# HS06 -m32 Vs –m64

D. Giordano (CERN)

HEPiX Benchmarking Working Group

Friday 30 Jun 2017

# Aim

- Compare HS06 scores when HS06 compiled at 32 bits (-m32) and 64 bits (-m64)
- Approach
  - Compare HS06 scores on different HW models and also different OS (SLC6 and CC7)
    - In order to access resources with different OS used VMs when possible
      - Standard WN configuration at CERN: bare-metal installed with CC7
      - Notable exceptions (old CPU models, such as Sandy Brige and Ivy Bridge)
  - Compare scale factors respect to other benchmarks (DB12, KV)
  - Modified HS06 output to report not only the overall score, but also the individual score per parallel process and benchmark

# Used Resources

Each server (physical or virtual) has been benchmarked fully loading the available cores

– CPU frequency typically auto-set to the nominal working value (no turbo-boost)

| Abbr. | Family | Model | nodes | OS |
|---|---|---|---|---|
| E5 0 | Sandy Bridge | E5-2690 0 @ 2.90GHz | 1 ph. node | SLC6 |
| E5 v2 | Ivy Bridge | E5-2650 v2 @ 2.60GHz | 2 ph. nodes | SLC6 |
| E5 v3 | Haswell | E5-2630 v3 @ 2.40GHz | VMs | SLC6 , CC7 |
| E5 v4 | Broadwell | E5-2630 v4 @ 2.20GHz | VMs | SLC6 , CC7 |

# VM configuration

VMs have been used for Haswell and Broadwell CPU models

- Physical Servers from CERN cell compute nodes
- SMT enabled
- Use NUMA aware node flavors
  - VMs confined to a single CPU socket
    - 2x $VM_{16}$ in a 32 logical cores ph. node
    - 2x $VM_{20}$ in a 40 logical cores ph. node
  - Single VM matching a ph. node
    - 1x $VM_{32}$ in a 32 logical cores ph. node
    - 1x $VM_{40}$ in a 40 logical cores ph. node

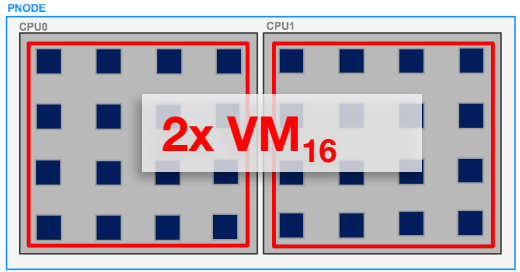


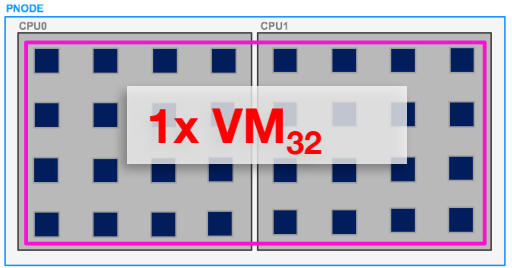

# HS06 score computation

- For each core (vCPU) the sequence of benchmarks runs 3 times
  - Each core sequence is independent (potential **misalignment**)
  - Multiple-Speed approach

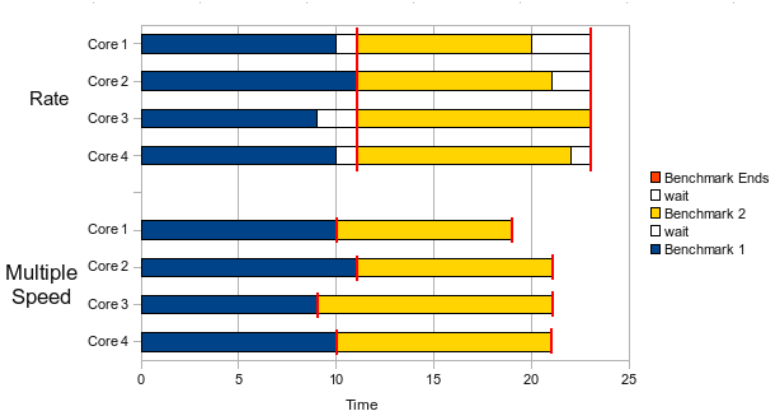


- For each core and benchmark, the median value of the 3 measurements is taken, and a ratio respect to a reference value is computed
  - NB: the 3 measurements are not consecutive

- Compute the geometric mean of the ratio values (per core)

