HEP Benchmark Suite: Enhancing Efficiency and Sustainability in Worldwide LHC Computing Infrastructures

Natalia Szczepanek, Domenico Giordano (CERN)
Alessandro Di Girolamo, Ewoud Ketele, Gonzalo Menendez Borge, Ladislav Ondris (CERN)
Ivan Glushkov (University of Texas at Arlington (US)), Emanuele Simili, David Britton (University of Glasgow (UK))

ACAT, 11th March 2024

Role of WLCG and Benchmarking

- WLCG counts around 1.4 million CPU cores spread over 170 data centres
- Increasing demand for computing resources (CPU, ARM, GPU)
- Due to the responsibility that comes with it, we adhere to best practices, including benchmarking of computing resources, which is being used for:
  - Site pledges
  - Accounting reports
  - Procurement procedures
  - Performance studies

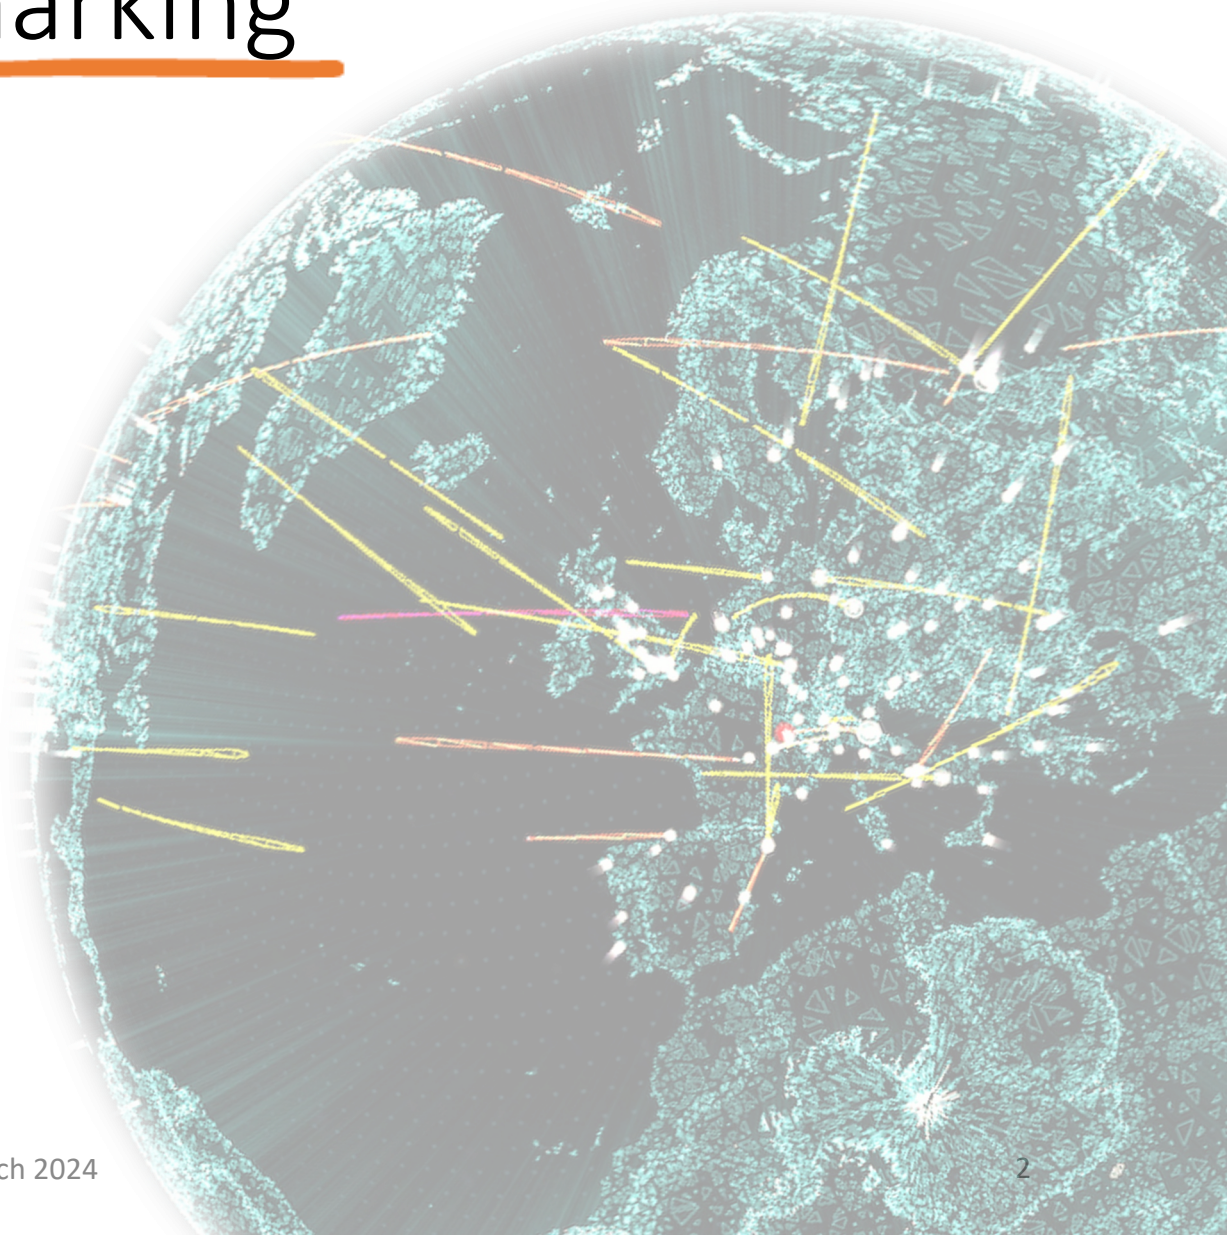

Main Objective Of This Work

- Enhance the efficiency of the server performance in production environment
  - Probe job slot performance on WLCG with HS23 and correlate it with usage metrics (load, frequency, power consumption, memory usage, etc.)
  - Monitor the performance in correlation with different metrics and understand how different configuration aspects affect the performance

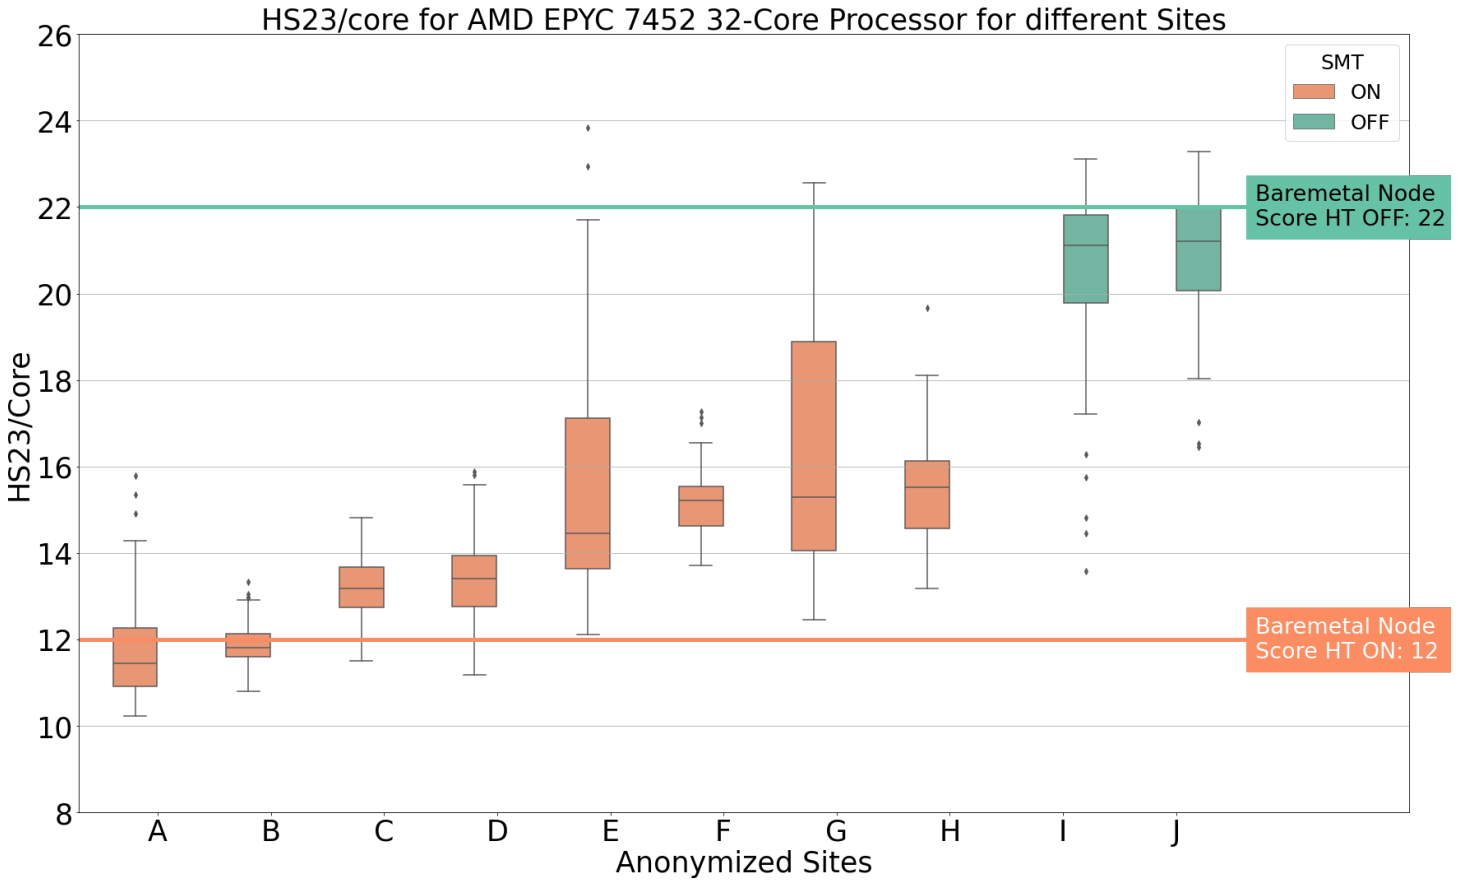

Approach

HEP Benchmark Project

HEP Benchmark Suite

- Orchestrator of multiple benchmark (HS23, HS06, SPEC CPU2017)
- Central collection of benchmark results

