


# HLT plans 2017

M.Krzewicki, D.Rohr for the HLT  
FIAS, CERN

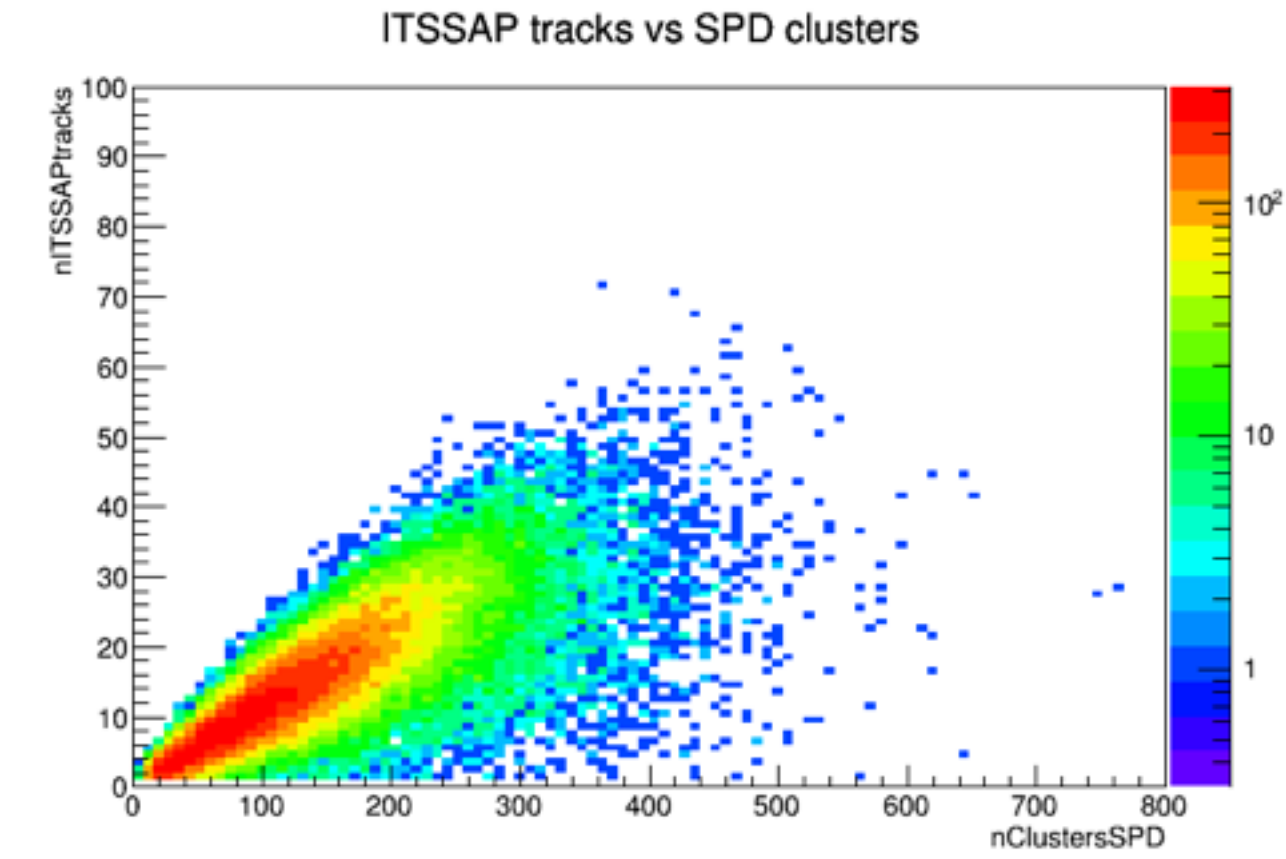
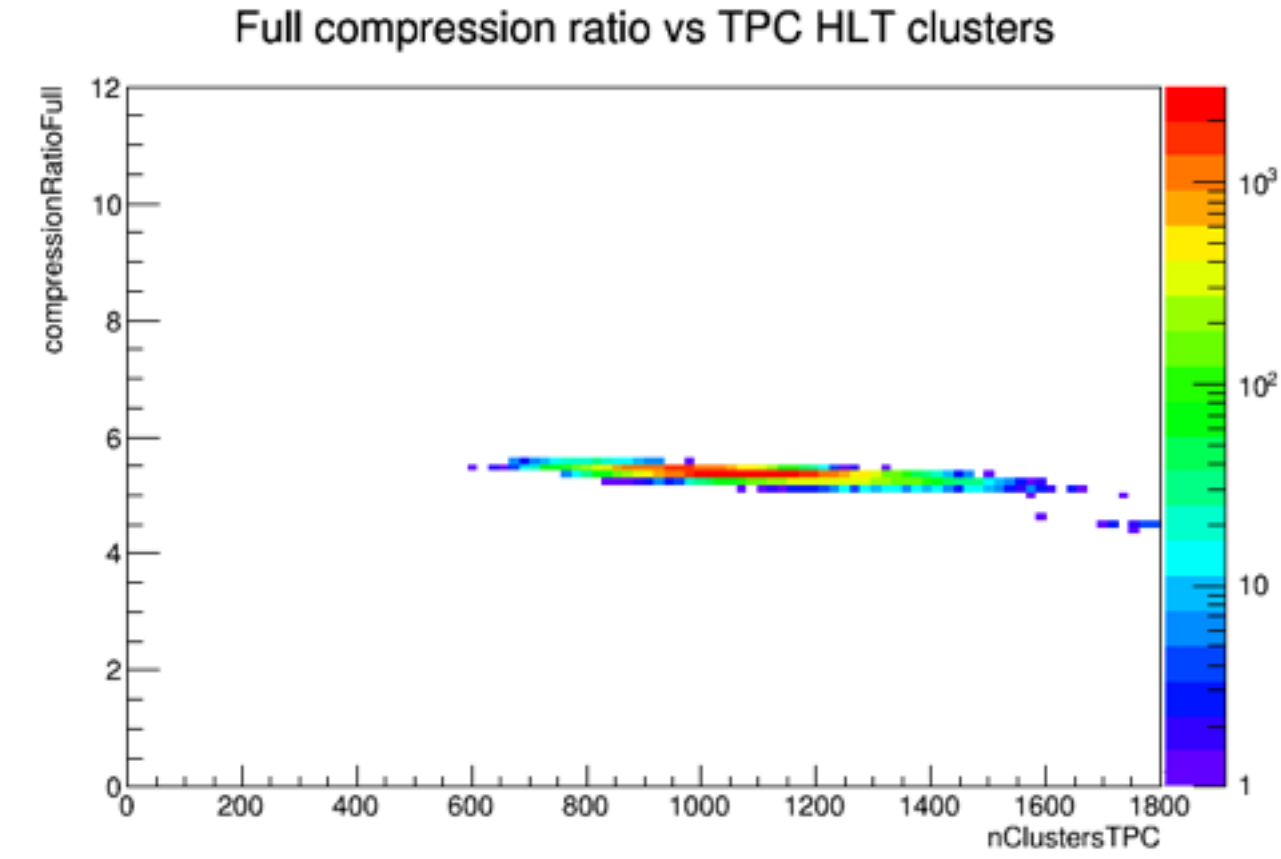
