Meeting with LHCC Referees
ATLAS updates

Simone Campana, Torre Wenaus
for the ATLAS collaboration

Tier-0 Status & News

- Smooth processing of new collisions data
  - Conventional workflow (RAWtoESD, ESDtoAOD, ESDtoDPD, DQM, …)
  - Reco wall-time: 9-10s/event. Major improvements wrt 2015:
    - HW configuration: no hyperthreading => more memory/core
    - Metadata handling: faster file merging

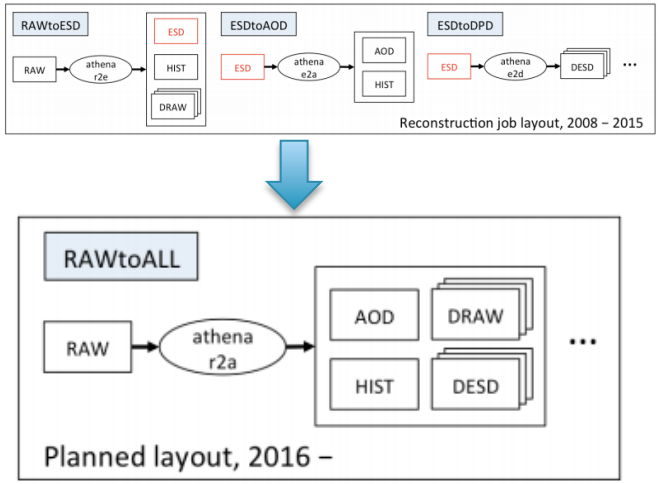

- RAWtoALL workflow commissioning
  - Extensive, large-scale performance tests
  - Performance: expected wall-time ~6s/event
  - Goal: validated by MD1 & TS1 (June 12)

- Resources
  - ~7k slots, 1600(+) more to be added soon
  - Back-fill by Grid Production, if not used by Tier-0

