Karlsruhe Institute of Technology

Benchmarking Update
GDB 2017-12-13

Manfred Alef

STEINBUCH CENTRE FOR COMPUTING (SCC)

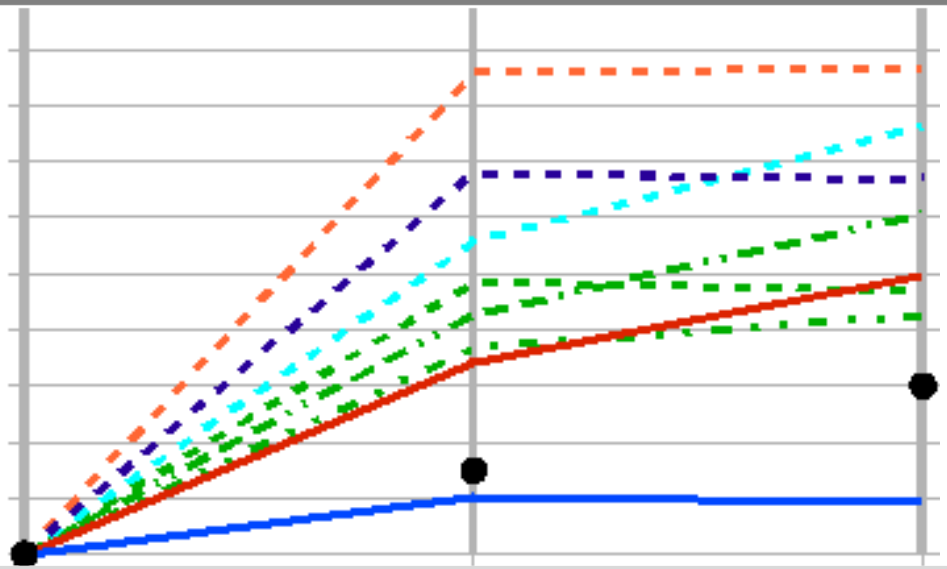

Analysis of Scaling Issues

- Investigating the impact on the number of job slots per core
  - ➔ Alice, Atlas, and LHCb done (see GDB Oct 11 2017) ...