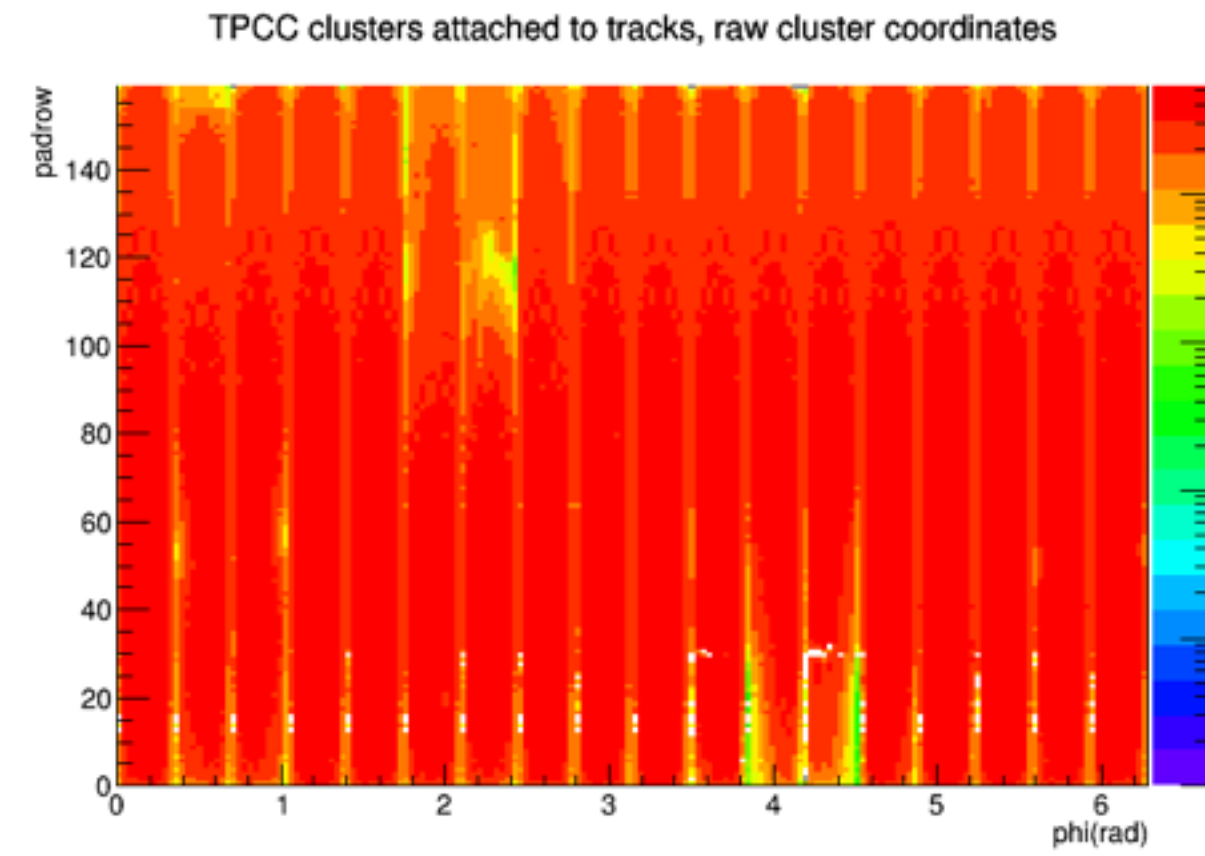
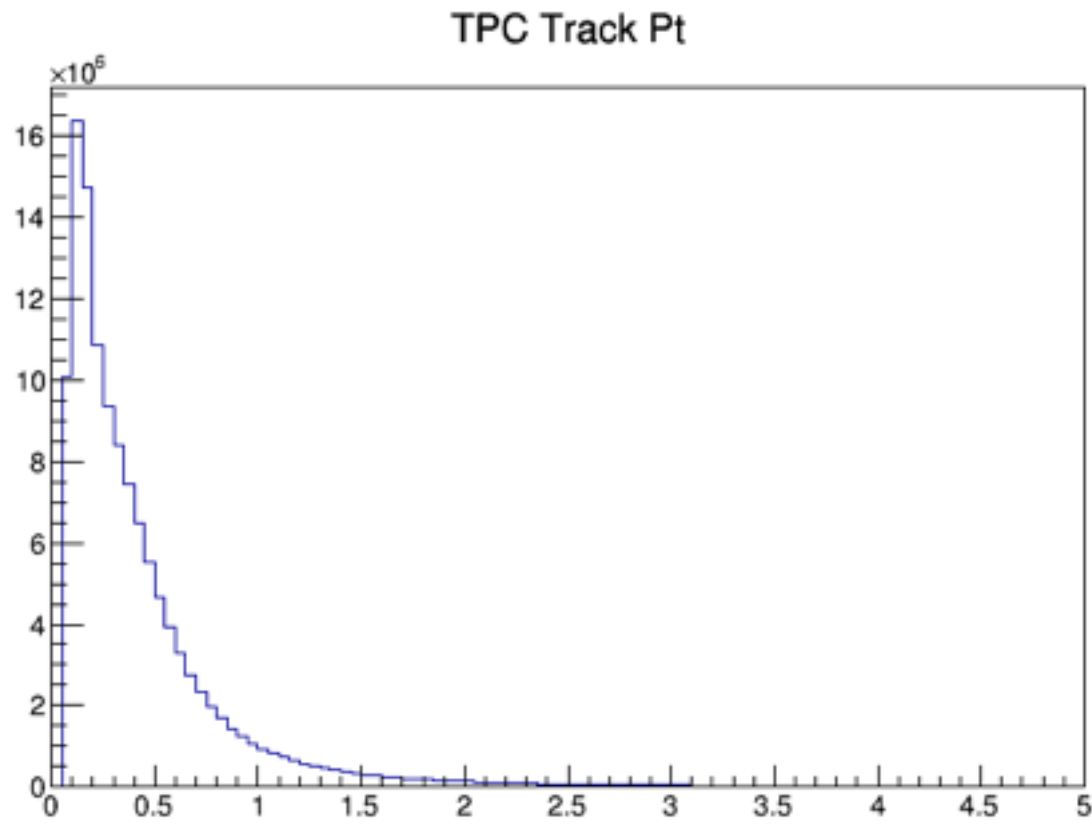
# Compression status 2016