- HS06 score = sum of the geometric means across cores

# HS06 benchmarks

| Bmk | Int vs Float | Description |
|---|---|---|
| 444.namd | CF | 92224 atom simulation of apolipoprotein A-I |
| 447.dealII | CF | Numerical Solution of Partial Differential Equations using the Adaptive Finite Element Method |
| 450.soplex | CF | solves a linear program using the Simplex algorithm |
| 453.povray | CF | a ray-tracer. Ray-tracing is a rendering technique that calculates an image of a scene by simulating the way rays of light travel in the real world |
| 471.omnetpp | CINT | discrete event simulation of a large Ethernet network. |
| 473.astar | CINT | derived from a portable 2D path-finding library that is used in game's AI |
| 483.xalancbmk | CINT | XSLT processor for transforming XML documents into HTML, text, or other XML document types |

# (score @ 64)/(score @32)

| CPUmodel | OS | isVM | vm | |
|---|---|---|---|---|
| E5 0 | SLC6 | 0 | olsnba02 | 1.10 |
| E5 v2 | SLC6 | 0 | p05496706a96156 | 1.13 |
| | | | p05496706h02393 | 1.13 |
| E5 v3 | CC7 | 1 | bmk16-0c06b8b8-0516- | 1.14 |
| | | | bmk16-ab111284-fc44- | 1.14 |
| | | | bmk32-03-centos | 1.13 |
| | | | bmk32-04-centos | 1.13 |
| | | | bmk32-07-centos-26e2 | 1.13 |
| | | | bmk32-07-centos-fffb | 1.13 |
| | SLC6 | 1 | bmk16-slc-7d1696e4-9 | 1.18 |
| | | | bmk16-slc-e839aa4f-c | 1.18 |
| | | | bmk32-05 | 1.18 |
| | | | bmk32-06 | 1.18 |
| | | | bmk32-server | 1.19 |
| E5 v4 | CC7 | 1 | bmk20-centos-3f9d1a5 | 1.15 |
| | | | bmk20-centos-51102ea | 1.15 |
| | | | bmk40-cc7-asd | 1.15 |
| | SLC6 | 1 | bmk40-67ba2527-0a46- | 1.16 |
| | | | bmk40-fc870633-5ce7- | 1.17 |

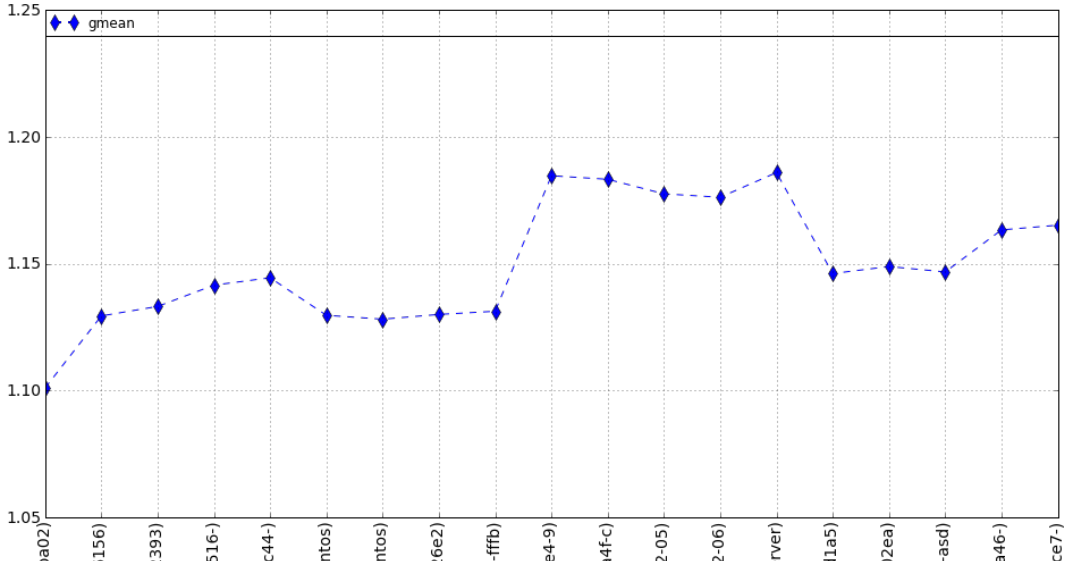


- Ratio of HS06 scores for each benchmarked server
  - Removes the CPU frequency dependency
  - Results for similar servers (or VMs) are comparable => can aggregate by CPU model classes in the next plot
  - Values range from 10% to 19%

# Ratio

| | | | bits | 32 | 64 | ratio |
|---|---|---|---|---|---|---|
| **CPUmodel** | **OS** | **isVM** | **ncores** | HS06 score | | |
| **E5 0** | **SLC6** | **0** | **32** | 334.443039 | 368.200709 | 1.10 |
| **E5 v2** | **SLC6** | **0** | **32** | 339.945662 | 384.547059 | 1.13 |
| **E5 v3** | **CC7** | **1** | **16** | 167.027239 | 190.819932 | 1.14 |
| | | | **32** | 336.416957 | 380.143415 | 1.13 |
| | **SLC6** | **1** | **16** | 167.284242 | 198.056291 | 1.18 |
| | | | **32** | 330.745914 | 389.659848 | 1.18 |
| **E5 v4** | **CC7** | **1** | **20** | 204.118525 | 234.214985 | 1.15 |
| | | | **40** | 404.375505 | 463.695030 | 1.15 |
| | **SLC6** | **1** | **40** | 398.016795 | 463.375925 | 1.16 |

