

HEPScore status update

D. Giordano, Randy Sobie

HEPiX Benchmarking WG

07 June 2023

About the TF mandate

Objectives accomplished

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

At the WLCG MB on May 16 agreed to close the Task Force activity

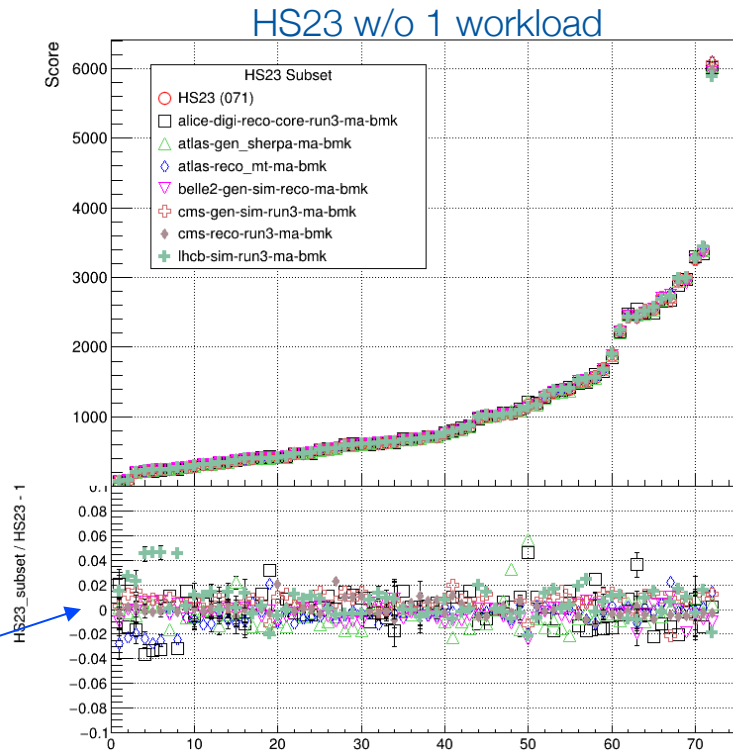
Remaining activities are under the role of other bodies

- Support & development: role of the **HEPiX Benchmarking WG**
- Monitor the adoption of HEPScore: more an **Accounting WG** task


HEPScore23 validation campaign

Since March 1st, measurements from a variety of servers and sites

- ~15 sites
- 46 distinct CPU models (Intel, AMD, ARM)
 - Including recent ARM nodes from vendors
- Small spread, ranging from ‰ to few ‰
- Several studies started
 - Performance Vs Energy
 - Performance in grid job slots
 - Effect of removing from HS23 one of the 7 benchmarks