Full compression ratio vs TPC HLT clusters

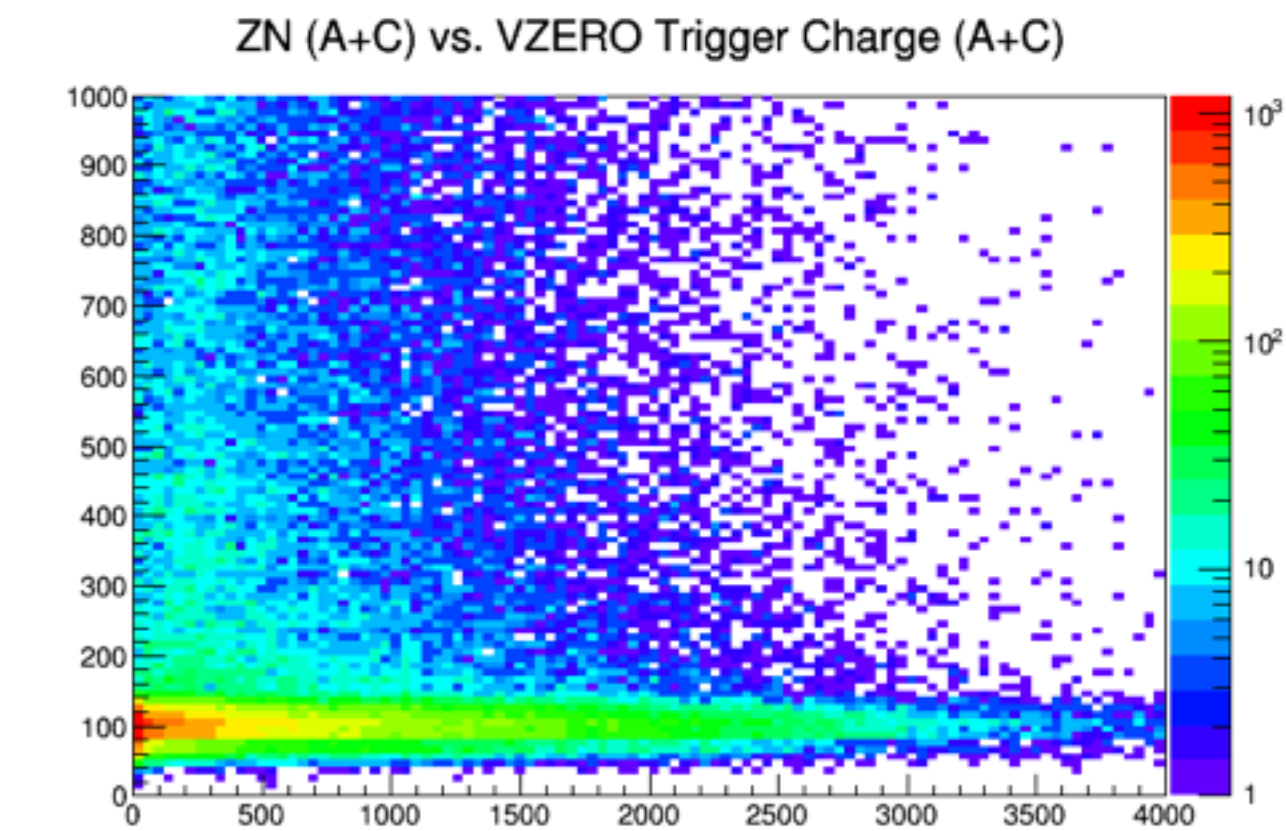


- Online TPC cluster compression: improved from factor  $\sim 4.3$  to  $\sim 5.5$ .
- Differential Huffman compression, tuned to 2016 data conditions.
- 20% more efficient raw data storage.
- 2017: reject noise clusters, better handling split clusters - later slides.
- Ongoing effort, compression studies important for run 3 upgrade.
- Under study (both run 2 and run 3): track model compression, smarter cluster charge encoding, junk removal.

