
# ARM software usage in ATLAS

Johannes Elmsheuser (BNL)

18 March 2024, Workshop on ATLAS Computing and Software Activities at BNL

# Introduction I

arm

- **What is ARM ?** (From Wikipedia)
  - Family of reduced instruction set computer (RISC) instruction set architectures for computer processors
  - Arm Ltd. develops the architectures and licenses them to other companies for their products like system on a chip (SoC) and system on module (SOM) designs
  - Low costs, minimal power consumption, and lower heat generation than their competitors

  $\rightarrow$ Mobile devices but also at computing centers (Bernd Panzer (CERN): ”ARM server market share is about 8% with vast majority AWS Graviton”)

# Introduction (II) - Access to ARM resources

- Access to ARM resources
  - CERN:
    - Public cluster lxplus-arm.cern.ch and build nodes with Ampere Altra
  - AWS:
    - Graviton2 and 3 used in Athena physics validation (see backup slides)
  - GoogleCloud:
    - Tau T2A machine series (Ampere Altra)
  - WLCG:
    - several sites provide/plan PanDA queues (see later)
  - HPCs:
    - Fugaku HPC at Riken Center for Computational Science, Japan is #4 in TOP500 supercomputer list of November 2023 (link)
    - EuroHPC Deucalion and Jupiter with ARM partitions (link)