HEPScore

- New High-Energy-Physics specific benchmark
- It has replaced HEP-SPEC06 in 2023
- Support for non-x86 architectures, open-source
- Uses the workloads from the HEP experiments and combines them in a single benchmark score

HEPScore23 (HS23)

- The official HEPScore configuration composed by 7 workloads from 5 experiments (3 ST, 4 MT)
- 1:1 normalization with HS06 for the reference CPU Intel® Xeon® Gold 6326 CPU @ 2.90 GHz (HT=On)
- Runtime ~4h

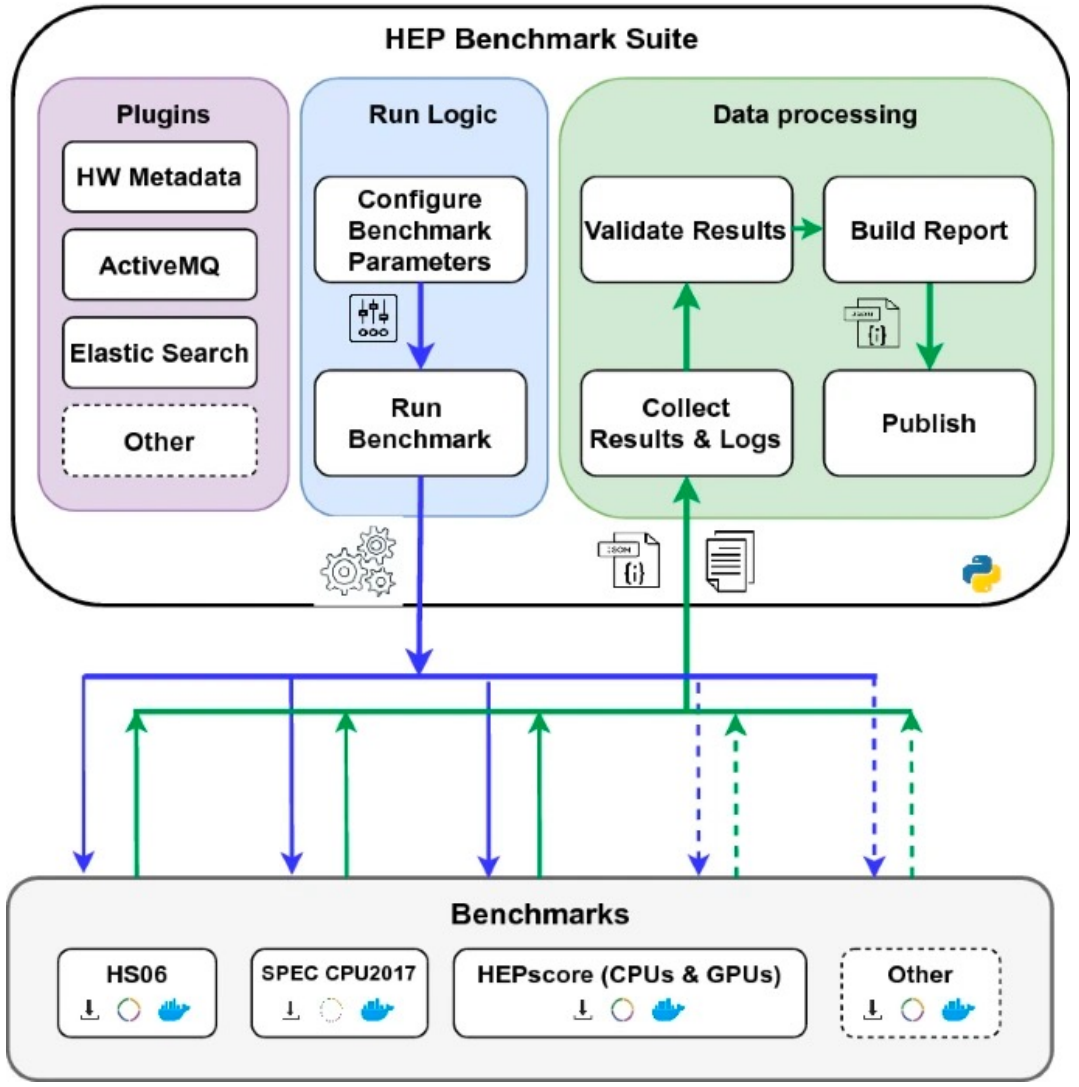

Source: https://link.springer.com/article/10.1007/s41781-021-00074-y/figures/4

The enhanced HEP Benchmark Suite

- The suite incorporates metrics such as machine load, memory usage, memory swap, and notably, power consumption
- Suite Plugins:
  - Configurable collection of various system metrics (e.g. load, memory usage, power consumption)
    - Run alongside benchmarks
    - Flexible modification and addition of collected metrics

Suite configuration

```
plugins:
  CommandExecutor:
    metrics:
      cpu-frequency:
        command: cpupower frequency-info -f
        regex: 'current CPU frequency: (?P<value>\d+).*'
        unit: kHz
        interval_mins: 1
      power-consumption:
        command: >
          sudo ipmitool sensor get 'PS1 Power In' ; sudo ipmitool sensor get
          'PS2 Power In'
        regex: 'Sensor Reading\s+:\s*(?P<value>\d+).*'
        unit: W
        interval_mins: 1
      load:
        command: uptime
        regex: 'load average: (?P<value>\d+.\d+),'
        unit: ''
        interval_mins: 1
```

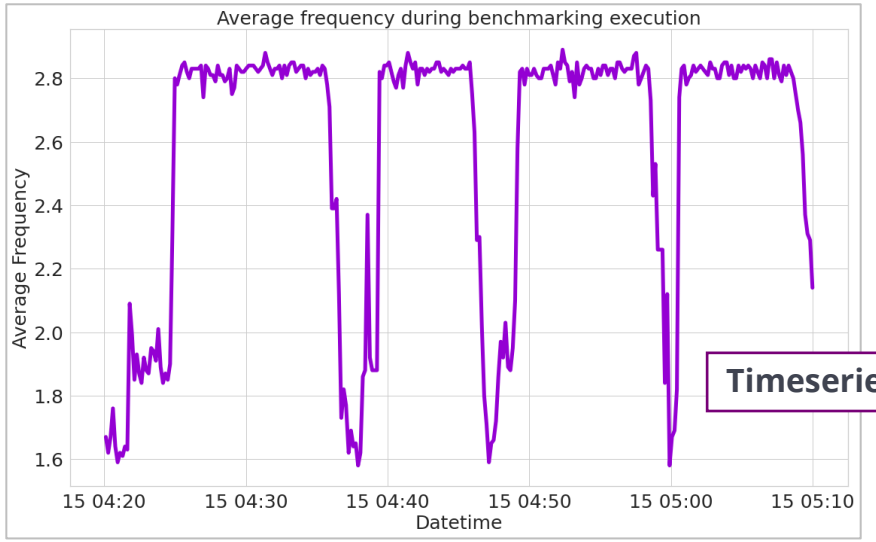

Timeseries data

Probing job slots

- Each site has servers with a variety of CPU models and number of cores (256, 128, 64…)
- We are running the benchmark injecting the HEP Suite script as a normal experiment job running inside the PILOT Apptainer
- We probe multi-core job slots (8/4/1 cores)