Analysis of Scaling Issues

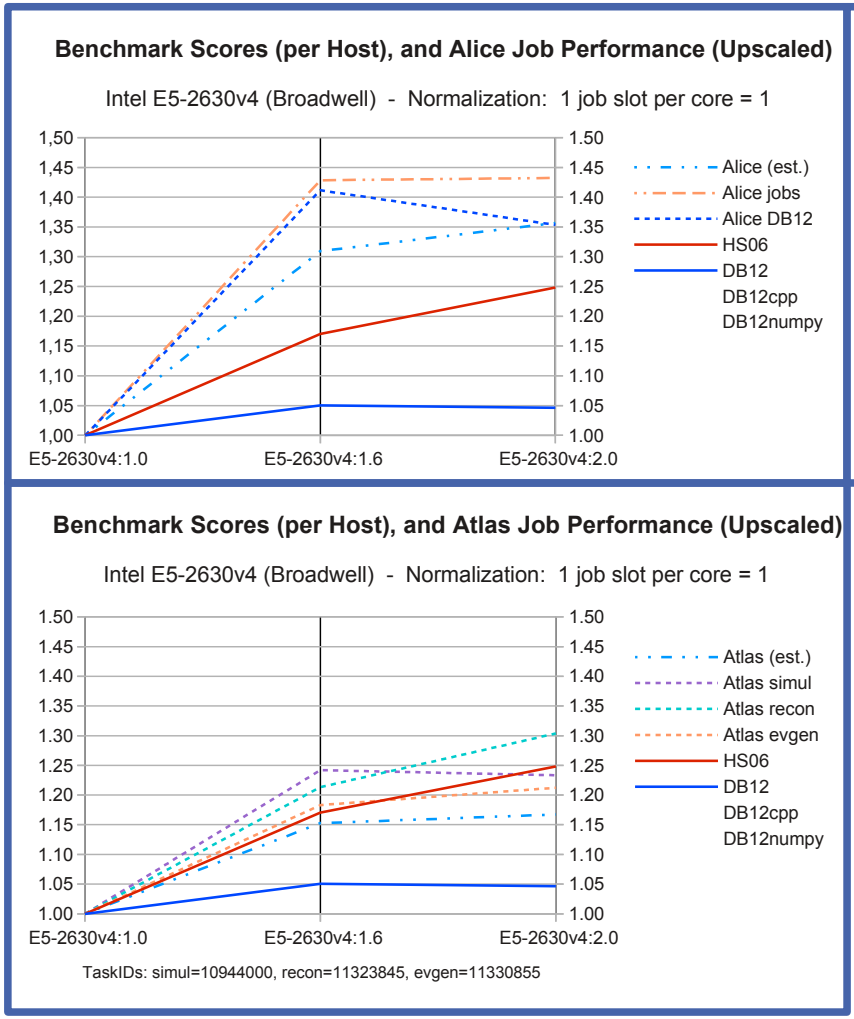

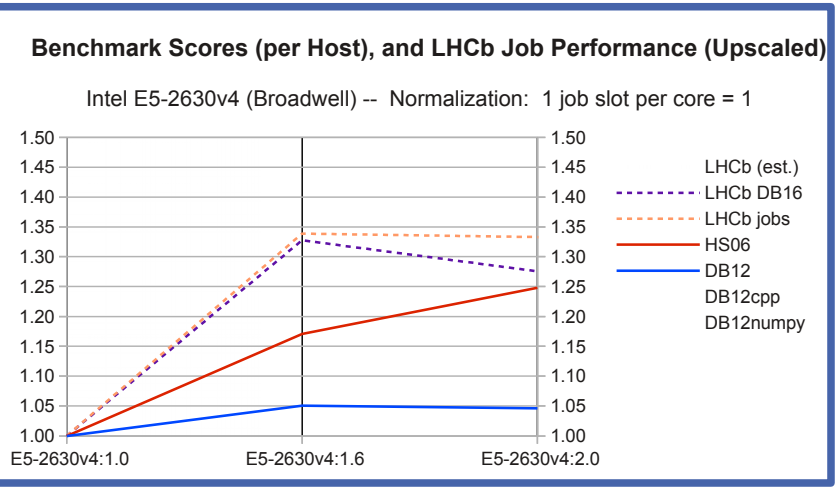

At least 1 rack of each WN configuration available at GridKa, running regular job mix (no dedicated benchmarking nodes!):

HW model	Job slots	Memory p. slot
E5-2630v4:1.0	20	3200 MB
E5-2630v4:1.6	32	4000 MB
E5-2630v4:2.0	40	≥2400 MB

Thanks to Costin Grigoras and Philippe Charpentier for providing the performance measurements and in-job benchmark scores of Alice and LHCb

Analysis of Scaling Issues

- Investigating the impact on the number of job slots per core
  - ➔ Alice, Atlas, and LHCb done (see GDB Oct 11 2017) ...
  - ➔ ... and first CMS results now available too
    (many thanks to Valentin Kuznetsov / CMS!)

* https://indico.cern.ch/event/595296/contributions/2558270/attachments/1440926/2225285/HEPiX_April_2017_benchmarking_giordano

Analysis of Scaling Issues

Benchmark Scores, and CMS Job Performance (Upscaled)