# ATLAS SOFTWARE SCHEMATIC OVERVIEW

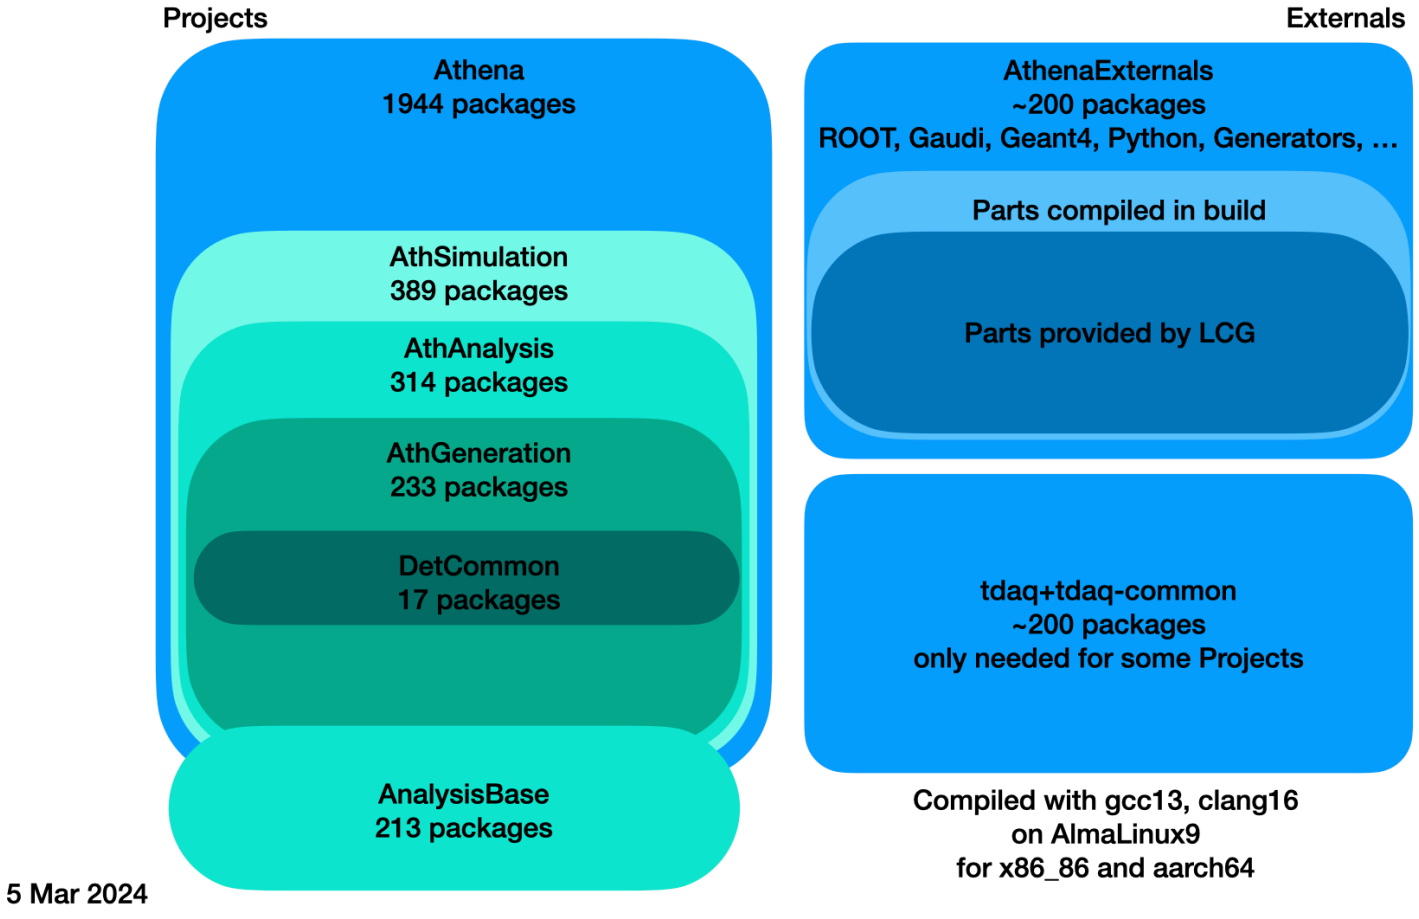


- 3 major blocks build by different teams:
  - Projects: ATLAS offline software
  - Externals: CERN EP/SFT through LCG layers and ATLAS offline software
  - TDAQ: ATLAS trigger/DAQ

# Nightly builds on ARM/aarch64

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| ARM | 24.0_AthSimulation_aarch64-el9-gcc13-opt | 2024-03-05T0001 | 05-MAR 01:09 | 0 (0) | 0 (0) | N/A | N/A | 05-MAR 01:41 |
| ARM | main_AnalysisBase_aarch64-el9-gcc13-opt | 2024-03-05T0220 | 05-MAR 07:05 | 0 (0) | 0 (0) | N/A | N/A | 05-MAR 08:11 |
| ARM | main_Athena_aarch64-el9-gcc13-opt | 2024-03-04T2101 | 05-MAR 03:45 | 0 (0) | 0 (0) | N/A | N/A | 05-MAR 06:21 |
| ARM | main_AthSimulation_aarch64-el9-gcc13-opt | 2024-03-04T2101 | 04-MAR 22:05 | 0 (0) | 0 (0) | N/A | N/A | 04-MAR 22:31 |
| ARM | main_DetCommon_aarch64-el9-gcc13-opt | 2024-03-04T2001 | 04-MAR 20:06 | 0 (0) | 0 (0) | N/A | N/A | 04-MAR 20:21 |

- 6 nightly builds for Athena, AthSimulation, AnalysisBase and DetCommon (link) projects fully integrated in standard ATLAS build system and available on CVMFS
- Selected stable Athena releases like 23.0.3 and 23.0.14 installed on CVMFS - used in physics validation (see later)
- Can easily build stand-alone docker/podman container for e.g. AthSimulation

- **Apple M1/M2/M3:** nightlies can easily be used for development/execution on Apple Silicon (documentation)

# ARM/aarch64 ATLAS builds and caveats


- Build flags
  - Using Armv8 defaults (gcc 13.1 allows up to armv9.3-a, gcc docu link)
  - Athena/Geant4 tested to work only up to `CXXFLAGS="-march=armv8.3-a -mtune=neoverse-n1"` on CERN ARM machines
  - No special `arch` option and no special compilation flags set apart from different linker flag in max-page-size
  - Test builds of Athena/AthenaExternals with clang16 work as well
- Floating Point Exception (FPE)
  - Athena FPE auditor code not working on ARM/aarch64 since it uses x86 specifics
- Potential numerical differences
  - E.g. due to different floating point libraries used - see StackOverflow link
  - Some fluctuations in physics objects at the level of $(10^{-4} - 10^{-6})$
  - Passed detailed physics validation for Geant4 simulation and reconstruction (see backup slides)
  - N.B. small numerical differences or Intel vs. AMD if IntelMathFunction used