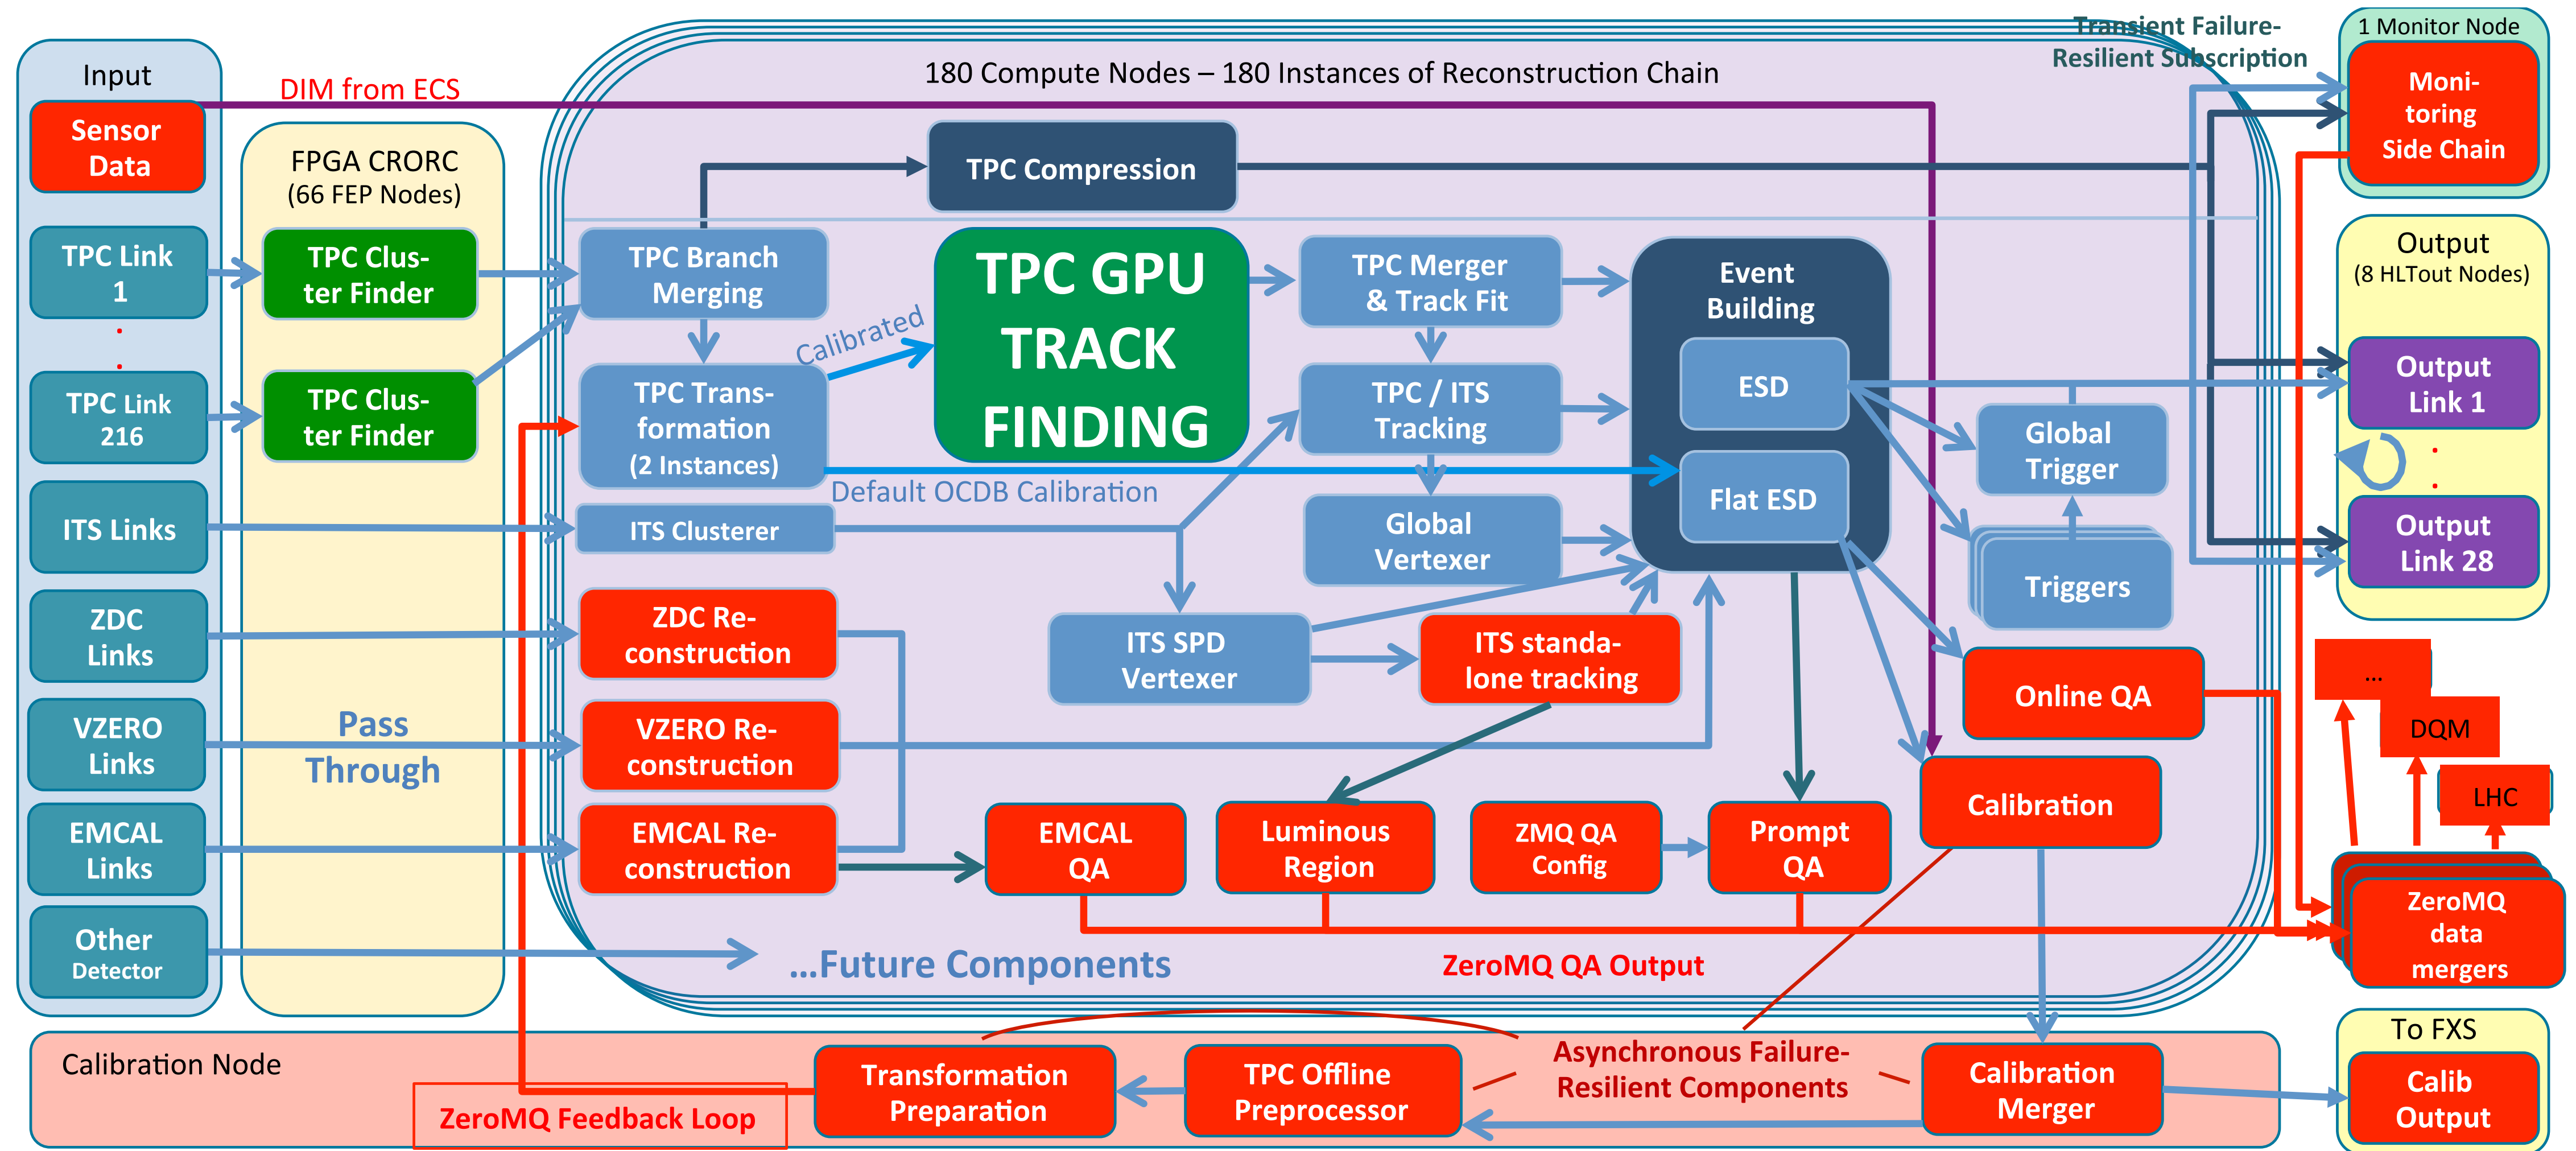
# Online physics QA and monitoring



- Utilising the online reconstruction of many detectors, a new monitoring scheme was developed to allow real-time monitoring of the physics performance of the ALICE detector.
- Includes slow out-of-chain, fast in-chain, and asynchronous running of offline QA and physics analysis code.
- Simple external interface.
  - Data, metadata, ROOT streamers etc. easily added to a single ZeroMQ message efficiently for use e.g. off-site.
- TPC Offline QA code ported to support online operation. Uses the same infrastructure developed for online calibration.
- TPC, HLT and EMCAL QA available in real-time online on OVERWATCH:
  - <https://aliceoverwatch.physics.yale.edu/monitoring>




# HLT functionality update



- HLT framework supports unidirectional, data synchronous flow only.
- Asynchronous processing, calibration feedback loop, out-of-chain processes (QA, LHC monitoring, etc.).
  - ZeroMQ used as additional transport supplementing the native HLT flow, providing the feedback loop and asynchronous capabilities.

