

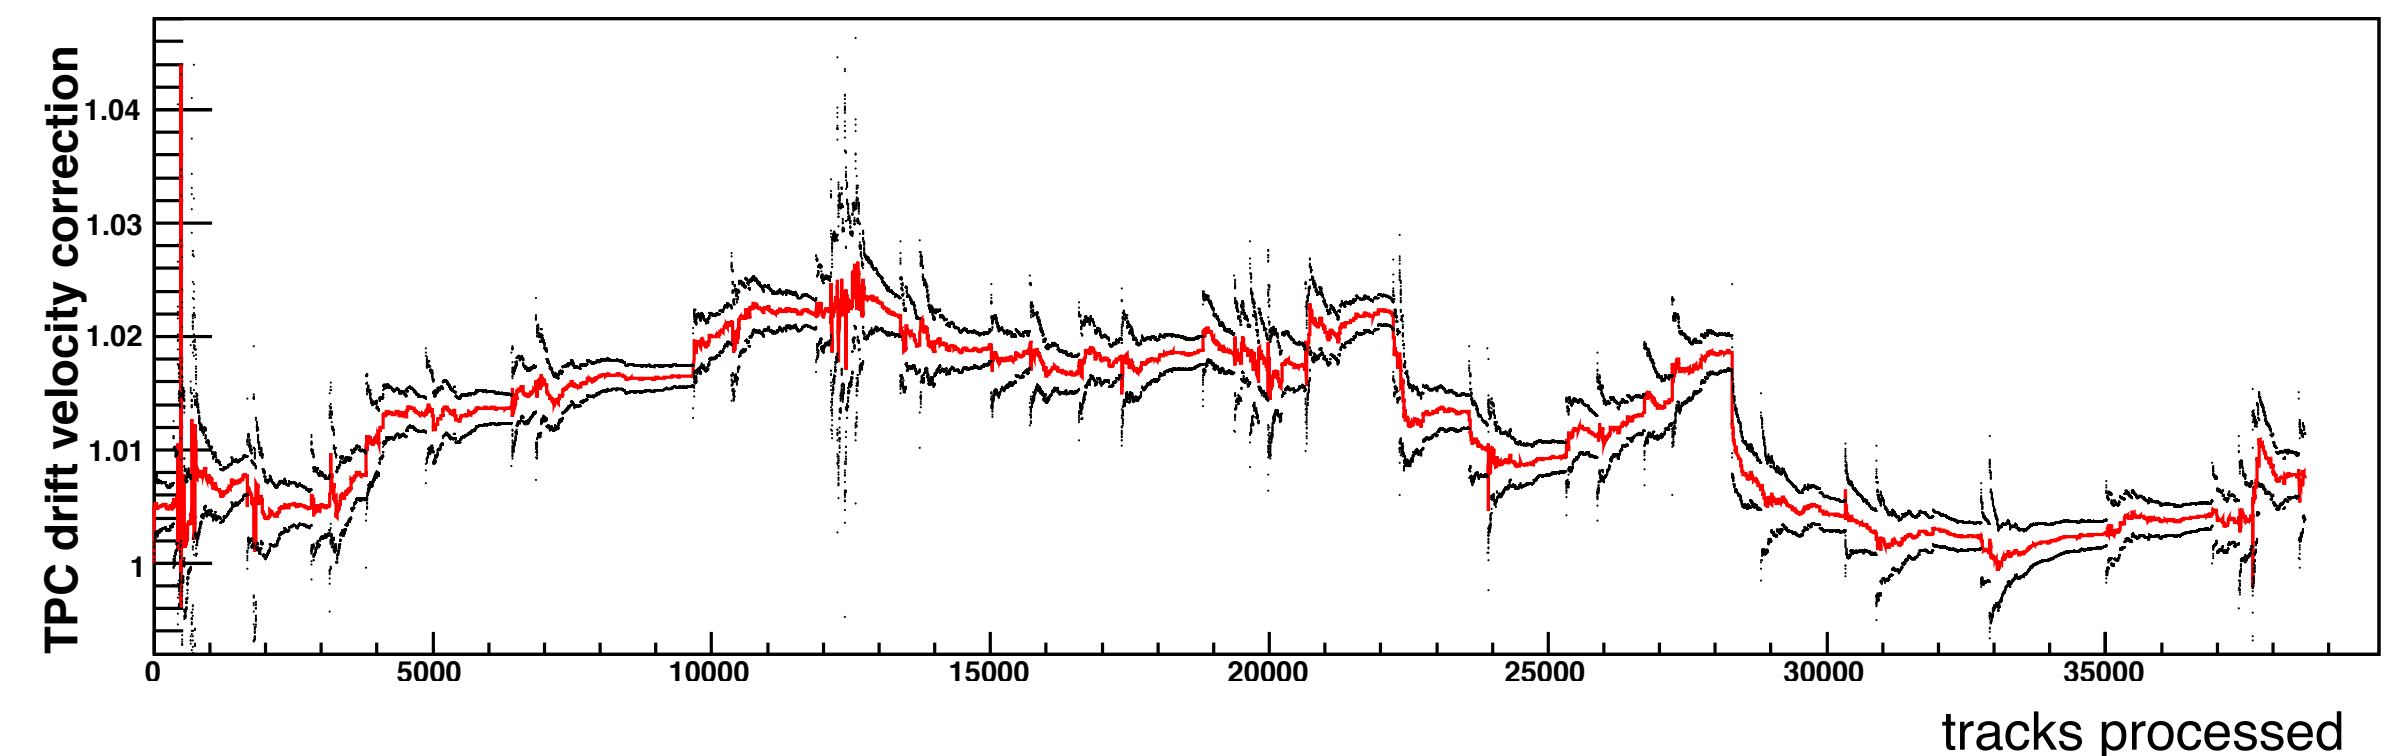