RAW Recording
May 9–18

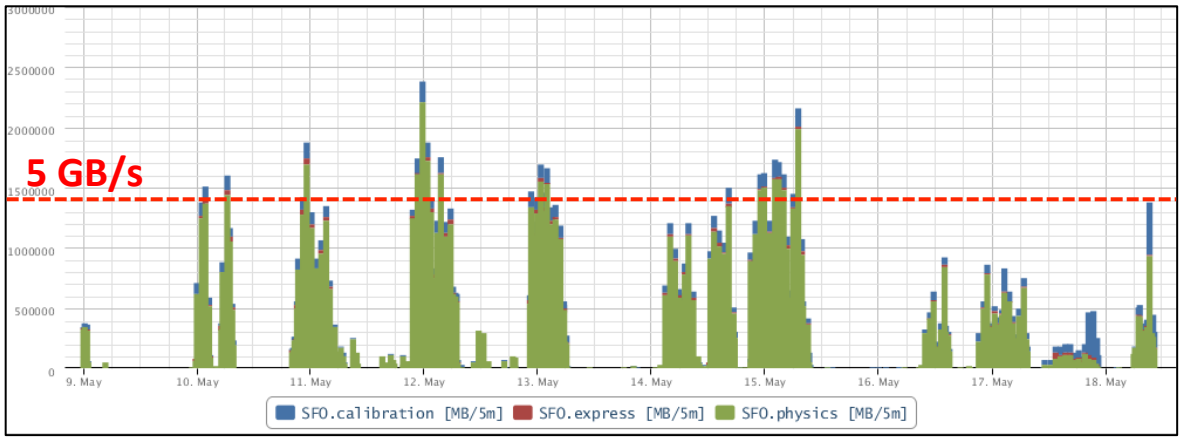

Processing at Tier-0
May 9–18

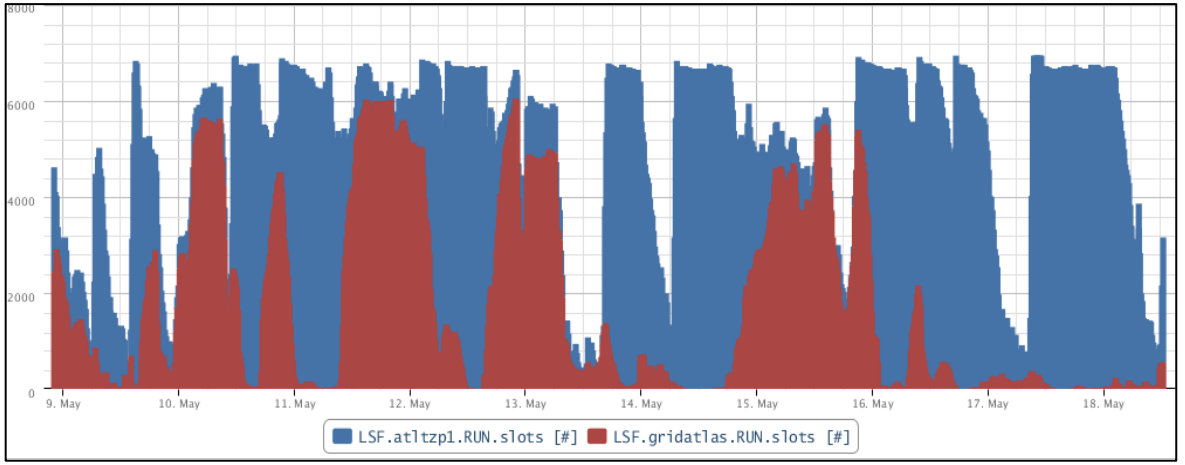

Tier-0 Processing Backlog
May 9–18

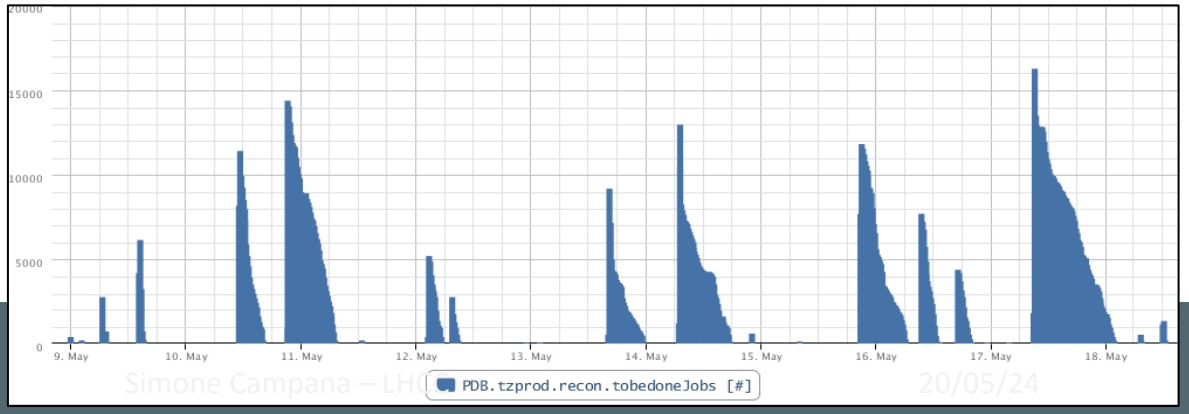

Reconstruction Software

Athena 20.7 offline software delivered in Q1 2016, will be used in 2016 and early 2017

2015 data and MC simulation (HITS) reconstructed with 20.7

Physics improvements: combined performance recommendations updated by respective groups

Same CPU/memory consumption of Athena 20.1 used in 2015

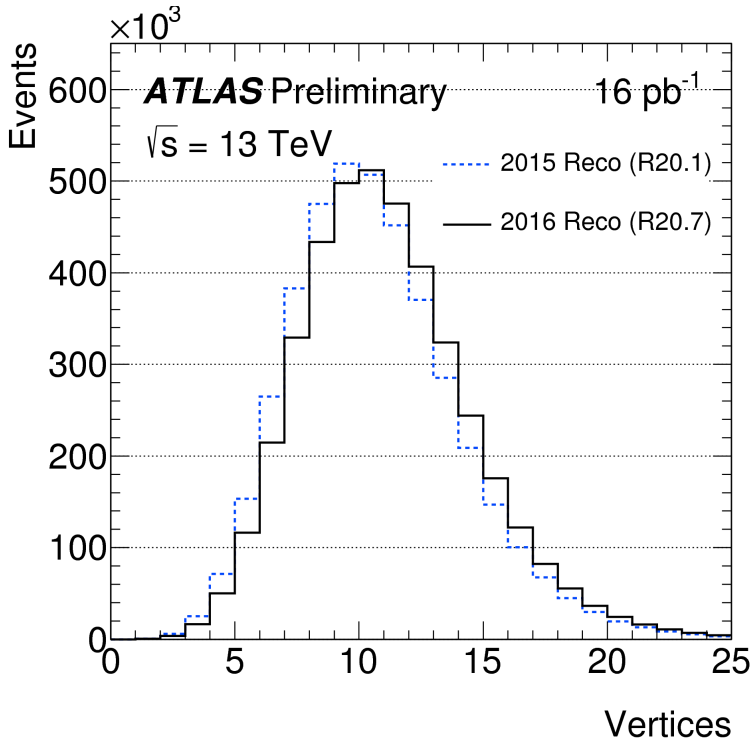

Comparison of the number of reconstructed vertices for the same run, using 2015 and 2016 software releases, showing improvement in efficiency for reconstructing pile-up vertices. The average number of pp interactions expected per beam crossing during this run (μ) was approximately 15.

MC and data processing (no HPCs)