# HLT production cluster as a WLCG site for ALICE jobs

- HLT production cluster running as a WLCG site during EYETS
- Improved setup using Docker containers instead of OpenStack VMs
- Containers are spawned for one job and destroyed afterwards
- Each container uses 1 CPU core and can use up to 5GB RAM
- 4244 containers running in parallel (24 per server)
- Performance last week: 122k ALICE jobs finished successfully only 1252 jobs failed (~1%)
- Job efficiency above 90 % - MC jobs.
- (CPU time / wall time)
- Total wall time contributed for ALICE jobs:
- 475k kSI2K hours in one day



J.Lehrbach, A.Gomez, D.Berzano, M.Concas

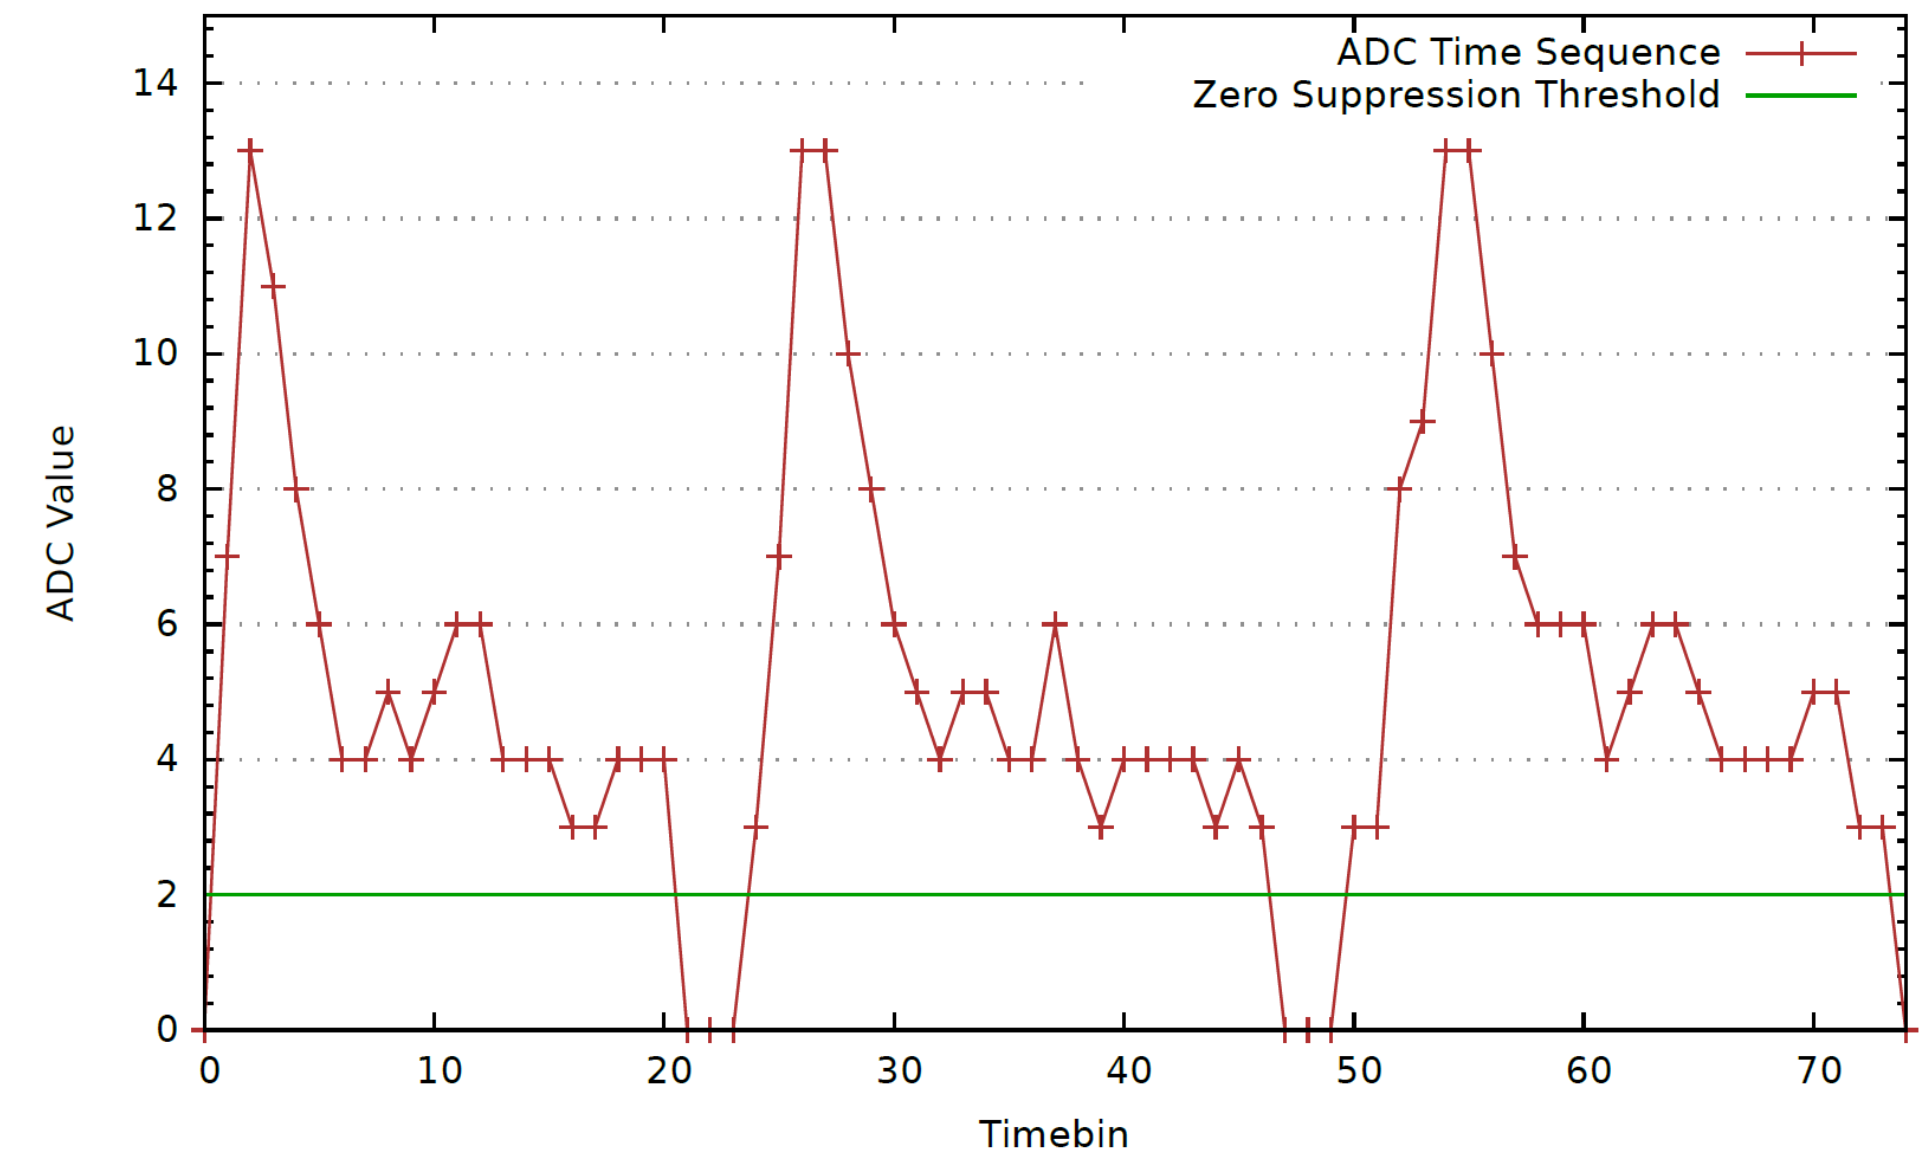
- What we have:
  - EMCAL, VZERO, ZDC, ITS, TPC online reconstruction
  - TPC real-time drift velocity calibration.
- New online TRD tracker (Ole Schmidt et al.)
  - Code ready, under investigation.
  - Intergration in the HLT framework OK, disabled by default.
  - Improve online QA.
  - Possibility to start exercising new online TPC calibration procedures (with O<sub>2</sub> in mind).
- TOF reconstruction and monitoring.
  - work in progress.
- T0 calibration monitoring.
  - work in progress
  
