




# HEP Workloads as a Benchmark: Benchmark Validation (Status: pre-GDB 2019-10-08)

**Manfred Alef** (on behalf of the HEPiX Benchmarking Working Group)

STEINBUCH CENTRE FOR COMPUTING (SCC)



# Checkpoints

- Robustness
  - ➔ Error detection, error handling
- Spread of results (total scores) when repeating the benchmark
  - ➔ Spread:  $(\text{score\_max} - \text{score\_min}) / \text{score\_mean} \leq 5\%$ ?
  - ➔ Possible option: median of 3 or more consecutive runs
- Runtime
- Memory consumption:
  - ➔ Benchmark must run on default WLCG WN
    - 2 GB RAM (physical memory) per job slot

---

*Remark: HS06 – runtime (median of 3 iterations per benchmark): ~3h, spread:  $\leq 2\%$ , memory footprint (32bit): 1 GB per copy*

# Test Systems

## ■ KIT:

- 2+x test systems running workloads continual since 2018
  - Intel Xeon E5-2660v3 (Broadwell, 10-core)
    - ◆ 2 systems – SL6/SL7, turbo on/off
  - Temporary test systems (Sandy Bridge, Haswell)
    - ◆ Up to 6 systems – hard disks/SSD
- All systems with at least 2 GB RAM per logical processor (WLCG default WNs)
- Hyperthreading enabled
- Turbo boost enabled
  - Disabling turbo boost doesn't reduce spread of the results

■ Test systems also at BNL, IN2P3, INFN, University Freiburg, ...


# Acknowledgement

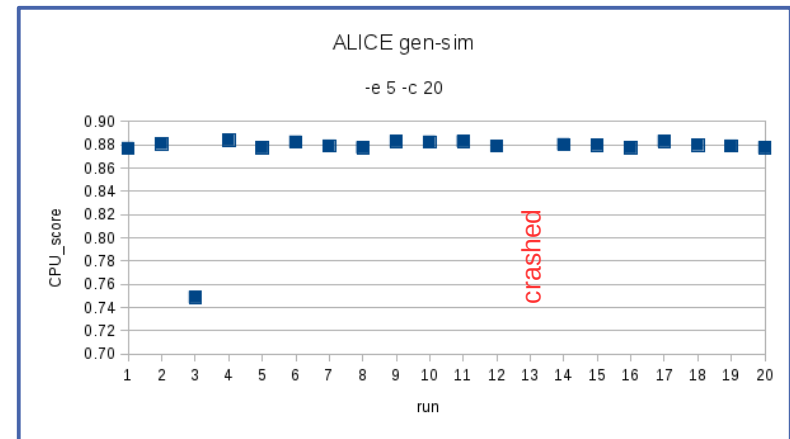
- Many thanks to
  - Chris Hollowell
  - Jean-Michel Barbet
  - Martina Javurkova
  - Michele Michelotto
  - Costin Grigoras
  - Olga Vladimirovna Datskova
  - Lorenzo Rinaldi
  - Martina Javurkova
  - David Lange
  - Andrea Sciabà
  - Andrea Valassi
  - Domenico Giordano

for improvements, debugging, bug fixes, and very valuable complementary studies!

# ALICE

## ■ Gen-sim-bmk:

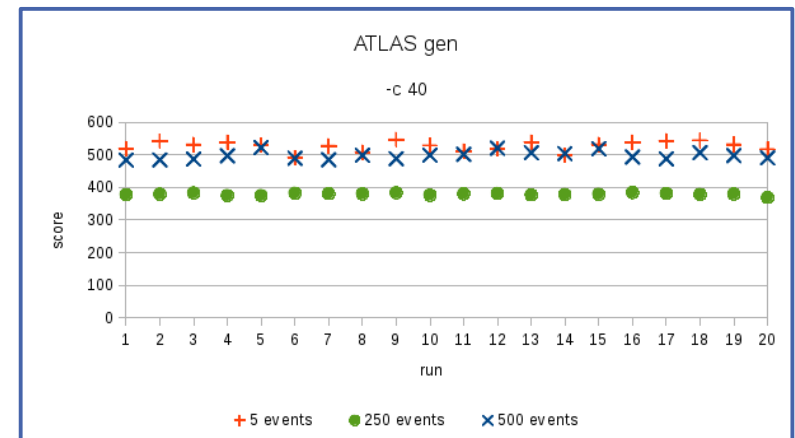
- Runtime: ~5'
- Robustness: frequent crashes (bug in Geant 3 + other issue) (JIRA: BMK-218), to be tackled by ALICE 
- Memory usage: ✓
- Individual outliers in the scores of repeated runs
  - Median of 3 or more runs: ✓



# ATLAS

## ■ Gen-bmk:


- ➔ Runtime: ~5...10'
- ➔ Robustness, memory usage: ✓
- ➔ Spread (5 ev.):  $\geq 5\%$ 
  - Even if Turbo mode off
  - Still under investigation
  - Optimum ~300 events?



