Benchmarking Cloud Resources

M. Alef[3], C. Cordeiro[1], A. De Salvo[2], A. Di Girolamo[1],
L. Field[1], D. Giordano[1], F. Schiavi[1], A. Wiebalck[1]

[1] CERN IT Department, [2] Università & INFN Roma I, [3] Karlsruhe Institute of Technology

Scenario

Performance measurements and monitoring are essential for the efficient use of computing resources. In a commercial cloud environment an exhaustive resource profiling has additional benefits due to the intrinsic variability of the virtualised environment. In this context resource profiling via synthetic benchmarking quickly allows to identify issues and mitigate them. Ultimately it provides information about the actual delivered performance of invoiced resources.

A Benchmarking Workflow [1]

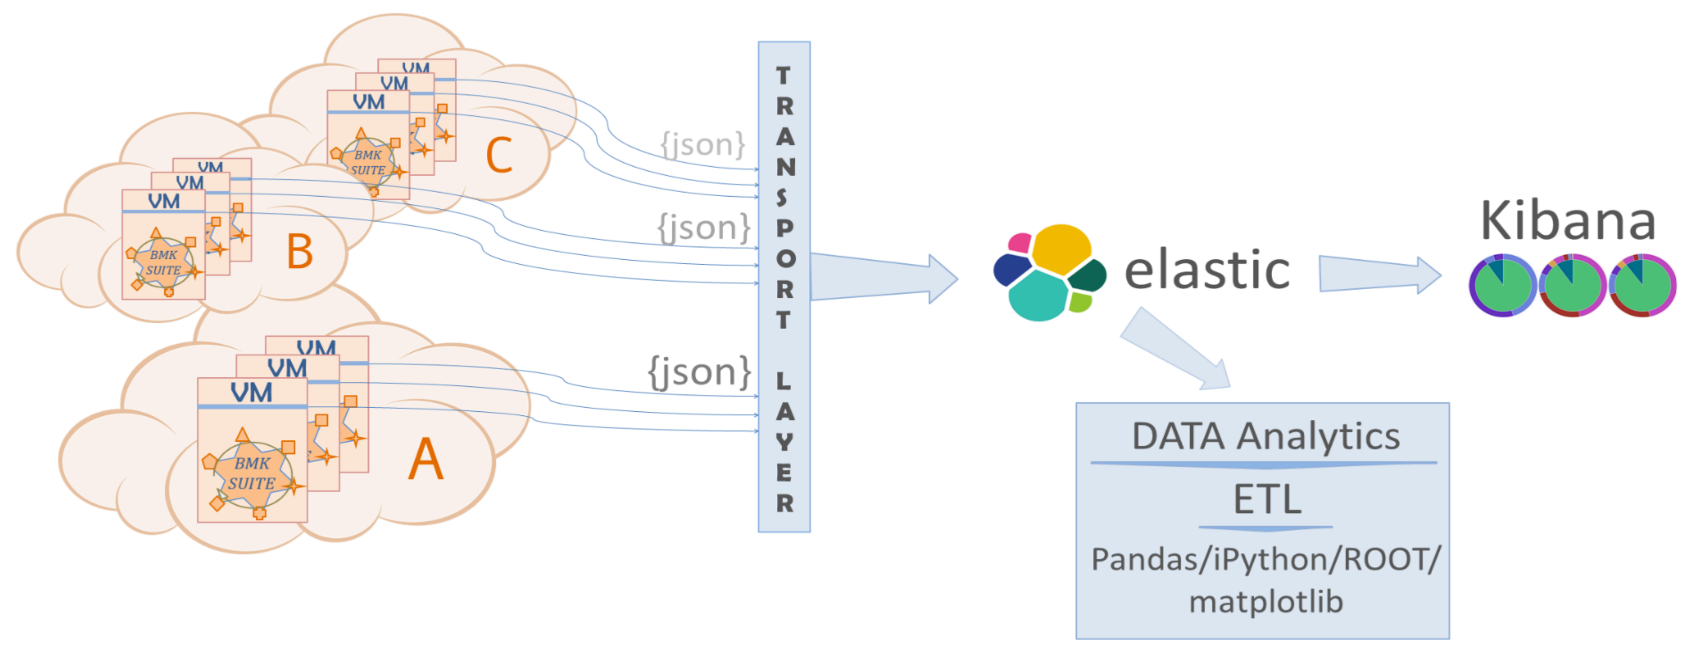

- A **configurable** sequence of benchmarks to run directly in the computing resources
- The results are pushed into a transportation layer, and later consumed into analytics platforms
- Kibana and ElasticSearch, for example, allow for analysis and prompt feedback of the collected results
- Detailed analysis performed with iPython analysis tools
- Up to now, a total amount ~3.6M results were collected from ~77k different machines, 89 different CPU architectures and 15 unique OS distributions

