

HEPiX Benchmarking Working Group Report

Chris Hollowell <hollowec@bnl.gov> - Scientific Data and Computing Center (SDCC)
On Behalf of the HEPiX Benchmarking Working Group

Domenico Giordano and Michele Michelotto (co-chairs)

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

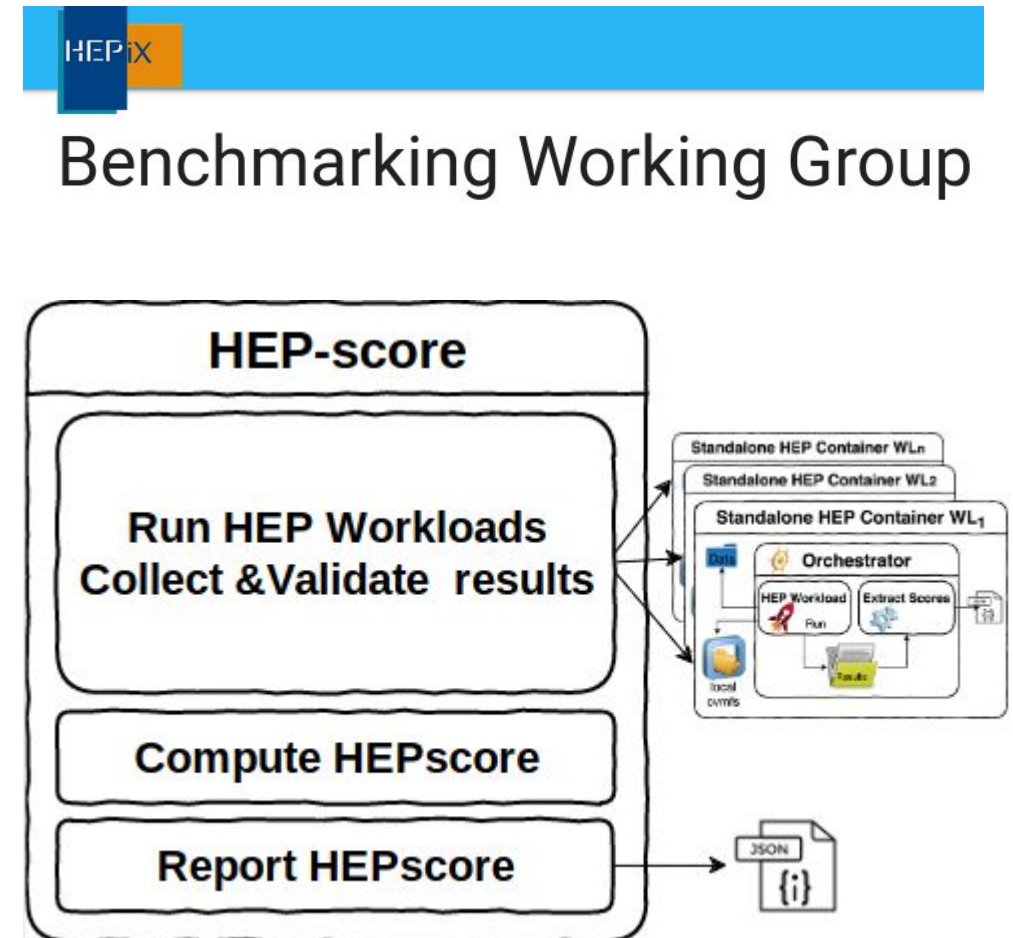
HEPiX Fall 2023 - University of Victoria
October 16, 2023