- More focus on Run 3 (O<sub>2</sub>).

# Update on TPC Cluster Finding / Compression in the HLT

- HLT is currently evaluating improvements to the cluster finding and data compression for 2017.
- These studies are anyway mandatory, in order to achieve the desired compression for Run 3!
  
- The new features include:
  - Improvements of the HLT TPC Cluster Finder to reject increased amount of noise seen in 2016.
  - Improvements of the Compression algorithm
    - Includes usage of track model compression instead of differential compression, data format improvements, arithmetic encoding.
  - Development and validation of the compression is much easier, because it is pure software and lossless.
    - Like ZIP, decompressed data is bit level identical (effect on physics performance need not be checked), error is seen by checksum check, several TB can be checked easily in software.
- These new developments bring additional benefits:
  - Split and edge cluster flags available, to be used for improved TPC dE/dx calculation later on.
  - HLT tracks can be used as seeds for offline, reducing memory footprint and compute time.

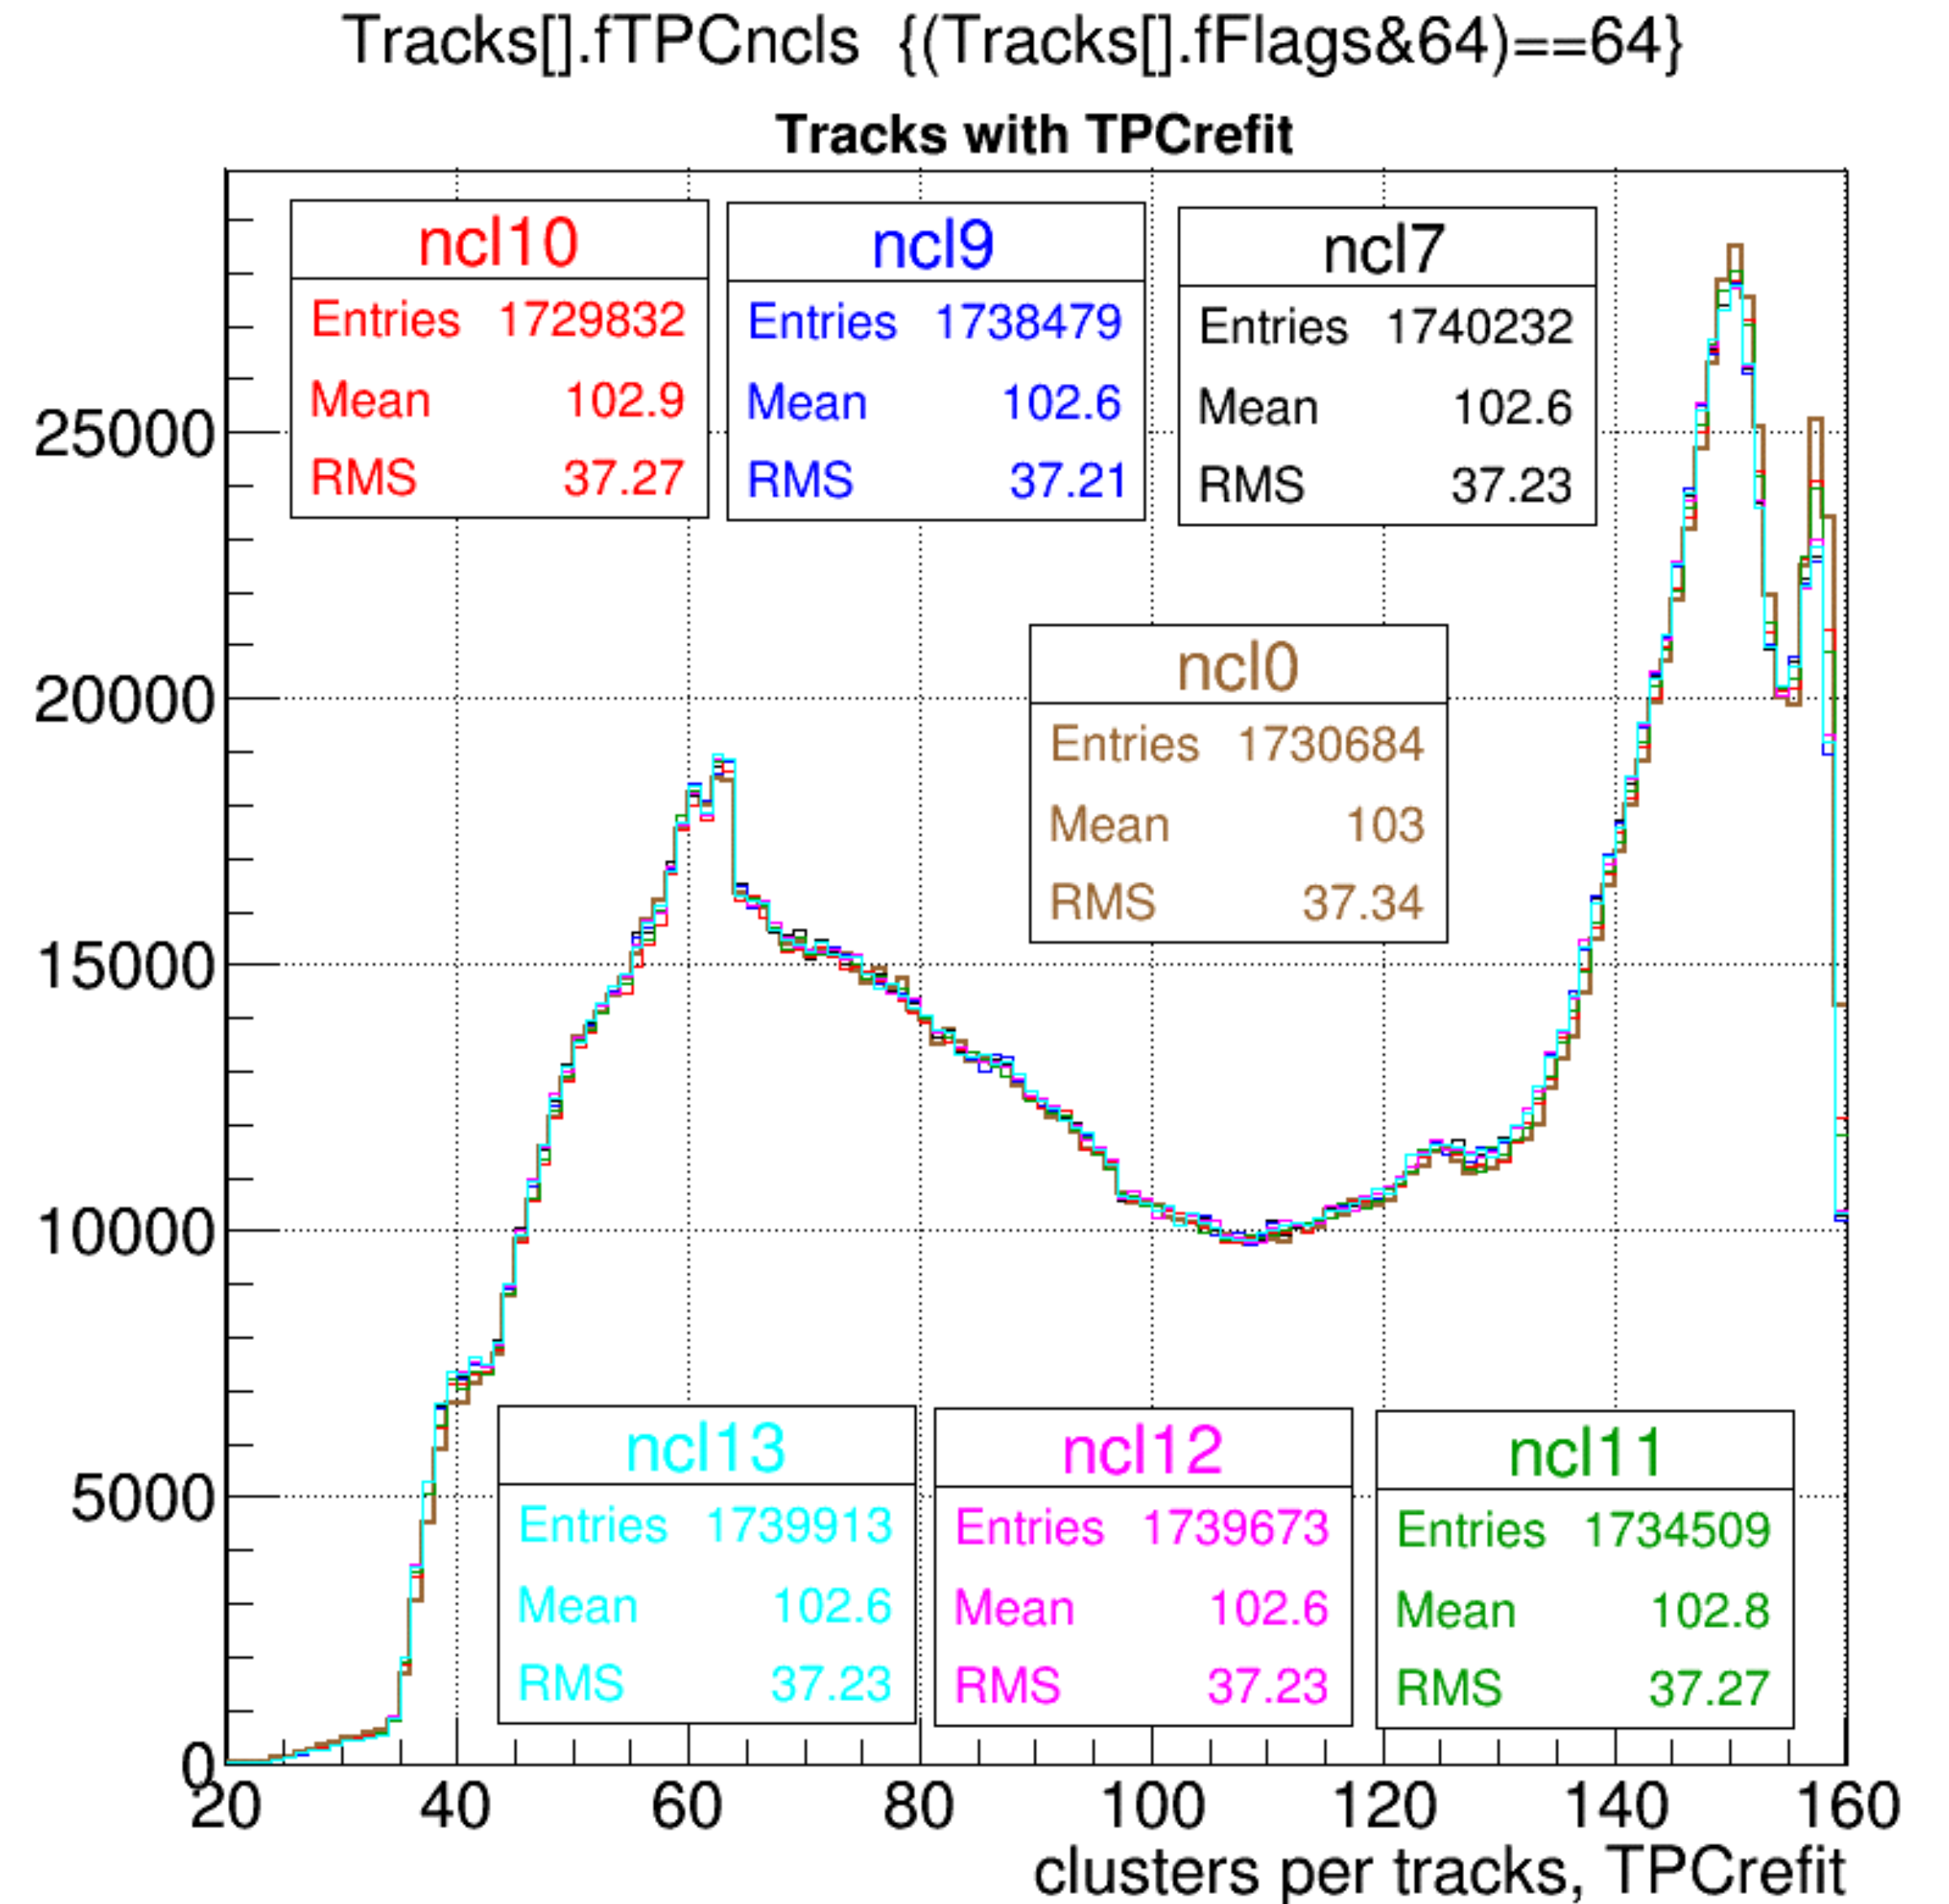
# Noise Cluster Rejection during Cluster Finding

- Noise in cluster tails above zero-suppression threshold poses a problem for HLT cluster finder.
- We are evaluating an improved more robust cluster finding algorithm to reject the noise.
- Evolution rather than revolution.
- Improve the peak detection heuristic.
- Goal:
  - Reject noise clusters.
  - Maintain the current physics performance.
  - (In fact we are rejecting clusters that should be ignored by the tracking anyway).
- Performance could improve slightly, because noise clusters could disturb tracking.
- Possibly slightly better cluster resolution by improved heuristics for corner cases.
- Problem more pronounced since 2015 with new gas, but already present during run 1.