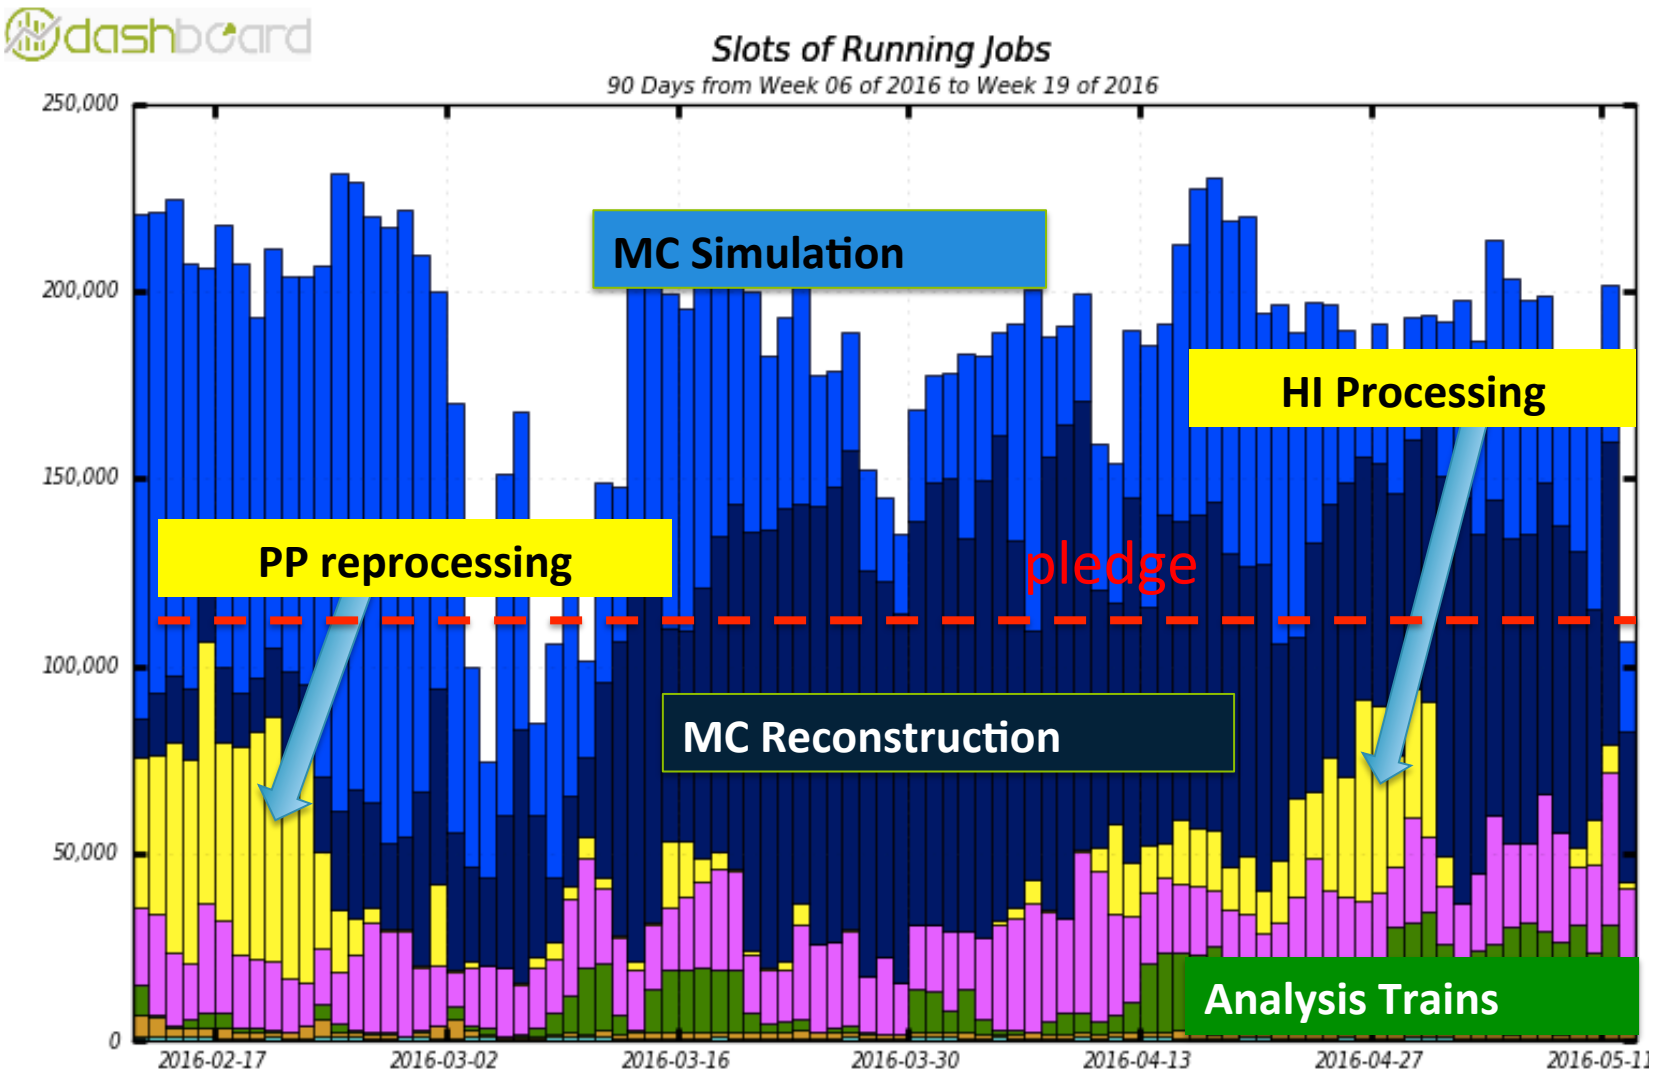

Heavy Ion processing

2015 HI software, selection cuts and workflow: 30GB PSS on 8 cores MP job and 15 seconds/ event

Decided not to process at T0 and improve performance

2016 HI processing: 24GB PSS on 8 cores MP job and 10 seconds/event

HI processing completed in 2 weeks on Grid. Data available to be presented in summer conferences