Events	Spread	Spread(medians)
5	10.6%	6.6%
250	4.0%	1.6%
500	7.6%	4.0%

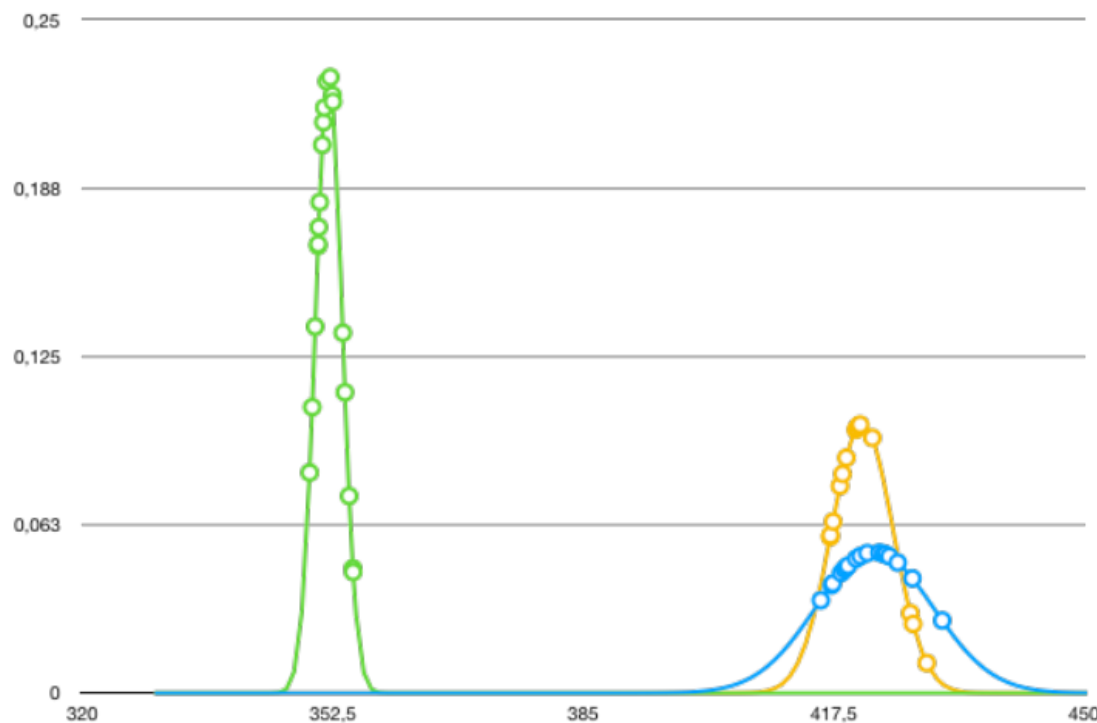
Spread Atlas Gen - score in light blue - median in orange

**Michele Michelotto**



# atlas-gen-bmk: reproducibility test

- **5 events**: spread=8% and stdev/mean=2% (~7 min, <cpu>: 64.375+/-9.710 ms)
- **500 events**: spread=2% and stdev/mean=0.5% (~15 min, <cpu>: 79.186+/-1.537 ms)
- **5 events with `—skipEvents="0"`**: spread = 3% and stdev/mean = 1%





# ATLAS

## ■ Sim-bmk (multi-threaded):

- Runtime: ~20...40'
- Spread, robustness, memory: ✓

## ■ Digi-reco-bmk:

- Runtime: ~30'
- Spread: ✓
- Robustness: ✓
- **Memory usage: up to ~3.4 GB RSS**

- ATLAS has understood what is happening, and they are working hard to make the container ready for testing very soon.

# CMS

- **Gen-sim-bmk (multi-threaded):**
  - ➔ Runtime: ~20'
  - ➔ Spread, robustness, memory: ✓