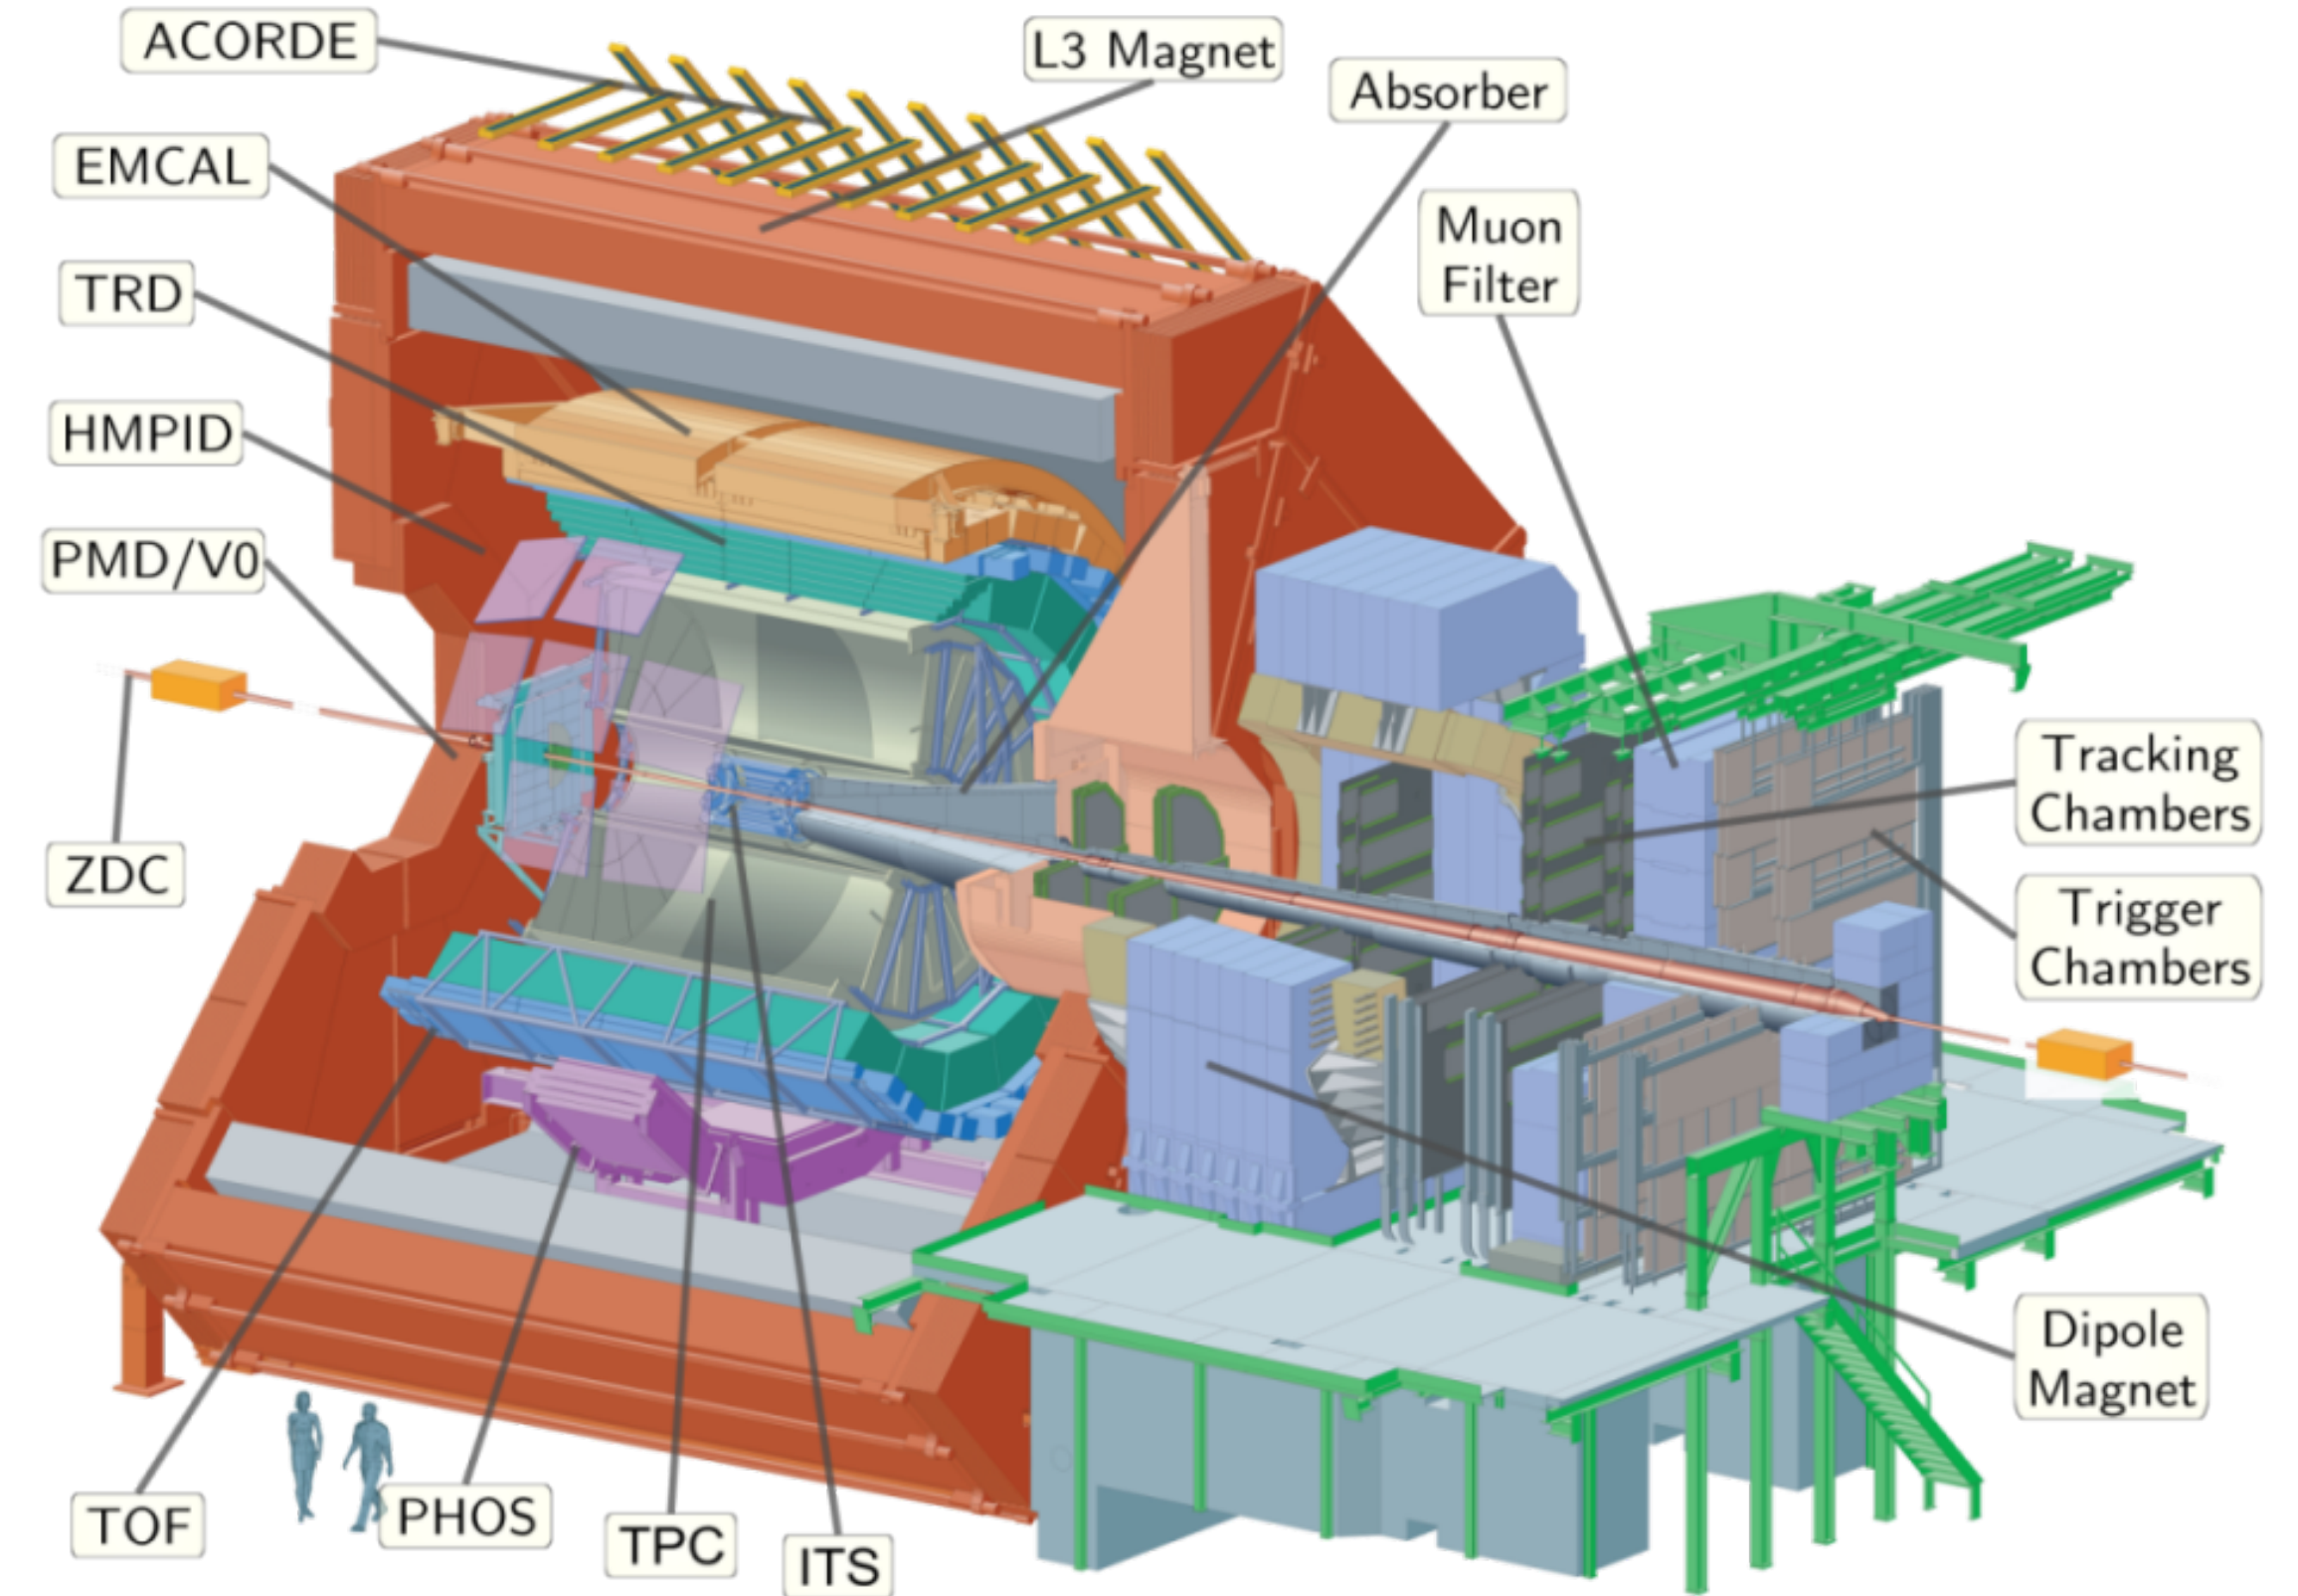