- Evaluation is ongoing...
- Checked with 2015 low intensity RCU1 data and 2016 high intensity RCU2, similar results, both OK.
- With the new version, we obtain similar cluster per tracks statistics at **35%** reduced number of clusters (fakes rejected) and even slightly improved  $\chi^2$ .
  - 20-25% improved compression (noisy clusters compressed well).
- Work in progress.
  - (<https://alice.its.cern.ch/jira/browse/ATO-366>)



# Study on total compression ratios

- *A mockup test in idealized conditions* with all new features achieved a maximum compression ratio of 9.1 for both pp and Pb-Pb (with limited statistics, average in 2016: 5.5, for Pb-Pb 5.8)
- Contributing factors:
  - **New Cluster Finder (20 – 25%, we are confident to have it ready).**
  - Track model compression and format optimizations (ongoing).
  - Arithmetic encoding (only ~4%, won't be ready, partially done by DAQ ROOT compression.)
- The individual compression factors from these contributions do not multiply directly.
- E.g., changes to the cluster finder change the entropy affecting Huffman compression.
- Work in progress, actual performance to be seen, mostly Run3 related.


- Do we need to rerun the HLT mode C validation for current HWCF?
  - With respect to the gas change:
    - The cluster finder was running with the new (old) gas already in 2011 – 2013.
    - We assume less distortions.
    - The fake cluster problem from 2015 / 2016 is very unlikely to increase going back to the old gas.
  - From the HLT side, there is technically no reason to rerun the validation.
    - We can still do it to double-check.
- The new cluster finder algorithm DOES need to be validated.
  - If the validation of the new cluster finder takes long, we can run in mode C with the old cluster finder
    - Fall-back solution to avoid the situation from last year.
- The primary problem in 2016 was: we could not run the validation in the low-intensity ramp-up phase at the beginning due to problems with RCU2 data ordering → Delayed until summer.
  - This is fixed → We can start the validation as soon as we have data.

# Summary cluster finder and compression developments

- New cluster finder implemented in software – Study and prototyping done in software.
- FPGA implementation ongoing.
- Data size:
  - ~35% reduction in number of clusters compared to old HLT, ~8% reduction compared to offline.
  - Data size (after compression) reduced by 20% - 25%, speedup during tracking due to less clusters.
- Performance:
  - No significant difference in Nclusters/track, dE/dx (MIP resol., PID sep.),  $\chi^2$  (indeed minimally better)
  - Slight improvement in TPC ITS matching at sector edge, might be able to improve further.
- Additional benefits:
  - Split/edge cluster flag available for improving dE/dx calculation.
  - HLT tracks available as seeds for offline (if new data compression scheme used).
- Revalidate of mode C at start of 2017 data taking, use old clusterer if data size grows too much.

