

# HEP Benchmark Suite

## The centralized future of WLCG benchmarking

G. Menéndez Borge<sup>1</sup> – On behalf of the HEPiX Benchmarking Working Group

<sup>1</sup> CERN IT Department



### The Suite

The HEP Benchmark Suite [1] is a toolkit developed for benchmark orchestration. It allows benchmarks to be managed from a single application, thus characterizing the performance of individual and clustered heterogeneous hardware.

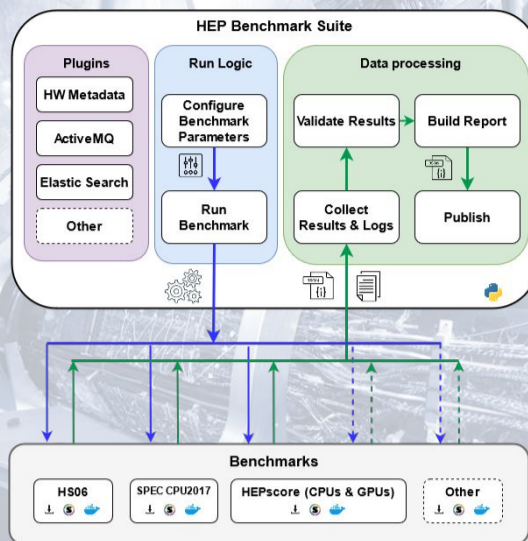
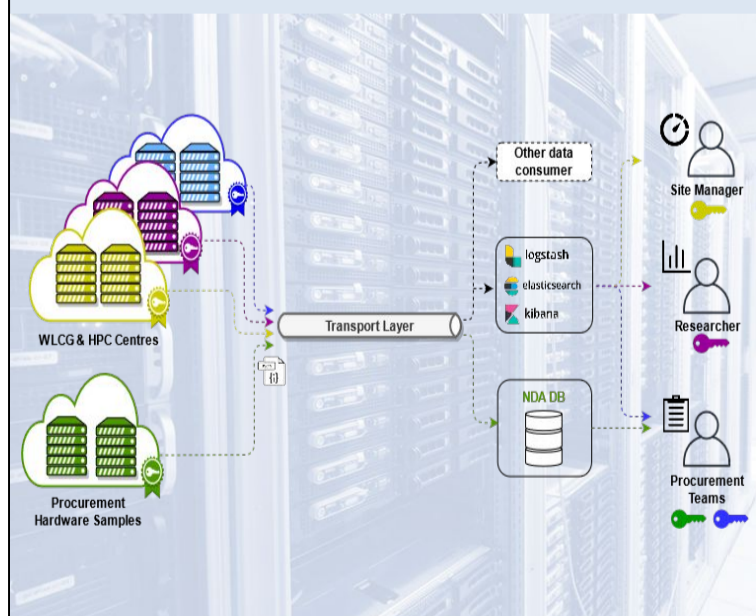
Given the *modular fashion* in which the Suite has been built, it allows the user to configure and run any number of benchmarks. While the list could be expanded indefinitely, a few remarkable examples include:

- HS06 [2]
- SPEC2017 [3]
- HS23 [4]
- DB12 [5]

Following the containerized nature of the benchmarks, the *HEP Benchmark Suite* has been developed with modularity in mind for this area as well.

**Docker**, as the most popular container platform in the market, has been fully integrated into the *Suite*, maximizing its compatibility.

Further, **Apptainer**, previously known as *Singularity*, the container system specialized in secure high performance computing is also fully compatible, and more may follow in the future.