# Support for Online Calibration in the ALICE HLT Framework

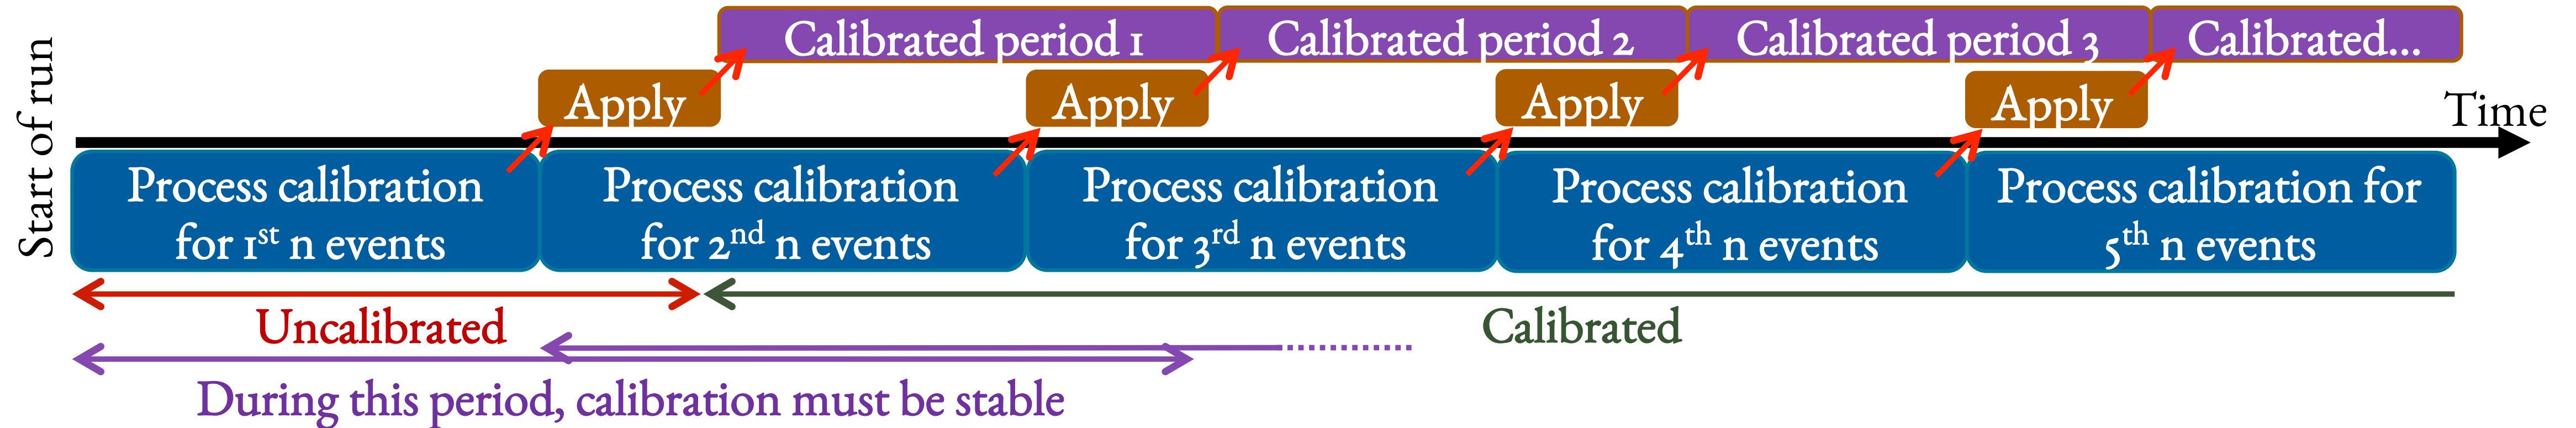
M.Krzewicki for the ALICE collaboration

# Online calibration

- Main tracker (TPC) is a gas drift chamber.
- Need exact drift velocity calibration to relate measurement (time) to position (z-coordinate).
- (Online) tracking quality depends heavily on calibration.
- Track finding (largely) unaffected.
- Matching to inner/outer detectors affected.
- Electron drift velocity is time dependent due to environmental conditions change (atmospheric pressure).