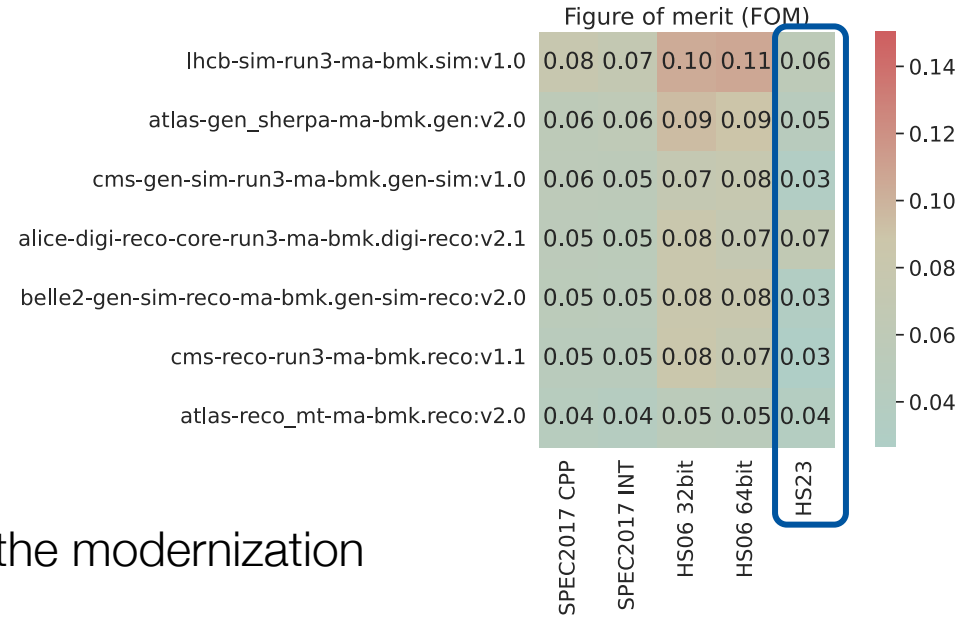


CPU models by year



HS23 is a more accurate representation of the modernization that has taken place in HEP applications compared to HS06

- Figure of merit (FOM) is a high-level measure of that



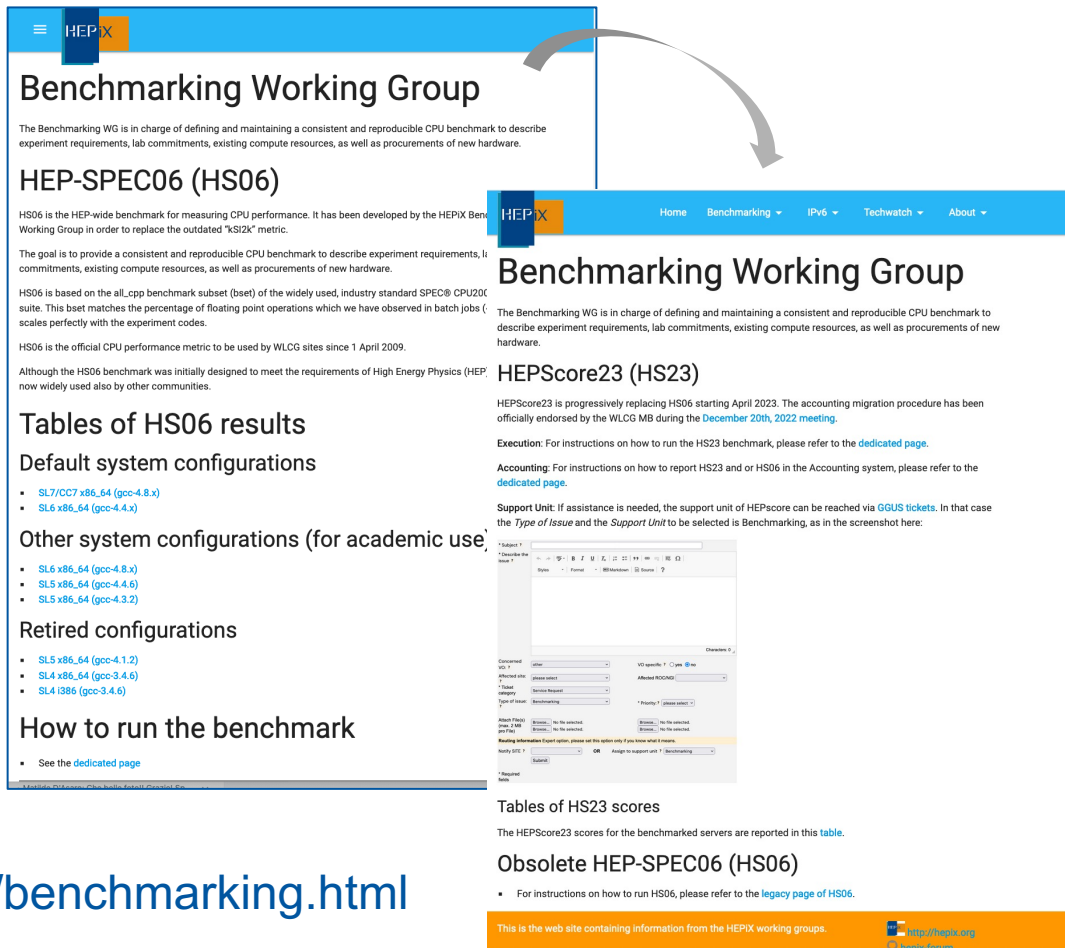
Documentation

Update the official HEPiX benchmarking page

- Report new instructions already reviewed by the TF
- Collect HS23 results as done in the past for HS06
 - Extracted from the benchmarking DB