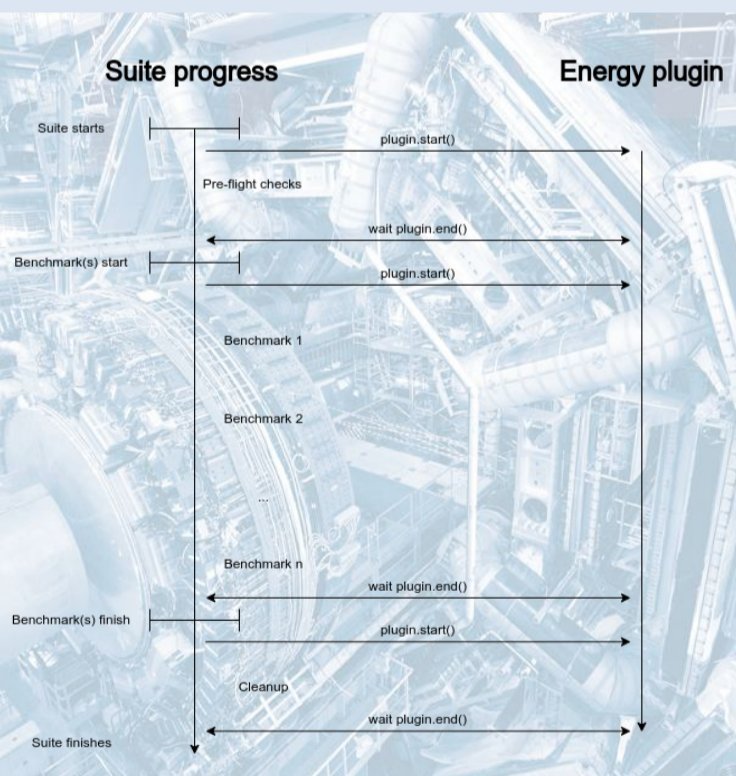


### Add-ons

The *modularity* of the suite does not only concern the benchmarks themselves, but also its add-ons. That way the suite can be configured with a number of *plugins* to report additional data aside from the benchmark results.

The **Extractor Plugin** retrieves hardware and software metadata which gives a deeper insight of the benchmark's running conditions. This ranges from the hostname or IP, to the CPU cache, the swap memory or the OS version.

An **Energy Plugin**, devised to report the consumption of the benchmarks, is being actively developed.

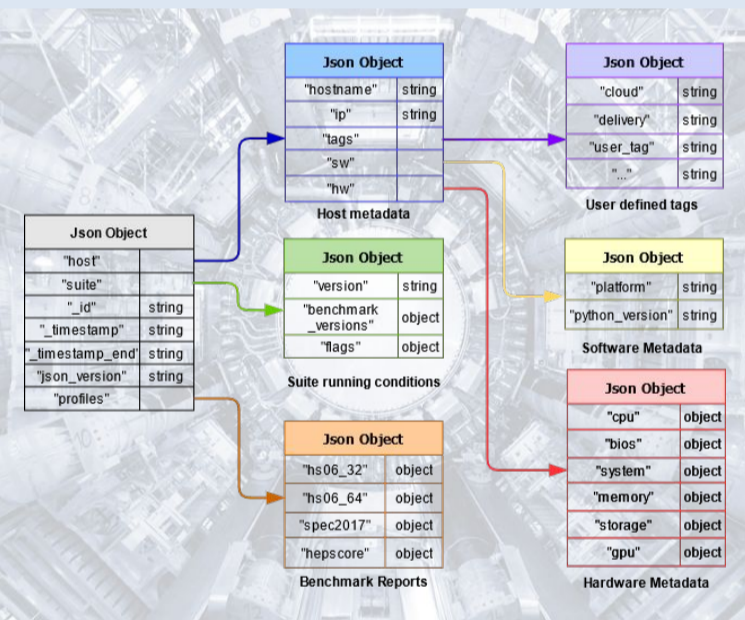


### A Central Database

The reports produced by the Suite can be published to an *Apache ActiveMQ* broker for further processing. By using CERN's dedicated AMQ benchmarking cluster, your data will be stored in our **central benchmarking Opensearch database** [6].

This will allow you to query your results alongside *over 80.000 existing entries* for reference and comparison. Thanks to the **Kibana** instance deployed on top of the DB, users may also create visualizations and dashboards to plot their results.

Furthermore, Opensearch/Elasticsearch being well-established databases, many tools include integrations to further extend the capabilities of this DB. As an example, **Grafana** can easily be integrated with such a data source, enabling further analysis, plotting, and monitoring.



### Support and License

The *HEP Benchmark Suite* is distributed under the terms of the *GNU General Public Licence v3* and is therefore **free software**. This means the community can both examine the code they execute and contribute to improve it if so desired.

Furthermore, the *Benchmarking Working Group* at CERN and the *WLCG* back this project, guaranteeing its long-term support.

References