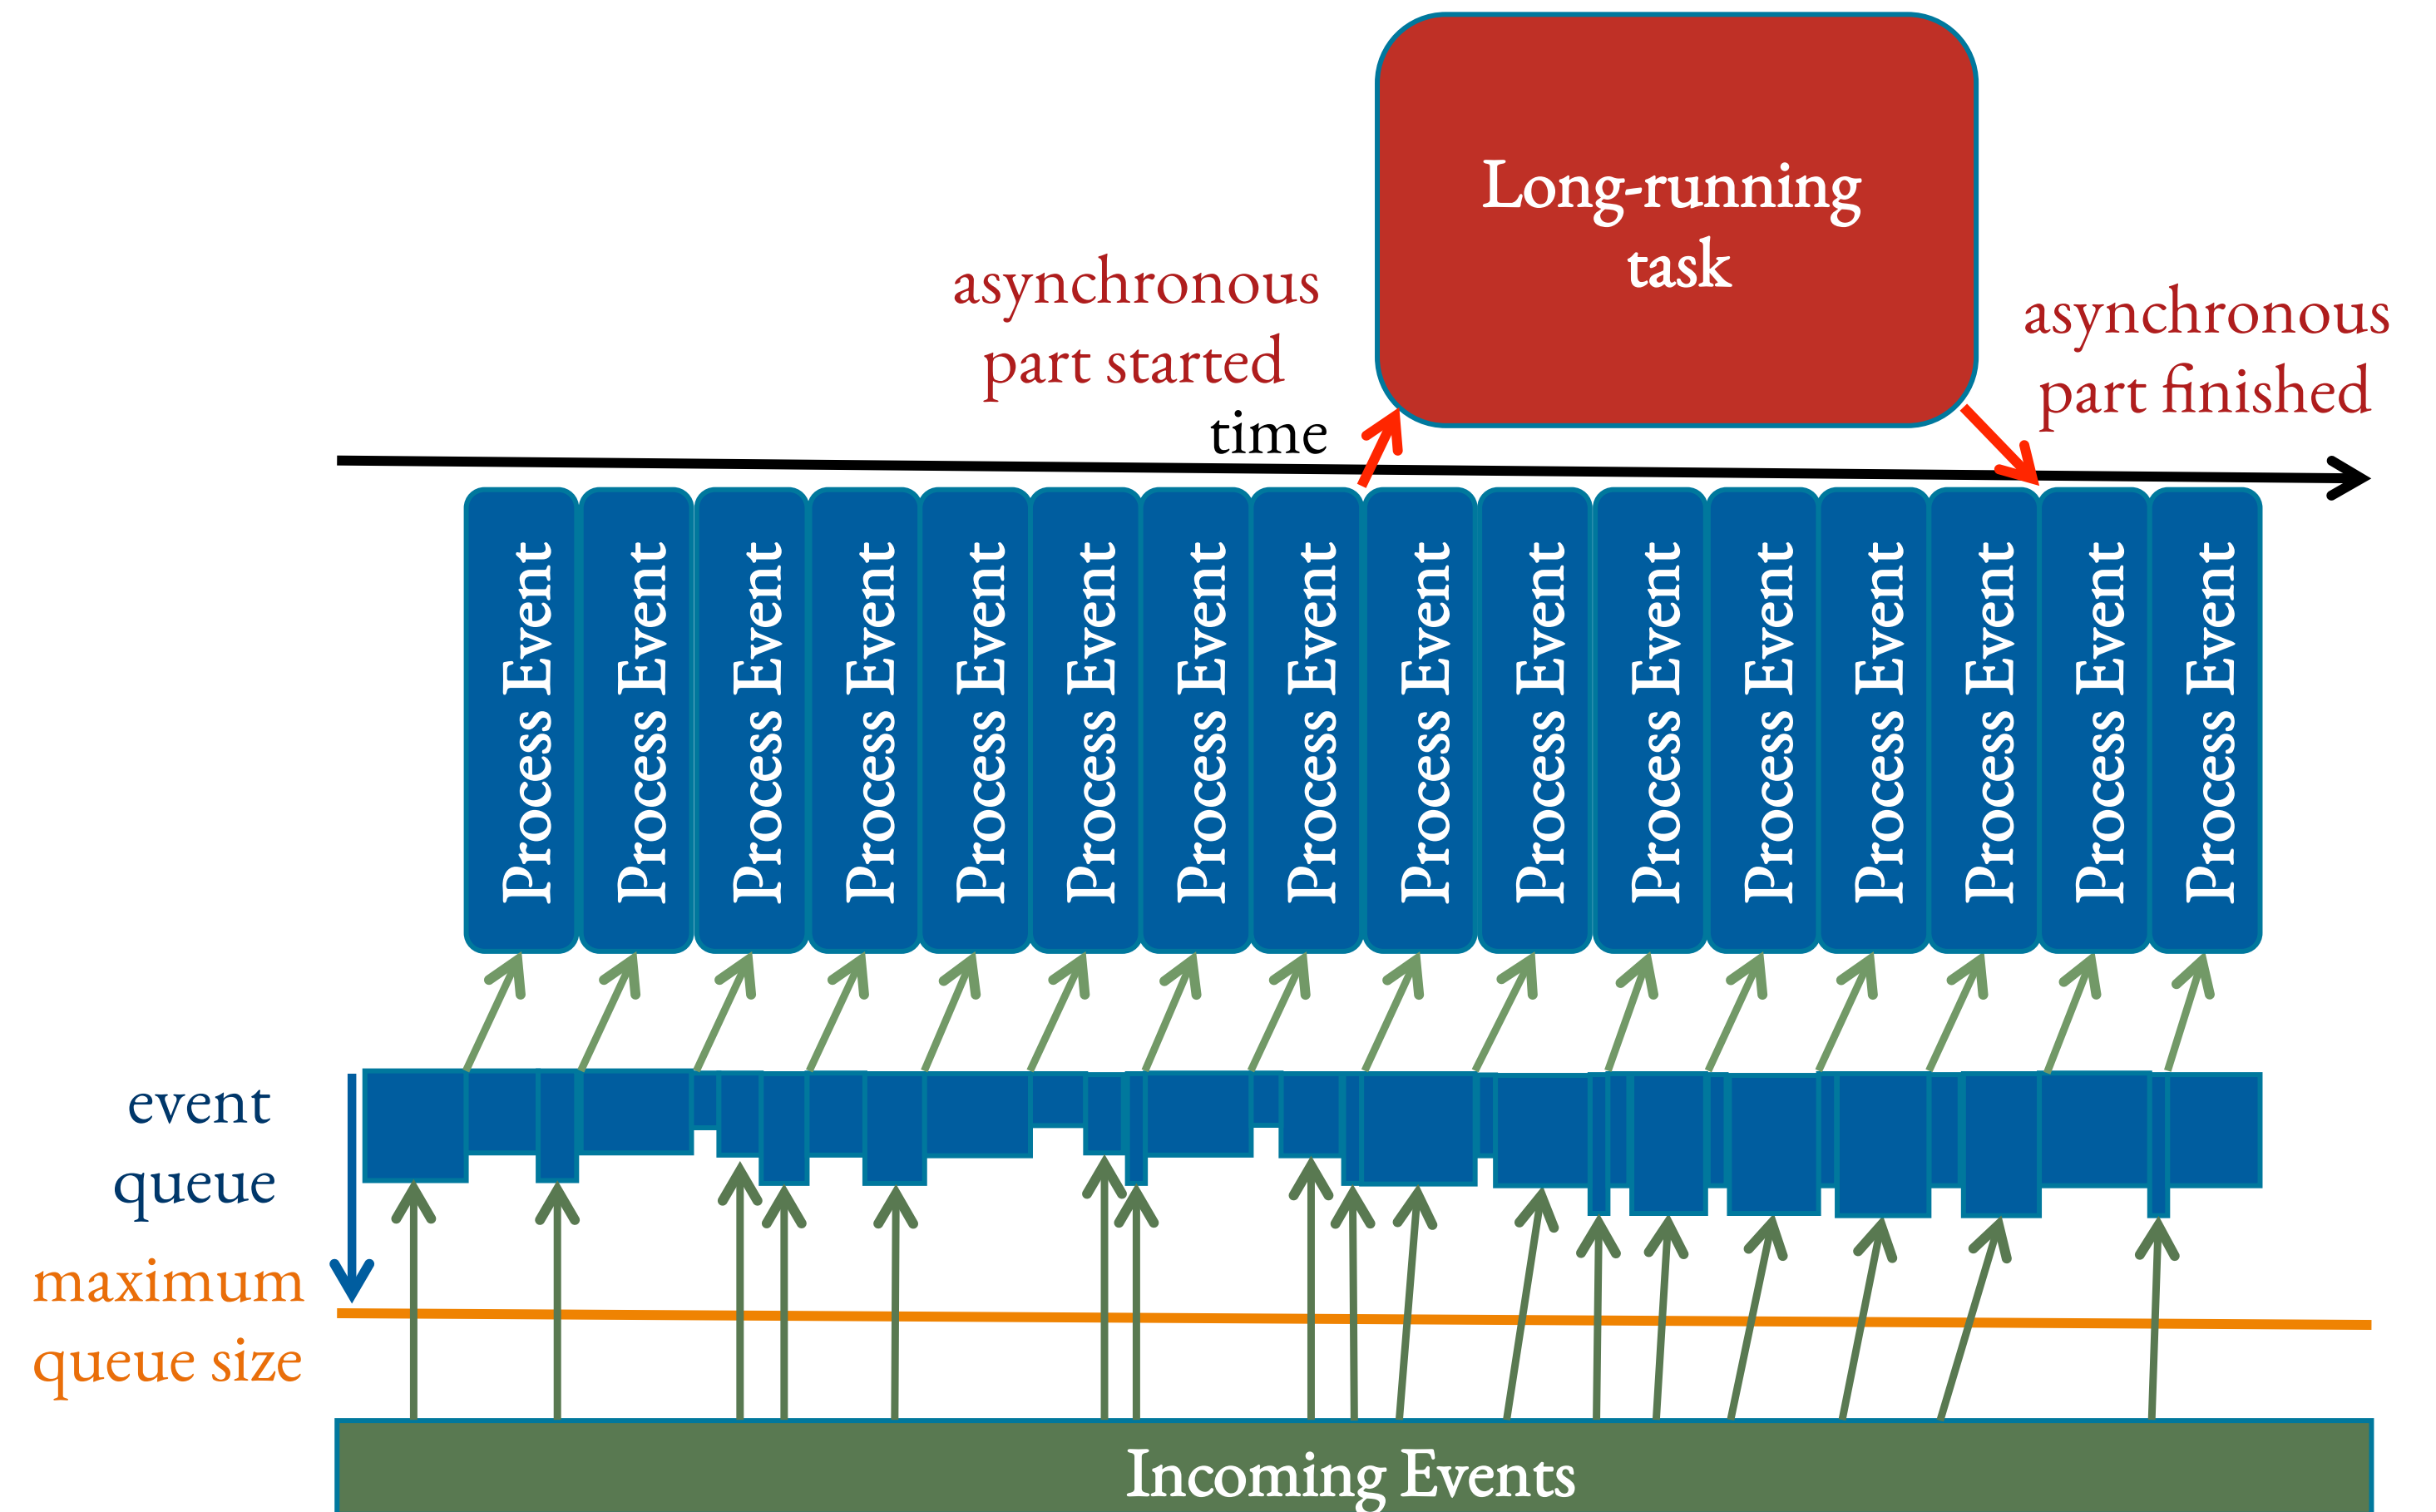
# Online TPC calibration challenges



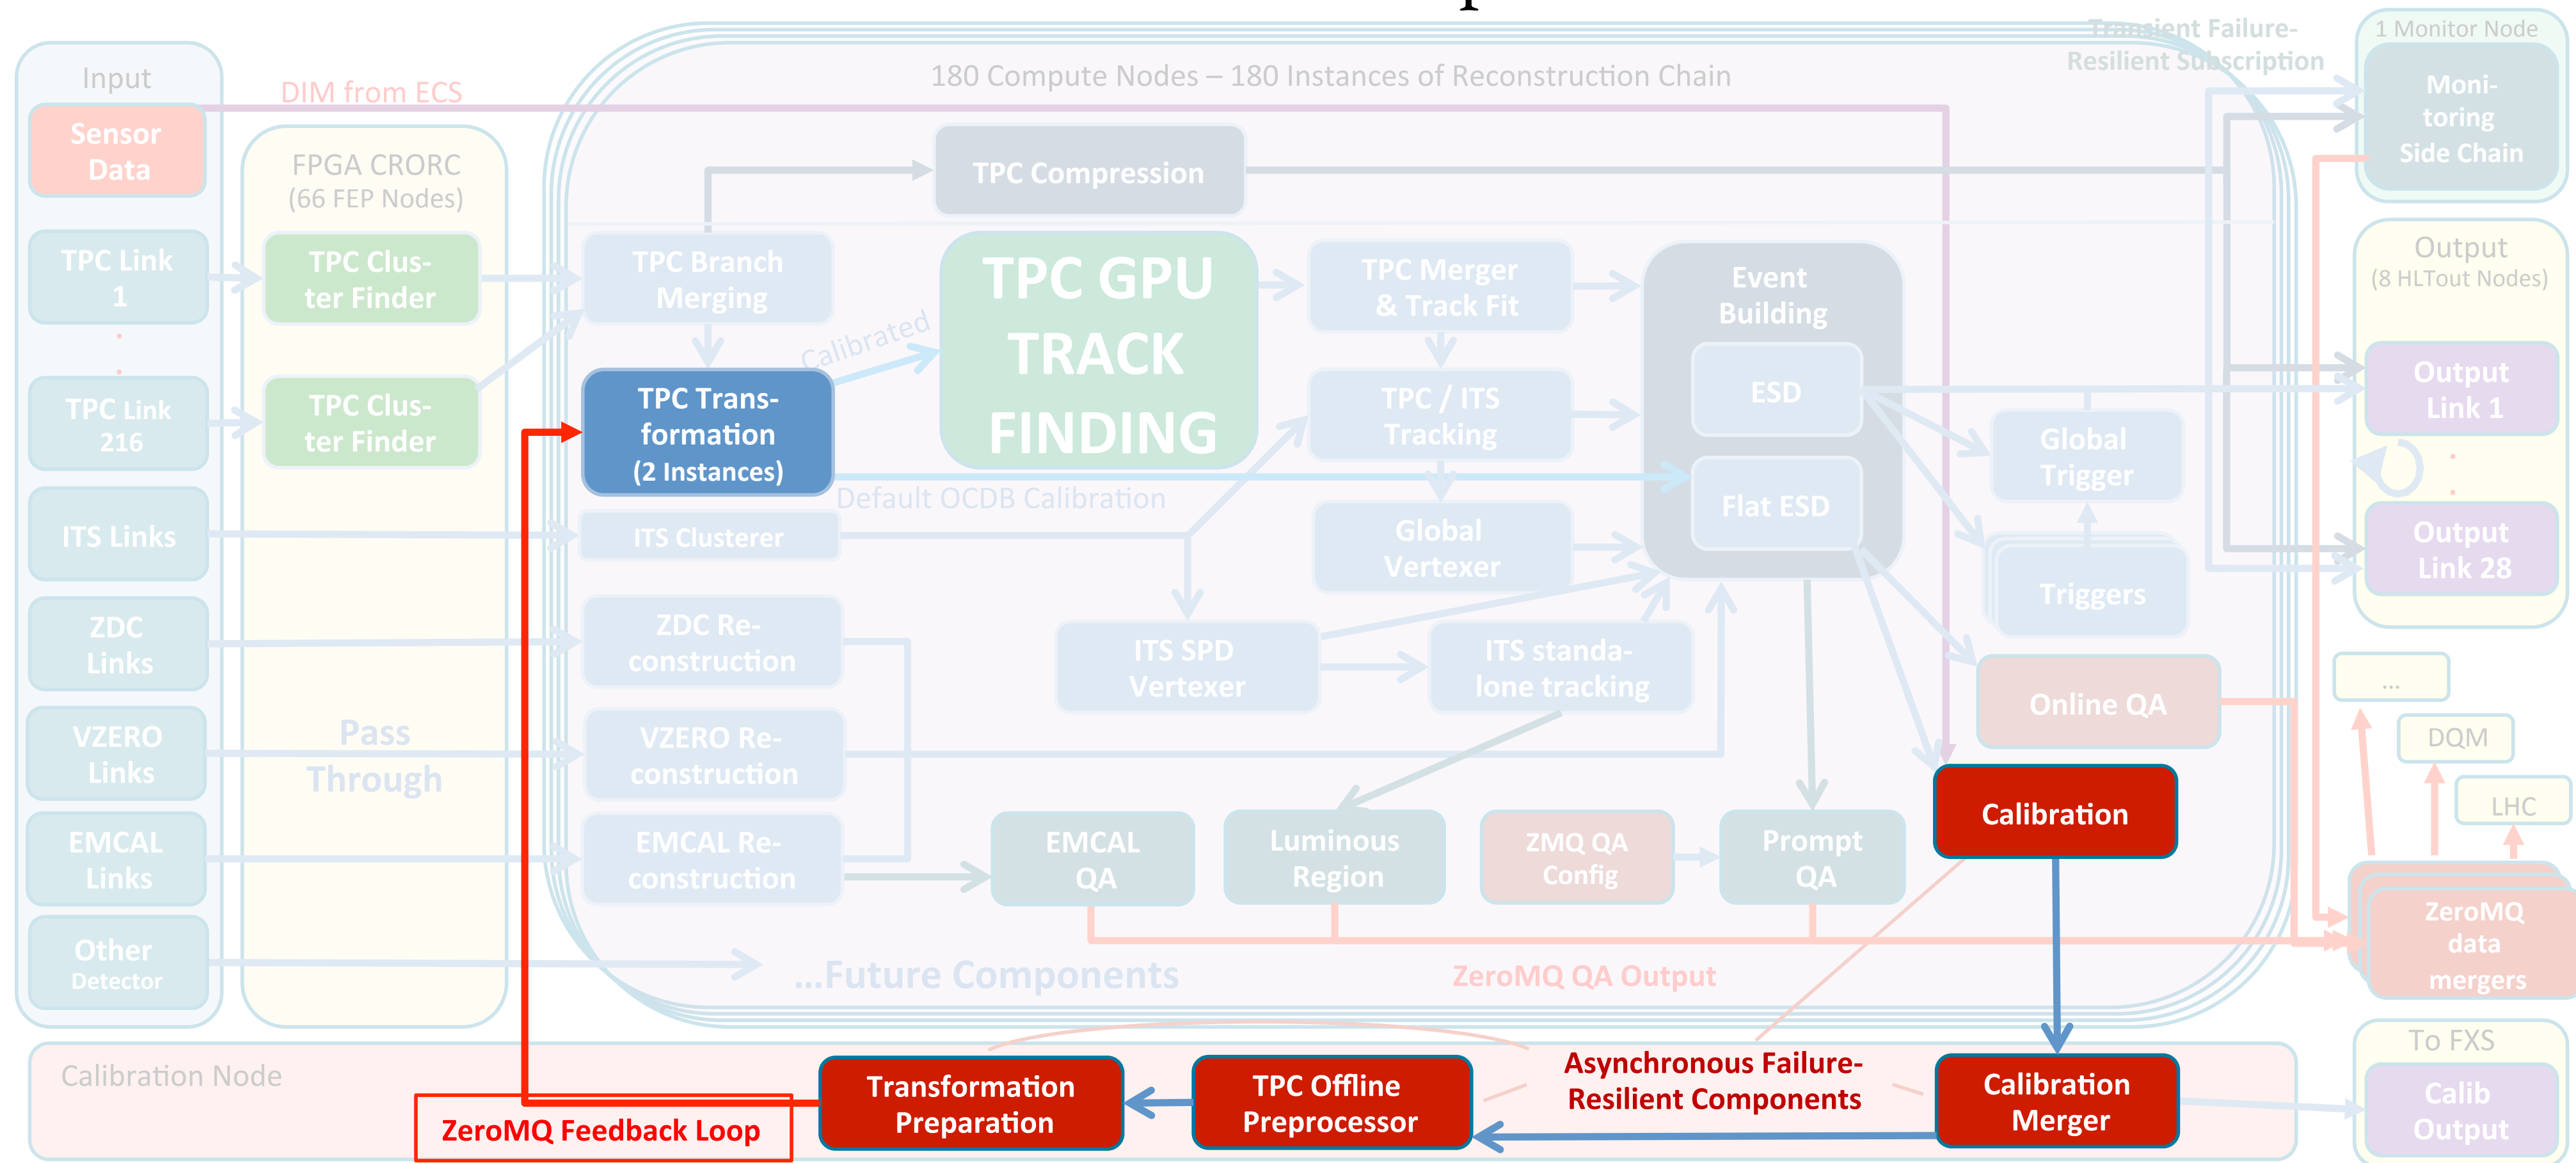
- TPC calibration parameters change slowly- semi stable within a window of ~15 minutes.
- Calibration needs to be (re-) calculated periodically and fed back to the reconstruction.
- Need about 3000 PbPb events to calibrate a time window (period).
- First period uncalibrated online, cannot go back to apply calibration.