HS06 (i.e. sum of geom. mean)

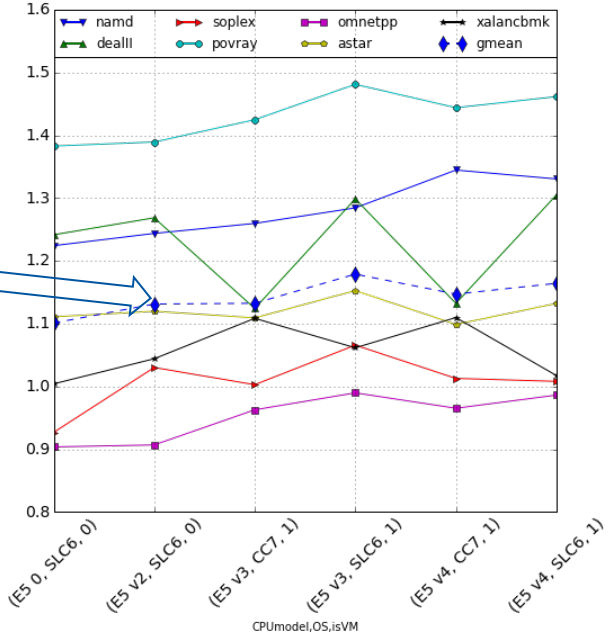


Ratio $<HS06_{64bits}>/<HS06_{32bits}>$ (average per CPU model, OS, cores)

- Ranges from 10% (S.B.) to 18% (Haswell)
- Is consistent for VM 16 and 32 on same ph. node
- Each individual benchmark in HS06 suite has completely different ratio values (remember: geom. mean)

# Ratio CC7/SLC6

Is there any difference in HS06 score when moving from SLC6 to CC7?

- Using the same HS06 scores, but doing a different ratio $<HS06_{CC7}>/<HS06_{SLC6}>$
- NB: **VMs** on same ph node => ratio removes CPU frequency e
- For VMs on a CC7 ph. node the difference is within few %
- Difference of ratio for –m32 and -m64 is within <1%

| | | | OS | CC7 | SLC6 | ratio_c_s |
|---|---|---|---|---|---|---|
| **CPUmodel** | **isVM** | **bits** | **ncores** | | | |
| **E5 v3** | **1** | **32** | **16** | 167.027239 | 167.284242 | 0.998 |
| | | | **32** | 336.416957 | 330.745914 | 1.017 |
| | | **64** | **16** | 190.819932 | 198.056291 | 0.963 |
| | | | **32** | 380.143415 | 389.659848 | 0.976 |
| **E5 v4** | **1** | **32** | **40** | 404.375505 | 398.016795 | 1.016 |
| | | **64** | **40** | 463.695030 | 463.375925 | 1.001 |

- Is this only valid for VMs running on a CC7 ph. host?
  - Need to try on ph. hosts with CC7 and SLC6

# Ratio CC7/SLC6

- What about fast benchmarks?
  - Do fast benchmarks scale similarly from SLC6 to CC7?
- Running in the same VMs used for HS06 measurements
  - Number of parallel processes (mp_num) == num vCPUs
  - Here there is not concept of 32 and 64 bits, as it was in HS06 case
    - KV is compiled at 64 bits
- Results $Ratio_{CC7/SLC6}$
  - DB12: +5 to +12%
  - DB12cpp: No difference
  - KV: -17% ??
    - But here KV is used in a not orthodox way: cvmfs libraries for SLC6 used in CC7
      - SLC6: __ieee754_log from libm-2.12.so
      - CC7: __ieee754_log_avx from libm-2.17.so

| | | Ratio_CC7/SLC6 | kv_speed | DB12 | DB12cpp |
|---|---|---|---|---|---|
| **CPUmodel** | **isVM** | **mp_num** | | | |
| **E5 v3** | **1** | **16** | 0.84 | 1.12 | 1.01 |
| | | **32** | 0.83 | 1.06 | 1.01 |
| **E5 v4** | **1** | **40** | 0.82 | 1.05 | 1.00 |

# Fast bmk Vs HS06 double ratio (PRELIMINARY)

- For each CPU model in VM or ph host with SLC6 OS
  - Measure
    - HS06 scores @ 32 and 64 bits
    - KV, DB12 and DB12cpp
  - Evaluate double ratio
    - Data in next slide
    - Ref: S.B.