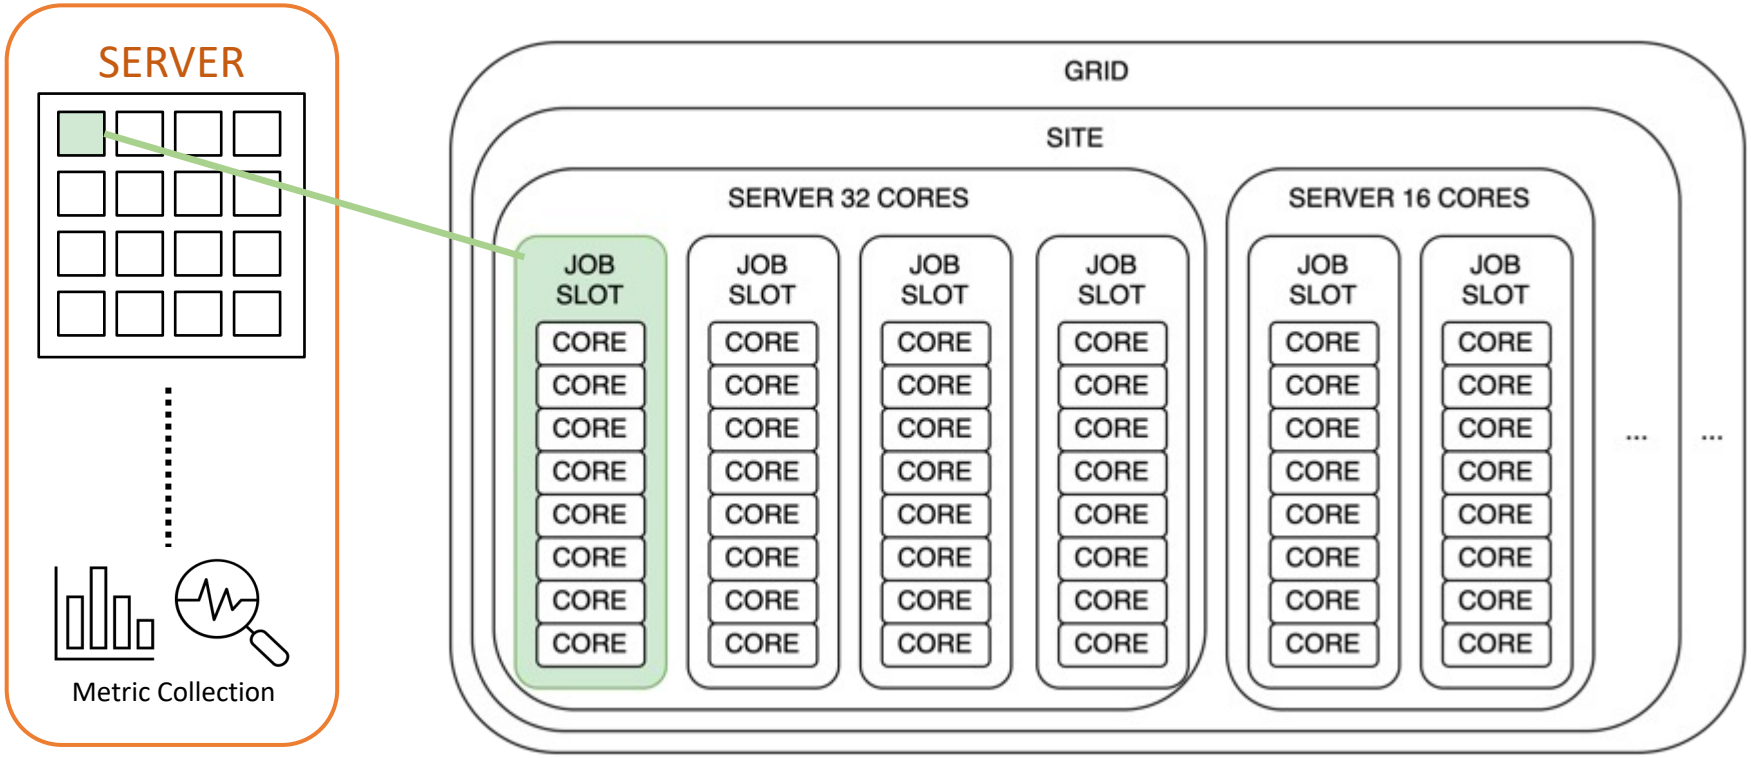

Submission Approach

- HEP Suite is adjustable for different WMS
- The benchmarking script is injected in the site job slot via the standard job submission system
- Successfully implemented and deployed the pipeline for:
  - **ATLAS:** Automated submission via HammerCloud, **over 200k jobs** finished
  - **LHCb:** Manual submission to DIRAC, **2.1k** jobs finished