# ARM resources through PanDA

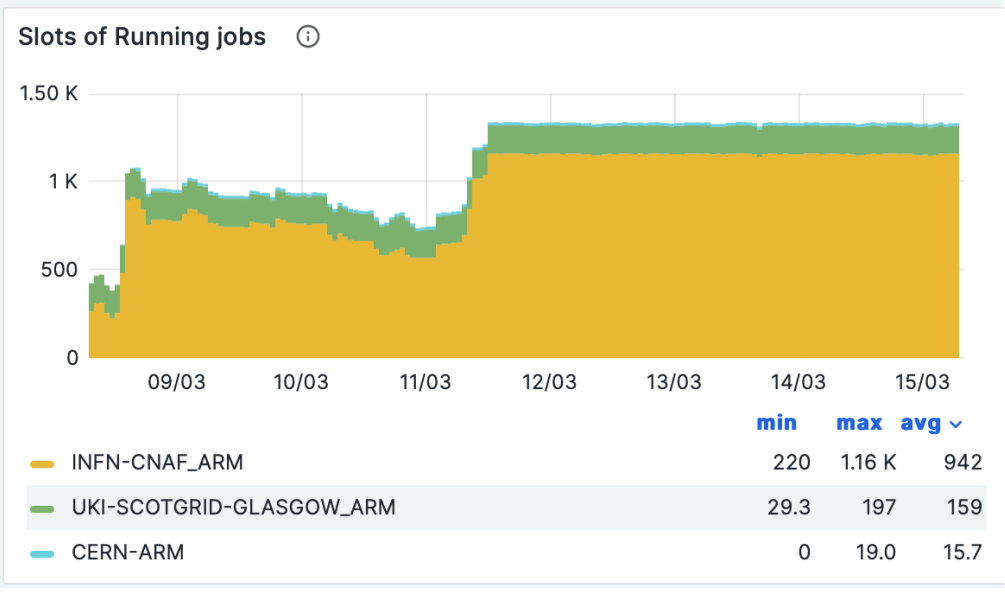


- Glasgow, INFN-T1/CNAF and CERN provide >1k job ARM job slots through PanDA
- Ampere Altra worker nodes
- Running MC simulation (Geant4) tasks right now