Benchmarking Commercial Clouds

In the context of its commercial cloud initiatives, CERN has acquired extensive experience in benchmarking commercial cloud resources, including Microsoft Azure, IBM, ATOS [2], T-Systems and the Deutsche Boerse Cloud Exchange.
Besides the performance validation, the thousands of benchmarks ran and collected from several distinct commercial cloud providers play a crucial rule on the analysis and correlation of data in several areas.

Performance vs Effective Cost

The exhaustive profiling across different HW models allows for a fine grained comparison of the effective cost of different cloud flavors

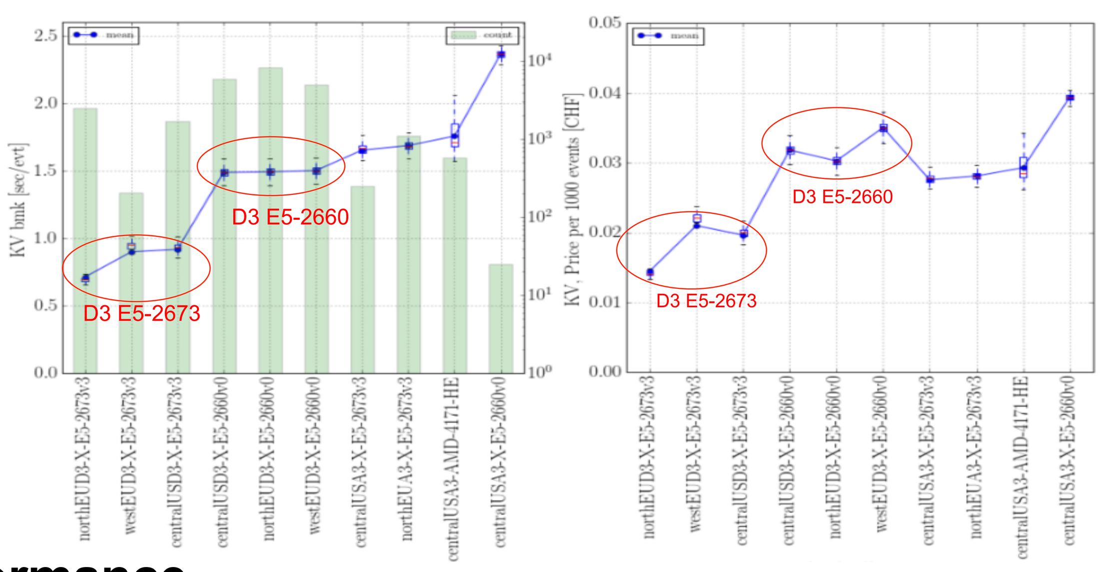

Estimation of WLCG jobs performance

By using benchmarks representative of the HEP workloads, it is possible to estimate the resource usage of the experiments jobs, anticipating therefore their runtime behavior.
In all WLCG sites, this compute power estimation is measured in units of HS06 [3], via a benchmark process that runs for about 6 hours on a baremetal host. This approach is not representative of the VMs performance during their lifetime. Faster benchmarks have been created and calibrated to scale with HS06 (e.g. the DIRAC spec [4]), allowing users to run them in just a few minutes.

