

Current HEP-SPEC and HEP-score Benchmarks at GridKa

Matthias J. Schnepf | 18. November 2021



Benchmarks with different Hardware Settings

- system
 - 2x Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz
 - results published on [CERN Kibana](#)
 - hostname: c01-010-173
 - Site: GridKa
 - power policy
 - performance
 - powersave
- Hyper-Threading
 - off
 - on
- active cores (siblings)
 - 16
 - 32

Benchmarks with different Hardware Settings: performance

■ HT off

Active Cores	Avg. HEP-score	Max HEP-score	Min HEP-score	Avg. HS06	Max HS06	Min HS06
16	310.335	310.582	310.203	293.88	294.067	293.783

■ HT on

Active Cores	Avg. HEP-score	Max HEP-score	Min HEP-score	Avg. HS06	Max HS06	Min HS06
16	307.795	308.758	306.831	288.634	288.791	288.478
32	376.914	377.119	376.709	360.697	360.864	360.529

■ without HT is 1% – 2% more performant than with HT and half of the cores

Benchmarks with different Hardware Settings: powersave

■ HT off

Active Cores	Avg. HEP-score	Max HEP-score	Min HEP-score	Avg. HS06	Max HS06	Min HS06
16	304.044	305.338	302.069	293.684	294.98	292.956

■ HT on

Active Cores	Avg. HEP-score	Max HEP-score	Min HEP-score	Avg. HS06	Max HS06	Min HS06
16	299.037	301.491	297.296	287.77	288.739	287.111
32	364.397	364.397	364.397	359.643	359.643	359.643

- negligible difference between powersave and performance in HS06
- powersave is 2% – 3% less performant than performance for HEP-Score

Benchmark Setup of new WNs

- benchmarks and BurnIn tests for about 200 new WN
 - 2x AMD EPYC 7742 64-Core
- reason for own data collection infrastructure
 - password authentication with service account file to publish automatically
 - results not reliable
 - different settings
 - insufficient power supply
 - could be also marked with tags for CERN DB
- collection infrastructure
 - Kibana and ElasticSearch already exist
 - activeMQ and logstash container used from hep-benchmark-suite test environment
 - adjust some configs

Results: 256 Copies

Avg. HEP-score	Max. HEP-score	Min. HEP-score	Avg. SPEC2017	Max. SPEC2017	Min. SPEC2017
3,262.721	3,441.979	3,116.183	430.826	458.332	421.981

Mode	Avg. HS06	Max. HS06	Min. HS06	No. of Benchmarks
64 Bit	3,320.977	3,358.532	3,286.904	11
32 Bit	2,920.971	2,979.005	2,876.71	684

- spread of $\pm 5\%$ in HEP-score and SPEC2017
- more than 10% in 64 Bit HS06 than 32 Bit HS06
- spread of $\pm 2\%$ in HS06

Results

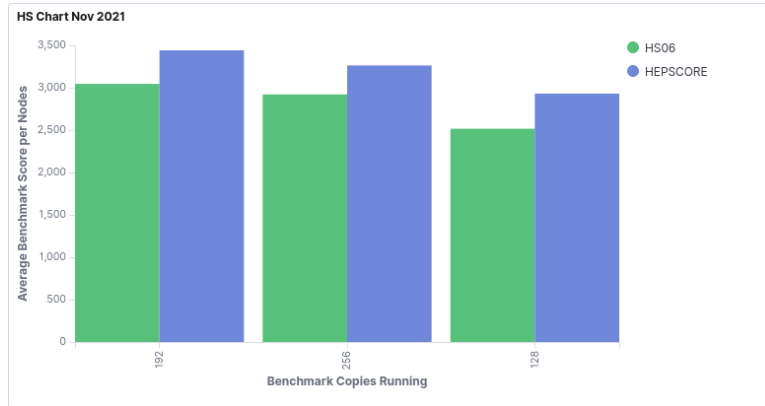
■ 128 copies (siblings)

Avg. HS06	Max. HS06	Min. HS06	Avg. HEP-score	Max. HEP-score	Min. HEP-score
2,516.593	2,534.347	2,498.434	2,930.634	2,943.25	2,925.529

■ 192 copies (siblings)

Avg. HS06	Max. HS06	Min. HS06	Avg. HEP-score	Max. HEP-score	Min. HEP-score
3,045.371	3,086.597	3,018.612	3,441.323	3,480.79	3,415.607

Result for Different Amount of Copies



- highest average HS06 and HEP-score score for 192 copies

Benchmark Results for WN with reduced power supply

- each box has 4 WN and two power supply
 - due to "sensitive" over current protection two boxes have only one power supply in use
- ⇒ 1.7kW instead of 2.8kW power consumption

Avg. HS06	Max. HS06	Min. HS06	Avg. HEP-score	Max. HEP-score	Min. HEP-score
2,544.204	2,554.435	2,531.007	2,956.181	3,046.549	2,872.651

- 87% of HS06 score of machines with two power supplies
- 90% of HEP-score score of machines with two power supplies

ATLAS-Gen Benchmark


- ATLAS-GEN benchmark consumed all available memory on several machines
- under investigation

Further Benchmarks

- repeat benchmarks with reduced amount of copies
- new machine at ETP
 - 2x AMD EPYC Milan 7713, 2GHz - 3,68GHz
 - OS: Ubuntu 20 or Ubuntu 21
 - jobs run in docker container

Backup

CPU Usage of Benchmarks



- DB12
- HS06
- SPEC2017
- HEP-score
- repeat