WallTime(s)/job

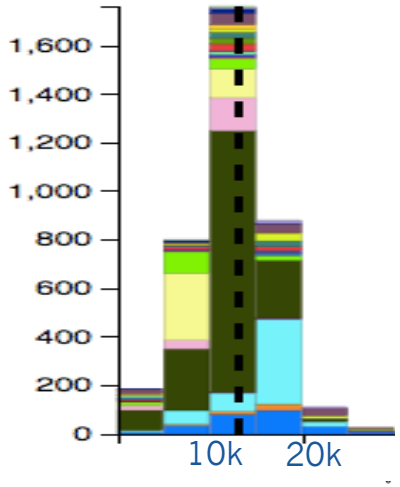

PSS(GB)/job

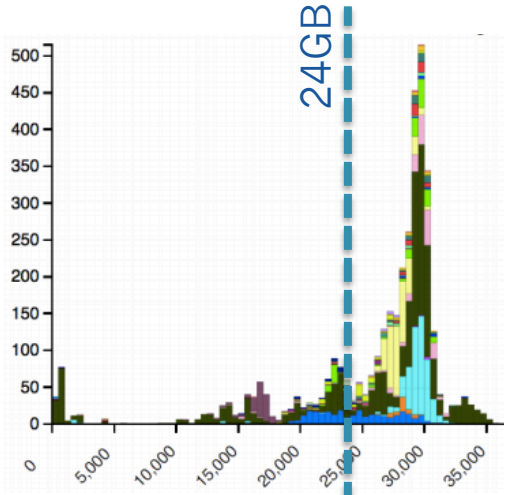

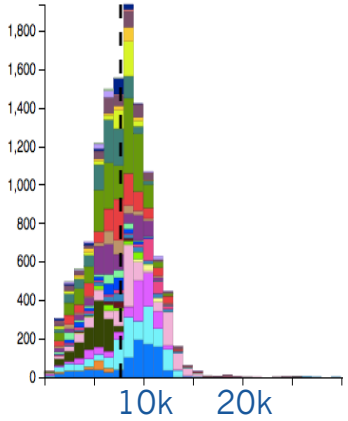

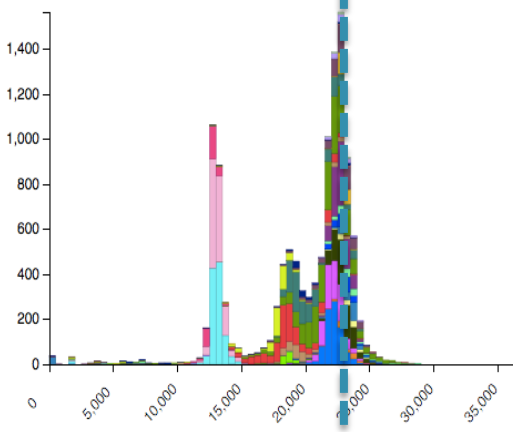

Non Grid Resources: HPCs

Grid and Cloud resources utilized at full capacity

- HPC resources are more bursty, depending on policies and allocations.
- Almost 400k++ running jobs in MIRA@Argonne for few days, played a crucial role in Event Generation for 2016 (ICHEP)
- HPCs will play an important role for ATLAS in the future, motivating the effort to integrate them