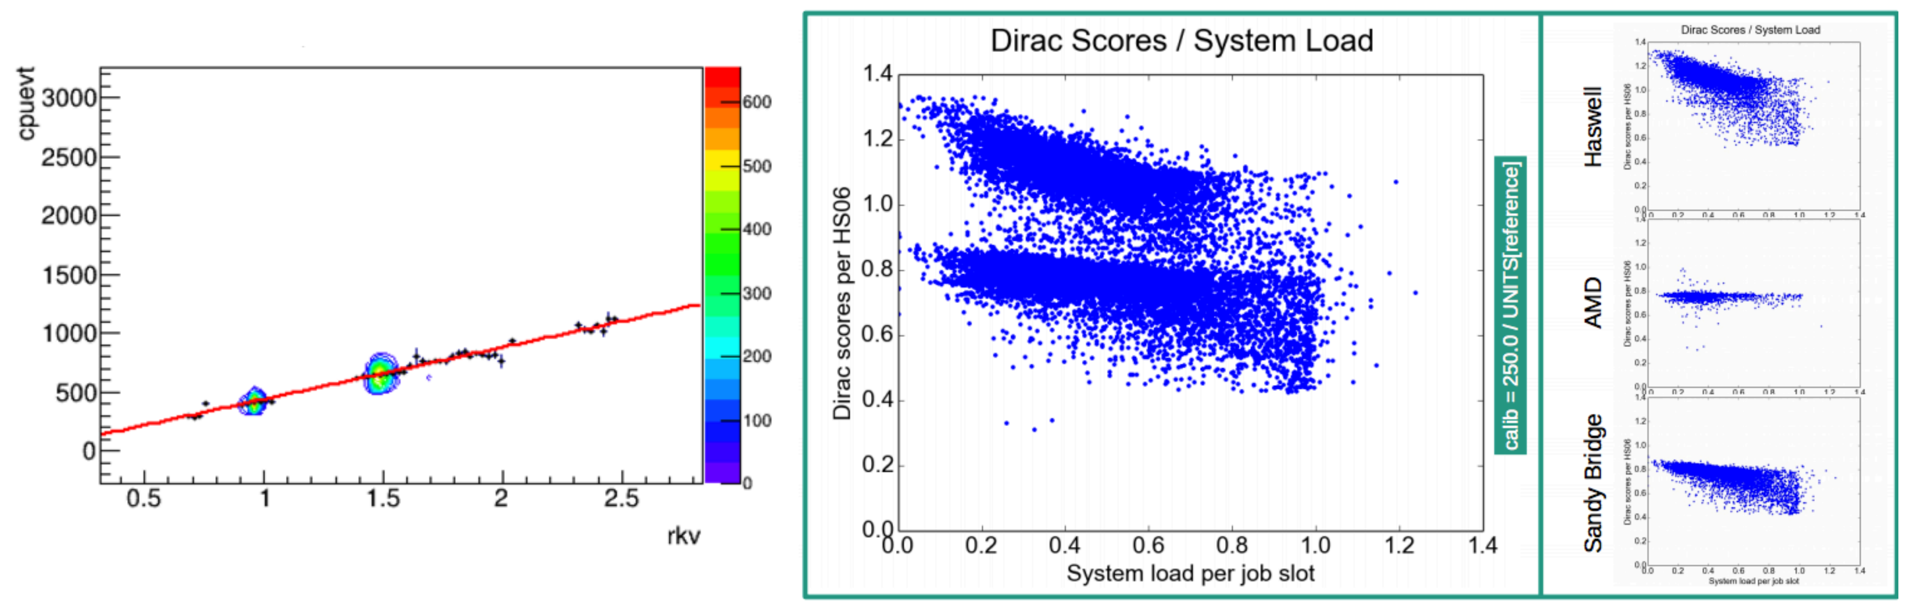

Cloud Procurement

The CERN cloud procurement process has greatly profited of benchmark measurements to assess the compliance of the bids with the requested technical specifications.

During the tendering phase, cloud providers are requested to benchmark their resources. The results collected by CERN are used not only to verify compliance with the Technical Requirements but also to collect statistics about cloud compute performance.

