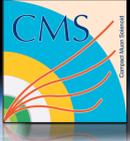


# The CMSSW Benchmarking Suite

## Using HEP code to measure CPU performance



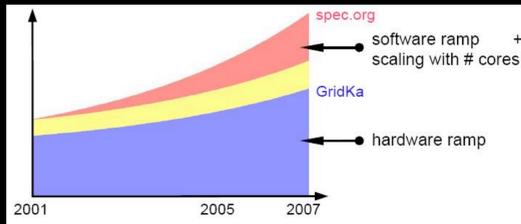
Gabriele Benelli  
CERN

On Behalf of the CMS Offline and Computing Projects



### HEP CPU benchmarking

In recent years a discrepancy has been observed between the CPU performance estimates given by the industry standard benchmarks (SPEC\*) used for HEP computing and actual performances of HEP code. The issue is due to multiple factors: compiler version, compilation flags, operating system, 32 vs. 64 bit architectures, the way the benchmark is run on multi-core architectures.



CPU benchmark divergence between published SPEC results and SPEC results obtained with HEP environment at GridKa [courtesy of M.Alef, FZK]

The results published at [www.spec.org](http://www.spec.org) include improvements to the environment in which the benchmarks are run, that do not reflect the HEP environment (for CMS Scientific Linux 4 64-bit, gcc3.4.5 32-bit, -O2 -pthread -fPIC flags). In this context the idea of using directly CMS experiment code as a realistic benchmarking tool for new architectures became important within the HEPiX CPU benchmarking working group.

The CMSSW benchmarking suite is based on centrally maintained tools used for Release Validation in CMS. The aim of the suite is to provide an out-of-the-box tool to benchmark architectures that ships with the CMS software release and to enable computing management and tier centers to access benchmarking information. The suite has been used to provide input to the HEPiX benchmarking group

\* SPEC®, SPECint®, and SPECfp® are registered trademarks of the Standard Performance Evaluation Corporation (SPEC).

### Running experience

The CMSSW performance suite has been used within the HEPiX CPU benchmarking working group to benchmark the lxbench cluster at CERN. The idea for this test was to investigate the performance of very different architectures with a variety of physics processes and processing steps to gain as much information as possible on the behavior of CMSSW as a benchmarking tool.

	lxbench01	lxbench02	lxbench03	lxbench04	lxbench05	lxbench06	lxbench07	lxbench08
Number of cores	2	2	4	4	4	4	8	8
Frequency (GHz)	2.8	2.8	2.2	2.66	3.0	2.6	2.33	2.33
Cache (MB)	1	2	2	4	4	2	8	12
Memory (GB)	2	4	4	8	8	8	16	16
Processor	Nocona	Irvingdale	Opteron 275	Woodcrest	Woodcrest	Opteron 2218 Rev.F	Clovertown	Xeon HarperTown E5410
Vendor	Intel	Intel	AMD	Intel	Intel	AMD	Intel	Intel

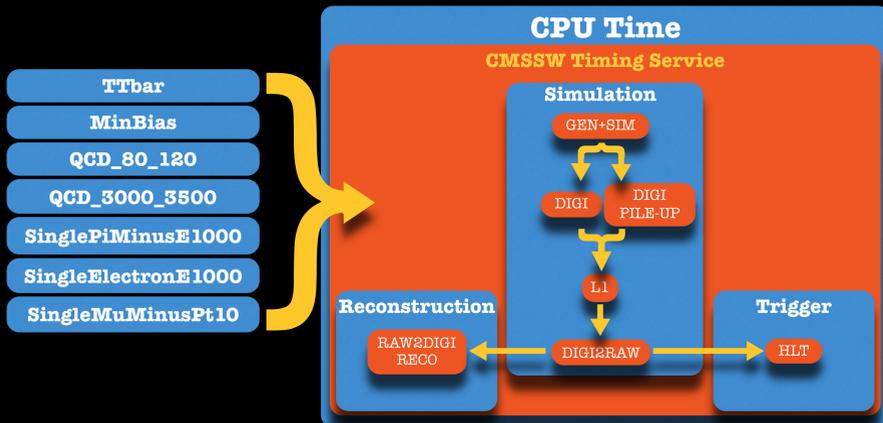
The lxbench cluster at CERN, used by the HEPiX CPU benchmarking working group.

On each lxbench machine CMSSW\_2\_2\_6 has been installed and the benchmarking server launched. The benchmarking client configuration file set the performance suite to run on all cores, on 6 physics processes and profile steps GEN-SIM, DIGI and RAW2DIGI-RECO. The client was launched on a CMS lxbuild machine running CMSSW\_2\_2\_6.

The execution time of the profiling ranged from 15 to 32 hours for this test, depending on the machine CPU power, and it resulted in writing on disk about 2.5 GB per core. All threads run independently keeping each core loaded, the cumulative execution time difference among cores were less than 1%.

### CMSSW performance suite

The CMSSW performance suite is designed to run at each (pre-)release as part of the Release Validation effort to provide detailed information about all aspects of software performance. The CMSSW benchmarking suite uses only the CPU time information.