@BrookhavenLab

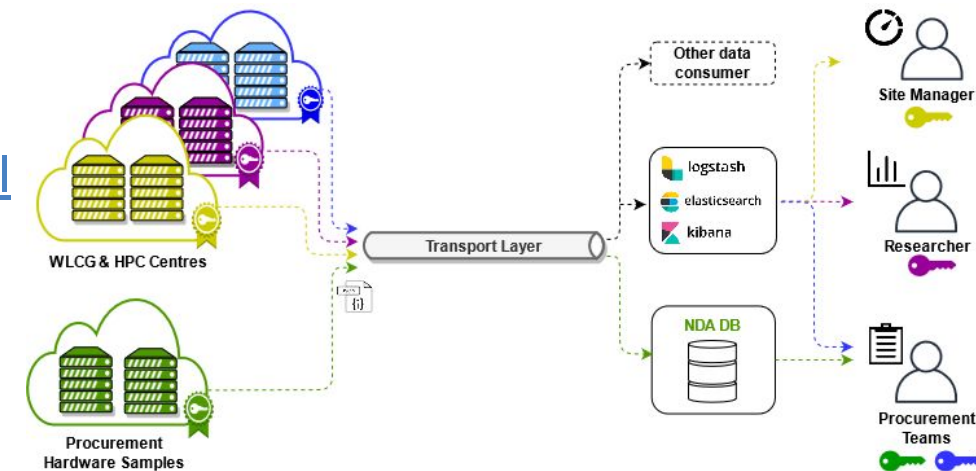
Introduction

- The [HEPiX Benchmarking Working Group](#) is tasked with managing the standard CPU benchmark used by the WLCG experiments and the HEPiX community
 - Utilized for accounting of WLCG compute resources
- For many years, this was HEPSPROC06, the “all_cpp” subset of the SPEC CPU2006 benchmark
- In 2019 the group started developing a new benchmark based on actual containerized workloads from the WLCG experiments: HEPscore
 - [HEPscore23](#) was ratified by the [WLCG HEPscore Deployment Task Force](#) as a replacement for HEPSPROC06 in April 2023, after an extensive and multi-site measurement/reproducibility validation campaign
- Also recently mandated by the WLCG Management Board to work on GPU benchmarking



Introduction (Cont.)

- Sites may continue to use HEPSPEC06 (HS06) scores for previously pledged resources
 - Resources purchased after April 2023 should be pledged with HEPscore23 instead
 - HEPscore23 is multiarch: supports transparent execution on both x86_64 and ARM
- The working group develops a tool to simplify the execution of HEPscore23, other benchmarks (HEPSPEC06, DB12, etc.) and the automated reporting of results
 - [The HEP Benchmark Suite](#)
- The group also maintains a central repository allowing sites and procurements teams to store (via the Suite) and search benchmark results for different hardware
 - https://w3.hepix.org/benchmarking/scores_HS23.html
- Consists of members from several HEPiX and WLCG institutions




HEPscore23

- 7 workloads included
- All utilize a recent version of the WLCG experiments' software
 - Run 3 based (for the LHC experiments)
- 3 Single process workloads and 4 multi-thread/process workloads
 - In default configuration, utilizes all logical cores on the host
- Workloads evenly weighted
 - Geometric mean of sub-benchmark scores used to calculate single overall score
- Reference server:
 - Intel® Xeon® Gold 6326 CPU @ 2.90 GHz with hyperthreading enabled

Exp	Workload	x86_64 / aarch64	Sw version
ALICE	digi-reco	✓	O2/nightly-20221215-1
ATLAS	gen_sherpa (SP) <small>Ⓢ SP: Single Process</small>	✓	Athena 23.0.3
	reco_mt	✓	Athena 23.0.3
Belle2	gen-sim-reco (SP)	✓	release-06-00-08
CMS	gen-sim	✓	CMSSW_12_5_0
	reco	✓	CMSSW_12_5_0
LHCb	sim (SP)	✓	v3r412

HEPscore23 (Cont.)




- HEPscore23 performs better on newer CPU microarchitectures than HEPspec06
 - It is a more accurate representation of the modernization that has taken place in HEP applications compared to HS06



HS23/HS06 Ratios by CPU Model Year

HEPscore23 (Cont.)