# HepScore integration

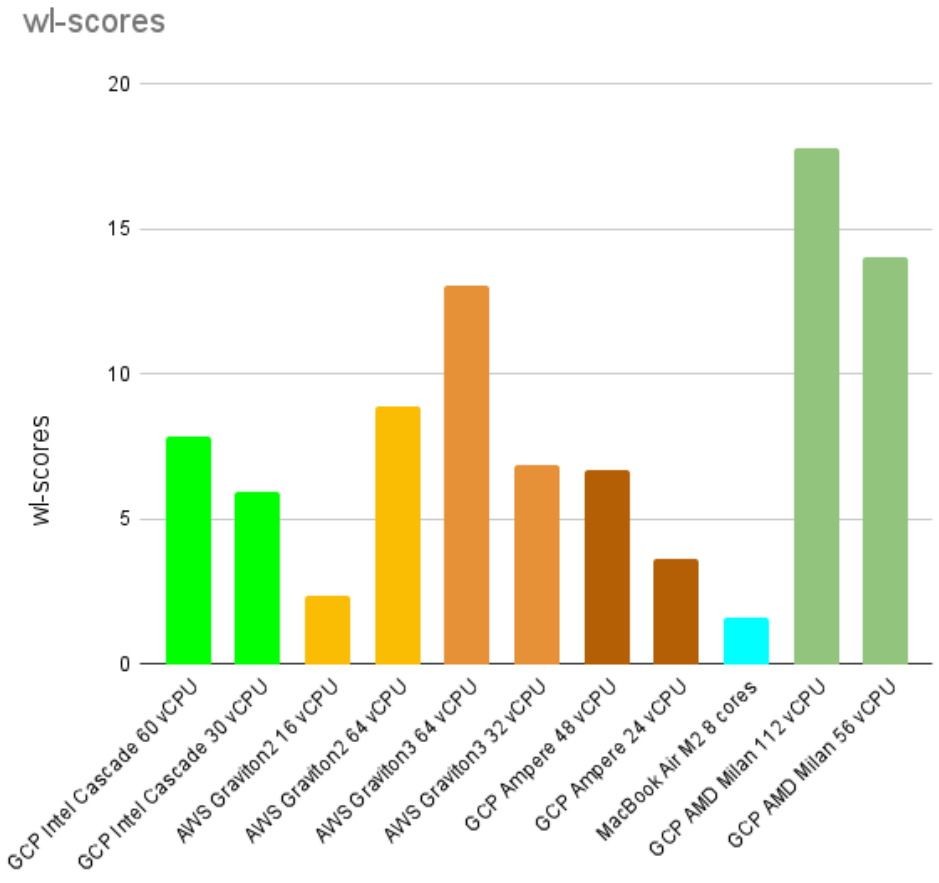


HepScore numbers for ATLAS reconstruction
Node fully packed with n×4 threads
Larger values are better

- WLCG replaced HepSpec06 benchmarks for resource pledging with 7 experiment workflows (ATLAS, CMS, LHCb, ALICE, Belle2) based HepScore23 (link)
- x86 and ARM
- ATLAS: single core Sherpa and AthenaMT data reconstruction
  → Different number of available cores explains different scales !
- Orange/Brown/Blue: ARM flavours
- Green: Intel/AMD flavours
- Some hyperthreading or IO related differences
- Similar trends for Event generation and simulation

# Further ideas, summary and Conclusions

- Further plans:
  - Derivation production tested to work fine
  - Some MC generators also work, but others need ATLAS MC expert attention due to complicated installation/setup on CVMFS
  - Power consumption requires bare-metal/super-user access to hardware
  - Looking forward to try out Nividia Grace
  - Try out G5g instances on AWS: Graviton2 CPU + NVIDIA T4G Tensor Core GPUs
- Summary:
  - Fully automated ARM ATLAS nightly builds available since more than 1.5 years
  - Geant4 and reconstruction physics/technical validation successfully passed using AWS through PanDA/Rucio production system
  - ATLAS data reconstruction and Sherpa event generation workflows for x86 and ARM in HepScore23 benchmark
  - ARM resources gradually appearing through PanDA for production usage
  - More details in CHEP 2023 talk and Proceedings paper

# BACKUP

# Geant4 physics validation with Athena,23.0.3

- Summer/Autumn 2022 successfully passed Geant4 10.6 physics/technical validation
- 1 million $t\bar{t}$ events (1k files, ≈700 GB output 1.5 days, 300*8 core job slots, 1800 USD) with Geant4 produced with stable Athena,23.0.3 release at AWS Graviton2 PanDA queue
- Compared with same events produced on x86_64 grid sites - all subsequent steps (reconstruction, merging histogramming) on x86_64
- Some fluctuations in different areas reported, but at the level as in other physics/technical validations