$$\frac{\frac{fbmk_{conf}}{HS^{x}_{conf}}}{\frac{fbmk_{ref}}{HS^{x}_{ref}}}$$

- Double ratio for
  - KV (GEANT) seems to have a speedup of 30-50% respect to S.B.
    - *Further investigation needed*
    - Slightly reduced @ 64bits
  - DB12cpp within 5%
  - DB12 python ~10% @ 32 bits

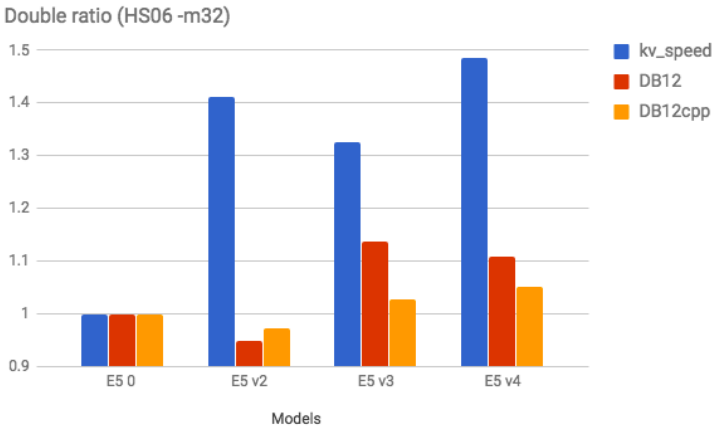


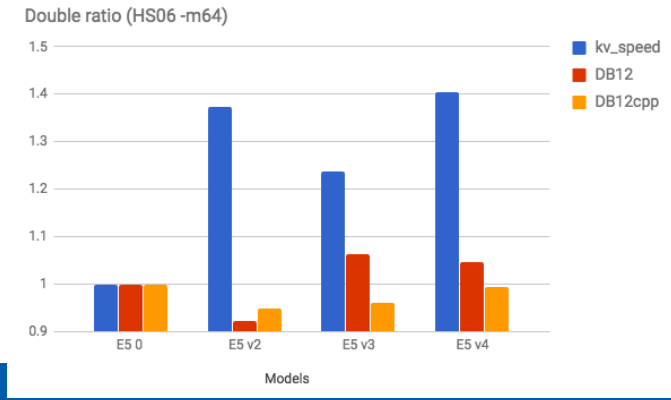

# Data

| | | | | kv_speed | DB12 | DB12cpp | HS06 -m32 | HS06 -m64 | ratio respect to E5 0 | | | | | double ratio HS06 -m32 | | | double ratio HS06 -m64 | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **CPUmodel** | **OS** | **isVM** | **mp_num** | | | | | | **kv_speed** | **DB12** | **DB12cpp** | HS06 -m32 | HS06 -m64 | **kv_speed** | **DB12** | **DB12cpp** | **kv_speed** | **DB12** | **DB12cpp** |
| **E5 0** | **SLC6** | **0** | **32** | 0.574719 | 10.039238 | 12.33 | 334.443039 | 368.200709 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1 | 1 | 1 | 1 | 1 | 1 |
| **E5 v2** | **SLC6** | **0** | **32** | 0.82482 | 9.678012 | 12.203 | 339.945662 | 384.547059 | 1.44 | 0.96 | 0.99 | 1.02 | 1.04 | 1.41 | 0.95 | 0.97 | 1.37 | 0.92 | 0.95 |
| **E5 v3** | | | **16** | 0.761516 | 10.257992 | 12.5165 | 167.284242 | 198.056291 | 1.33 | 1.02 | 1.02 | 1.00 | 1.08 | 1.32 | 1.02 | 1.01 | 1.23 | 0.95 | 0.94 |
| | **SLC6** | **1** | **32** | 0.752906 | 11.293019 | 12.526 | 330.745914 | 389.659848 | 1.31 | 1.12 | 1.02 | 0.99 | 1.06 | 1.32 | 1.14 | 1.03 | 1.24 | 1.06 | 0.96 |
| **E5 v4** | **SLC6** | **1** | **40** | 0.813137 | 10.589897 | 12.3545 | 398.016795 | 463.375925 | 1.41 | 1.05 | 1.00 | 0.95 | 1.01 | 1.49 | 1.11 | 1.05 | 1.41 | 1.05 | 1.00 |

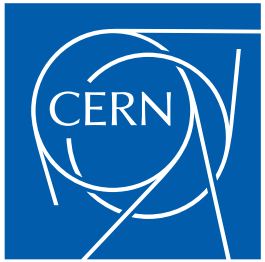
CERN