- Extensive studies on multiple CPU models and architectures have shown the benchmark to be highly reproducible
 - Scatter of results less than 0.5% on all systems in the CERN testbed



hep-score Application Updates

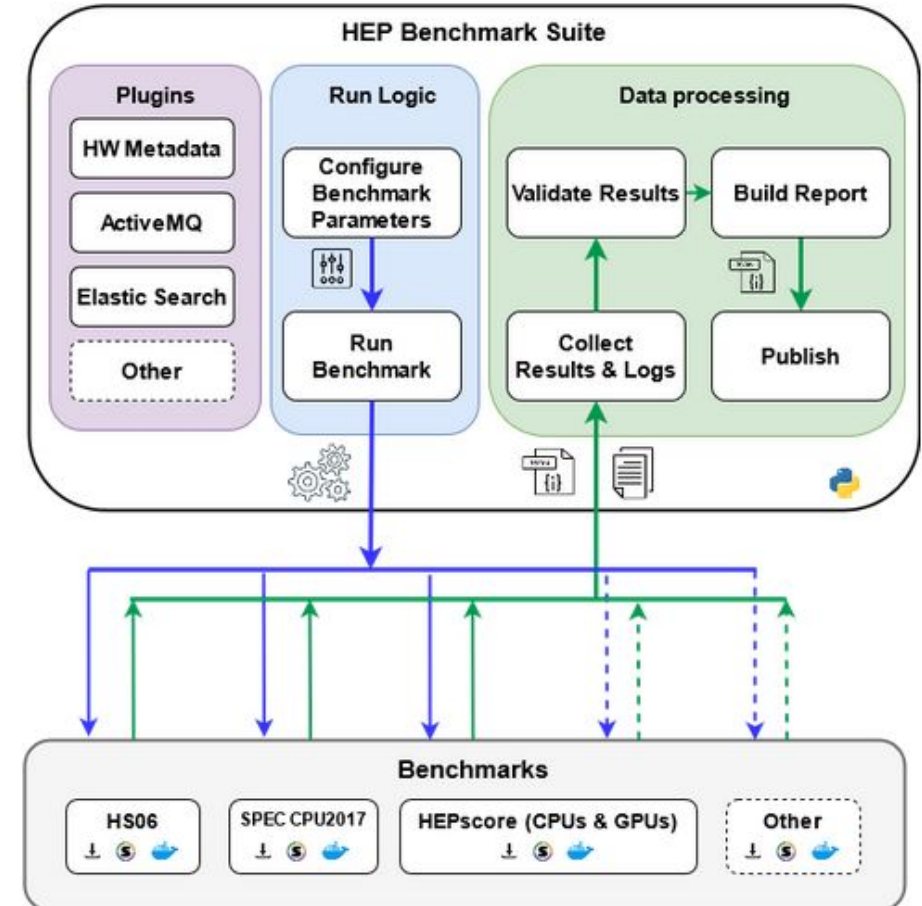
- **hep-score application v1.5 released in March 2023**
 - Current stable version, and default version used for suite execution
- **Expect v1.6 release in the near future**
 - Several v1.6 release candidates (rc) made available for testing
 - Many new features and bugfixes, including:
 - Allow registry to be overridden via the command line and options field
 - Add support for lists of registry types, with one to utilize selectable via the '-i' option
 - Configurable number of cores to be loaded (ncores) to be passed to the workloads, if the workloads implement this feature
 - Unset problematic Singularity/Apptainer environment variables

```
$ hep-score -i dir -n 2 /tmp
2023-10-11 23:14:23 hepscore [INFO] HEPscore23 Benchmark
2023-10-11 23:14:23 hepscore [INFO] Config Hash:                f5216b9cd1e2d5bd8c799ead4fa7c4ecd3f39c71d7cd624d6025056906204856
2023-10-11 23:14:23 hepscore [INFO] HEPscore version:        1.6.0.0rc4
2023-10-11 23:14:23 hepscore [INFO] System:                Linux test00.sdcc.bnl.gov x86_64
2023-10-11 23:14:23 hepscore [INFO] Container Execution:    singularity
2023-10-11 23:14:23 hepscore [INFO] Implementation:        aptainer
2023-10-11 23:14:23 hepscore [INFO] Registry:               dir:///cvmfs/unpacked.cern.ch/gitlab-registry.cern.ch/hep-benchmarks/hep-workloads
2023-10-11 23:14:23 hepscore [INFO] Output:                 /tmp/HEPscore_11Oct2023_231423
2023-10-11 23:14:23 hepscore [INFO] Date:                   Wed Oct 11 23:14:23 2023

2023-10-11 23:14:23 hepscore [INFO] Executing 3 runs of atlas-gen_sherpa-ma-bmk [v2.0_x86_64]
2023-10-11 23:14:23 hepscore [INFO] Enforcing run of each workload on only 2 cores
2023-10-11 23:14:23 hepscore [INFO] Starting run0
```