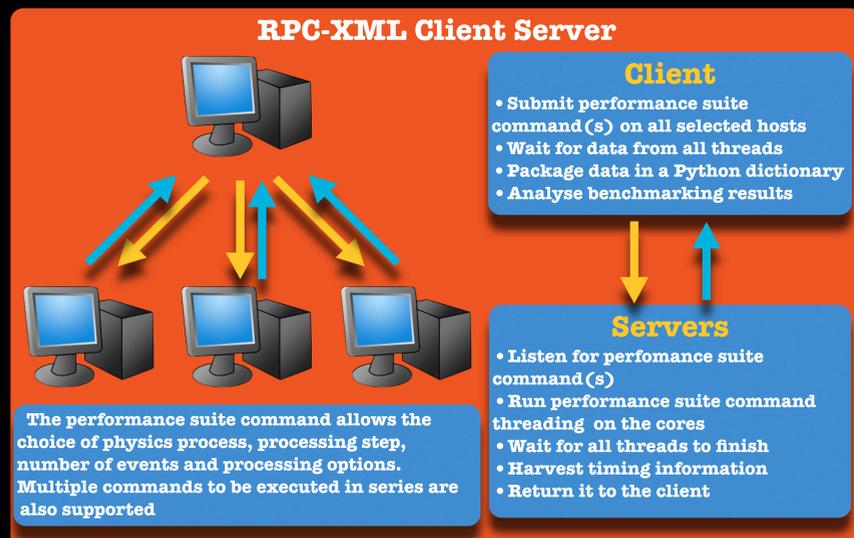


The CMSSW performance suite physics processes and processing steps for the CMSSW Timing Service profiling.

The suite is driven by a Python command-line interface that can be configured to run on any of the Release Validation candles above and for any of the processing steps. For the architecture benchmarking use case the suite can load all cores of a multicore machine, by threading itself in each core.

### CMSSW benchmarking suite

The benchmarking suite is based on a client/server approach to allow benchmarking of multiple machines at once. The only requisite is that the machines can run CMSSW.



### Results

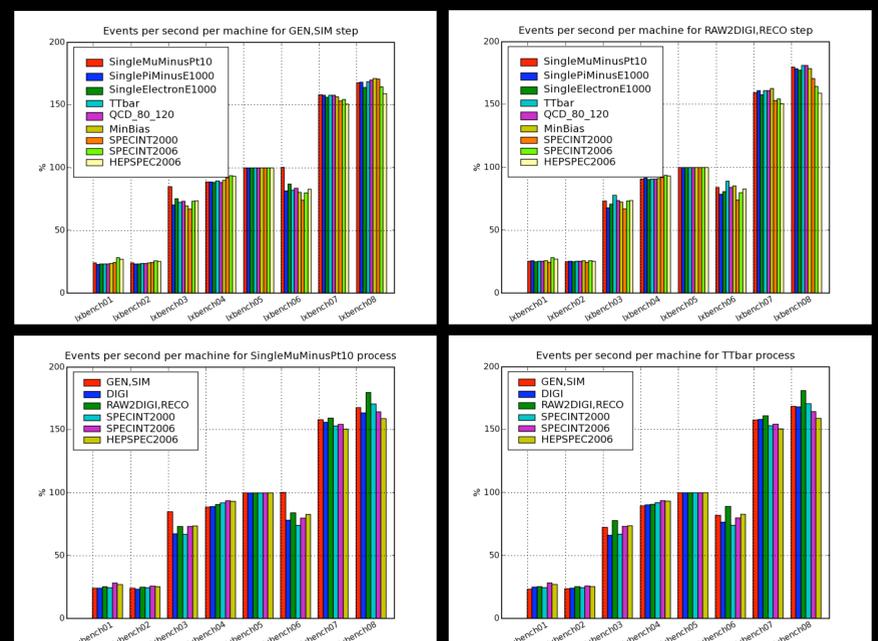
In order to provide a result comparable with the SPEC benchmarks that increase for increasing CPU performance, the results are reported in terms of number of events per unit of time. The timing information used by the CMSSW benchmarking suite is per event. On each core the average time per event is then inverted into the number of events per seconds and the sum over all cores is given as the machine performance.

The benchmarking results analysis provides an ASCII table output for each machine and a series of comparison plots (if more than one machine was profiled).

Step	SingleMuMinusPt10	SinglePiMinusE1000	SingleElectronE1000	TTbar	QCD_80_120	MinBias
GEN,SIM	12041.82	79.8	54.88	53.18	63.17	489.86
DIGI	8361.72	7365.57	8588.17	3514.57	3979.87	7714.21
RAW2DIGI,RECO	17535.64	17840.67	8597.12	1739.21	2496.56	19179.63

CMSSW benchmarking results in events/1000 seconds score for all physics processes and processing steps, for lxbench08

Comparison plots are produced per processing step showing all candles for all machines and per candle showing all processing steps for all machines. All plots shown have been normalized to lxbench05, a very popular Intel worker node architecture. The SPEC results run with HEP environment are also reported for comparison.



The choice of default representative candle(s) and processing steps is to be determined based on the experiment needs. By construction the tool allows detailed studies of CMSSW software performance on different architectures.