backup

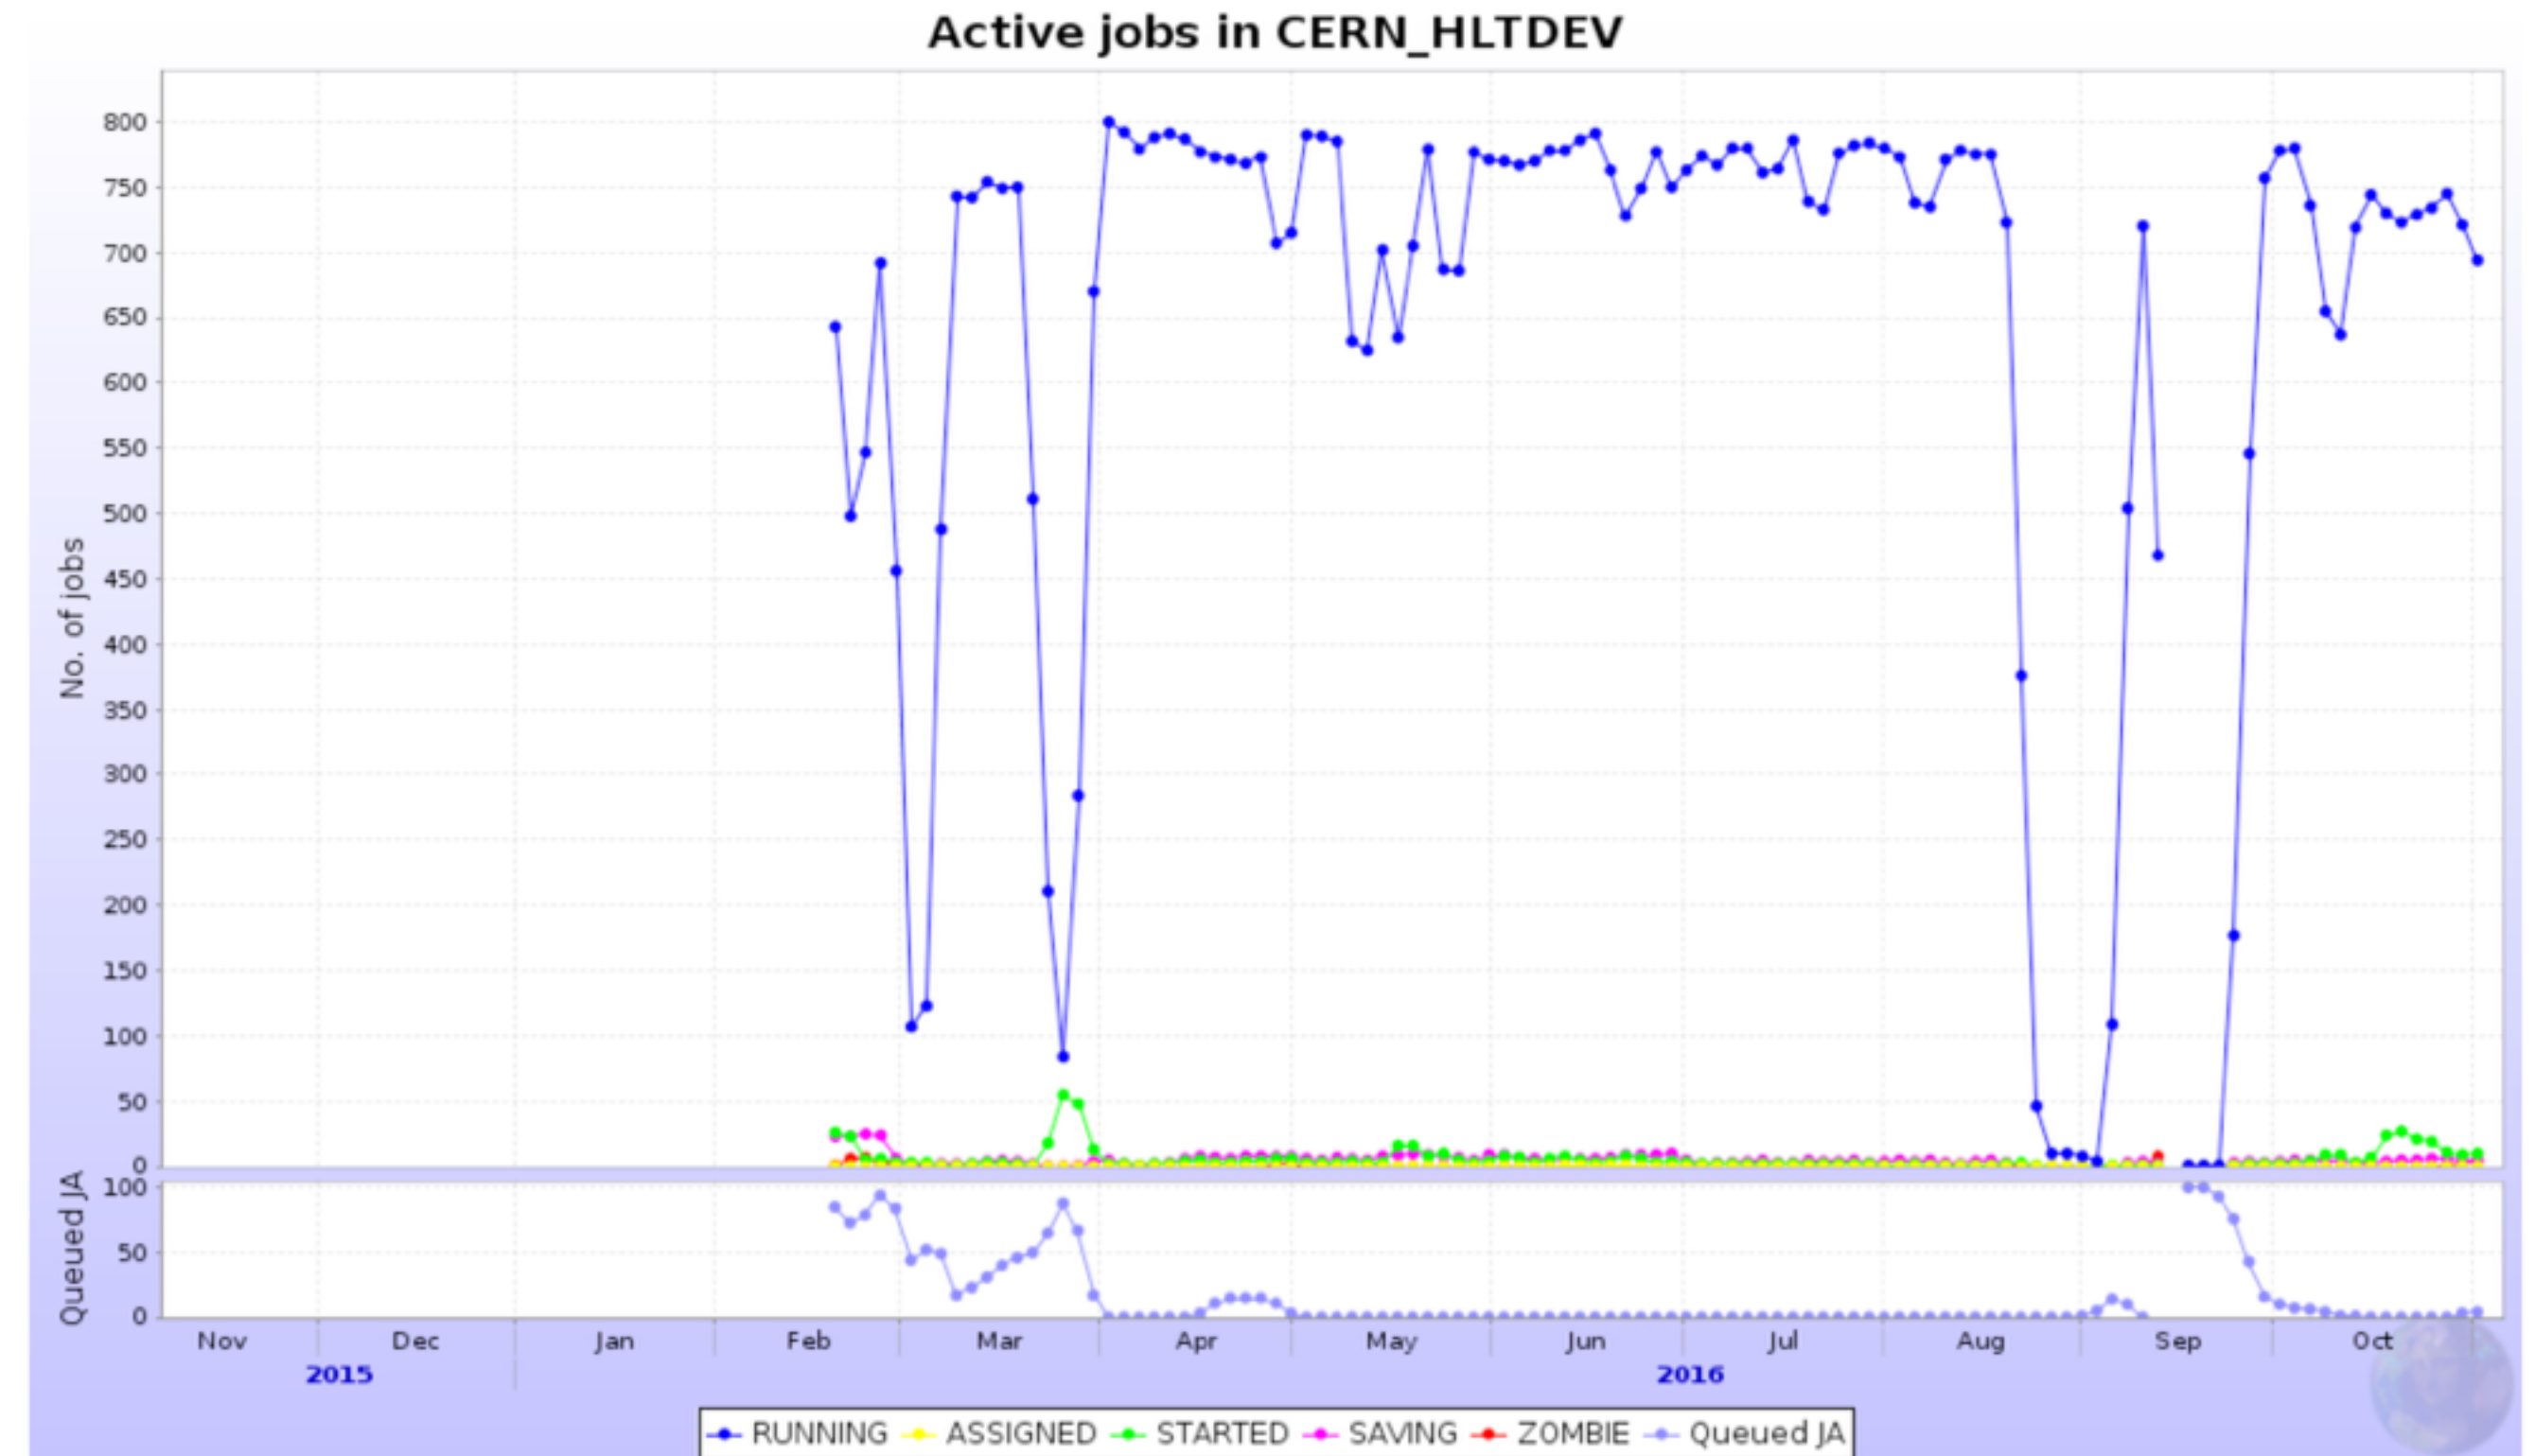
Total wall time kSI2K hours for ALICE jobs in 2016



- The spare compute resource, the development cluster consisting of older HLT infrastructure is run as a tier-2 GRID site using an Openstack-based setup, contributing as many resources as feasible depending on the data taking conditions.
- In periods of inactivity during shutdowns, also the production cluster is used.
- only MC jobs.
- 650k jobs done, ~2.5% of MC load (as of september).

# Development cluster

- Used as a GRID site continuously.
- MC jobs, few hundred running all the time since February.



- used for O<sub>2</sub> development:
  - nodes equipped with GPU.
  - InfiniBand network.
  - All necessary software installed (moving target...).
  - In use by many groups/people at request.

- PHYSICS runs with beam with HLT in (not mode A): 1392.
- runs stopped due to HLT problems: ~2% of total.
- Problem list (as of 2016-11-07)
  - GPU driver issue (3 runs, solved by upgrading the kernel driver).
  - 1 broken GPU (2 runs, replaced).
  - 1 broken CPU (2 runs, replaced).
  - Generic hardware failure of one node (MCE) - (3 runs, node excluded).
  - Online TPC reconstruction code problem (1 run, fixed).
  - Network communication problem (1 run).
  - GPU processing stuck (partial software workaround found, rate lowered).