HEP Benchmark Suite Application Updates

- Suite v2.2 released in June 2023
 - Current stable version, and default version used for suite execution with `run_HEPscore.sh`
 - Numerous changes, including:
 - Multi-python version and platform wheel builds
 - Support for builtin hepscore configurations
 - The inclusion of a script (`populate_cache.py`) to pre-populate Singularity/Apptainer/Docker caches with workload images
- Expect v3.0 release in the near future
 - Release candidates (rc) available for testing
 - Includes support for time series plugin extensions



HEP Benchmarking Suite Time Series Plugins

- HEP Benchmark Suite extended to support executing commands while benchmarks are running (and also before and after) and record time series data: CommandExecutor plugin
 - Can be used to record energy consumption, load, memory use, CPU frequency
 - via ipmitool, uptime, free, etc.
- Can be correlated with the benchmark score measurements
 - Increasing interest in energy/performance studies in WLCG


Example: Plugin Configuration

```
{
  "name": "power-consumption",
  "description": "Retrieves power consumption of the system. Requires elevated p",
  "command": "ipmitool dcmi power reading",
  "regex": "Instantaneous power reading:\\s*(?P<value>\\d+) Watts",
  "unit": "W",
  "example-output": "\n  Instantaneous power reading:          124 Wa",
  "expected-value": 124
},
{
  "name": "load",
  "description": "Retrieves the one minute system load average. Note that lo",
  "command": "uptime",
  "regex": "load average: (?P<value>\\d+\\.\\d+),",
  "unit": "",
  "aggregation": "sum",
  "example-output": " 11:02:47 up  3:03,  1 user,  load average: 0.18, 0.38,",
  "expected-value": 0.18
}
```

Plugins' Report


```
114+   "plugins": {
115+     "CommandExecutor": {
116+       "hepscore": {""},
305+       "pre": {
306+         "load": {
307+           "start_time": "2023-09-03T11:19:08.772095Z",
308+           "config": {
309+             "command": "uptime",
310+             "interval_mins": 1,
311+             "aggregation": "sum",
312+             "regex": "load average: (?P<value>\\d+\\.\\d+),",
313+             "unit": ""
314+           },
315+           "values": [11.05,10.47,10.49,10.22,10.45,10.54],
316+           "end_time": "2023-09-03T11:24:08.774851Z",
317+           "statistics": {
318+             "min": 10.22,
319+             "mean": 10.536,
320+             "max": 11.05
321+           }
322+         },
323+         "status": "success",
324+         "used-memory": {""},
325+         "used-swap-memory": {""}
326+       },
327+       "post": {""}
328+     }
329+   }
330+ }
```

