



Benchmarking ATLAS applications

Franco Brasolin - INFN Bologna - <Franco.Brasolin@bo.infn.it>

Alessandro De Salvo - INFN Roma1 - <Alessandro.DeSalvo@roma1.infn.it>

May, 08 -2008

Outline

- **The ATLAS benchmark suite**
- **Benchmark execution**
- **Accessing the test results**
- **Conclusions and next steps**



The ATLAS benchmark suite

- **The goal is to define an ATLAS benchmark system, based on the experiment software**
 - **To have more realistic measurements of the CPU power, based on the experiment code and real jobs**

- **The ATLAS tests are based on an existing infrastructure**
 - **KitValidation**
 - **Global KitValidation portal**
 - <https://pc-ads-01.roma1.infn.it/KV>



The test infrastructure (1/2)

- **KitValidation (KV)**
 - KitValidation has been developed to test the ATLAS software
 - It can execute custom tests, defined by the software developers and included in the experiment software releases
 - Used by the community since 2003 to test the ATLAS software installation
 - Grid installation
 - End-user installation on desktops, laptops, ...
 - Generic pre-release testing

- **Global KitValidation (GKV) portal**
 - KitValidation is able to send the test results to a web service (GKV), if the corresponding feature is enabled
 - Few informations are sent to the portal
 - Machine infos (CPU type, CPU speed, Memory size, ...)
 - Test status (OK, FAILED)
 - Test logfiles and xml job info
 - Through the GKV web interface search engine, the test results can be selected
 - <https://pc-ads-01.roma1.infn.it/KV>



The test infrastructure (2/2)

- **The benchmark results are obtained by dynamically parsing the logfiles of the KV tests**
 - Timings are obtained from the ATLAS software framework (Athena)
 - Using the Athena ChronoStatSvc timing service for full job timing
 - Using specific timers from the Athena routines, currently available only for G4 simulation
 - Almost no bias is introduced when using the framework timing services
 - The ChronoStatSvc is measuring the actual job running time, excluding the basic initialization and finalization of the framework
 - Specific timing services can be added to fine tune the measurements

- **The benchmark results are collected in GKV**
 - All users can view the benchmark results
 - Only power users have full access to the informations contained in the GKV pages, using their personal certificate



Benchmarks execution

- **The test execution is performed by using a script**
 - <https://pc-ads-01.roma1.infn.it/KV/sw-mgr>
 - **The sw-mgr script is able to**
 - Install and test a specific ATLAS software release
 - Test an existing ATLAS software installation
 - **Multiple concurrent tests are also possible, with and arbitrary number of threads**
 - Useful when testing multi-core, multi-CPU machines at full load
 - **Example of a test of the existing installation of release 14.0.0 at CERN**

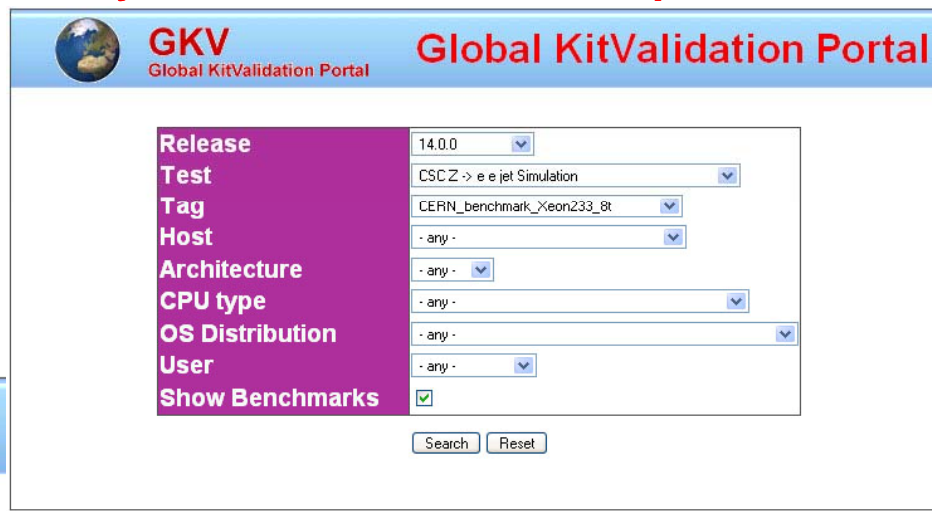
```
> VO_ATLAS_SW_DIR=/afs/cern.ch/project/gd/apps/atlas/slc3
> sw-mgr --test 14.0.0 -l $VO_ATLAS_SW_DIR/software/14.0.0 \
  -p $VO_ATLAS_SW_DIR/prod/releases/rel_14-0 -o -P AtlasOffline \
  -T release -t _i686_slc4_gcc34 -O opt -m 3.24 \
  --no-tag --tthreads 4 --kvpost --kvpost-tag "CERN_benchmark_AMD22_4t" \
  --kv-cache http://classis01.roma1.infn.it/pacman/test/cache
```