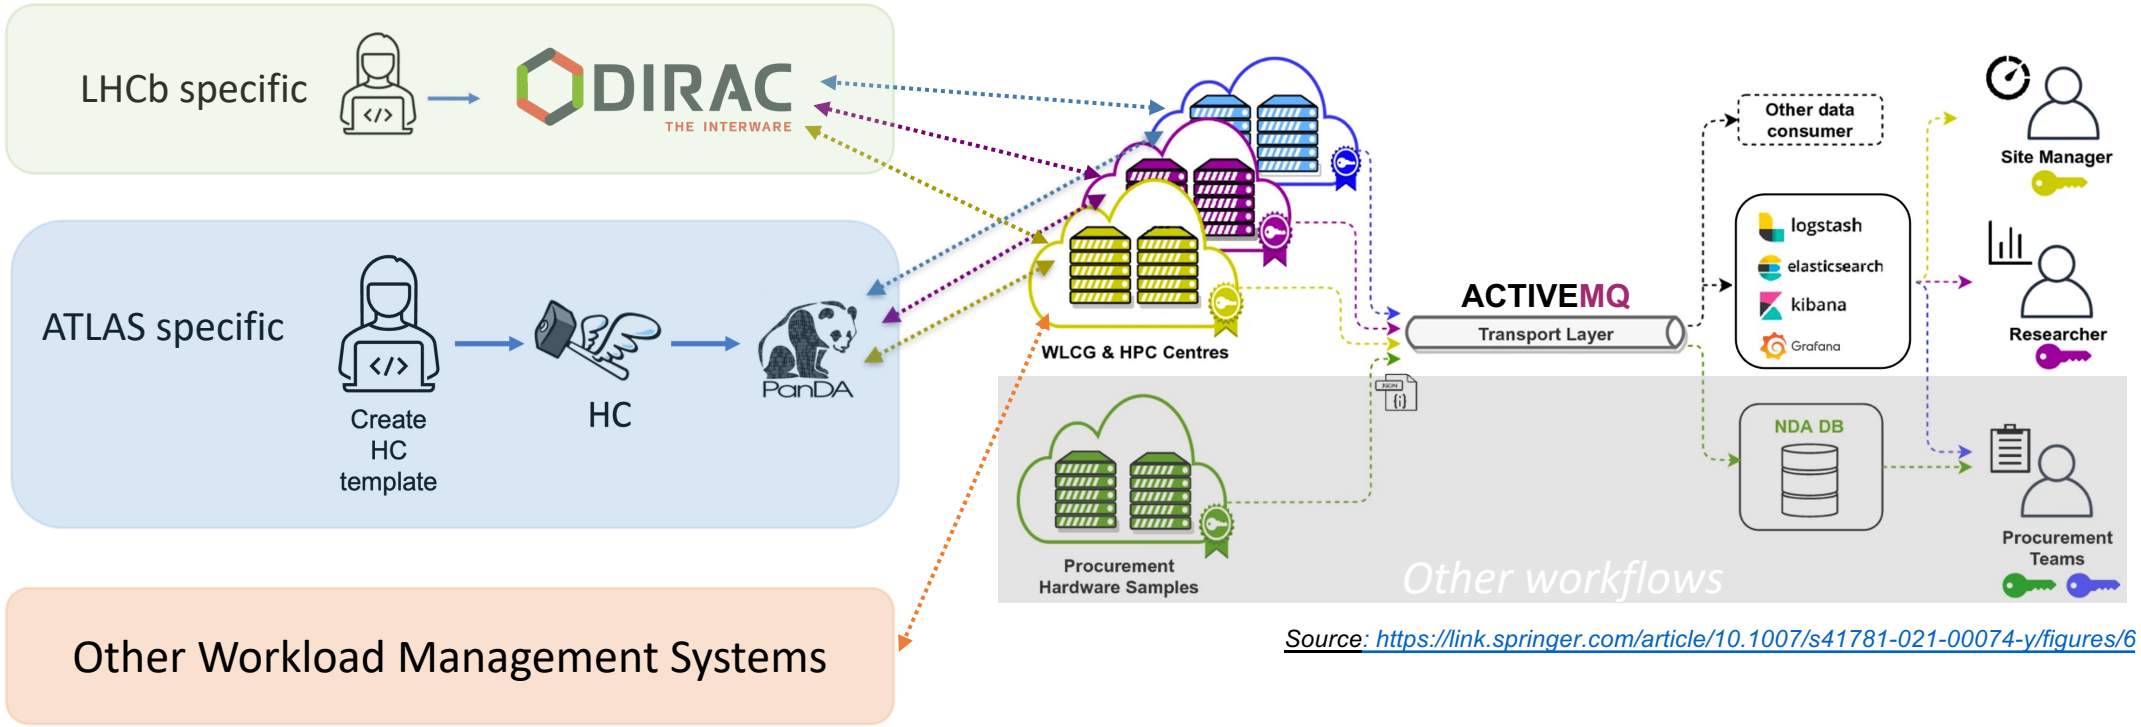

Source: https://link.springer.com/article/10.1007/s41781-021-00074-y/figures/6

Data Analysis

Expanding the feature space

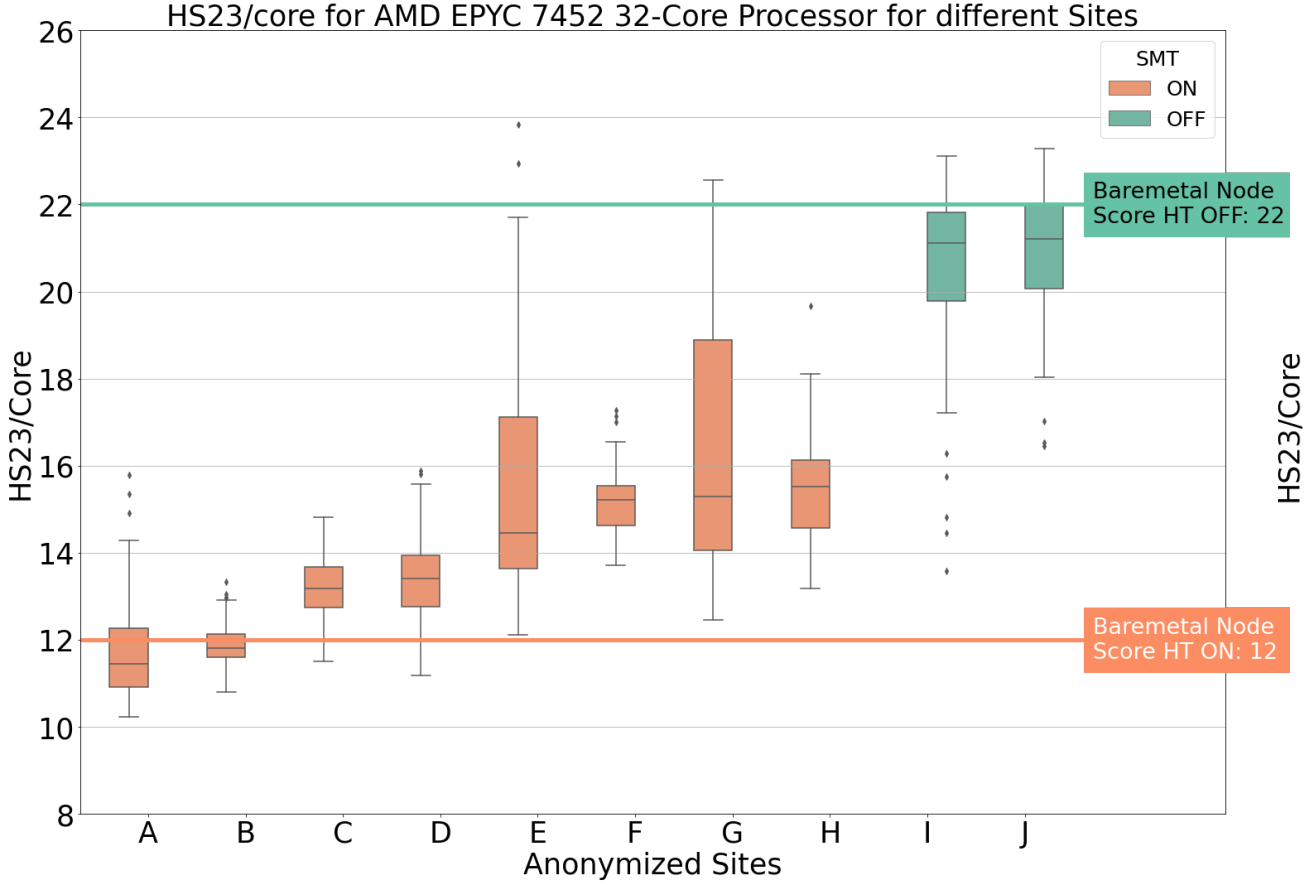

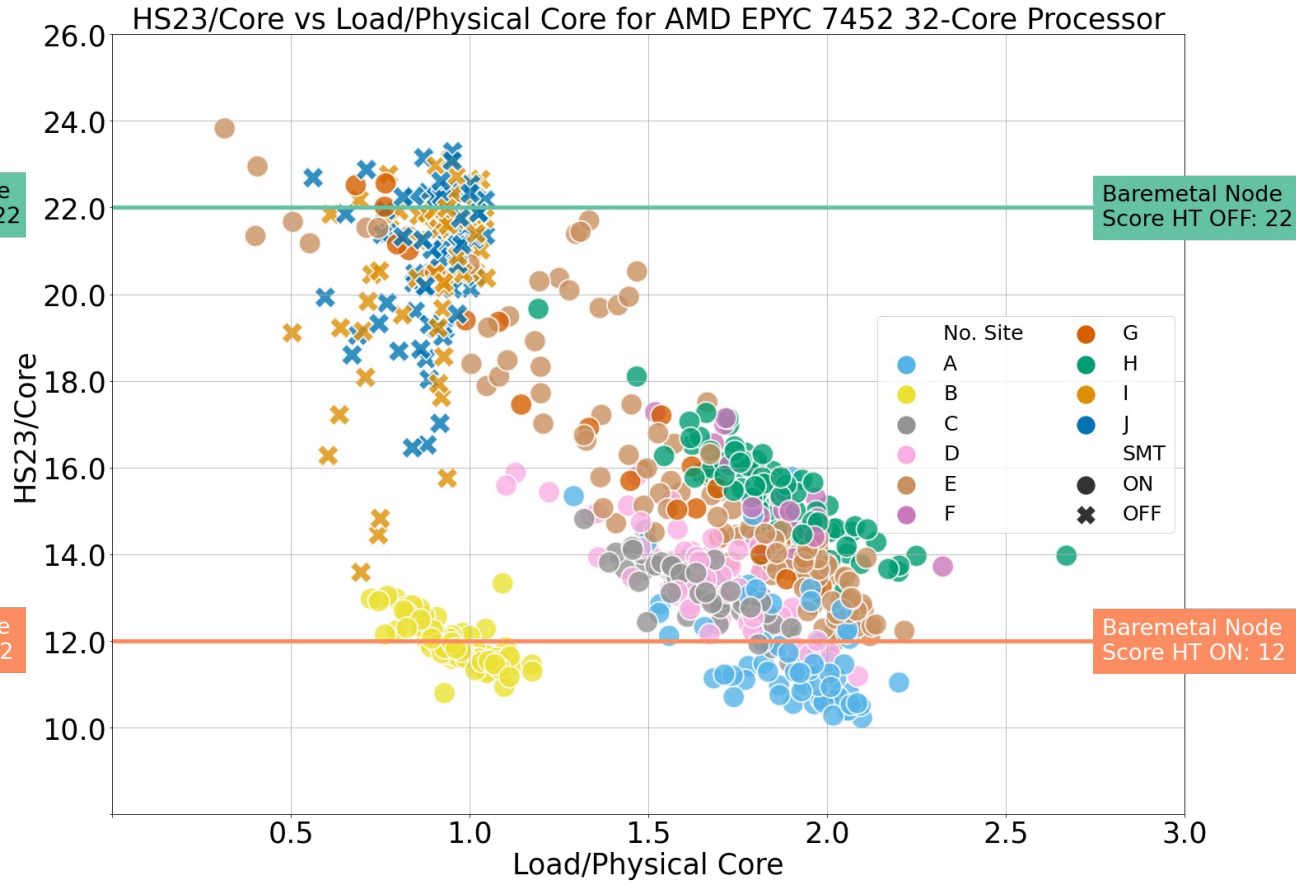

Discriminate normal trend and anomalies

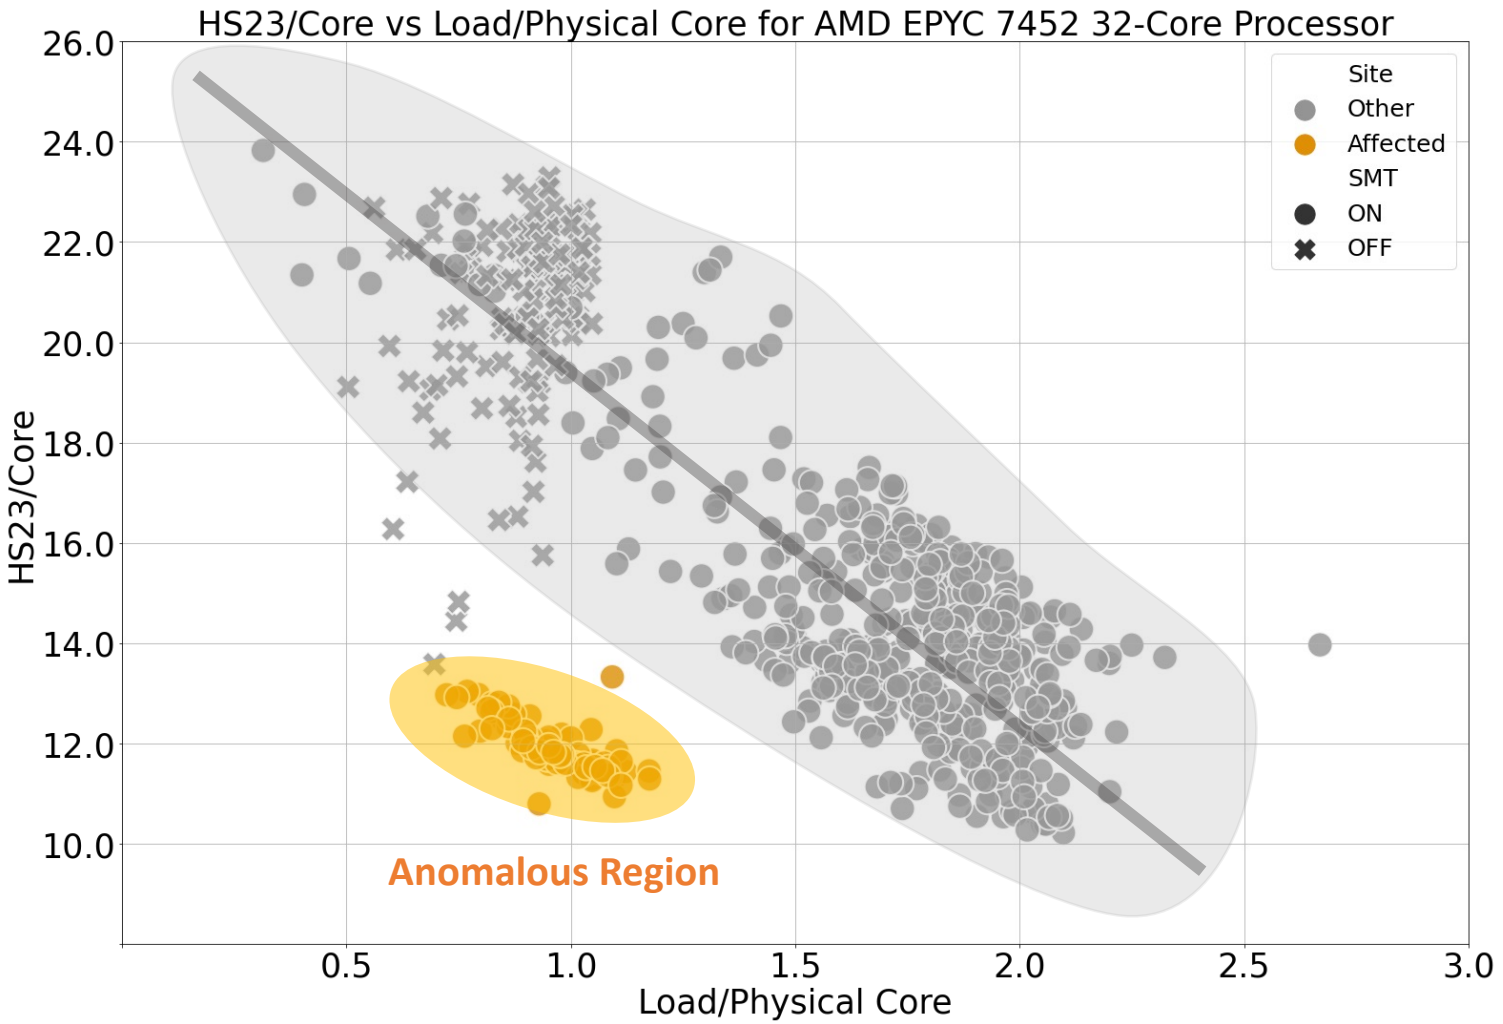

- Performance of affected site in terms of score per core is low and comparable to other sites that have double load in the servers
- After site admin investigation, it was found out that servers had a wrong masking of the disabled cores while SMT was ON
  - Then, fix has been applied