Rebuild time series from stored data



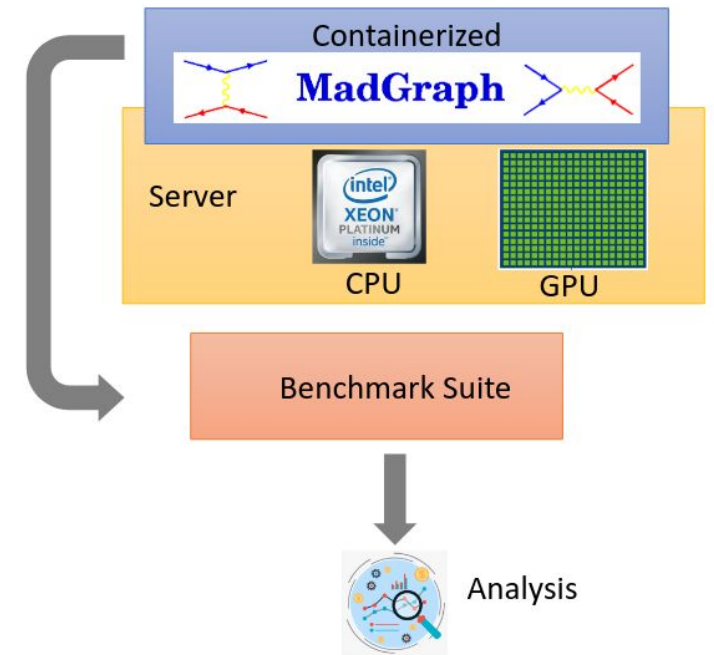
Preparation of HEPscore GPU Workloads

- CERN [student project](#) to measure energy efficiency on GPU workloads
- Utilized containerized Madgraph workload ported to CPU and GPU with the CommandExecutor HEP Benchmark Suite time series plugin



Keshvi Tuteja

- Server Specs:**
- **CPU**
Intel Platinum 8362 @ 2.80GHz
Cores per socket: 32
Sockets: 2
 - **GPU**
Model name: Nvidia L4
Total board power: 72 W
GPU clocks: 2 GHz (Boost)



Keshvi Tuteja

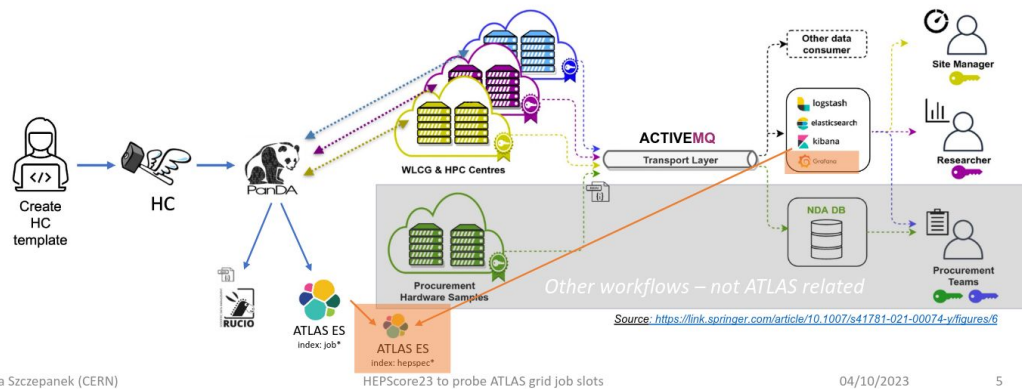
- Other GPU hep-workloads containers available: Simpletrack, HLT-CMS, MLPF

Studies on Grid Jobs

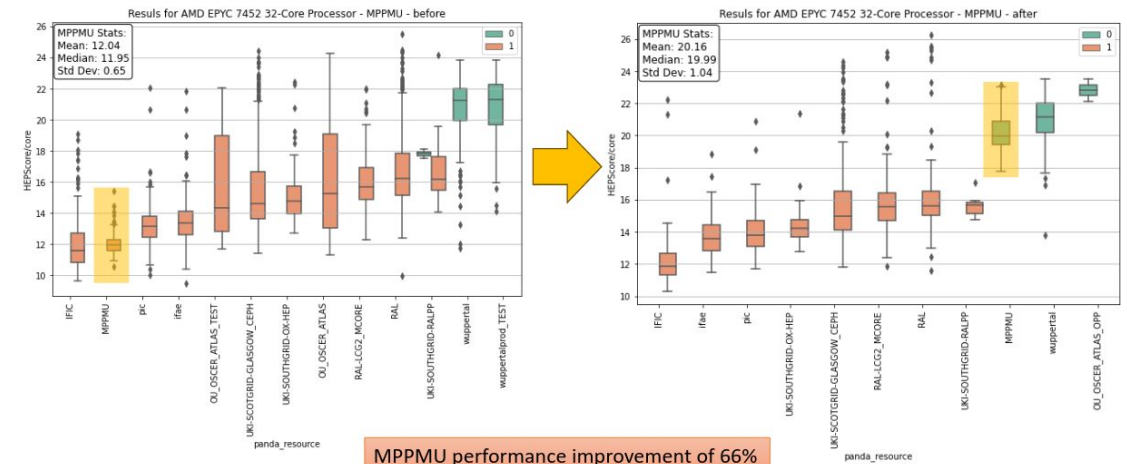
- Ran HEPscore23 via PanDA to continuously measure grid site performance
- Was utilized to determine some underperforming sites, including one where HT was incorrectly disabled
 - The group worked with site admin on correcting: resulted in a 66% performance improvement

