
ATLAS & Benchmarks

Andrej Filipčič
for the ADC

What do we need benchmarks for in ATLAS

- Accounting:
 - Precise benchmarking, eg HSPEC06
 - Consumption in spent CPU hours per site
- Job cputime/walltime estimate
 - Rough estimation, 10% accuracy is good enough
 - PanDA dynamic jobs (JEDI) – estimate cputime/event, shape the jobs to fit within given walltime
 - New simulation tasks already use HSPEC06 units for brokering and walltime estimate – site values stored in AGIS

ATLAS benchmark infrastructure

- Benchmarks installed in
 - `/cvmfs/atlas.cern.ch/repo/benchmarks/`
 - For now, HTCondor and KitValidation benchmarks
- Pilot wrapper can execute the given benchmark at the beginning of the job if requested in the job description
- The results are reported back to PanDA
 - Type and value, node parameters
- Not yet fully in place, but within a month we will start the collection
- The results will be available to both JEDI engine and to collect the site statistics on per node basis

Which benchmarks are useful for ATLAS

- HSPEC06: the standardized one, but close to impossible to run automatically
 - ➔ Full node reservations
 - ➔ Too long for comprehensive site evaluation
 - ➔ Licensing issues
 - ➔ ATLAS will likely not use it
- HTCondor: `condor_kflops`, `condor_mips`
 - ➔ Based on Whetstone, Dhrystone
 - ➔ Promising for few % accuracy, but current cvmfs binaries do not give consistent results, to be verified
- Boinc client runs Whetstone+Dhrystone at initialization
 - ➔ collected results in [ATLAS@Home](#) database show the values to be consistent within few % (benchmark vs MC Simulation) over wide range of CPUs
 - ➔ In general, Whetstone is reliable within few % and can run 10s to quickly estimate the node power
- KitValidation benchmark: muon MC simulation with ATHENA on fixed number of events
 - ➔ few minutes of runtime, but more accurate for ATLAS
 - ➔ Good candidate to automatically estimate the site power with a high precision

How to proceed

- HSPEC06 unusable
- Fast benchmarks will be very useful for ATLAS production system
 - ~10s runtime is acceptable
 - They are proportional to heavy benchmarks within 10%
- But there is no common WLCG benchmark which can be automatically executed – but it is crucial to have one
 - We expect from WLCG to define one
- What accuracy is needed? 1%?
 - Note that heavy benchmarks stressing memory, cpu and disk can be very inconsistent, and can depend on site architecture