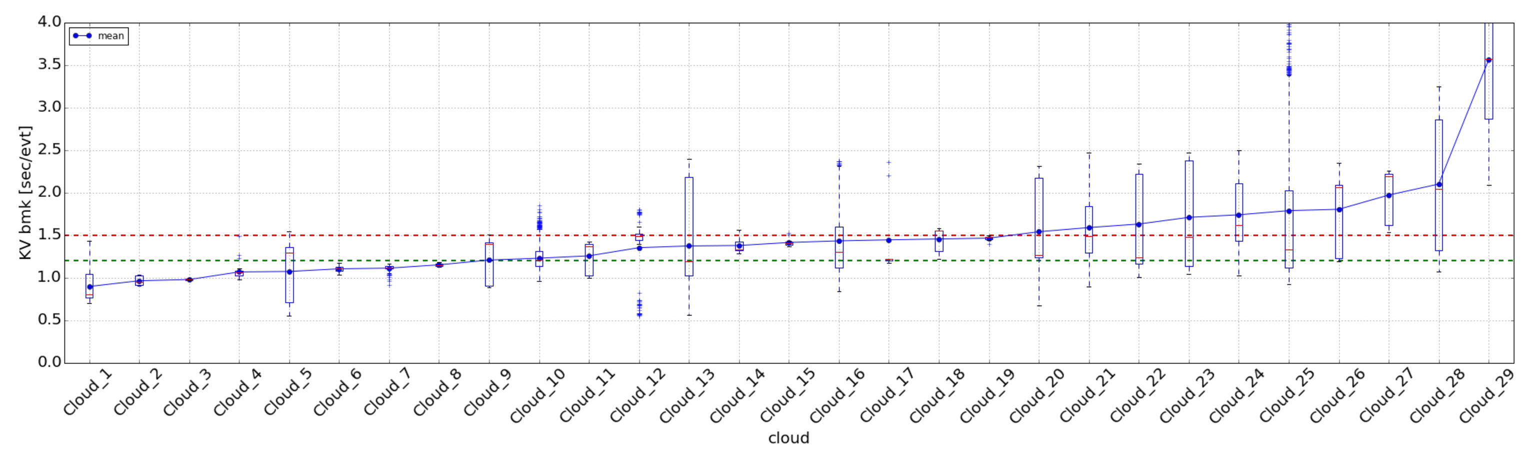

Continuous benchmarking during the production activity serves as an indicator of the delivered performance. In situations where the delivered performance falls under certain pre-defined benchmark reference values, Service Credit Compensation measures might be applied as a penalty.

Analysis of a Compute Cell at CERN

A recent benchmark study has examinated the compute performance of a new cell consisting of 240 physical servers located in the Wigner data center (Hungary). The cell runs Openstack and is configured to be an integral part of the HTCondor batch system at Tier-0.

Each baremetal server is equipped with two 8-core Intel Xeon Processors E5-2630 v3, summing up to 32 threads per machine with Simultaneous Multi-Threading enabled.

Different use cases were investigated, both directly on the baremetal servers and on VMs. The virtualized tests consisted in provisioning and sequentially benchmarking VMs, both fully and partially loading the hypervisors, and changing the flavors with each use case. The benchmarks were configured and scheduled to run synchronized for all the VMs.
Preliminary benchmark results show good agreement of the measured performance with respect to the expected scaling factors.