# Asynchronous processing

- Same calibration code offline & online.
- HLT analysis manager framework; can run any analysis task, used also for QA.
- Offline code not optimised for latency:
  - ➔ need async processing.
- Out of our control:
  - ➔ need failure resistance.
- Framework forks a process:
  - Asynchronous data exchange (shmem) - no latency issues for normal data flow.
  - Results ready when ready.
  - Resilient - in case of problem, we restart (or ignore), data taking continues.



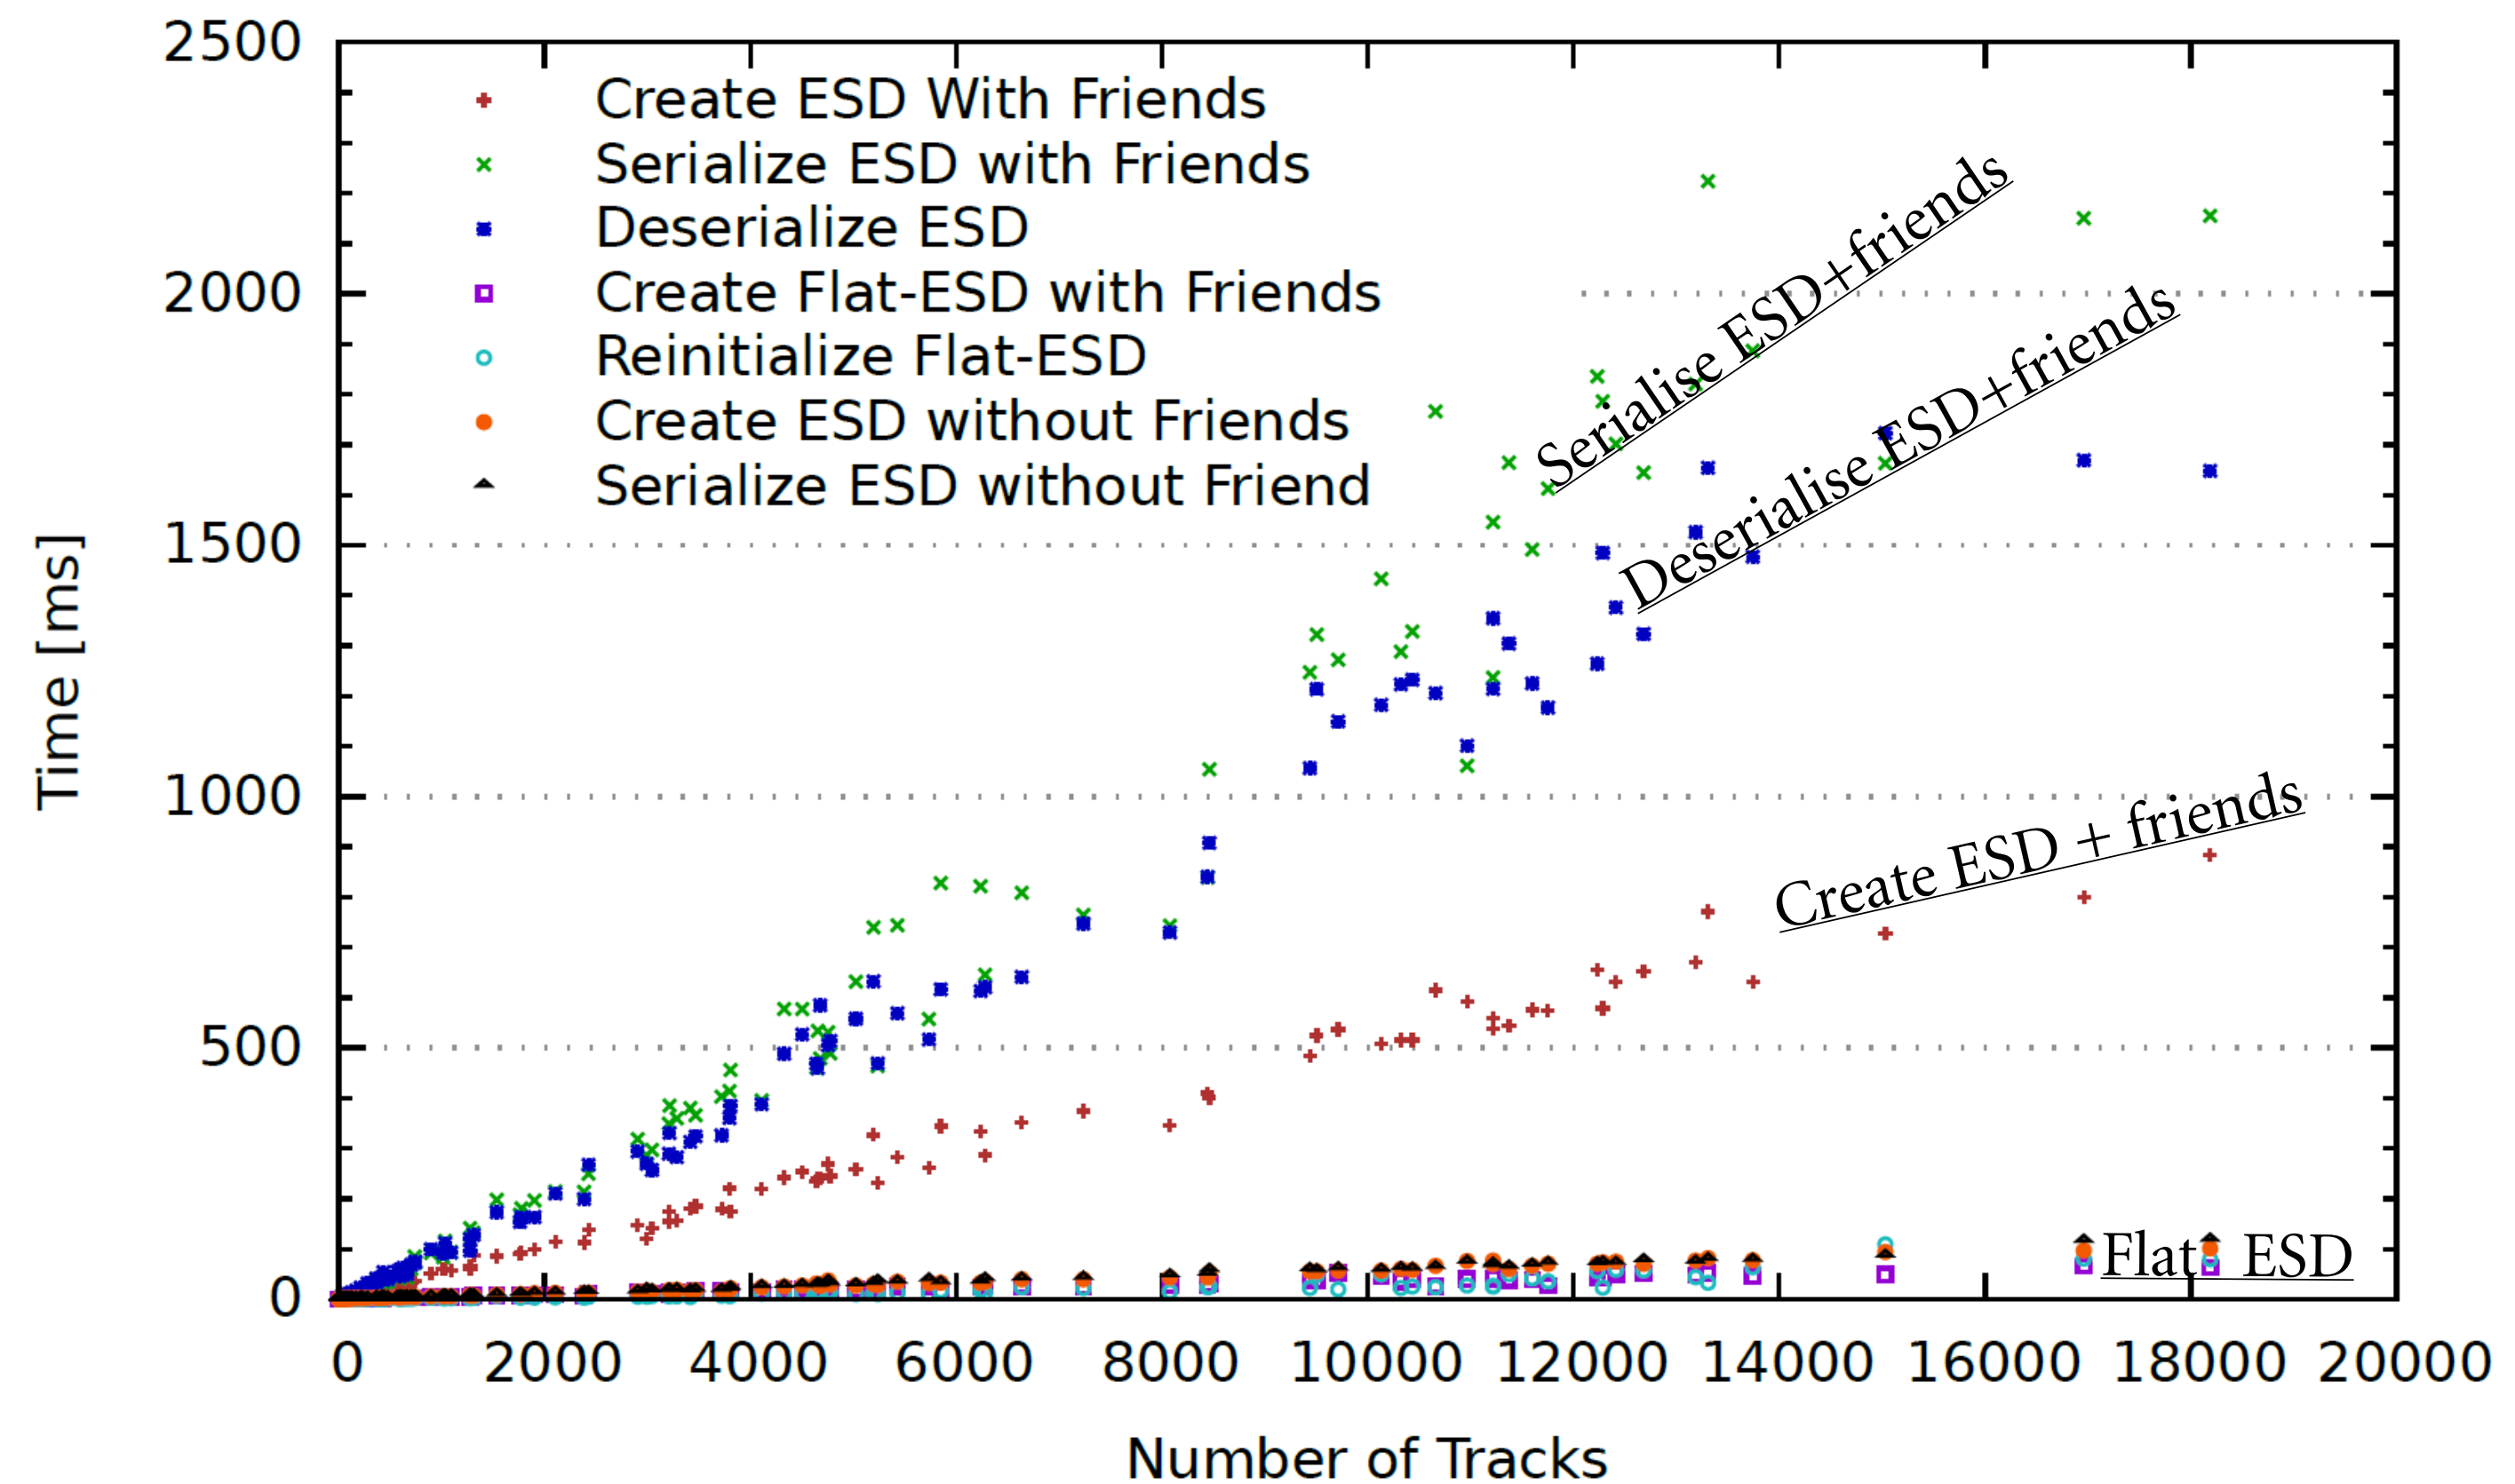
# Feedback loop



- Calibration components (on each node) process the data asynchronously.
- Single merger instance feeds the calculation of the calibration object (cluster transformation map).
- ZeroMQ handles the (asynchronous) distribution of the calibration object to all instances of the tracking.
  - data injected into HLT data stream, components unaware of the asynchronous channel.