- **Every user is allowed to perform such tests and publish the information to the GKV portal**



Accessing the test results

- The test results may be browsed via the GKV portal



The screenshot shows the GKV Global KitValidation Portal search interface. It features a search form with the following fields:

- Release:** 14.0.0
- Test:** CSC Z -> e e jet Simulation
- Tag:** CERN_benchmark_Xeon233_8t
- Host:** - any -
- Architecture:** - any -
- CPU type:** - any -
- OS Distribution:** - any -
- User:** - any -
- Show Benchmarks:**

Buttons for 'Search' and 'Reset' are located at the bottom of the form.



For comments or informations please send me a mail (Alessandro.DeSalvo@roma1.infn.it)

Benchmarks are shown in the last column.

Multiple values are allowed for composite jobs where more than 1 task is chained.

Values in black are cumulative timings extracted via the ChronoStatSvc, while values in red are average times per event extracted from the G4 event timers

Host	Session	Tag	Test	Date	Result	Log	Benchmark
lxbench07.cern.ch	2FF84948-B22F-27A2-60C7-6A21965E0FA1	CERN_benchmark_Xeon233_8t	CSC Z -> e e jet Simulation	23-04-2008 19:26	OK	logfile	320.5944 +- 47.7870 s 53.5 min
lxbench07.cern.ch	4EB57CEB-847A-AE77-A23B-82CB716B219B	CERN_benchmark_Xeon233_8t	CSC Z -> e e jet Simulation	23-04-2008 19:27	OK	logfile	322.1444 +- 48.0519 s 53.7 min
lxbench07.cern.ch	D4C0F8EB-9067-E50E-CFBA-B9C9AB75B96D	CERN_benchmark_Xeon233_8t	CSC Z -> e e jet Simulation	23-04-2008 19:29	OK	logfile	320.7489 +- 47.8209 s 53.5 min
lxbench07.cern.ch	6EB34C6B-B850-E701-AE38-0C9018E158C1	CERN_benchmark_Xeon233_8t	CSC Z -> e e jet Simulation	23-04-2008 19:30	OK	logfile	322.1811 +- 47.9859 s 53.7 min




lxbench01	Intel(R) Xeon(TM) CPU 2.80GHz, 2 GB RAM
lxbench02	Intel(R) Xeon(TM) CPU 2.80GHz, 4 GB RAM
lxbench03	Dual Core AMD Opteron(tm) Processor 275
lxbench04	Intel(R) Xeon(R) CPU 5150 @ 2.66GHz
lxbench05	Intel(R) Xeon(R) CPU 5160 @ 3.00GHz
lxbench06	Dual-Core AMD Opteron(tm) Processor 2218
lxbench07	Intel(R) Xeon(R) CPU E5345 @ 2.33GHz

Preliminary Benchmarking Results for ATLAS

- Preliminary results obtained by running with release 12.0.6 (SLC3, gcc 3.2.3, 32 bits) over different processors at CERN

ATLAS vs. SPEC





Conclusions

- **The ATLAS benchmarking suite is in place and functional**
 - All the results are available from the GKV portal
 - <https://pc-ads-01.roma1.infn.it/KV>
 - Every user may contribute to the tests, even if non ATLAS people, by simply wget-ing the test script and executing it
 - An initial set of tests has been defined and performed on several architectures, both in SLC3 compatibility mode (release 12.0.6) and in SLC4 native mode (release 14.0.0)
 - The preliminary results are encouraging but still some work needs to be done (see below)

- **Next steps and improvements**
 - Provide more user-friendly pre-compiled scripts and examples to run the benchmarks for non-ATLAS users
 - Add to GKV the feature of presenting the summary of the benchmark sessions
 - Define a final set of tests to be performed in the benchmark sessions
 - It would be highly desirable to have a compatible set of test (physics channels) among the experiments, if possible