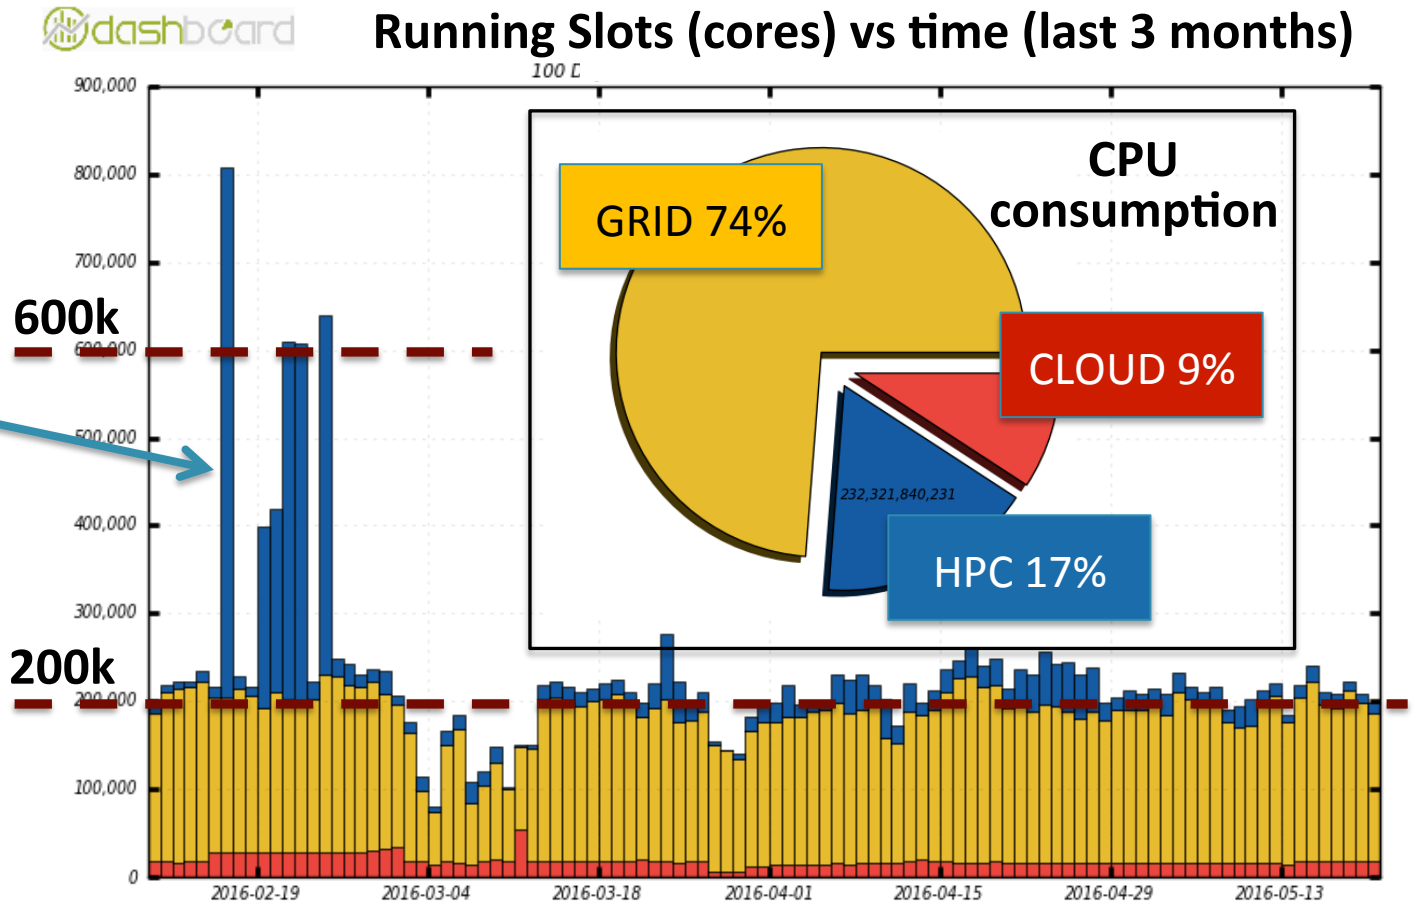

Non Grid Resources: Clouds

Reminder: we run MC production on AWS in September 2015. We used the ATLAS Event Service to optimize spot market resources utilization. The 40k running processed target was achieved and the AES benefits wrt "classical" workflow demonstrated

Trying to repeat the exercise at 100k scale we hit a limitation in VM provisioning (we use Condor). The Condor team is actively working on it and USATLAS (BNL) is actively working on a new procurement on AWS

We will repeat the 100k scale production in the summer

In contact with Google for a possible similar exercise but leveraging Google specific features (and API). Important to demonstrate we do not have a "vendor lock"

Meanwhile, we gathered more experience with AES on Grid resources and HLT. Its utilization is growing (staged rollout)

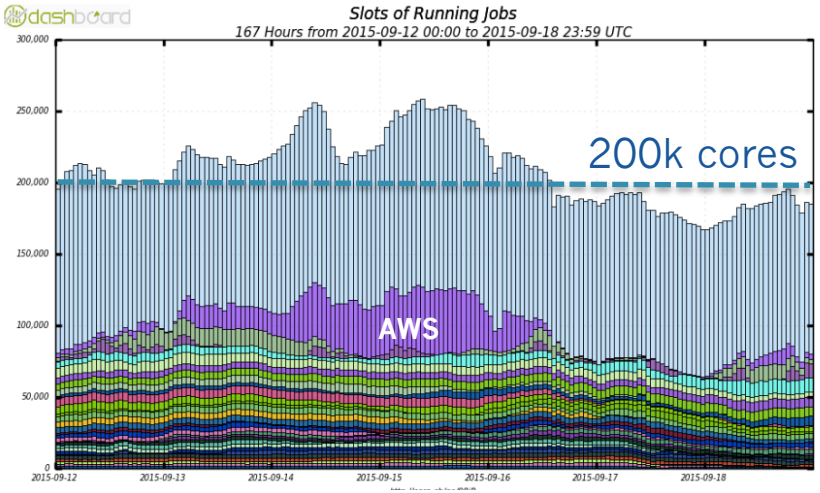

The Event Service 2016

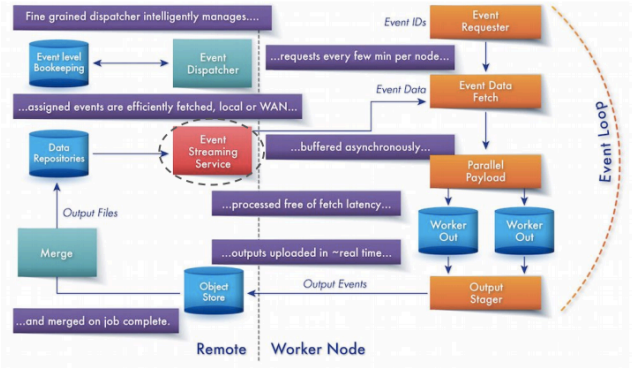

Event Index

A complete catalogue of ATLAS events (including MC)

- **Event Picking**
- Trigger counts
- Detection of data corruption
- Monitoring of event overlaps

Many improvements in the architecture as recommended by the 2015 review

- Consolidation and evolution of Data Production and Data Collection
- HDFS backend hosts the data and Hadoop ecosystem enables the client to consume them with very good performance
- Oracle backend hosts a subset of data in a schema optimized for data event picking. Excellent performance and very rich user interface

