

ATLAS WN

Status and Expectations


Ivan Glushkov
for ATLAS Distributed Computing team

Current status

- ATLAS requirements (per score slot):
 - **Mcore Slots:** 8-core is optimal (more cores, less efficiency)
 - **HDD:** 20 GB (min. 10-15 GB)
 - **RAM:** 2 GB (better: 3-4 GB)
 - **Swap:** RAM+Swap \geq 4GB
 - **Network bandwidth:** min. 0.25 Gbit/s (CPU speed dependent)
 - **Software environment:** OS, CVMFS, all jobs run in Singularity containers
- CRIC:
 - Sites should enter values manually
 - Values are not kept up-to-date with the evolution of hardware
- Job mix:
 - Job distribution from the same task can have long tails in RAM and disk utilization i.e. hard limits are waste of resources
 - Lost walltime in job retries
 - Allocated resources which are not being used
 - Soft limits on job level works due to the mix of jobs with different resource requirements.

Software improvements

- MT (Multithreading)
 - Lower RSS per core (1GB for reconstruction)
 - Efficiency decreases with number of threads
- New pile-up
 - Pre-mixed overlay
 - Streaming of events instead of random events from multiple input files.
 - Reduced io load (important for sites with shared filesystem)
- DAOD_PHYS
 - Common format - potentially same file used by many analysis jobs reading directly from local storage



Memory usage of the MT ATLAS reconstruction (including Data Quality monitoring) as a function of number of threads in release 22 (blue triangles) compared to that of the MP reconstruction as a function of the number of workers in release 21 (open green circles). The solid lines are linear fits.

Experience

- Type of resources:
 - Grid:
 - x86
 - No hard limit on RAM or disk
 - Diskless sites:
 - Limited by network. Thus they run only simulation or generation.
 - HPC:
 - Preferably able to run all workflows (Example: Vega)
 - Anything else is with limited application and requires dedicated operational attention
 - Clouds:
 - Can be used by overlaying a batch system on them via cloudscheduler
 - GPUs:
 - Limited to ML workflows which represent a minimal portion of ATLAS production
- Resource description
 - If sites do not properly declare the parameters of the resources they provide, we have no way of guessing / fully utilizing them.

Expectations

- No significant requirement changes expected
- Run-3 workflows are not significantly different to Run-2