Fixing misconfiguration issues

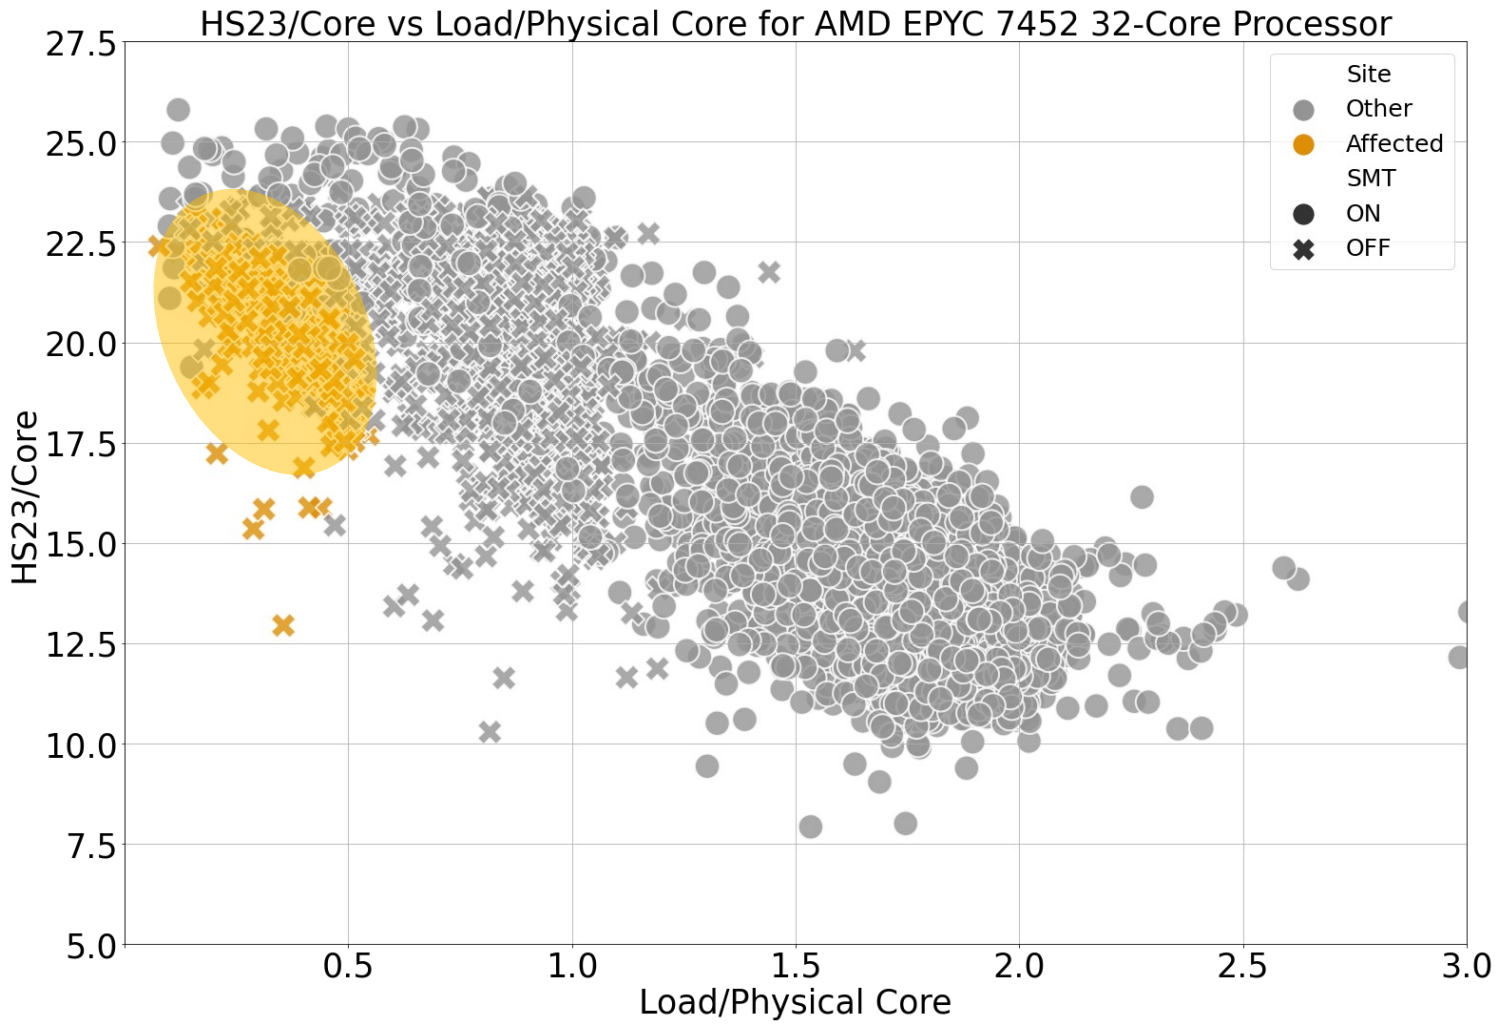

- Applied changes: SMT off fixed the problem with cores masking misconfiguration
- **Performance of affected site increased by 66%**
- **The analysis and issue resolution enhanced the site efficiency**
- Recursive studies for different CPU Models
  - Contacted > 10 sites
  - Solved misconfiguration issues
  - Or understood sites configuration choices

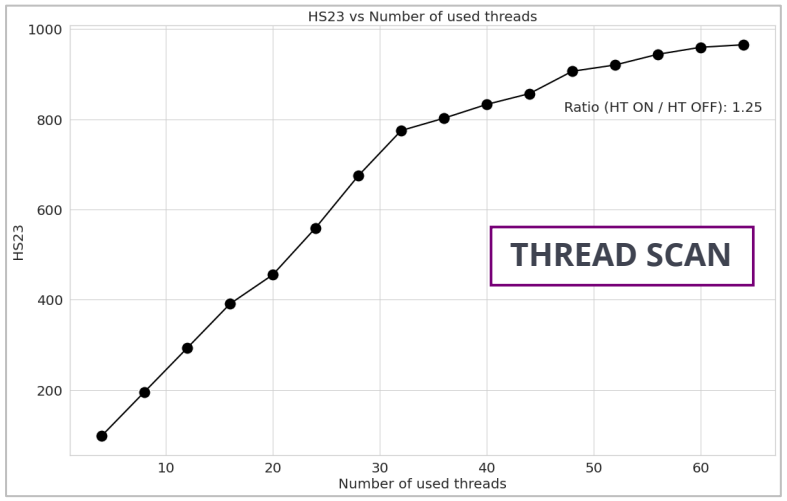

Modelize collected data

- Two main patterns:
  - Consistent with the thread scan
  - Hidden feature which cause the inconsistency with the thread scan