Harmonize the gitlab documentation of the Suite, HEPscore, HEP Workloads

<https://w3.hepix.org/benchmarking.html>



The screenshot shows the HEPiX Benchmarking Working Group website. The main content area includes sections for HEP-SPEC06 (HS06) and HEPscore23 (HS23). The HEP-SPEC06 section describes the benchmark's purpose and provides links to tables of results and default system configurations. The HEPscore23 section describes the benchmark's replacement of HS06 and provides links to execution instructions and support unit information. A sidebar on the right contains a navigation menu and a table of HS23 scores. A footer at the bottom of the page contains the text "This is the web site containing information from the HEPiX working groups." and the HEPiX logo.

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.

HEP-SPEC06 (HS06)

HS06 is the HEP-wide benchmark for measuring CPU performance. It has been developed by the HEPiX Benchmarking Working Group in order to replace the outdated "SIZK" metric.

The goal is to provide a consistent and reproducible CPU benchmark to describe experiment requirements, lab commitments, existing compute resources, as well as procurements of new hardware.

HS06 is based on the `all_cpp` benchmark subset (`bset`) of the widely used, industry standard SPEC@ CPU2000 suite. This `bset` matches the percentage of floating point operations which we have observed in batch jobs (scales perfectly with the experiment codes).

HS06 is the official CPU performance metric to be used by WLCG sites since 1 April 2009.

Although the HS06 benchmark was initially designed to meet the requirements of High Energy Physics (HEP), now widely used also by other communities.

Tables of HS06 results

Default system configurations

- [SL7/CC7 x86_64 \(gcc-4.8.x\)](#)
- [SL6 x86_64 \(gcc-4.4.x\)](#)

Other system configurations (for academic use)

- [SL6 x86_64 \(gcc-4.8.x\)](#)
- [SL5 x86_64 \(gcc-4.4.6\)](#)
- [SL5 x86_64 \(gcc-4.3.2\)](#)

Retired configurations

- [SL5 x86_64 \(gcc-4.1.2\)](#)
- [SL4 x86_64 \(gcc-3.4.6\)](#)
- [SL4 i386 \(gcc-3.4.6\)](#)

How to run the benchmark

- See the [dedicated page](#)

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.

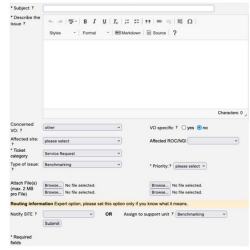
HEPscore23 (HS23)

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:



Tables of HS23 scores

The HEPscore23 scores for the benchmarked servers are reported in this [table](#).

Obsolete HEP-SPEC06 (HS06)

- For instructions on how to run HS06, please refer to the [legacy page of HS06](#).

This is the web site containing information from the HEPiX working groups. <http://hepix.org>

New documentation

Menu to access old HS06 pages, accounting instructions, and how to run

HEP-X Home Benchmarking IPv6 Techwatch About

HEP-SPEC06 (5)

THIS PAGE IS PROVIDED FOR LEGACY REASONS. THE PREVIOUSLY OFFICIAL BENCHMARK, HS06, REPLACED BY HEPScore23 SINCE 2022. HEPScore23 is the new official benchmark.

- WG Activities
- HS06 legacy page
- Accounting
- How to run HS23
- HS23 score table

HS06 is the HEP-wide benchmark for measuring CPU performance. It has been developed by the HEP-X Benchmarking Working Group in order to replace the outdated "kSI2k" metric.

The goal is to provide a consistent and reproducible CPU benchmark to describe experiment requirements, lab commitments, existing compute resources, as well as procurements of new hardware.

HS06 is based on the all_cpp benchmark subset (bset) of the widely used, industry standard SPEC® CPU2006 benchmark suite. This bset matches the percentage of floating point operations which we have observed in batch jobs (~10%), and it scales perfectly with the experiment codes.

HS06 is the official CPU performance metric to be used by WLCG sites since 1 April 2009.

