



# Benchmarking Working Group

---

## Introduction

Kickoff Meeting  
2016-06-03

Manfred Aef (KIT)  
manfred.aef@kit.edu



# Benchmarking Working Group

---

**Welcome to the HEPiX Benchmarking Working Group!**



# Benchmarking Working Group

---

## Mandate:

- **Fast benchmark**  
to estimate the performance of the provided job slot (in traditional batch farms) or VM instance (in cloud environments)
  - Job matching / masonry  
(e.g. “can a pilot run another payload with the resources left?”)
  - Accounting if HS06 score not available
  - ...
  
- **Next generation of long-running HEP-SPEC benchmark**  
for installed capacities, accounting, procurements aso. in WLCG  
(successor of HS06)



# Benchmarking Working Group

---

## Volunteers:

### → Experiment representatives:

- **Alice:** Costin Grigoraş
- **Atlas:** Alessandro Di Salvo, Franco Brasolin
- **CMS:** ?
- **LHCb:** Ben Couturier

### → Many site admins



# Benchmarking Working Group

---

## Sharing of information, meetings:

### → Mailing list:

~~hepixoncpubenchmark@cern.ch~~

hepixoncpubenchmark@hepixon.org

**NEW**

### → WIKI:

<https://twiki.cern.ch/twiki/bin/view/HEPIX/CpuBenchmark>



# Benchmarking Working Group

---

## Sharing of information, meetings:

- **Regular Vidyo meetings**
  - **Regular time slot:**  
**Fri 14:00 CEST (12:00 UTC)?**
  - **Next meeting:**  
**Fri June 17th?**
- **Face-to-Face meeting at HEPiX Fall 2016 (October 17-21, Berkeley)?**
- **Email to mailing list** for brief news or information updates (archive available at <http://listserv.in2p3.fr>)



# Benchmarking Working Group

---

**First tasks:**



# Benchmarking Working Group

---

## First tasks:

- **Fast benchmarks**





# Benchmarking Working Group

---

## First tasks:

### → **Fast benchmarks:**

- Short run time
- Results most probably less accurate than HS06 (e.g. because of varying system load!)
- Must scale well with representative applications and/or with HS06




# Benchmarking Working Group

## First tasks:

### → Fast benchmarks:

Examples (fast benchmark candidates in GridKa batch farm, runtime < 4 min)





# Benchmarking Working Group

---

## First tasks:

### → Fast benchmarks:

- Candidates:
  - ◆ Dirac fast benchmark (LHCb, Belle)
  - ◆ KitValidation/Geant4 (Atlas)
    - ✓ Single muon events
    - ✓ Full chain
  - ◆ ROOT stress test (Alice)
  - ◆ Other fast benchmarks (HTCondor, Boinc)
    - ✓ Whetstone
    - ✓ Dhrystone



# Benchmarking Working Group

---

## First tasks:

### → Fast benchmarks:

- **Run several candidates in various batch farms (single and multi-core jobs), or in cloud environments**



# Benchmarking Working Group

---

## First tasks:

### → Fast benchmarks:

- Suggested benchmark harness:
  - ◆ KitValidation (KV) to run fast benchmarks and collect results (see talk by Domenico Giordano)



# Benchmarking Working Group

---

## First tasks:

### → **Fast benchmarks:**

- How to run the fast benchmarks:
  - ◆ Site admins: please run fast benchmarks and HS06
  - ◆ Experiment representatives: please report the performance of your applications (events/s) as well as the fast benchmark scores
  - ◆ Repeat fast benchmarks from time to time to find out how the benchmark score depends on the system load



# Benchmarking Working Group

---

## First tasks:

### → Fast benchmarks:

- Scaling with HS06:
  - ◆ Sites are encouraged to publish HS06 score of physical hosts and of long running VMs via MJF store
    - ✓ \$MACHINEFEATURES/hs06:  
total HS06 score of WN
    - ✓ \$JOBFEATURES/hs06\_job:  
HS06 score available for job

---

MJF:  
<https://github.com/HEP-SF/documents/raw/master/HSF-TN/2016-02/HSF-TN-2016-02.pdf>



# Benchmarking Working Group

---

## First tasks:

- Long running benchmark





# Benchmarking Working Group

---

## First tasks:

### → Long running benchmark:

- Scaling issues with HS06 have been reported:
  - ◆ LHCb: performance on Haswell hosts around 40% better than on Sandy Bridge [1]
  - ◆ Atlas [2]:
    - ✓ Simulation: similar to LHCb
    - ✓ Reconstruction, merge, event generation: still scaling well with HS06
  - ◆ Alice, CMS: ?

---

[1] <https://indico.cern.ch/event/319751/contributions/739785/attachments/1153280/1656518/150909-MJFandBenchmarking-LHCb.pdf>

[2] [https://indico.cern.ch/event/433164/contributions/1930250/attachments/1220374/1783838/Results\\_of\\_HS06\\_Scaling\\_Studies\\_at\\_GridKa\\_2016-02-01.pdf](https://indico.cern.ch/event/433164/contributions/1930250/attachments/1220374/1783838/Results_of_HS06_Scaling_Studies_at_GridKa_2016-02-01.pdf)



# Benchmarking Working Group

---

## First tasks:

### → Long running benchmark:

- Possible causes:
  - ◆ It seems that the odd scalings are not influenced by the general purpose application performance of the hosts but by extensions of new generations
  - ◆ Enhanced hardware features of recent processor generations:
    - ✓ AVX2 (Advanced Vector Extension V2)
      - Haswell and newer generations
    - ✓ Hardware random number generator
      - some Ivy Bridge models (?), all Haswell and newer



# Benchmarking Working Group

---

## First tasks:

### → Long running benchmark:

- It's highly recommended to look into the scaling issues and to understand the causes before we start to develop the next generation benchmark (successor of HS06)
  - ◆ Influence of AVX2, HWRNG, ...?
  - ◆ Can we check scaling of Alice or CMS applications with HS06?



# Benchmarking Working Group

---



# Benchmarking Working Group

---

## Miscellaneous:

- SPEC has published the **Cloud IaaS 2016** benchmark. It is designed to stress provisioning as well as runtime aspects of a cloud using I/O and CPU intensive cloud computing workloads. This benchmark is available at [www.spec.org/cloud\\_iaas2016](http://www.spec.org/cloud_iaas2016) (500 USD non-profit pricing).