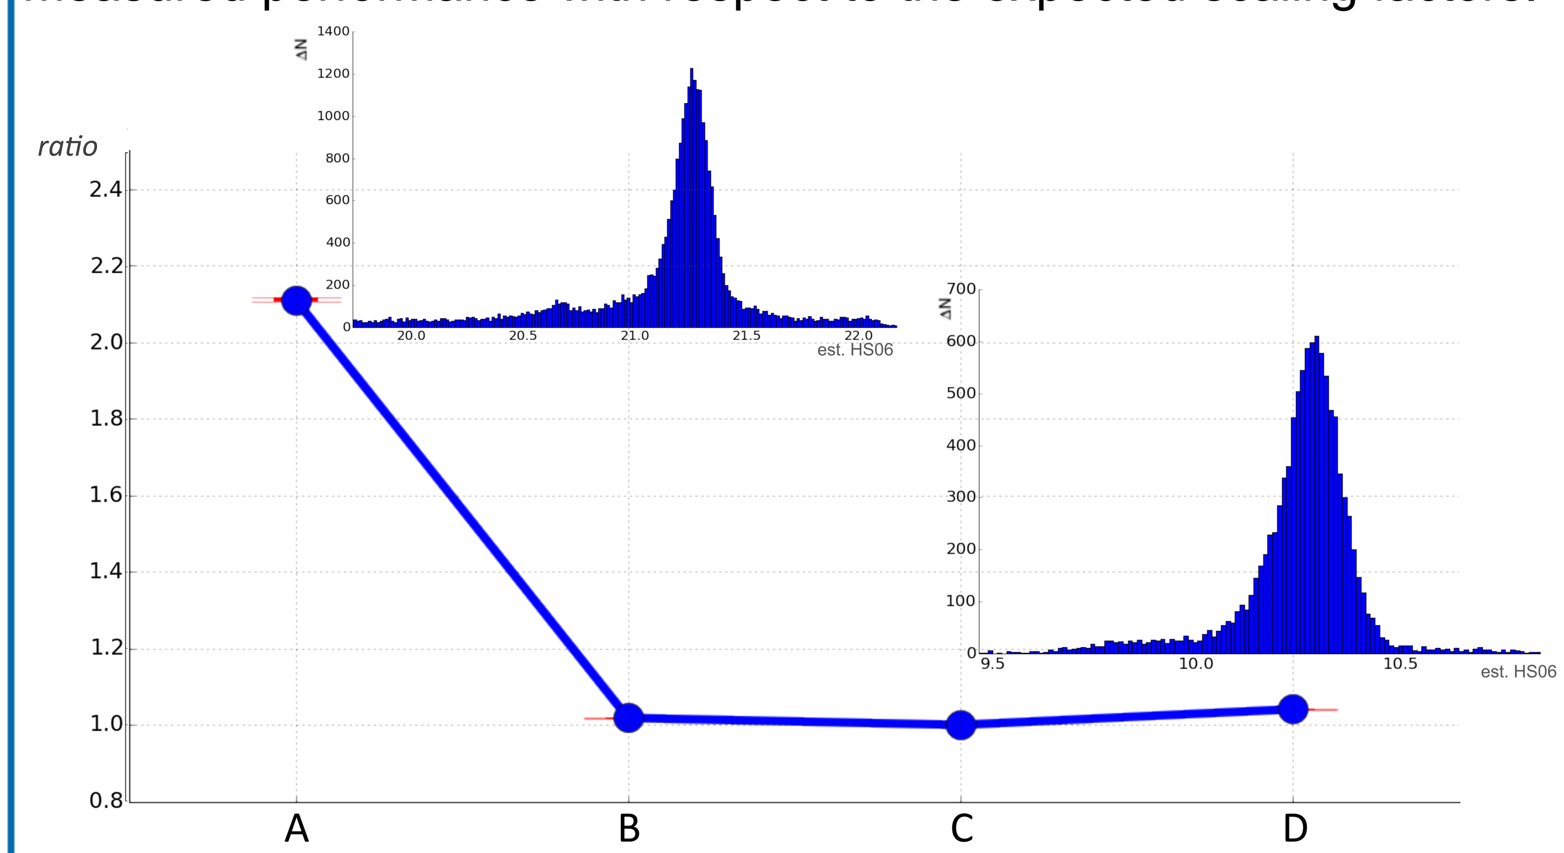

- **A:** *1x* 8-core VM, on a physical CPU
- **B:** *2x* 8-core VMs, on a physical CPU
- **C:** *4x* 8-core VMs, **two** per physical CPU
- **D:** *2x* 16-core VMs, **one** per physical CPU

References

[1] https://bmkwg.web.cern.ch/bmkwg/
[2] "Accessing commercial cloud resources within the European Helix Nebula cloud marketplace", D. Giordano et al, 2015 Journal of Physics: Conference Series 664 022019
[3] "A comparison of HEP code with SPEC benchmarks on multi-core worker nodes", Michele Michelotto et al, 2010 J. Phys.: Conf. Ser. 219 052009 doi:10.1088/1742-6596/219/5/052009
[4] https://diracgrid.org/

Contact: domenico.giordano@cern.ch