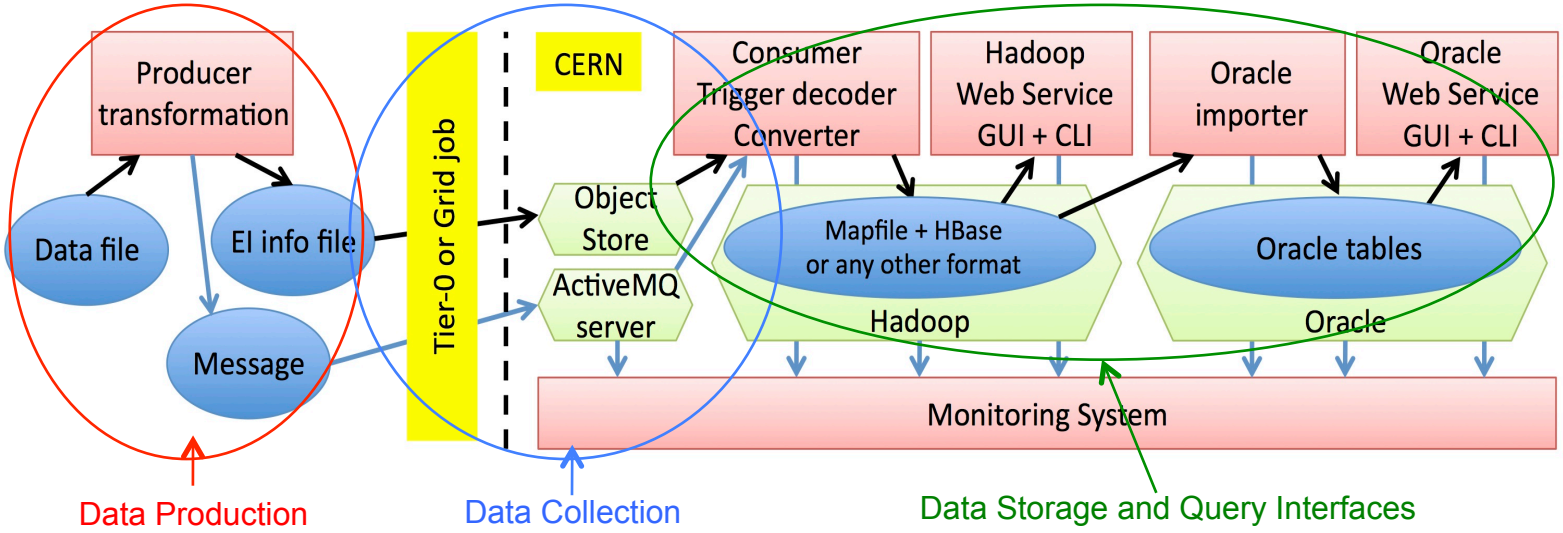

Software Infrastructure

- **Considerable progress after the review in Fall 2015**
  - Migration from CMT to CMake **accomplished:** -40% build time and margin to improve
  - RPM distribution of the ATLAS SW available
  - Soon nightly builds will be available in CVMFS and therefore on the Grid
  - Resources needed for Real Time Tests reduced by 50% while improving the quality of the tests
- **Milestones for 2016/2017**
  - Integration of Jenkins in the Build System
  - Migrate SW repository to Git