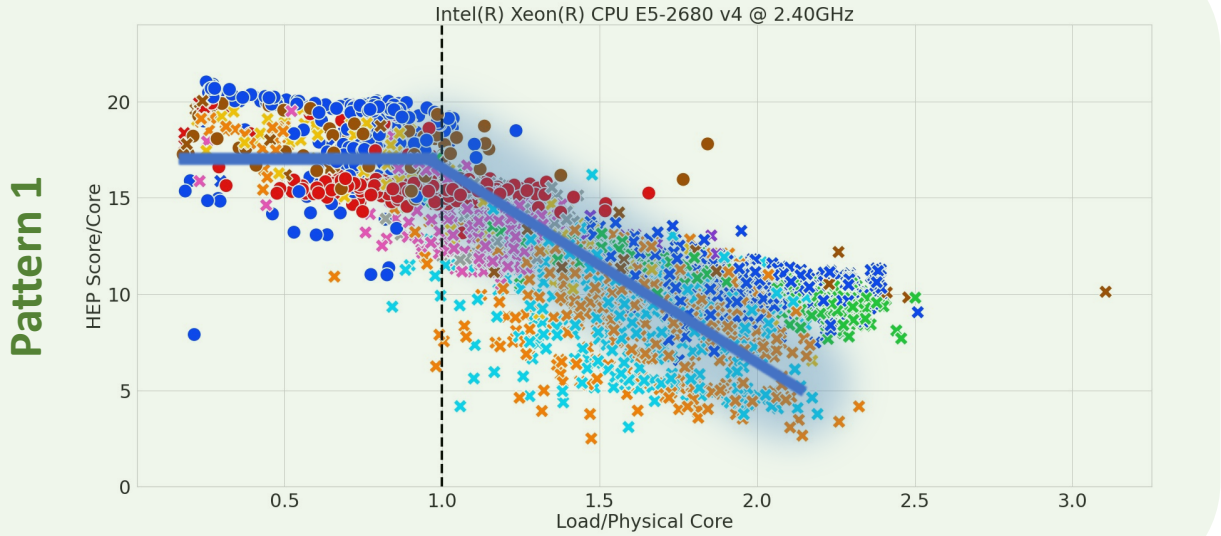

DIFFERENT MARKER COLOR = DIFFERENT SITE

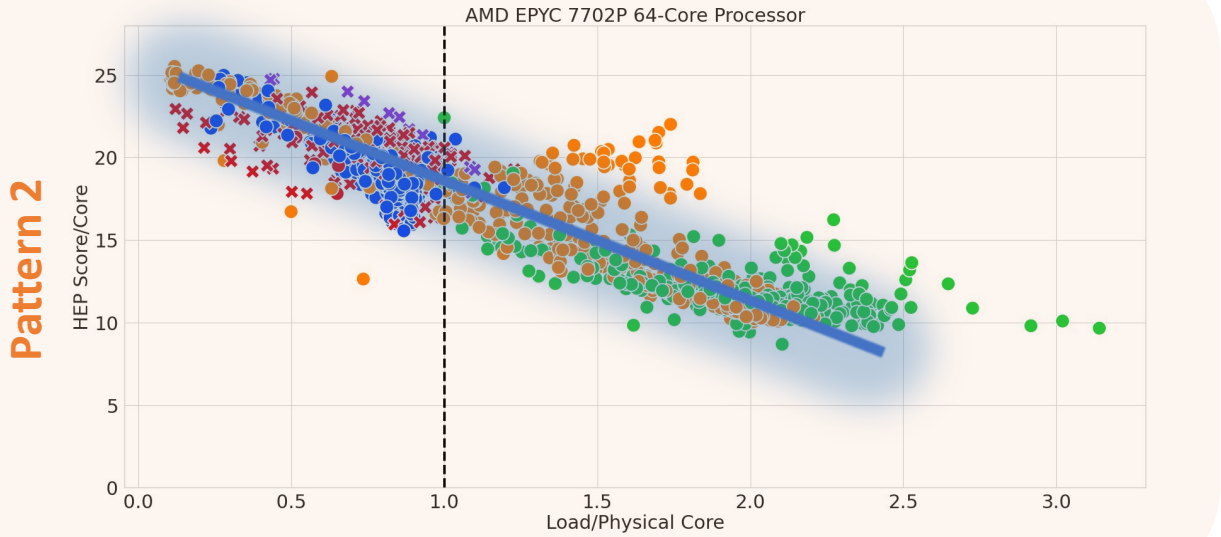

Governors Study on baremetal node

Dedicated tests on baremetal node to check the impact of different governors on performance:

- Run benchmark in a thread scan with modularity of 4
- Repeated for different governors on AMD EPYC 7302 16-Core Processor

Findings:

- Powersave governor performance is low compared to other governors, but it reduce the power consumption by working on low frequencies
- All governors represent the first visible pattern
- Results on baremetal node are consistent with results achieved from grid

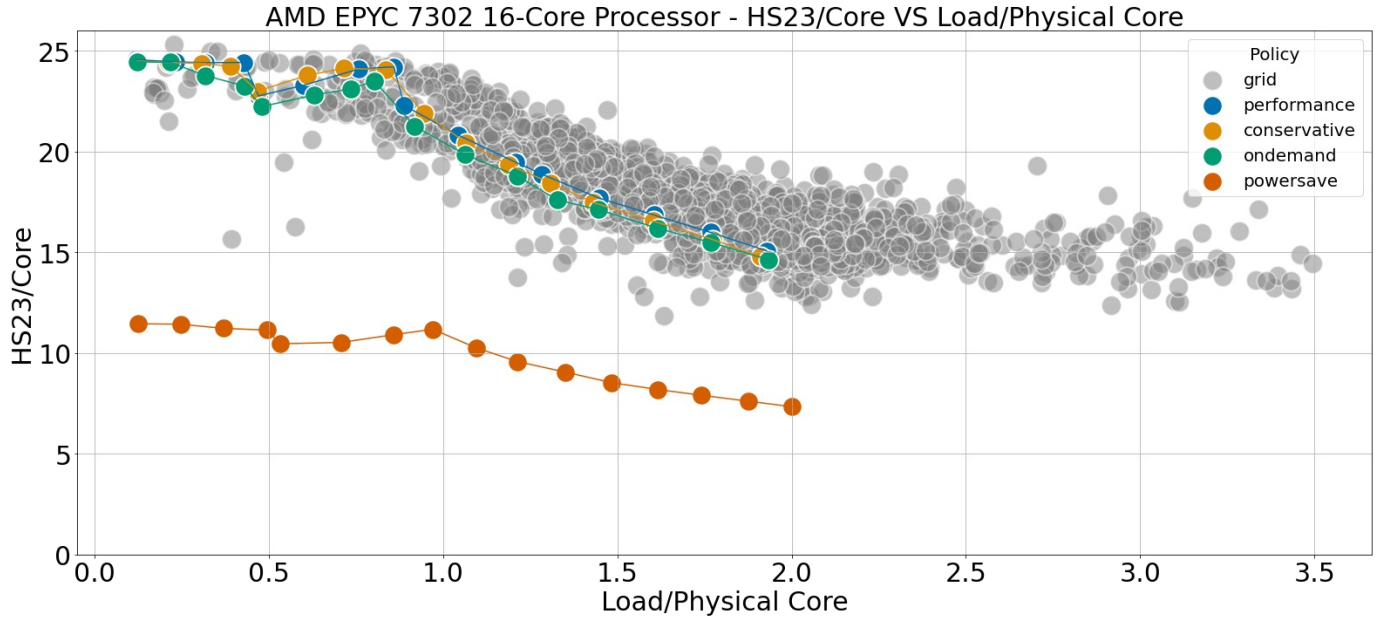

Governors Study on baremetal node

- Tests repeated for the following baremetal node: **Intel® Xeon® CPU E5-2630 v4 @ 2.20GHz** which represent the second visible pattern
- We can see similar tendency on baremetal node, there is increase of performance in the range of load (0, 1) for performance, conservative, ondemand and userspace governors
- Powersave governor looks consistent with first pattern on baremetal node, which is not the case for results achieved from grid (second pattern)
- **We need to perform more tests to find the feature responsible for this behaviour**

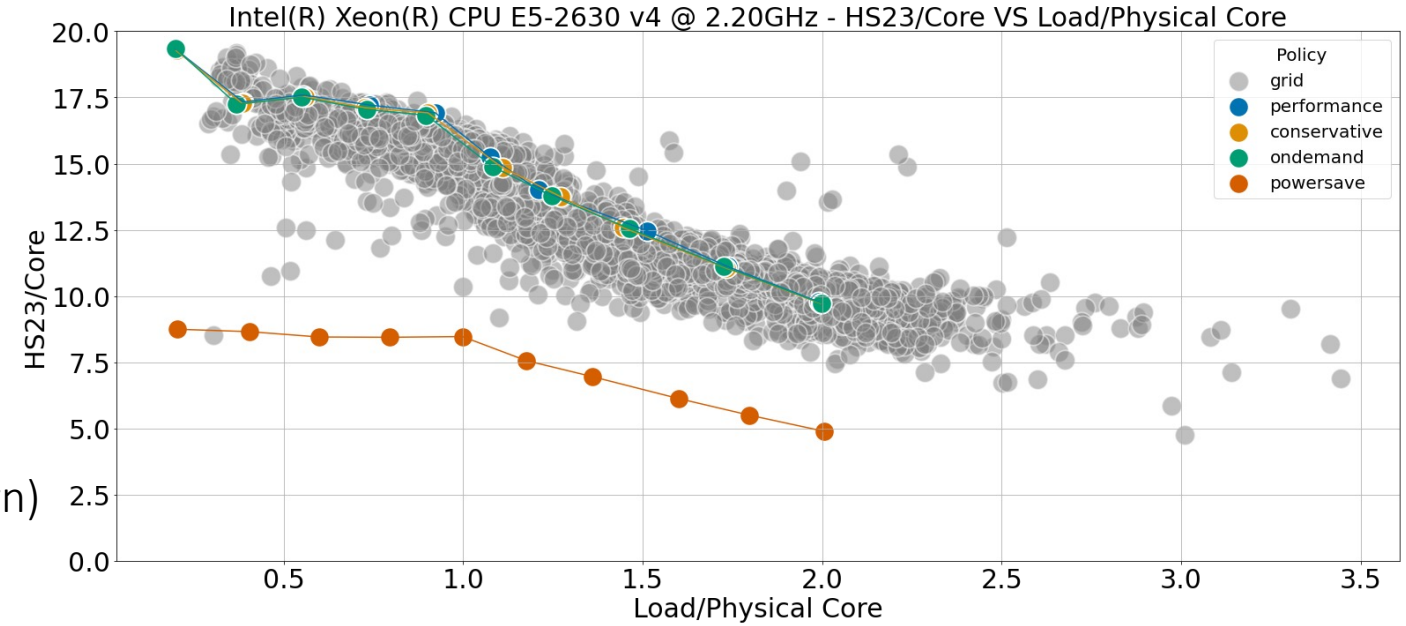

Data Model (so far)