Although the HS06 benchmark was initially designed to meet the requirements of High Energy Physics (HEP) labs, it is by now widely used also by other communities.

Tables of HS06 results

Default system configurations

- SL7/CC7 x86_64 (gcc-4.8.x)
- SL6 x86_64 (gcc-4.4.x)

Other system configurations (for academic use)

- SL6 x86_64 (gcc-4.8.x)
- SL5 x86_64 (gcc-4.4.6)
- SL5 x86_64 (gcc-3.4.2)

Retired configurations

- SL5 x86_64 (gcc-4.1.2)
- SL4 x86_64 (gcc-3.4.6)
- SL4 i386 (gcc-3.4.6)

How to run the benchmark

- See the [dedicated page](#)

HEP-X Home Benchmarking IPv6 Techwatch About

How to Run HEPscore23 benchmark

- WG Activities
- HS06 legacy page
- Accounting
- How to run HS23
- HS23 score table

- Requirements
- Run the HEP Benchmark Suite
 - Script mandatory parameters
 - Publish results (Optional)
 - DN extraction
- Run HEPscore23 standalone
- Troubleshooting
 - ulimit configuration on CENTOS7 (reason and procedure)
 - DVMFS (as image repository) configuration

This document provides instructions on how to execute the HEPscore23 benchmark.

Requirements

It is crucial that the server is fully dedicated to the benchmarking activity during the run, to ensure accurate measurements and prevent potential errors.

The server must have a minimum hardware configuration (see requirements below) and include the following packages:

- Container engine Apptainer (version 1.1.0 or higher);
- Python version 3.9 or higher;
- python3-pip;
- git

The user will need pip and git to install HEPscore23 as a Python package.

Hardware requirements:

- A disk space proportional to the number of available cores on the server (about 1 GB per logical core) is necessary to temporarily store the results;
- The server must have at least 2GB of RAM per logical core;
- ulimit configuration (see details below)

Run the HEP Benchmark Suite

While it is possible to install HEPscore23 standalone (see later), it is recommended to use the HEP Benchmark Suite alongside HEPscore23 to include in the benchmark report metadata about the server's running conditions. The metadata includes details about the server's CPU, RAM, disks, IP addresses, and other relevant information. In future versions of the suite, there will be the capability to configure additional measurement plugins, expanding the functionality beyond its current state. These new plugins will include options such as an energy consumption plugin or a load and memory usage plugin.

HEP-X Home Benchmarking IPv6 Techwatch About

Accounting

- WG Activities
- HS06 legacy page
- Accounting
- How to run HS23
- HS23 score table

- Example1: Site with a different cluster per CPU
- Example2: A site with a single cluster

The migration strategy for the accounting side is detailed by the Accounting TF in this [document](#). This strategy involves implementing software changes on the site side as well as APEL, EGI portal, WAU sides. To streamline the process and minimize the number of changes, several strategic approaches have been discussed within the WLCG collaboration, in particular at the Lancaster Workshop. These approaches have been endorsed by the WLCG Management Board during the [December 20th, 2022 meeting](#).


To summarize, the transition from HS06 to HEPscore23 should be gradual and seamless. This will be achieved through the following measures: The HEPscore23 benchmark will use the same scale factor as HS06, which is fixed on a reference server.