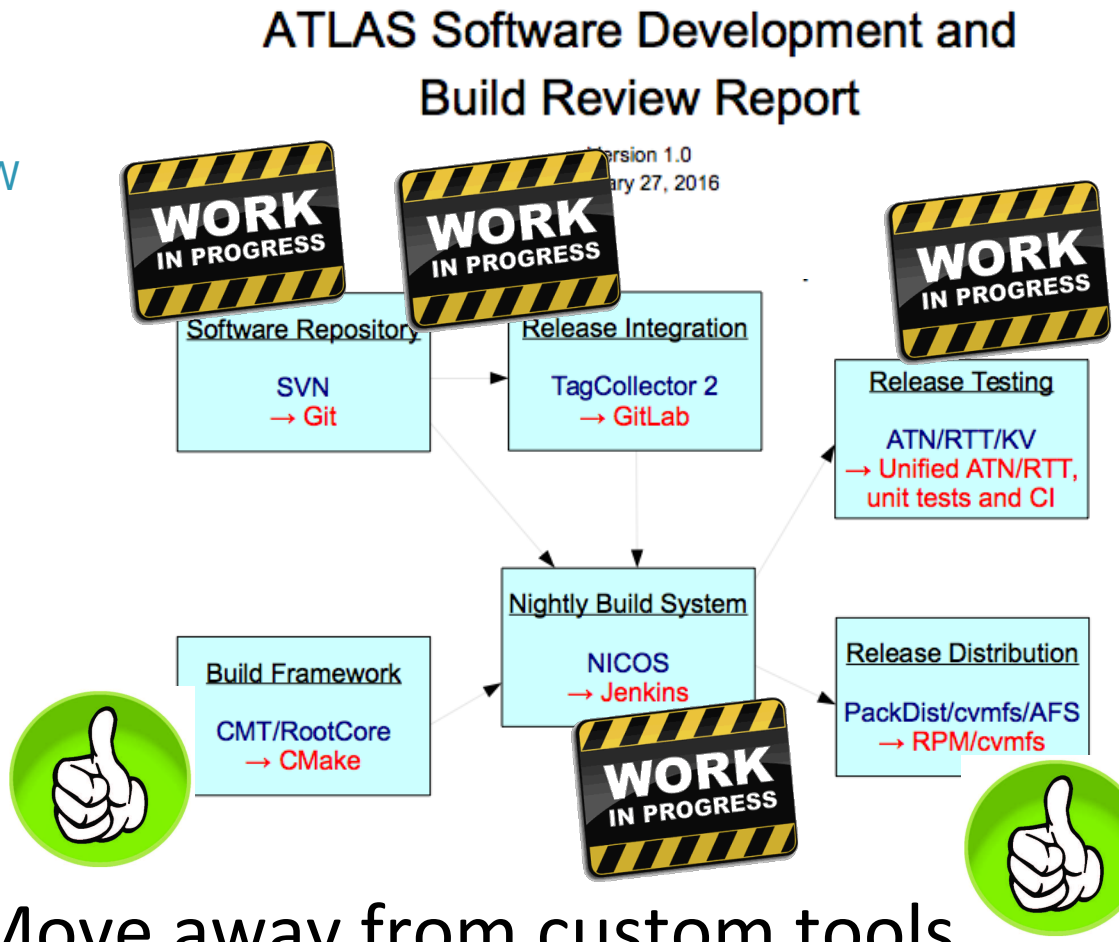

Move away from custom tools

Software: Infrastructure

- ATLAS simulation software now builds on Aarch64 (ARM)
- Physics Validation is ongoing but looks promising
- Next target is PowerPCs
- The process improves the overall quality of the ATLAS software
- A lightweight release for simulation was a good time investment

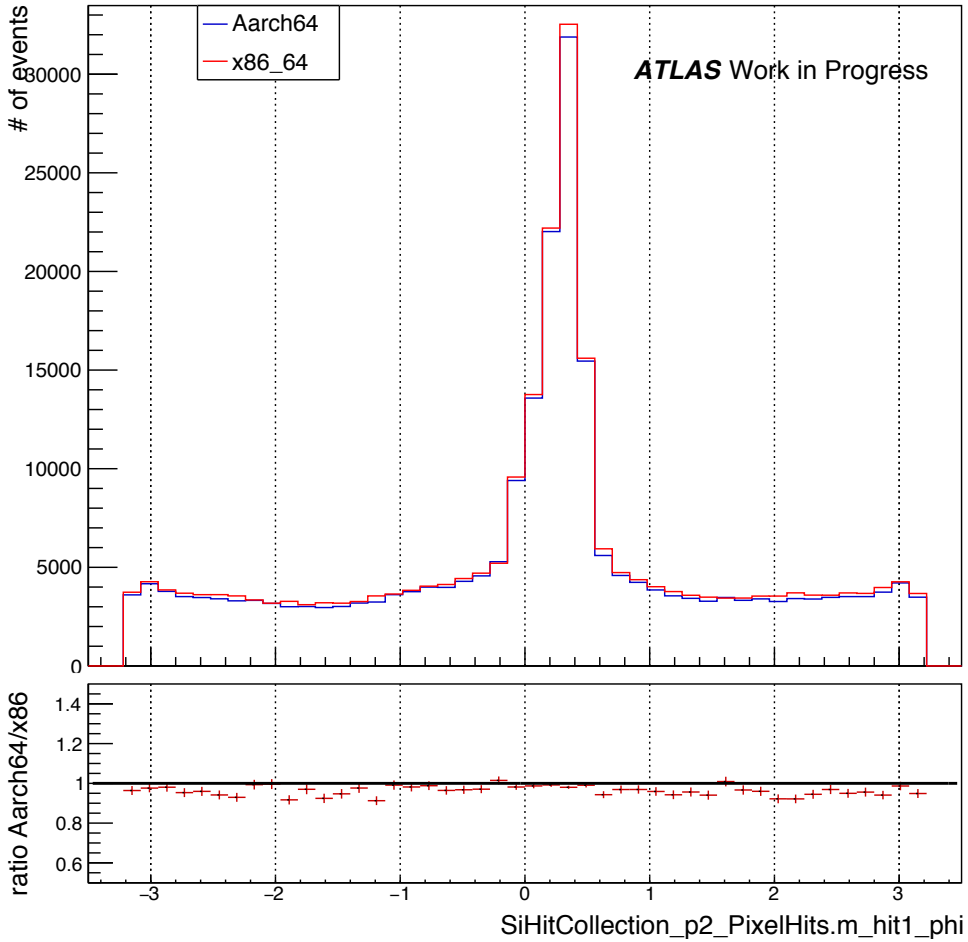

Software milestones for 2016/2017

- Simulation working on release for the MC16 campaign
  - Improved geometry based on 2015 experience (particularly in describing new elements/sub-detectors)
  - Based on Geant 4.10, which supports multithreading
  - Plan to launch MC16 in June
- Reconstruction software working on R21, to be used from 2017 until LS2
  - Plan to validate the software by Fall 2016
- AthenaMT (multithreading) being actively developed
  - MT version of Gaudi provides the foundation: originated at CERN and is widely used but needs more and long term support
  - Framework planned to be ready by 2018, still many challenges to face

