

WLCG HEP-SCORE Deployment Task Force

Helge Meinhard / CERN-IT
WLCG Management Board
19 January 2021

Future HEP Benchmarking (1)

- Need to replace HEP-SPEC06 (EOL of SPEC CPU 2006)
- WG of sites formed under HEPiX umbrella
 - Worked for ~ three years with participation by experiments
 - Considered SPEC CPU 2017 – rejected
 - Developed flexible scheme for benchmarking
 - Listed policy decisions to be taken by WLCG
 - Recent reports:
 - Domenico Giordano's report to MB 26-May-2020:
<https://indico.cern.ch/event/917098/contributions/3855129/attachments/2045174/3426154/WLCG-MB-26-05-2020-giordano.pdf>
 - Domenico Giordano's report to HEPiX 13-Oct-2020:
<https://indico.cern.ch/event/898285/contributions/4034096/attachments/2121862/3571531/HEPiX-Workshop-13-10-2020-giordano.pdf>



Future HEP Benchmarking (2)

- Result of the HEPiX benchmarking working group:
 - “HEP Benchmark Suite”: Framework itself (automatising benchmark runs, ensuring structured delivery and storage of results)
 - “HEP Workloads”: Collection of (mostly) HEP workloads for which it is desirable to obtain performance information
 - Rather dynamic – add improved workloads, new compilers/flags/OS, ...
 - Sometimes called “the matrix”
 - “HEP Score”: Single number based on a defined, stable combination of defined, stable reference workloads; prototype implementation (“HEP-SCORE20”) using various LHC experiment workloads (from Run 2)
 - Good consistency with HEP-SPEC06



HEP-SCORE Deployment Task Force

- WLCG Management board discussed in May and July; decided to launch a task force
- Following summer and an intense autumn with workshop etc. activities, task force started in November; bi-weekly meetings since then
- Task force membership:
 - Experts on pledge etc. process / procurements
 - Experiment experts
 - Four LHC experiments
 - Belle 2, DUNE, LIGO/Advanced VIRGO(/KAGRA), JUNO/BES III etc.
 - Site experts
 - Some MB members



Task Force Members (2)

Name	Function	Name	Function
Tommaso Boccali	CMS	Jeff Templon	Nikhef
Simone Campana	WLCG	Andrea Valassi	LHCb
Domenico Giordano	Benchmarking WG	Ian Collier	STFC-RAL
Michel Jouvin	Tier-2s	Gonzalo Merino	PIC
Walter Lampl	ATLAS	Fazhi Qi	JUNO, BES etc.
Andrew McNab	DUNE	Oxana Smirnova	NDGF
Helge Meinhard	WLCG and Convener	Tony Wong	US Tier-1s
Bernd Panzer	CERN	Josh Willis	LIGO/Adv. VIRGO
Stefano Piano	ALICE	Manfred Alef	KIT and Benchmarking WG
Randy Sobie	Belle 2		



Task Force Topics

- Compute facilities at WLCG sites
 - Still very much x86 dominated (mostly Intel, some AMD)
- Compute facilities used (quasi-)opportunistically
 - Various processors (x86, POWER, ARM), e.g. at HPC sites
 - Various GPUs in various relations with CPUs, e.g. at HPC sites
 - May see some of this soon at WLCG sites, too – even as part of the pledges
 - More may be coming, e.g. FPGAs



Discussions and Agreements So Far

- Start with CPU benchmarking on x86-based systems; look at other CPUs and/or GPUs later
 - Aim: single benchmark with a stable definition for at least (a typical CPU server lifetime cycle | a complete LHC machine cycle)
- Framework by benchmarking WG is very attractive
 - Use it to record behaviour of (wide) selection of workloads (not limited to HEP) across machines
- No “final” candidate for new benchmark yet
 - Intense discussions whether HEP-SCORE20 is a suitable starting point, and whether HEP workloads should be used at all
 - Decision deferred to after we have got hold of actual experiment workloads (LHC: Run 3?) and have collected data of selection of workloads (see above)



Reactions and What Next?

- Reporting and collecting feedback
 - Presentation and discussion at GDB last week: Questions on TF, discussion mainly on impact on accounting
 - CPU accounting infrastructure must be prepared to support more than one unit
 - Presentation and discussion at WLCG MB
- Task force plans
 - Round of status reports on workloads (ATLAS and ALICE to start with this week)
 - Implement workloads, once stable, into the “HEP Workloads” framework for data collection
 - Once ready, run on a variety of CPU server configurations to obtain behaviour of all workloads
 - Define representative sample of workloads for the final benchmark (HEP-SCORE2x)
 - Including relative weights
 - Propose steps for transition from HEP-SPEC06 to HEP-SCORE2x



(Personal) Conclusions

- Defining a HEP-SPEC 06 successor is a complex, ambitious task
 - Will take many months rather than weeks
 - ... and then we will have to tackle non-x86 CPUs and GPUs
- Quite a number of different views on how exactly this should be done
 - Good representation in the task force
- Nonetheless, discussions in a constructive and collaborative spirit
 - Thanks to all task force members!



Questions? Comments?