HEPScore23 via PanDA

- We are running on 134 different Panda Resources (Queues)
 - BNL, CERN, CA-VICTORIA, DESY-HH, JINR, Vega...
- Infrastructure:
 - PanDA, HammerCloud, Rucio, ActiveMQ, Elasticsearch, Grafana, Kibana...



HS23 statistics before and after changes



Support via GGUS



- Benchmarking support via [GGUS](#) available since April 2023
 - Have received a number of tickets
 - Preferred mechanism for support due to its inherent history tracking
 - Still receiving some requests via email
- Suite result JSON that users have any issues submitting through the application can be sent via GGUS as an attachment

A screenshot of the GGUS web application's ticket submission form. The form is titled '* Subject ?' and '* Describe the issue ?'. It includes a rich text editor with a toolbar containing icons for undo, redo, bold, italic, underline, strikethrough, bulleted list, numbered list, quote, link, unlink, and insert link. Below the editor are dropdown menus for 'Concerned VO: ?' (set to 'other'), 'Affected site: ?' (set to 'please select'), '* Ticket category' (set to 'Service Request'), and 'Type of issue: ?' (set to 'Benchmarking'). There are also radio buttons for 'VO specific ?' (set to 'no') and a dropdown for 'Affected ROC/NGI'. A 'Priority: ?' dropdown is set to 'please select'. The form has four 'Attach File(s)' sections, each with a 'Browse...' button and the text 'No file selected. (max. 2 MB pro File)'. A yellow banner highlights the 'Routing information' section with the text 'Expert option, please set this option only if you know what it means.' Below this, there are dropdowns for 'Notify SITE ?' and 'Assign to support unit ?' (set to 'Benchmarking'), separated by 'OR'. A 'Submit' button is at the bottom. A legend at the bottom left indicates '* Required fields'.

Contributors Needed

- Developing, maintaining and supporting HEPscore23, the HEP Benchmark Suite, and our results table requires significant effort
 - Many of the working group's participants are contributing part time
 - A number of members are students with temporary appointments/internships
- We are looking for additional members of the HEPiX community to contribute
 - Testing benchmarks
 - Development of benchmark workload containers
 - Development of the hep-benchmark-suite and hep-score applications
- If you are interested, please contact the co-chairs:
 - Domenico Giordano <Domenico.Giordano@cern.ch>
 - Michele Michelotto <Michele.Michelotto@pd.infn.it>

Conclusions

- HEPscore23 was adopted as the standard WLCG CPU benchmark in April 2023
 - Support available via GGUS if needed
- The hep-score and hep-benchmark-suite applications continue to be updated with new features
 - Expect new releases in the coming weeks
- A new CommandExec plugin for the hep-benchmark-suite application allows for execution of commands and recording of time-series data while benchmarks are executing
 - Permits the collection of power usage data (i.e. via ipmitool), and other system characteristics such as memory use over time
- Studies of numerous grid site performance via HEPscore23 have successfully been completed via PanDA
 - Allowed the working group to help improve the performance of one of the sites measured
- Development and testing of GPU hep-workload containers is continuing to progress
- We are looking for additional members of the HEPiX community to join our efforts