Work in progress

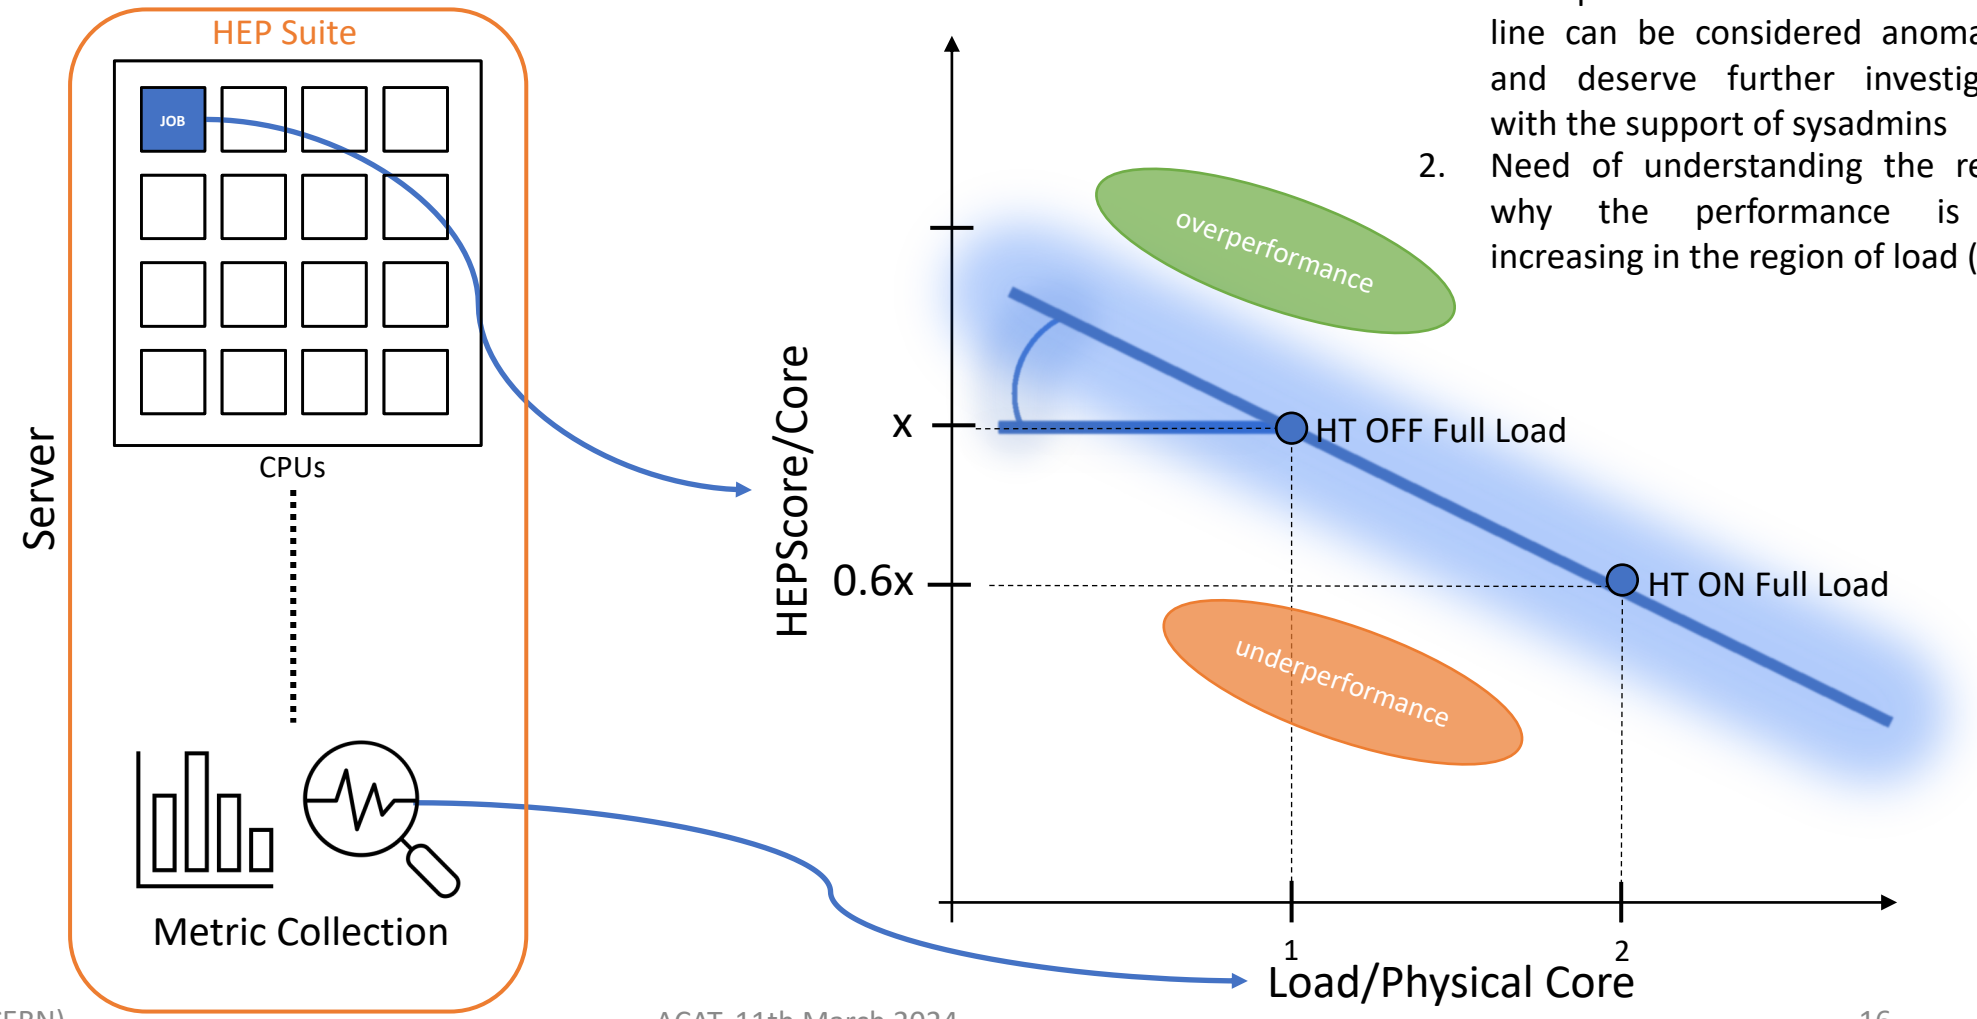

1. Data points outside the main trend line can be considered anomalous, and deserve further investigation with the support of sysadmins
2. Need of understanding the reason why the performance is still increasing in the region of load (0, 1)