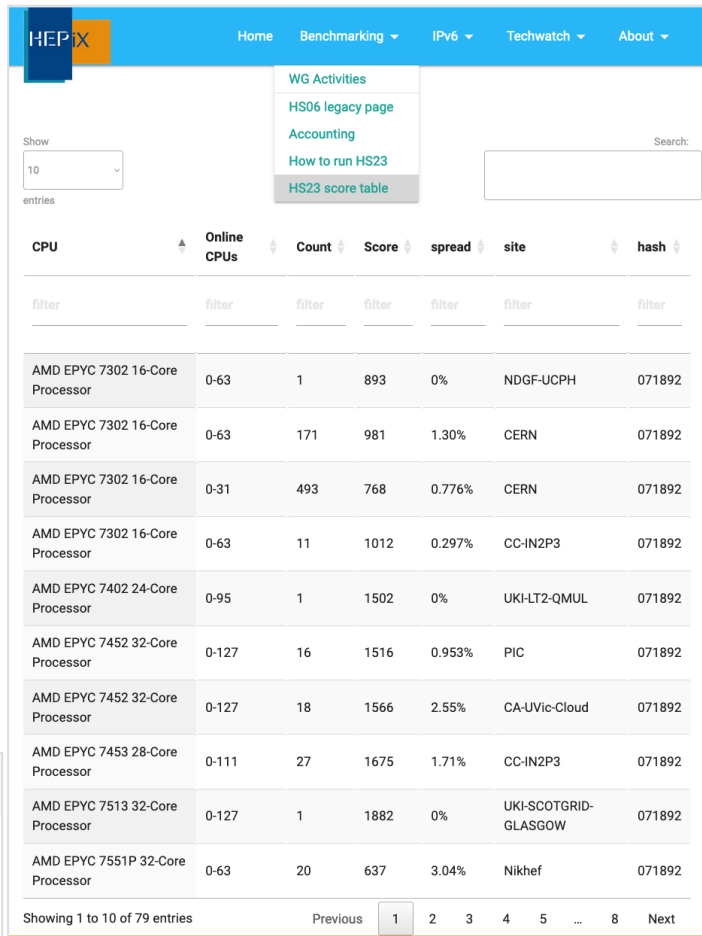
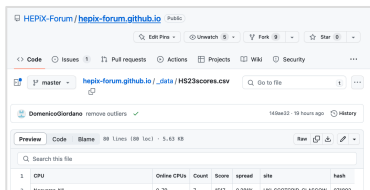
Sites are only expected to benchmark new resources with HEPscore23. Old servers do not need to be re-benchmarked for accounting purposes. This ensures that the installed capacity pledged by the sites will remain unchanged. Sites are free to re-benchmark their servers if they wish, but they are not required to submit this information to the accounting portal. However, they can still use the HEP Benchmark Suite to publish their results in the benchmark database, which is separate from the accounting infrastructure.

How do these procedures reflect what is done in a given WLCG site? Below we describe how to calculate the benchmarking factor depending on site configuration and how the report would look like in accordance with the new specification.

Example1: Site with a different cluster per CPU

HEPScore23 results table

- ❑ Open access as the whole website. No need of ACL as for the Benchmarking DB 
- ❑ Leverage DataTable to filter, order, paging
 - Not a static table as for HS06
 - Reporting CPU model, number of online CPUs, Number of measurements, score, spread, site and hash of the HEPscore config. Can be extended to other metadata (RAM, OS, etc)
- ❑ Data extracted from the Benchmarking DB
 - Semi-automatic procedure under preparation:
 - Perform analysis of the new data via Jupyter notebooks (investigate outliers, fit values, etc)
 - Inject new table rows, as csv, in the github repo hosting the Benchmarking website



CPU	Online CPUs	Count	Score	spread	site	hash
AMD EPYC 7302 16-Core Processor	0-63	1	893	0%	NDGF-UCPH	071892
AMD EPYC 7302 16-Core Processor	0-63	171	981	1.30%	CERN	071892
AMD EPYC 7302 16-Core Processor	0-31	493	768	0.776%	CERN	071892
AMD EPYC 7302 16-Core Processor	0-63	11	1012	0.297%	CC-IN2P3	071892
AMD EPYC 7402 24-Core Processor	0-95	1	1502	0%	UKI-LT2-QMUL	071892
AMD EPYC 7452 32-Core Processor	0-127	16	1516	0.953%	PIC	071892
AMD EPYC 7452 32-Core Processor	0-127	18	1566	2.55%	CA-Uvic-Cloud	071892
AMD EPYC 7453 28-Core Processor	0-111	27	1675	1.71%	CC-IN2P3	071892
AMD EPYC 7513 32-Core Processor	0-127	1	1882	0%	UKI-SCOTGRID-GLASGOW	071892
AMD EPYC 7551P 32-Core Processor	0-63	20	637	3.04%	Nikhef	071892


Publication

The plan is to prepare a paper that describes HS23 and the process to get here

- CSBS is the candidate journal: already published the previous paper
- Randy and I as editors
- Authors: all contributors from the Working group and the Task force