**ttbar_SingleLep**

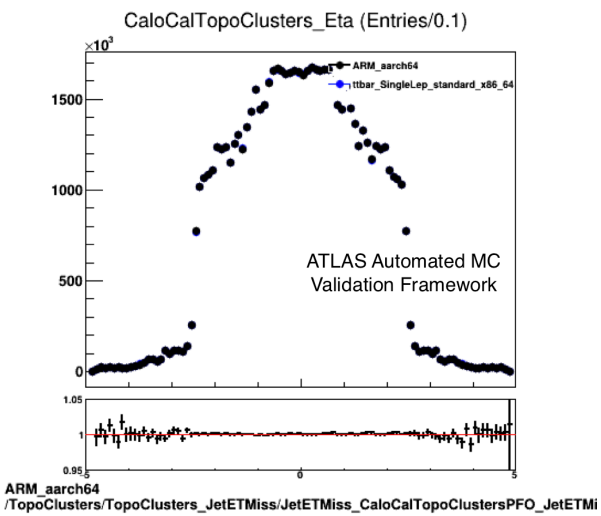


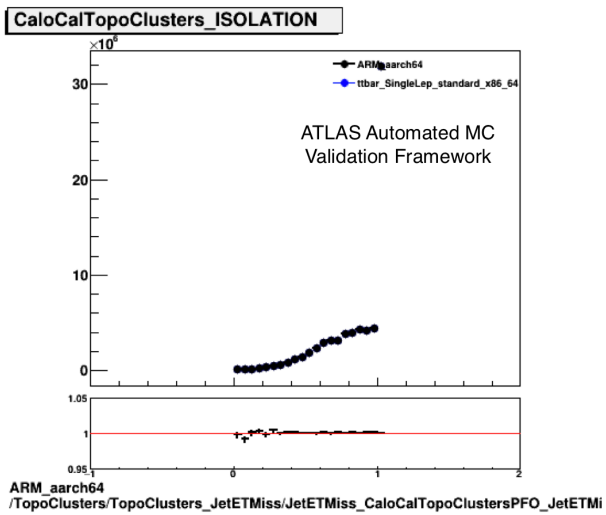


Example plots from calorimeter cluster validation of ARM/aarch64 (test) vs. x86_64 (ref)

# Reconstruction physics validation with Athena,23.0.14

- March 2023 successfully passed reconstruction physics validation
- Standard procedure of 13 MC physics processes (100k events each), no pile-up
- Digitization+Reconstruction step on AWS Graviton2 PanDA queue (130 jobs, 215 GB output, 1.5 days, 450 USD in total)
- Compared with same events produced on x86_64 grid sites - all subsequent steps (merging, histogramming) on x86_64 - task speed comparison in the backup

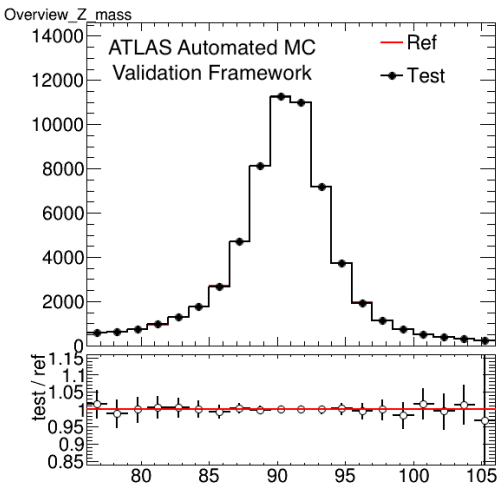


Reconstructed MC $Z \to \mu\mu$ invariant mass

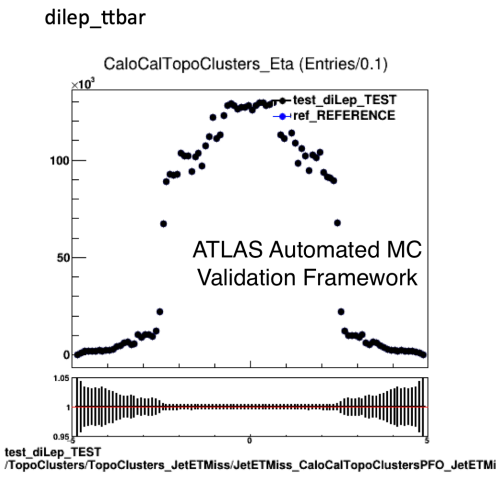


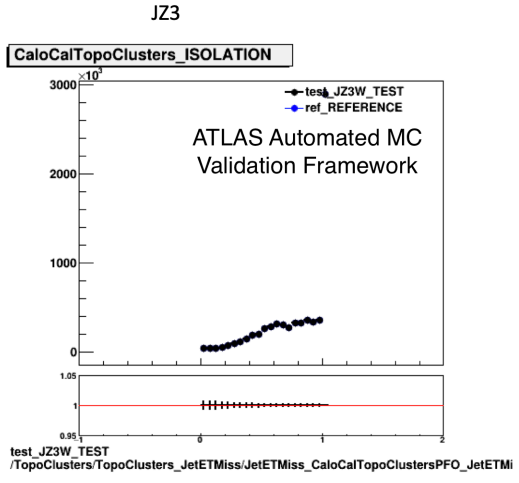


**Perfect agreement**

Example plots from calorimeter cluster validation