Power Measurements

Power Consumption Study

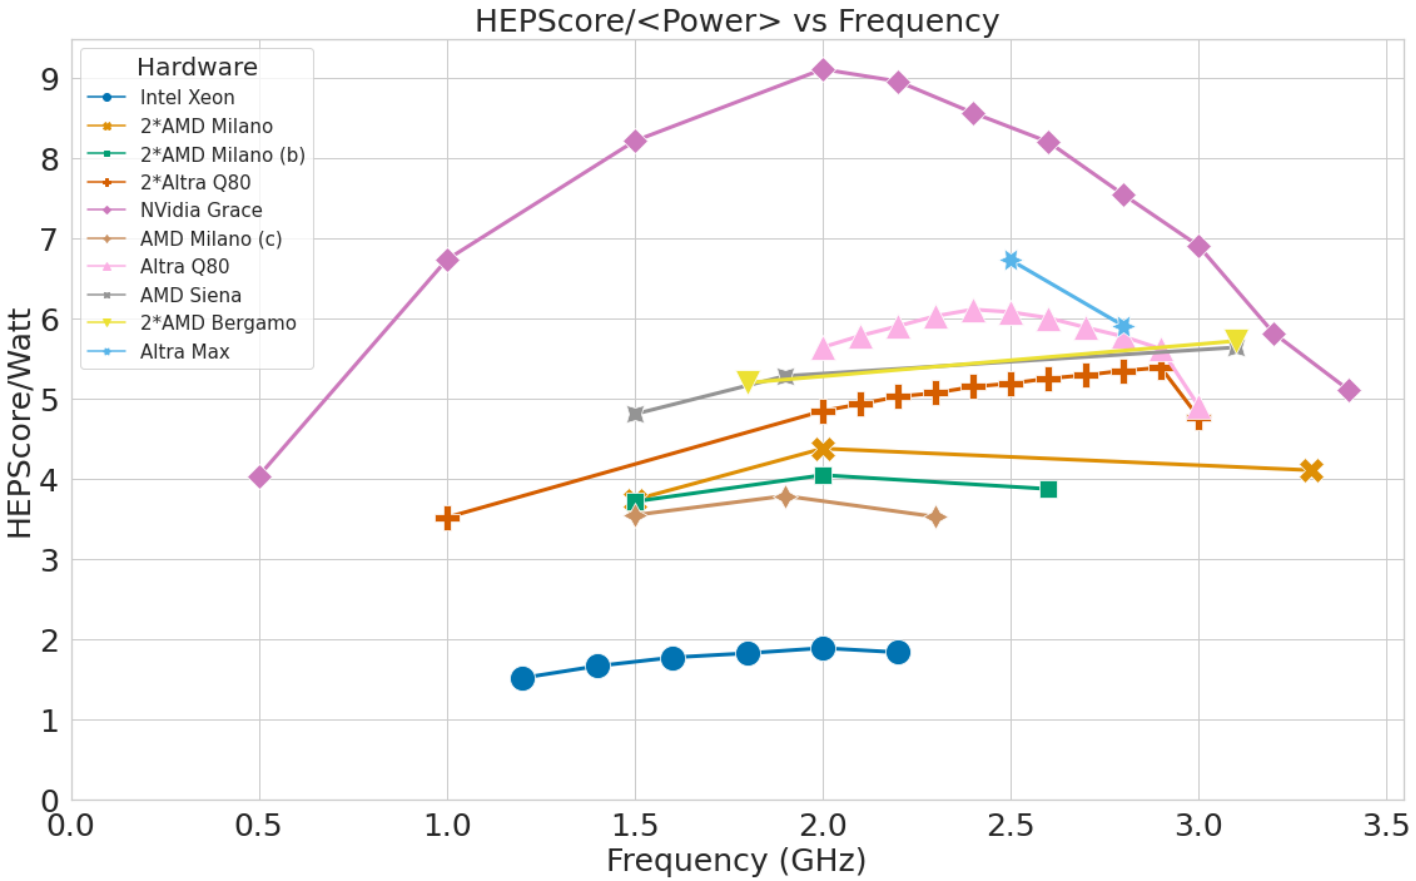

- Another use case of the Enhanced HEP Benchmark Suite is the power measurement together with frequency and HEPScore
- Measurements done by the ScotGrid Glasgow group on different servers show that almost all servers have a frequency that maximises the HEPScore/Watt
- This study also allows to measure the machine sustainability of different architectures and hardware

HEPScore Documentation

- How to run HEPScore
- Table with HS23 scores declared by sites

HS23 table: *http://w3.hepix.org/benchmarking/scores_HS23.html*

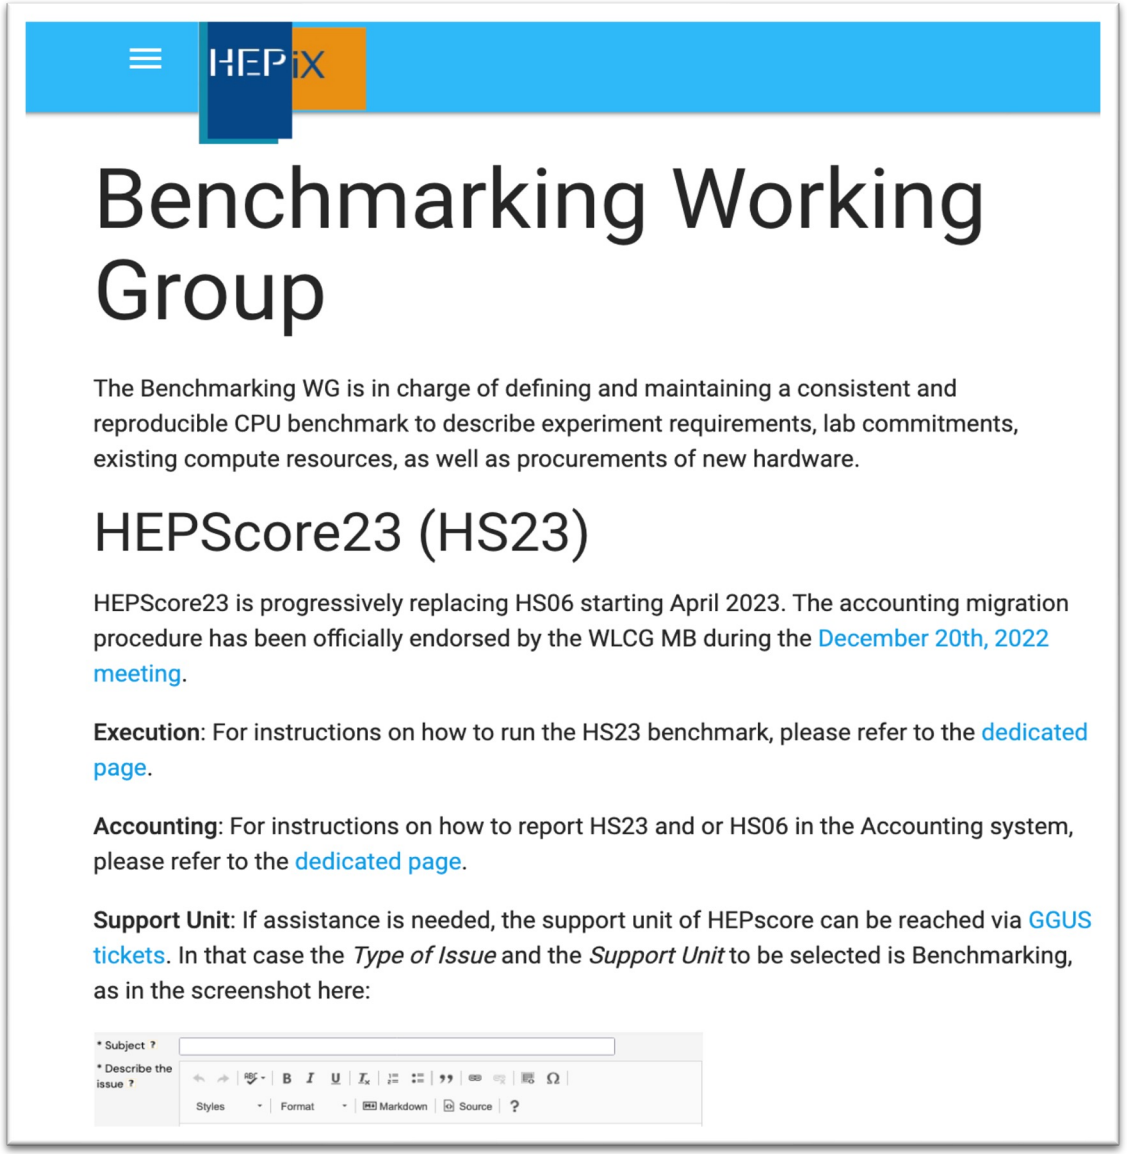

HEPiX

Benchmarking Working Group

The Benchmarking WG is in charge of defining and maintaining a consistent and reproducible CPU benchmark to describe experiment requirements, lab commitments, existing compute resources, as well as procurements of new hardware.

HEPScore23 (HS23)

HEPScore23 is progressively replacing HS06 starting April 2023. The accounting migration procedure has been officially endorsed by the WLCG MB during the December 20th, 2022 meeting.

Execution: For instructions on how to run the HS23 benchmark, please refer to the dedicated page.

Accounting: For instructions on how to report HS23 and or HS06 in the Accounting system, please refer to the dedicated page.

Support Unit: If assistance is needed, the support unit of HEPscore can be reached via GGUS tickets. In that case the *Type of Issue* and the *Support Unit* to be selected is Benchmarking, as in the screenshot here:

Source: *http://w3.hepix.org/benchmarking/*

Summary

Infrastructure

- HEP Benchmark Suite is a powerful tool, ready to use by other WMS thanks to its flexibility
- Enhancements were included to measure metrics such as machine load, frequency, memory usage, power consumption etc.
- We can use enhanced toolkit to perform power consumption studies

Analysis

- Model of HS23 vs Load allow us to fix misconfiguration issues and to understand how different configuration aspects affect the performance
- Comparison of results gathered from job slots with bare metal nodes disclose useful insights
- We need to perform additional tests to understand the cause of increasing performance in low load region

Q&A

Appendix

ATLAS

data from: 07/04/23 – 08/01/24

- Automated job submission every 3 hours on each panda resource
  - 139 Panda Resources
  - 227 CPU Models
  - 28246 unique hosts
- Over 200k jobs finished
- Each job: 8 core slot
- Median of job's walltime: 81minutes
  - HEPScore23 configuration with 1 repetition
  - 0.06% of total walltime_x_core

LHCb

data from: 01/08/23 – 01/11/23

- Manual job submission
  - 48 Sites
  - 110 CPU Models
  - 1650 unique hosts
- 2.1k jobs finished
- Each job: 1 or 4 core slot (most 1core)
- Median of job's walltime: 43minutes
  - lhcb-sim-run3-ma-bmk with 3 repetitions