Initial studies on Computing for HL-LHC

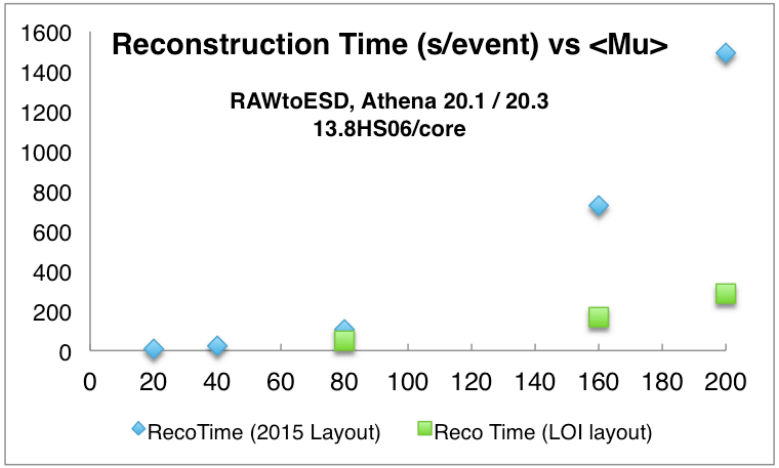

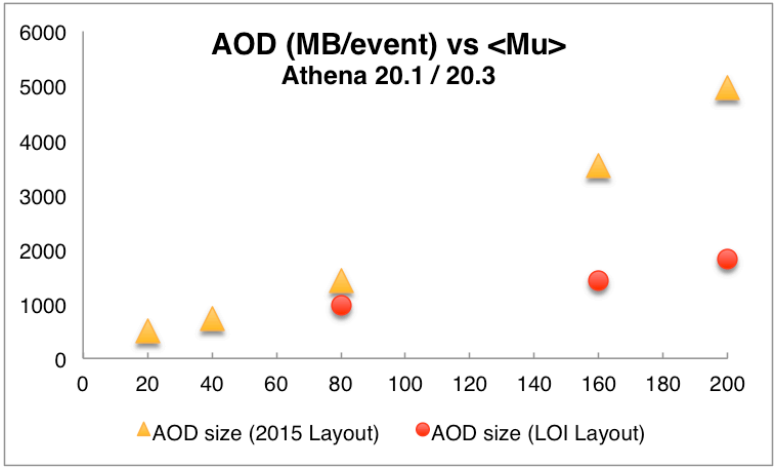

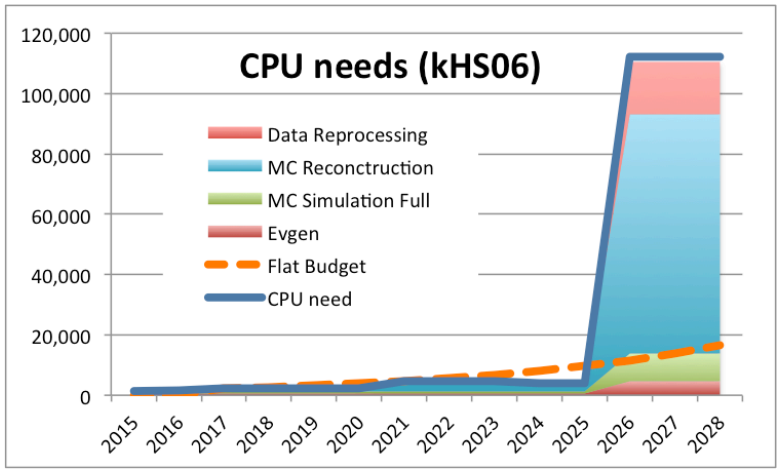

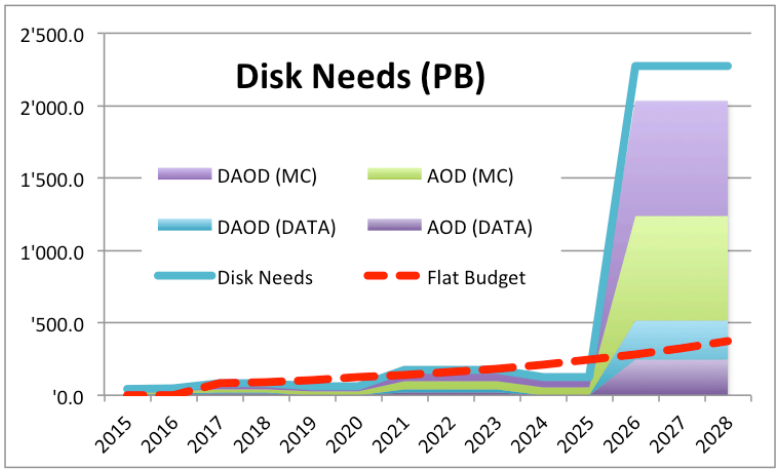

Computing for HL-LHC: disclaimers

- The plots are based on the ATLAS layout in the Upgrade Phase 2 Letter of Interest. Could change considerably with the TDR layout
- Plots are based on today's computing model: number of replicas, ration data/MC, ration Full/Fast simulation.
- We consider 10kHz of HLT output for Run4, while in the LOI we mention 5 to 10kHz
- We consider the legacy from Run1/2/3 in terms of data volume negligible for Run4 and we do not include it
- In conclusion, what presented are ballparks and not absolute numbers

Conclusions

- We are supporting Run-2 data taking and physics while evolving the Software and Computing systems
- Constantly utilizing our pledged resources and leveraging heterogeneous opportunistic ones
- Scenario for Run-2 and Run-3 is rather clear now, we started looking at Run-4 as well, in concert with WLCG and the other experiments
- The Long Term Planning document together with ongoing focused reviews guide the evolution
- Substantial improvements in software performance and capability: new efficient workflows, new platforms, better physics performance and resource utilization