# Cms on C7 increasing # of events

	A	B	C	D	E	F	G	H	I	J	K
1	HEP-Benchmarking Group : ALICE Workload										
2	Date Report 26/04/19										
3	Jean-Michel BARBET, Subatech IN2P3, FR										
4	Image : cms-gen-sim										
5	Image Tag	V0.4		V0.4		V0.4		V0.4		V0.4	
6		Singularity		Singularity		Singularity		Singularity		Singularity	
7	Host	nanpcw61 (with tboost)		nanpcw61 (with tboost)		nanpcw61 (with tboost)		nanpcw61 (with tboost)		nanpcw61 (with tboost)	
8	Operating System	CentOS7		CentOS7		CentOS7		CentOS7		CentOS7	
9	CPU Model	E5-2620 v4 @ 2.10GHz		E5-2620 v4 @ 2.10GHz		E5-2620 v4 @ 2.10GHz		E5-2620 v4 @ 2.10GHz		E5-2620 v4 @ 2.10GHz	
10	Results file	10ev-24Avr2019.results		20ev-23Avr2019a.results		30ev-23Avr2019a.results		100ev-25Avr2019a.results		300ev-28Avr2019a.results	
11	20 Runs (4threads)	CPU_score	throughput_score	CPU_score	throughput_score	CPU_score	throughput_score	CPU_score	throughput_score	CPU_score	throughput_score
12		0,1692	0,6138	0,1682	0,6422	0,1657	0,6424	0,1702	0,6707	0,1725	0,6866
13		0,1671	0,6141	0,1678	0,6341	0,1651	0,6398	0,1664	0,6575	0,1738	0,6926
14		0,1668	0,6083	0,1694	0,6473	0,1644	0,6327	0,1689	0,6672	0,1731	0,689
15		0,1678	0,6064	0,1679	0,634	0,1638	0,6323	0,1696	0,6715	0,1705	0,6784
16		0,1687	0,6177	0,1656	0,6287	0,166	0,6403	0,1682	0,6647	0,1728	0,6879
17		0,1671	0,6035	0,1679	0,6274	0,1645	0,6304	0,1665	0,6575	0,1739	0,6928
18		0,1681	0,604	0,1668	0,636	0,1643	0,6313	0,165	0,6524	0,1743	0,6943
19		0,1667	0,5923	0,1686	0,6407	0,1649	0,6384	0,1674	0,662	0,1731	0,689
20		0,1677	0,5845	0,1663	0,6349	0,1652	0,643	0,1674	0,663	0,173	0,6885
21		0,1664	0,6127	0,1637	0,6217	0,1661	0,642	0,1661	0,6564	0,173	0,6888
22		0,1685	0,6033	0,1687	0,632	0,1669	0,6473	0,1661	0,6572	0,1711	0,6806
23		0,1692	0,6041	0,1685	0,6334	0,1619	0,6294	0,1678	0,6637	0,1724	0,6865
24		0,1678	0,6195	0,1685	0,6425	0,162	0,6296	0,1665	0,6558	0,1778	0,6895
25		0,168	0,6171	0,1689	0,6424	0,1662	0,6384	0,1686	0,6667	0,1728	0,6886
26		0,1692	0,6159	0,1685	0,6462	0,1655	0,6353	0,1656	0,6558	0,171	0,681
27		0,1668	0,6048	0,1663	0,6346	0,1646	0,6372	0,1645	0,6514	0,1693	0,6741
28		0,1666	0,6119	0,1651	0,6154	0,1617	0,6244	0,1686	0,6666	0,1686	0,6713
29		0,1621	0,5933	0,168	0,6392	0,1616	0,6234	0,167	0,6606	0,1708	0,6801
30		0,1633	0,6027	0,164	0,6211	0,164	0,6381	0,1668	0,6589	0,1707	0,6801
31				0,1657	0,6302	0,1643	0,6333	0,1638	0,6484	0,1705	0,679
32	AVERAGE(B12:B31)	0,16721578947	0,606836842105	0,16722	0,6342	0,164435	0,63545	0,16705	0,6604	0,17225	0,684935
33	STDEV(B12:B31)	0,00184338552	0,009378948681	0,00167759226	0,008479697115	0,00157154801	0,006331001417	0,00167095938	0,006362058588	0,00203818805	0,006405447928
34	MIN(B12:B31)	0,1621	0,5845	0,1637	0,6154	0,1616	0,6234	0,1638	0,6484	0,1686	0,6713
35	MAX(B12:B31)	0,1692	0,6195	0,1694	0,6473	0,1669	0,6473	0,1702	0,6715	0,1778	0,6943
36	SPREAD(MAX-MIN/AVG)	4,25%	5,77%	3,41%	5,03%	3,22%	3,76%	3,83%	3,50%	5,34%	3,36%

# CMS

## ■ Digi-bmk (multi-threaded):

- Runtime: ~15'
- Spread, robustness, memory: ✓

## ■ Reco-bmk (multi-threaded):









































- Runtime: ~15'
- Spread, robustness, memory: ✓

# LHCb

## ■ Gen-sim-bmk:

- Runtime: ~30...40'
- Spread, robustness, memory: ✓

# Summary (Status: pre-GDB 2019-10-08)

	ALICE	ATLAS			CMS			LHCb
	gen-sim	gen	sim	reco-digi	gen-sim	digi	reco	gen-sim
Robustness	 <sup>1</sup>							
Spread	 <sup>2</sup>	 <sup>3</sup>						
Runtime								
Memory				 <sup>4</sup>				
<b>Readiness</b>								

 okay
  fine tuning
  to do / still in progress

Remarks:

<sup>1</sup> Frequent crashes

<sup>2</sup> Use median of 3 or more runs to reduce spread

<sup>3</sup> Fine tuning in progress (e.g. more events, `-skipEvents`)

<sup>4</sup> Work in progress, improved container version will appear soon