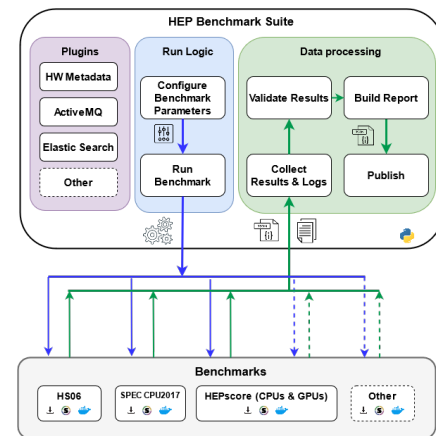
Ongoing work: Suite

- Extend the metadata plugin library to timeseries metrics: Energy consumption, load, CPU frequency,
- To be correlated with the score measurements



Benchmark report

```
"plugins": {  
  "CPUFrequencyPlugin": {  
    "pre": {"interval": 60...},  
    "hepscore": {"interval": 60...},  
    "post": {"interval": 60...}  
  },  
  "IpmiSdrEListPowerConsumptionPlugin": {  
    "pre": {"interval": 60...},  
    "hepscore": {  
      "interval": 60,  
      "tstart": "2023-05-30T14:53:32.982843Z",  
      "tend": "2023-05-30T16:58:33.708735Z",  
      "unit": "W",  
      "values": [  
        120,  
        120,  
        120,  
        ...  
      ]  
    }  
  }  
}
```



Plots from E. Simili et al.

Ongoing work: HEPScore

- Consolidate the python code
- Support multiple container registries
- Release a version v2.0

Ongoing work: HEP Workloads

- ❑ Allow to load a configurable number of cores
- ❑ Improve the validation of input params
- ❑ Study score stability Vs number of events/thread
- ❑ Progress on GPU workloads

THANK YOU!!!

HEP-SCORE deployment TF meeting #1 - launch

Wednesday 4 Nov 2020, 16:00 → 17:00 Europe/Zurich

Zoomland

Helge Meinhard (CERN)

Description Zoom link: <https://cern.zoom.us/j/93877918007>

16:00 → 16:05 **Welcome** ⌚ 5m

16:05 → 16:35 **Objectives and next steps** ⌚ 30m

[2020-11-04-HEPSco...](#) [2020-11-04-HEPSco...](#) [D Giordano's talk at ...](#) [D Giordano's talk at ...](#)

16:35 → 16:50 **Administrivia** ⌚ 15m

- Regular meeting slot
- Meeting frequency
- Note-taking
- Mailing list
- Web page
- Shared document space

16:50 → 17:00 **Any other business** ⌚ 10m

- Next meeting

HEPscore Workshop

19–20 Sept 2022
CERN
Europe/Zurich timezone

- Overview
- Timetable
- Contribution List
- My Conference
 - My Contributions
- Registration
- Videoconference

The goal of the workshop is to define the strategy for the adoption of HEPscore as replacement of HepSpec06.

The HEPspec06 benchmark has been a reliable estimate of CPU performance for many years, and is currently used by the WLCG for accounting and pledges. However, HEPspec06 is based on the SPEC2006 benchmark that is no longer supported. Further, it uses applications that do not reflect those used by the HEP community and will not provide benchmark for the new CPU+GPU systems.

The HEPscore Workshop will consist of several sessions:

Monday September 19

The **first session** will summarize the work of the WLCG HEPscore Task Force and the HEPiX Benchmark Working Group. During the session, potential candidates for the new HEPscore22 benchmark will be introduced.

The **second session** is devoted to the presentation of the current HEP-Workloads provided in the past year by the LHC experiments, Belle2, Juno and IWGN. Besides a description of the applications and of their performance, representatives of each experiment will highlight their expectations about the HEPscore22 composition and its lifetime.

The **last session** will introduce the framework developed to run HEP benchmarks and collect benchmark measurements ensuring traceability and monitoring. The HEP Benchmark Suite will be described and feedback about its usability will be provided by WLCG sites having used the suite.

Tuesday September 20

That has been a long, motivating, constructive journey!

