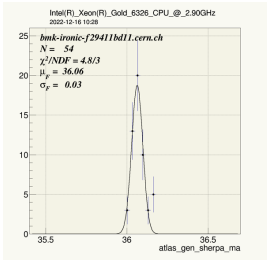

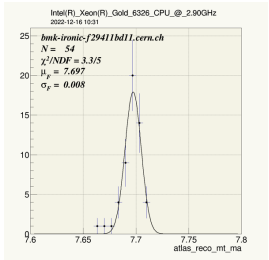

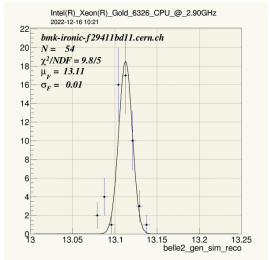

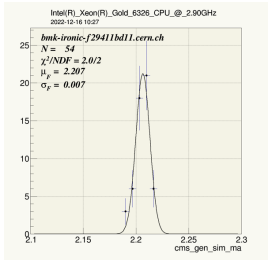

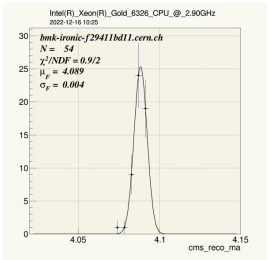

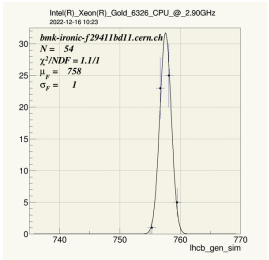

Reference Server

CPU model: Gold 6326 CPU @ 2.90GHz (Ice Lake)

- Dual socket, 32 processors (HT OFF), 257 GB RAM
- Server available at CERN. One of the most recent.
  - Request for long term maintenance to be addressed to the CERN facility team
  - Experiments can request via Openstack-Ironic similar servers for tests
- All available workloads have been already validated on this server

HS23(beta) in action

Testbed at CERN is running a beta version of HS23

- Reproducibility of HS23 score is of the order of 2‰
- This version will be update with the new Alice and LHCb workloads

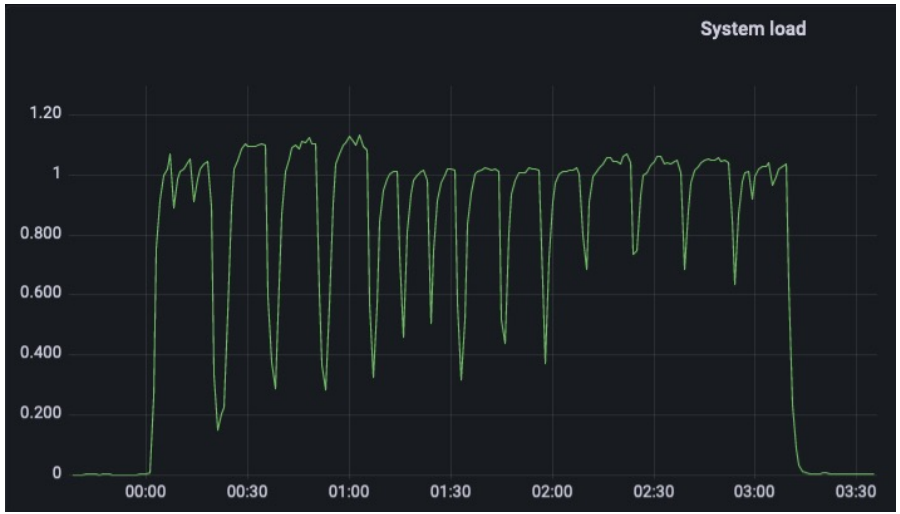

CPU	Online CPUs	Count	50th percentile of score	Spread [%]
Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz	0-31	18	306.733994	0.42
Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz	0-15	7	268.556885	0.59
Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz	0-47	18	458.076248	0.27
Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz	0-23	7	400.858795	0.35
Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz	0-55	18	592.281708	0.51
Intel(R) Xeon(R) CPU E5-2680 v4 @ 2.40GHz	0-27	7	522.627319	0.19
Intel(R) Xeon(R) Gold 6326 CPU @ 2.90GHz	0-31	9	859.123108	0.22
Intel(R) Xeon(R) Gold 6326 CPU @ 2.90GHz	0-63	1	909.185791	0.00
AMD EPYC 7302 16-Core Processor	0-63	1	960.730103	0.00

Milestones

- **1st April 2023** HEPscore23 in production
- **1st March 2023** HEPscore23 configuration frozen
  - Allows for 1 month for tests
- **14th Feb. 2023** HEP Workloads frozen
  - Allows for 2 weeks of tests/fixes
  - Latest date to have HEPscore23 for x86 and ARM
    - Otherwise ARM support will be added in a next version HEPscore2x, with x>3
  - In case a new workload does not pass the validation:

    (a) the corresponding current one can be used; (b) exclude from HEPscore23

Accounting perspective

- Migration procedure detailed by Julia in the last GDB (talk)
- The **1:1 normalization** of HS06 and HS23 simplifies the transition
  - Less changes to the code
- Sites will be expected to only benchmark new hardware with HEPScore. Old hardware does not need to be re-benchmarked
  - However, sites wishing to re-benchmark old hardware may do it (*outcome of WLCG workshop and the last GDB*)

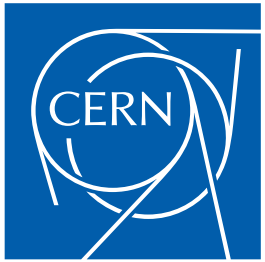
CERN

Why ALICE is still under preparation

- This workload is different from all the other workloads
  - The microservice design poses more challenges for the control of the resources
  - Because of that, CPU utilization and CPU load have very different pattern

A Large Ion Collider Experiment

Reminder of ALICE software framework in Run3

- ALICE developed new software framework for Run3 data processing, accompagnying major detector upgrade (DPL = data processing layer)
- DPL manages a topology of interacting **microservices = "devices"**, based on the concept of a (reactive) **data-flow architecture**
- "Devices" are long lived **process, transforming input data** and forwarding result via **messaging (sockets, shared memory)**
- **Common building blocks / model** for online and offline computing:
  - real-time reconstruction (data reduction)
  - async reconstruction
  - MC workflows + Analysis

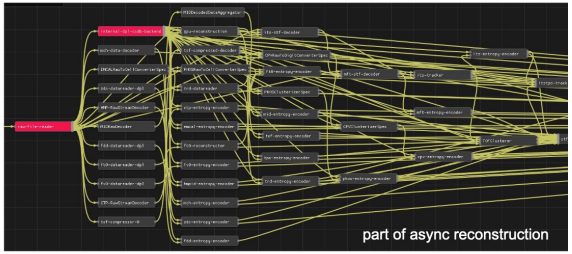

its-tracker "device" (a long-lived linux process)

Flow of data

HEPscore Workshop | September 19, 2022 | Sandro Wenzel

2