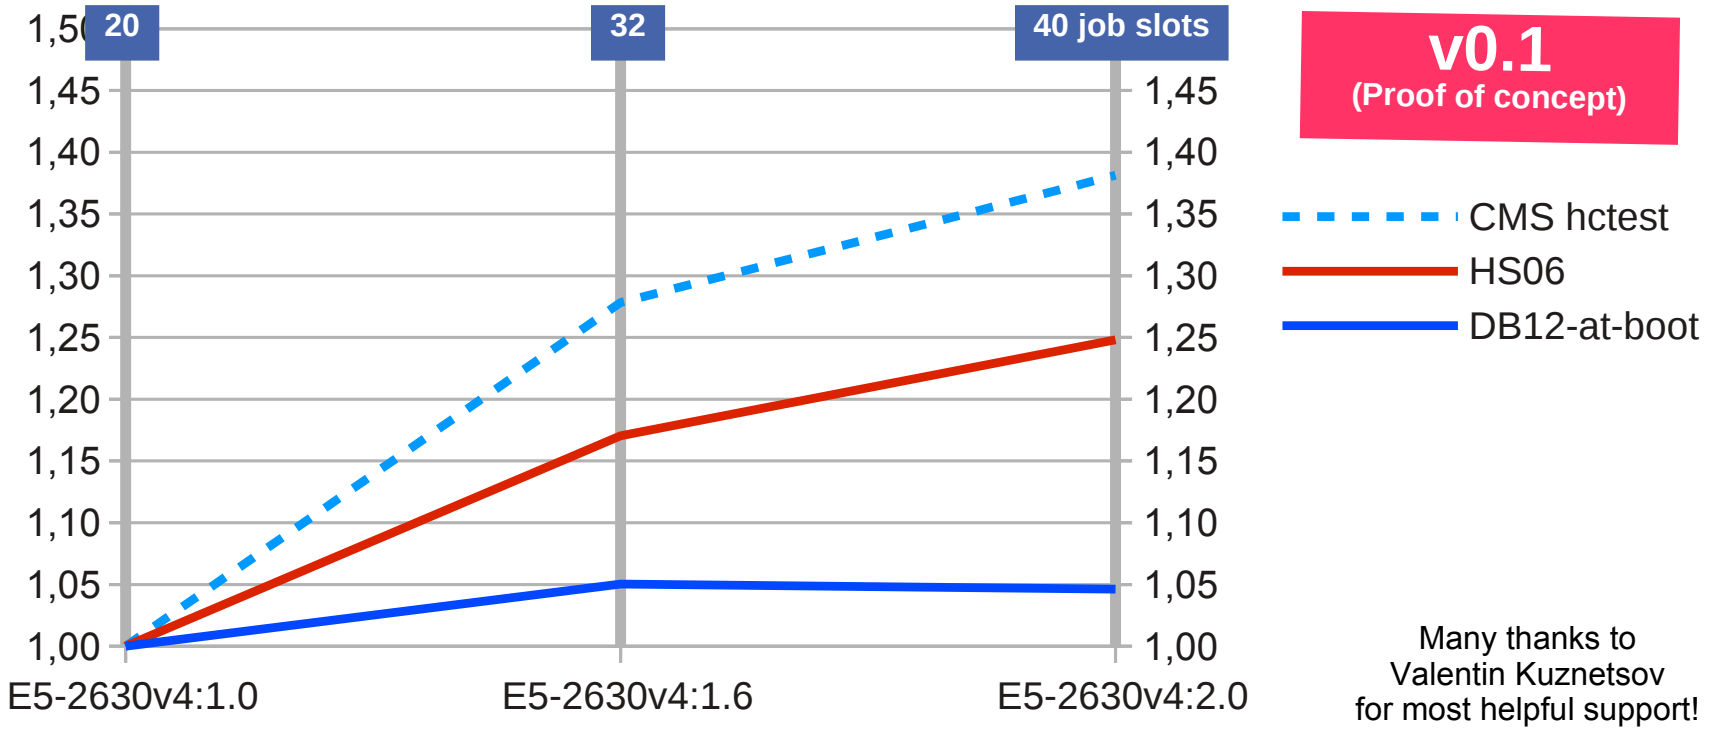

Many thanks to
Valentin Kuznetsov
for most helpful support!