# Offline calibration code/serialization

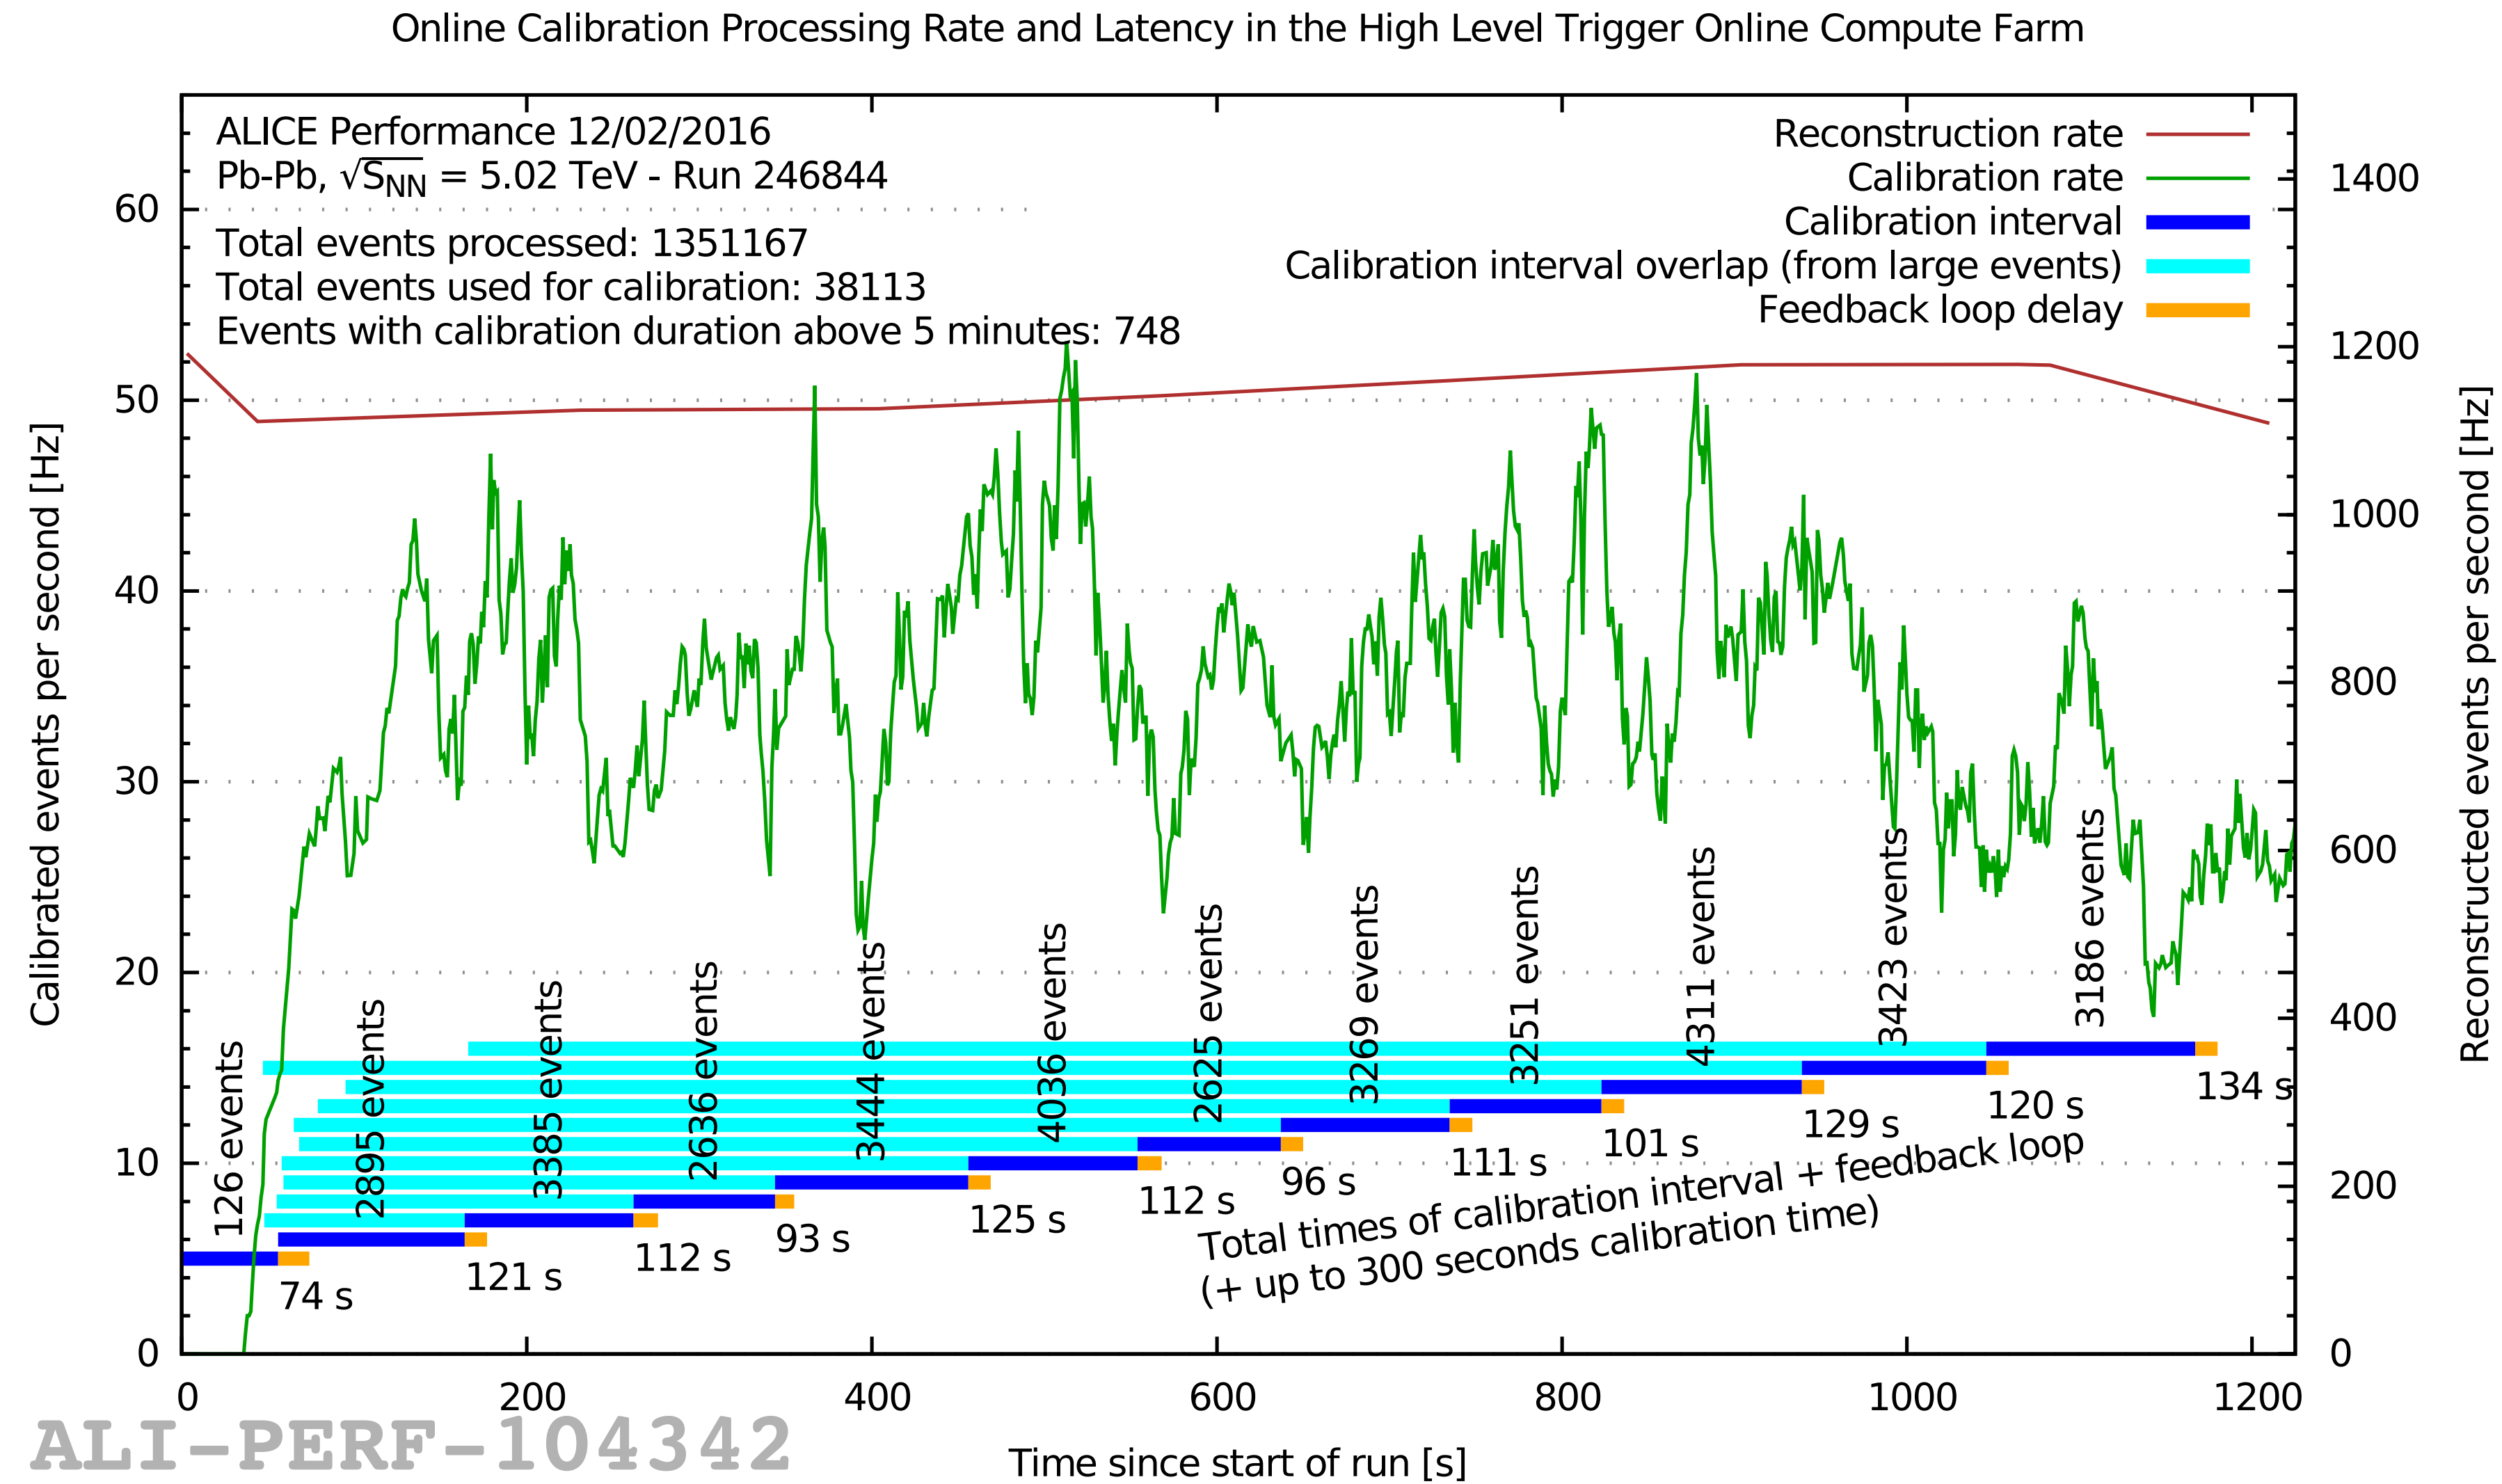
- Same calibration code offline & online.
- input is ESD.
- Online: distributed components exchanging data via network/shmem.
- Packing/unpacking overhead for ROOT objects large.
- Custom ESD representation: Flat-ESD:
  - Common virtual interface with ESD.
  - Flat buffer, zero serialisation overhead.
  - Small reinitialisation cost (restore vtable pointer).
  - Some stored data (track seeds, clusters) are stored as POD - special access logic.



- ESD: event summary data, ALICE format for reconstructed data.
- ESD Friends: additional data for calibration
  - ➔ track seeds + clusters attached to tracks.