Analysis of Scaling Issues

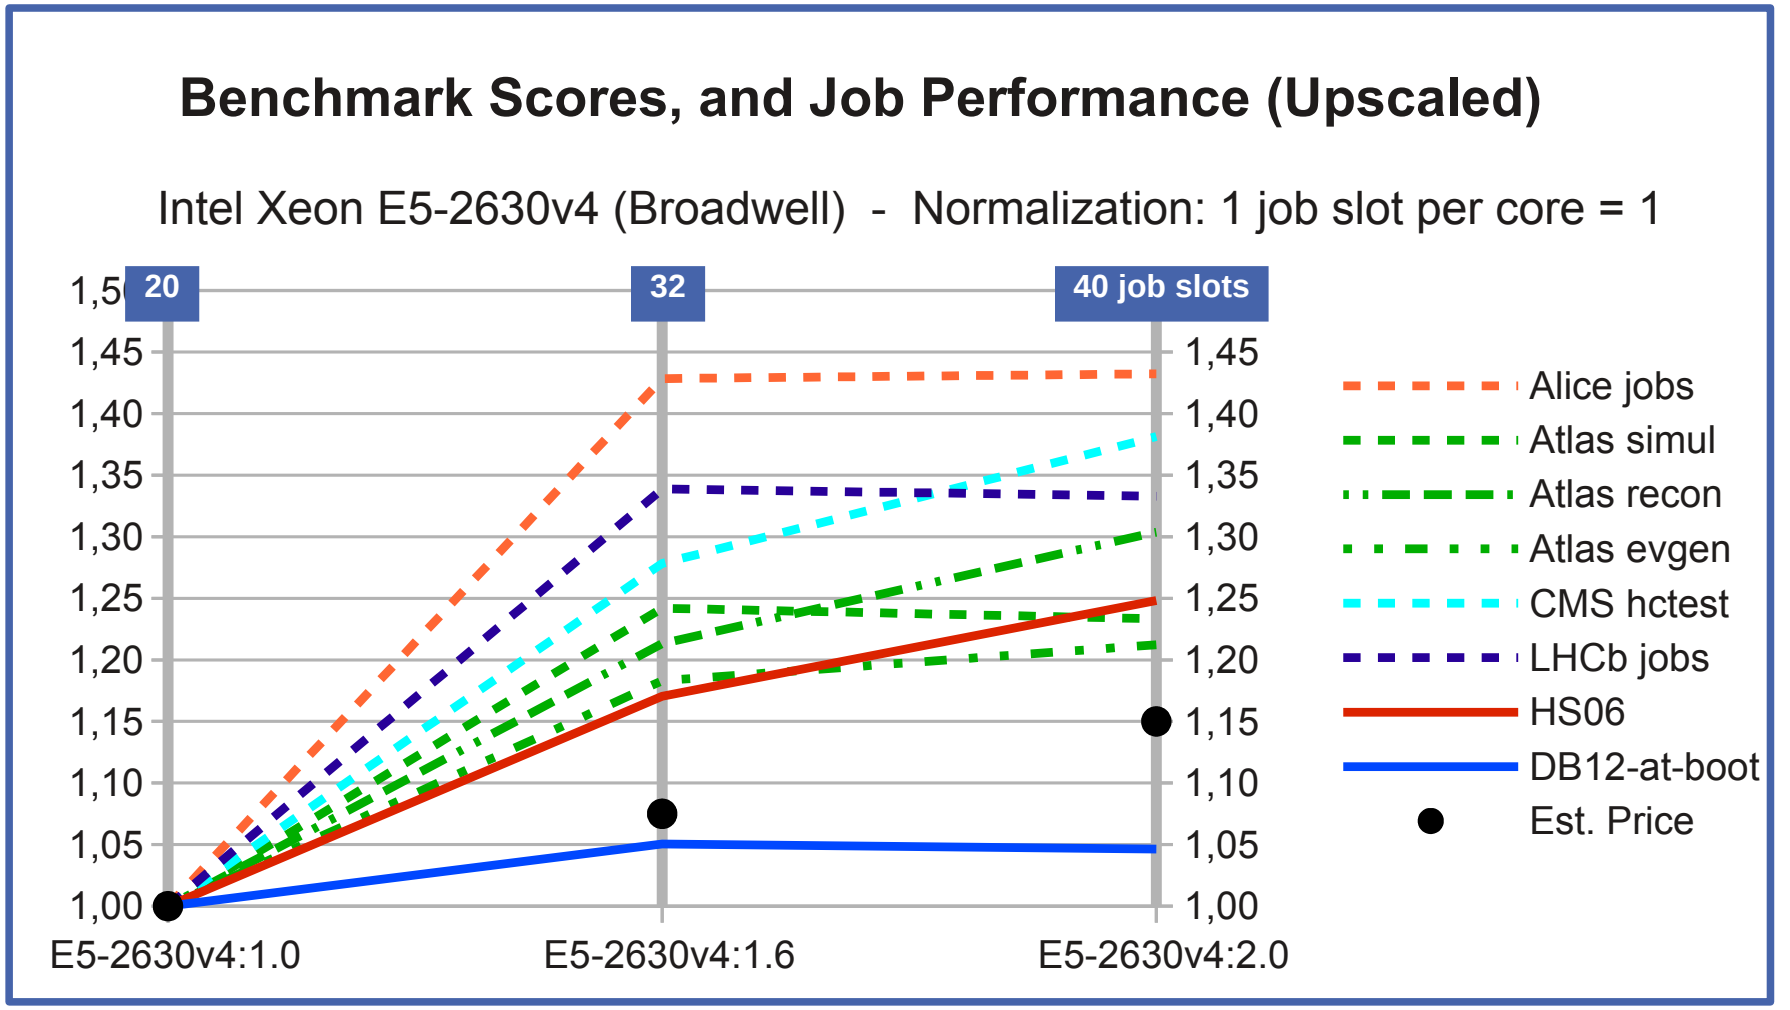

Analysis of Scaling Issues

- **Question to the working group:**
  - **Any chance to repeat this analysis in virtual environments?**

- **Questions to the sites:**
  - **Are there any statistics about WN configurations at WLCG sites?**
    - **Bare metal**
    - **VM**
  - **What's the most relevant configuration?**

Assessing SPEC CPU2017

- The HEPiX Benchmarking Working Group is assessing the SPEC CPU2017 as a possible HS06 replacement
  - ➔ Trying out compiler options
    - Remember HS06:
      - `-O2 -fPIC -pthread -m32`
    - SPEC CPU2017 default Linux config file (64bit default):
      - `-g -O3 -march=native \`
        `-fno-unsafe-math-optimizations \`
        `-fno-tree-loop-vectorize`
    - Testing:
      - `-O2 / -O3`
      - `-O2 -fPIC -pthread / -O3 -fPIC -pthread`

Assessing SPEC CPU2017

- **Questions to the experiments:**
  - **Which (common) compiler flags are in use?**
    - **Today**
    - **In the future (multithreading, ...?)**

Summary

- Analysis of scaling issues of HS06 versus HEP applications
  - ➔ Analysis of job accounting records:
    - Now all 4 WLCG experiments are involved
      - Significant differences depending on number of configured job slots per core
    - To double-check benchmark scores not only with reference workloads on dedicated benchmarking hosts but also with everyday jobs of all 4 WLCG experiments

Summary

- Analysis of scaling issues of HS06 versus HEP applications
  - ➔ Influence of the ratio of job slots per physical core:
    - Better performance (of Alice, Atlas, LHCb, and also CMS jobs) at ~1.5 or more job slots per core
      - Performance increase exceeds extra cost (RAM, disk)
      - In contrast to DB12-at-boot (i.e. running multiple DB12 copies on idle host) → not a suitable HS06 replacement
    - Statistics about WN configurations at WLCG sites?

Summary

- Assessing SPEC CPU2017 as a possible HS06 replacement:
  - ➔ Recommended compiler flags?

Announcement

- Pre-GDB about benchmarking in spring 2018

Addendum – How to Run the Analysis

- Example: CMS
  - I have got a CSV file (thanks to Valentin Kuznetsov/CMS) of jobs run at GridKa (i.e. WNHostName contains domain '.gridka.de')
    - Relevant rows:
      TaskMonitorId, WNHostName, Type, NEvProc, WrapWC
  - Grep the most relevant TaskMonitorId(s)
    - → 'hctest' jobs (around 2/3 of all included jobs)
    - Other tasks (analysis) with too few jobs each to plot lines
      - No outliers found, though
  - For each WN model calculate
    Average_Efficiency = sum(NEvProc)/sum(WrapWC)

Addendum – How to Run the Analysis

- Example: CMS
  - ➔ Results for Intel Xeon E5-2630v4 (Broadwell):

WN Model	Slots	Average_Efficiency	Upscaled	Normalized
E5-2630v4:1.0	20	3,96	79,2	**1,00**
E5-2630v4:1.6	32	3,12	99,8	**1,26**
E5-2630v4:2.0	40	2,72	108,8	**1,37**

    - Upscaled values: Average_Efficiency * Slots
    - Normalizing: dividing by the Upscaled value of E5-2630v4:1.0

  - ➔ Remark: HS06 scores of all WN hosts at GridKa are available from MJF (look into `$JOBFEATURES/hs06_job`), so experiments can consider to run their own analysis