# Performance

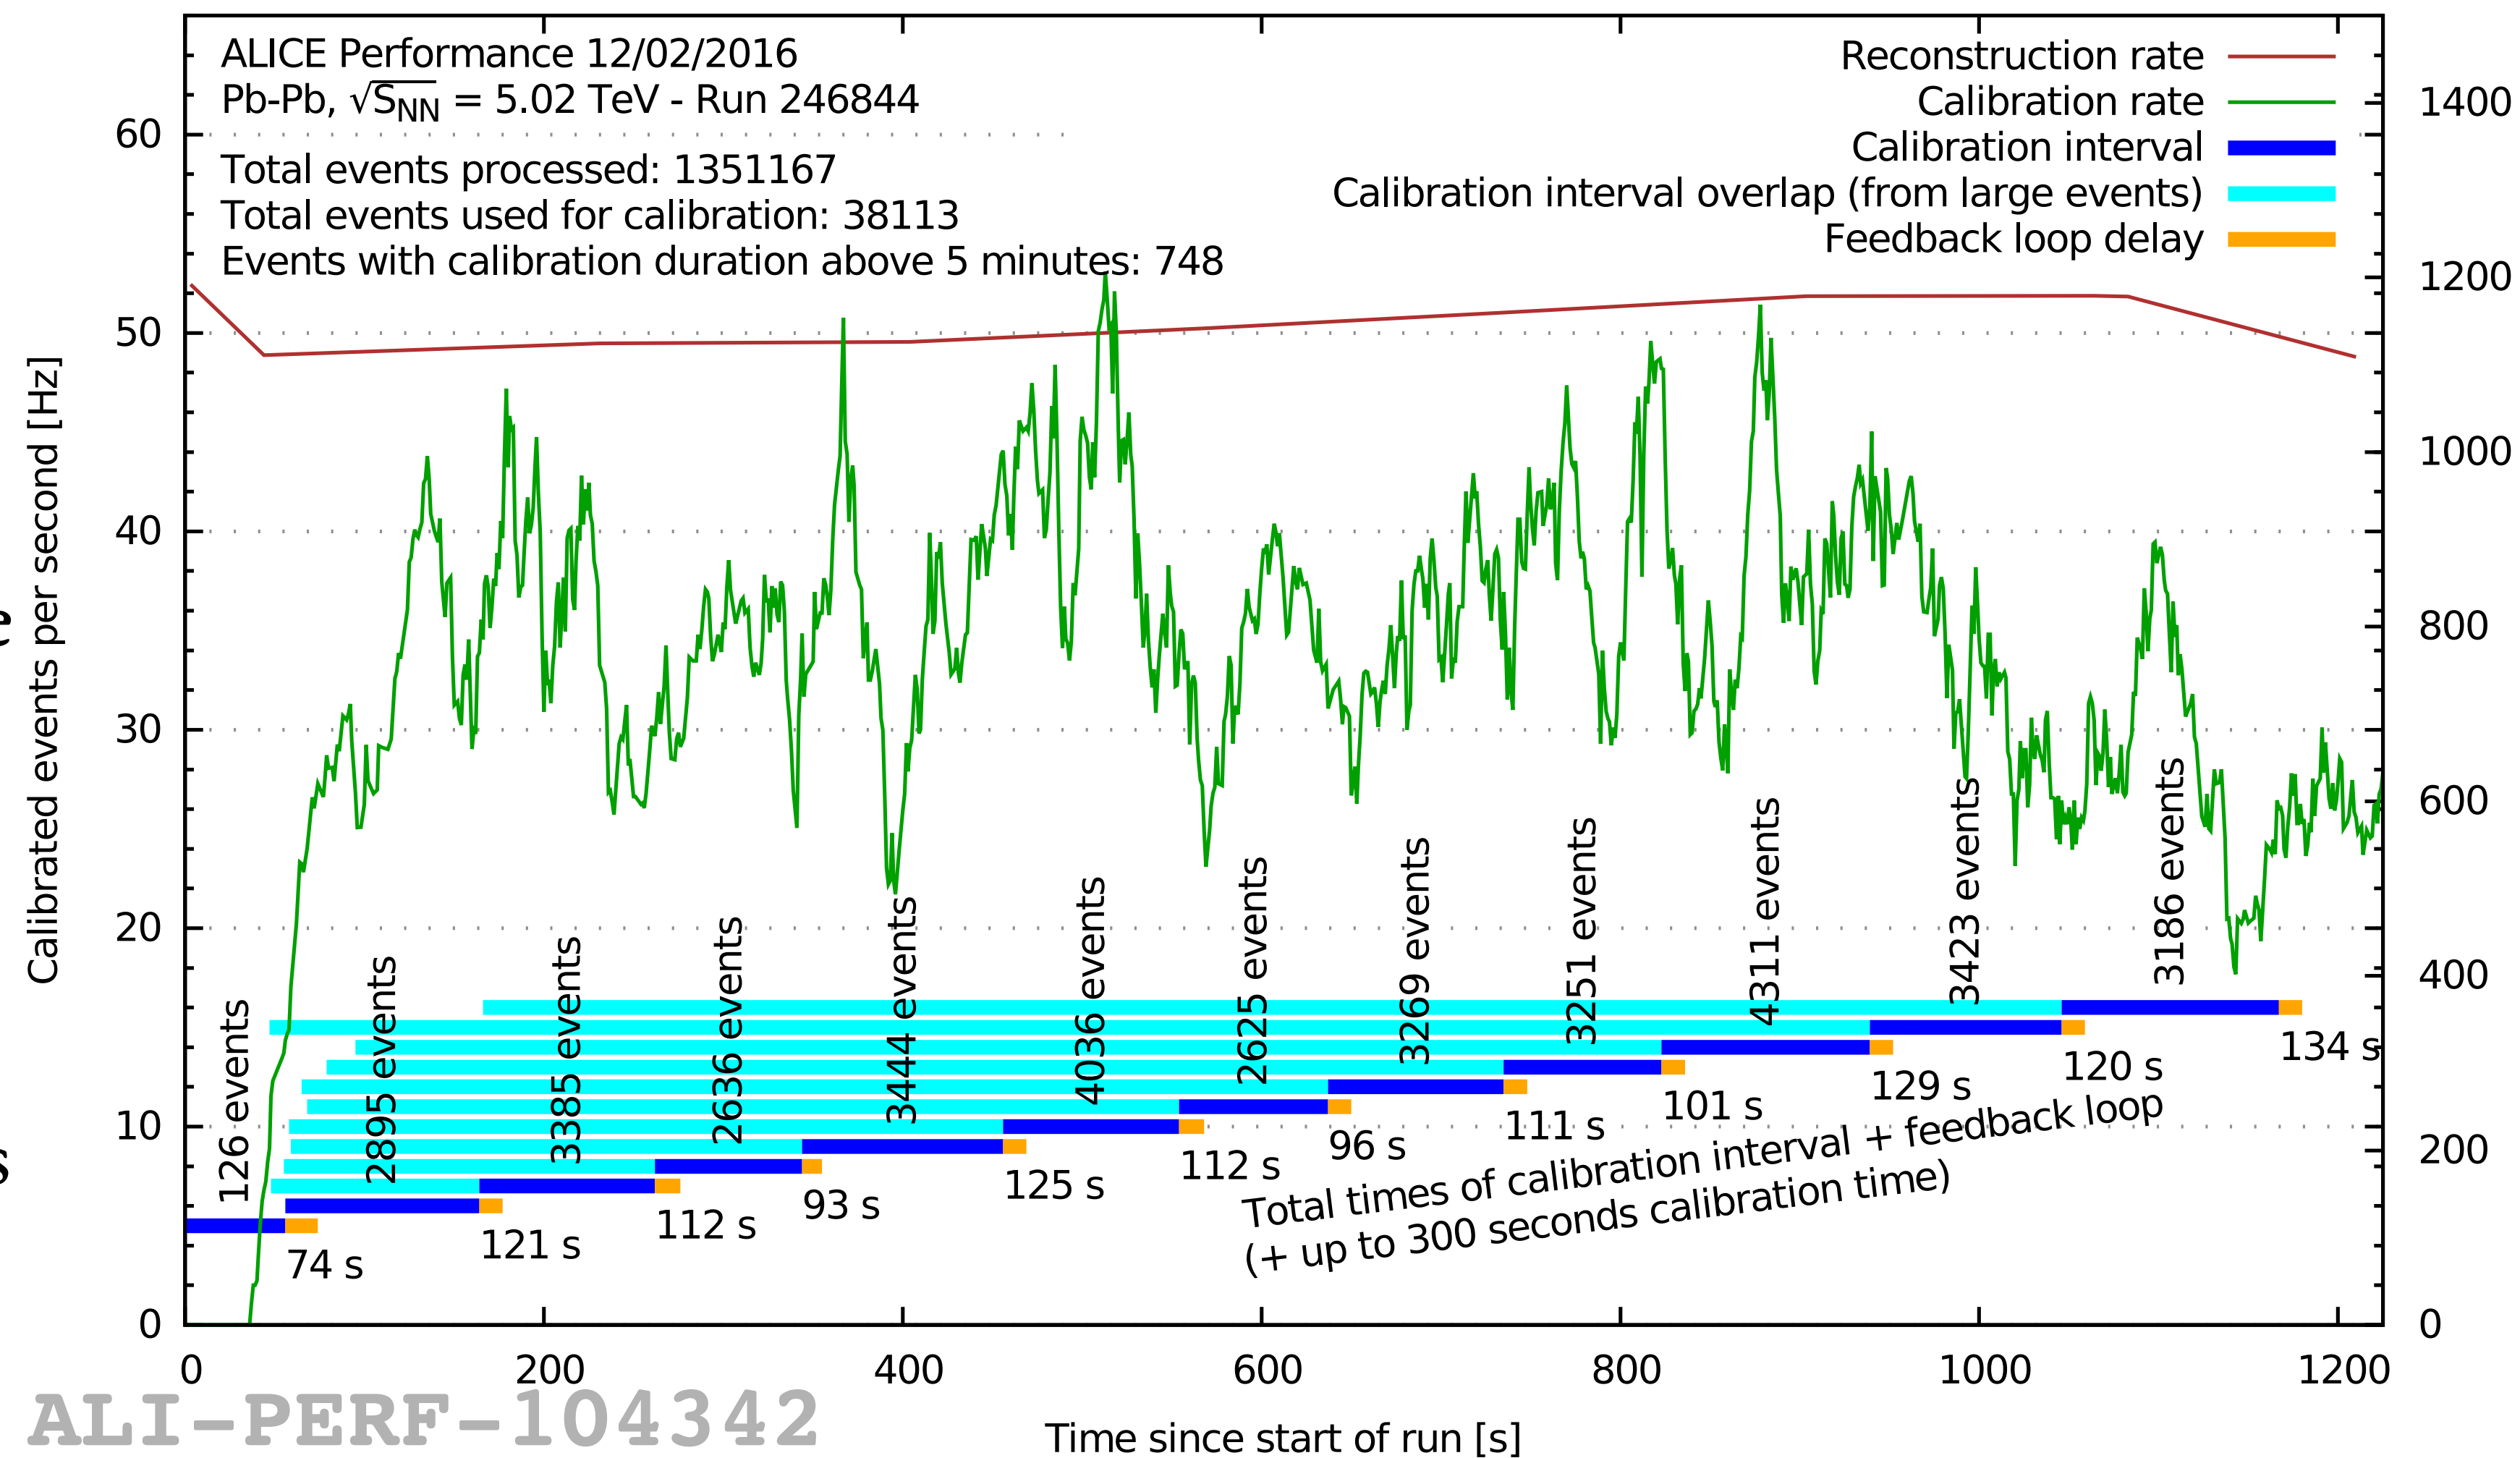
- Green line:
  - calibration rate aggregated over all HLT nodes.
- Dark blue bands:
  - duration of the calibration interval.
- Yellow extension: delay between calibration period ends and new calibration is used in tracking (~140s).



# Performance

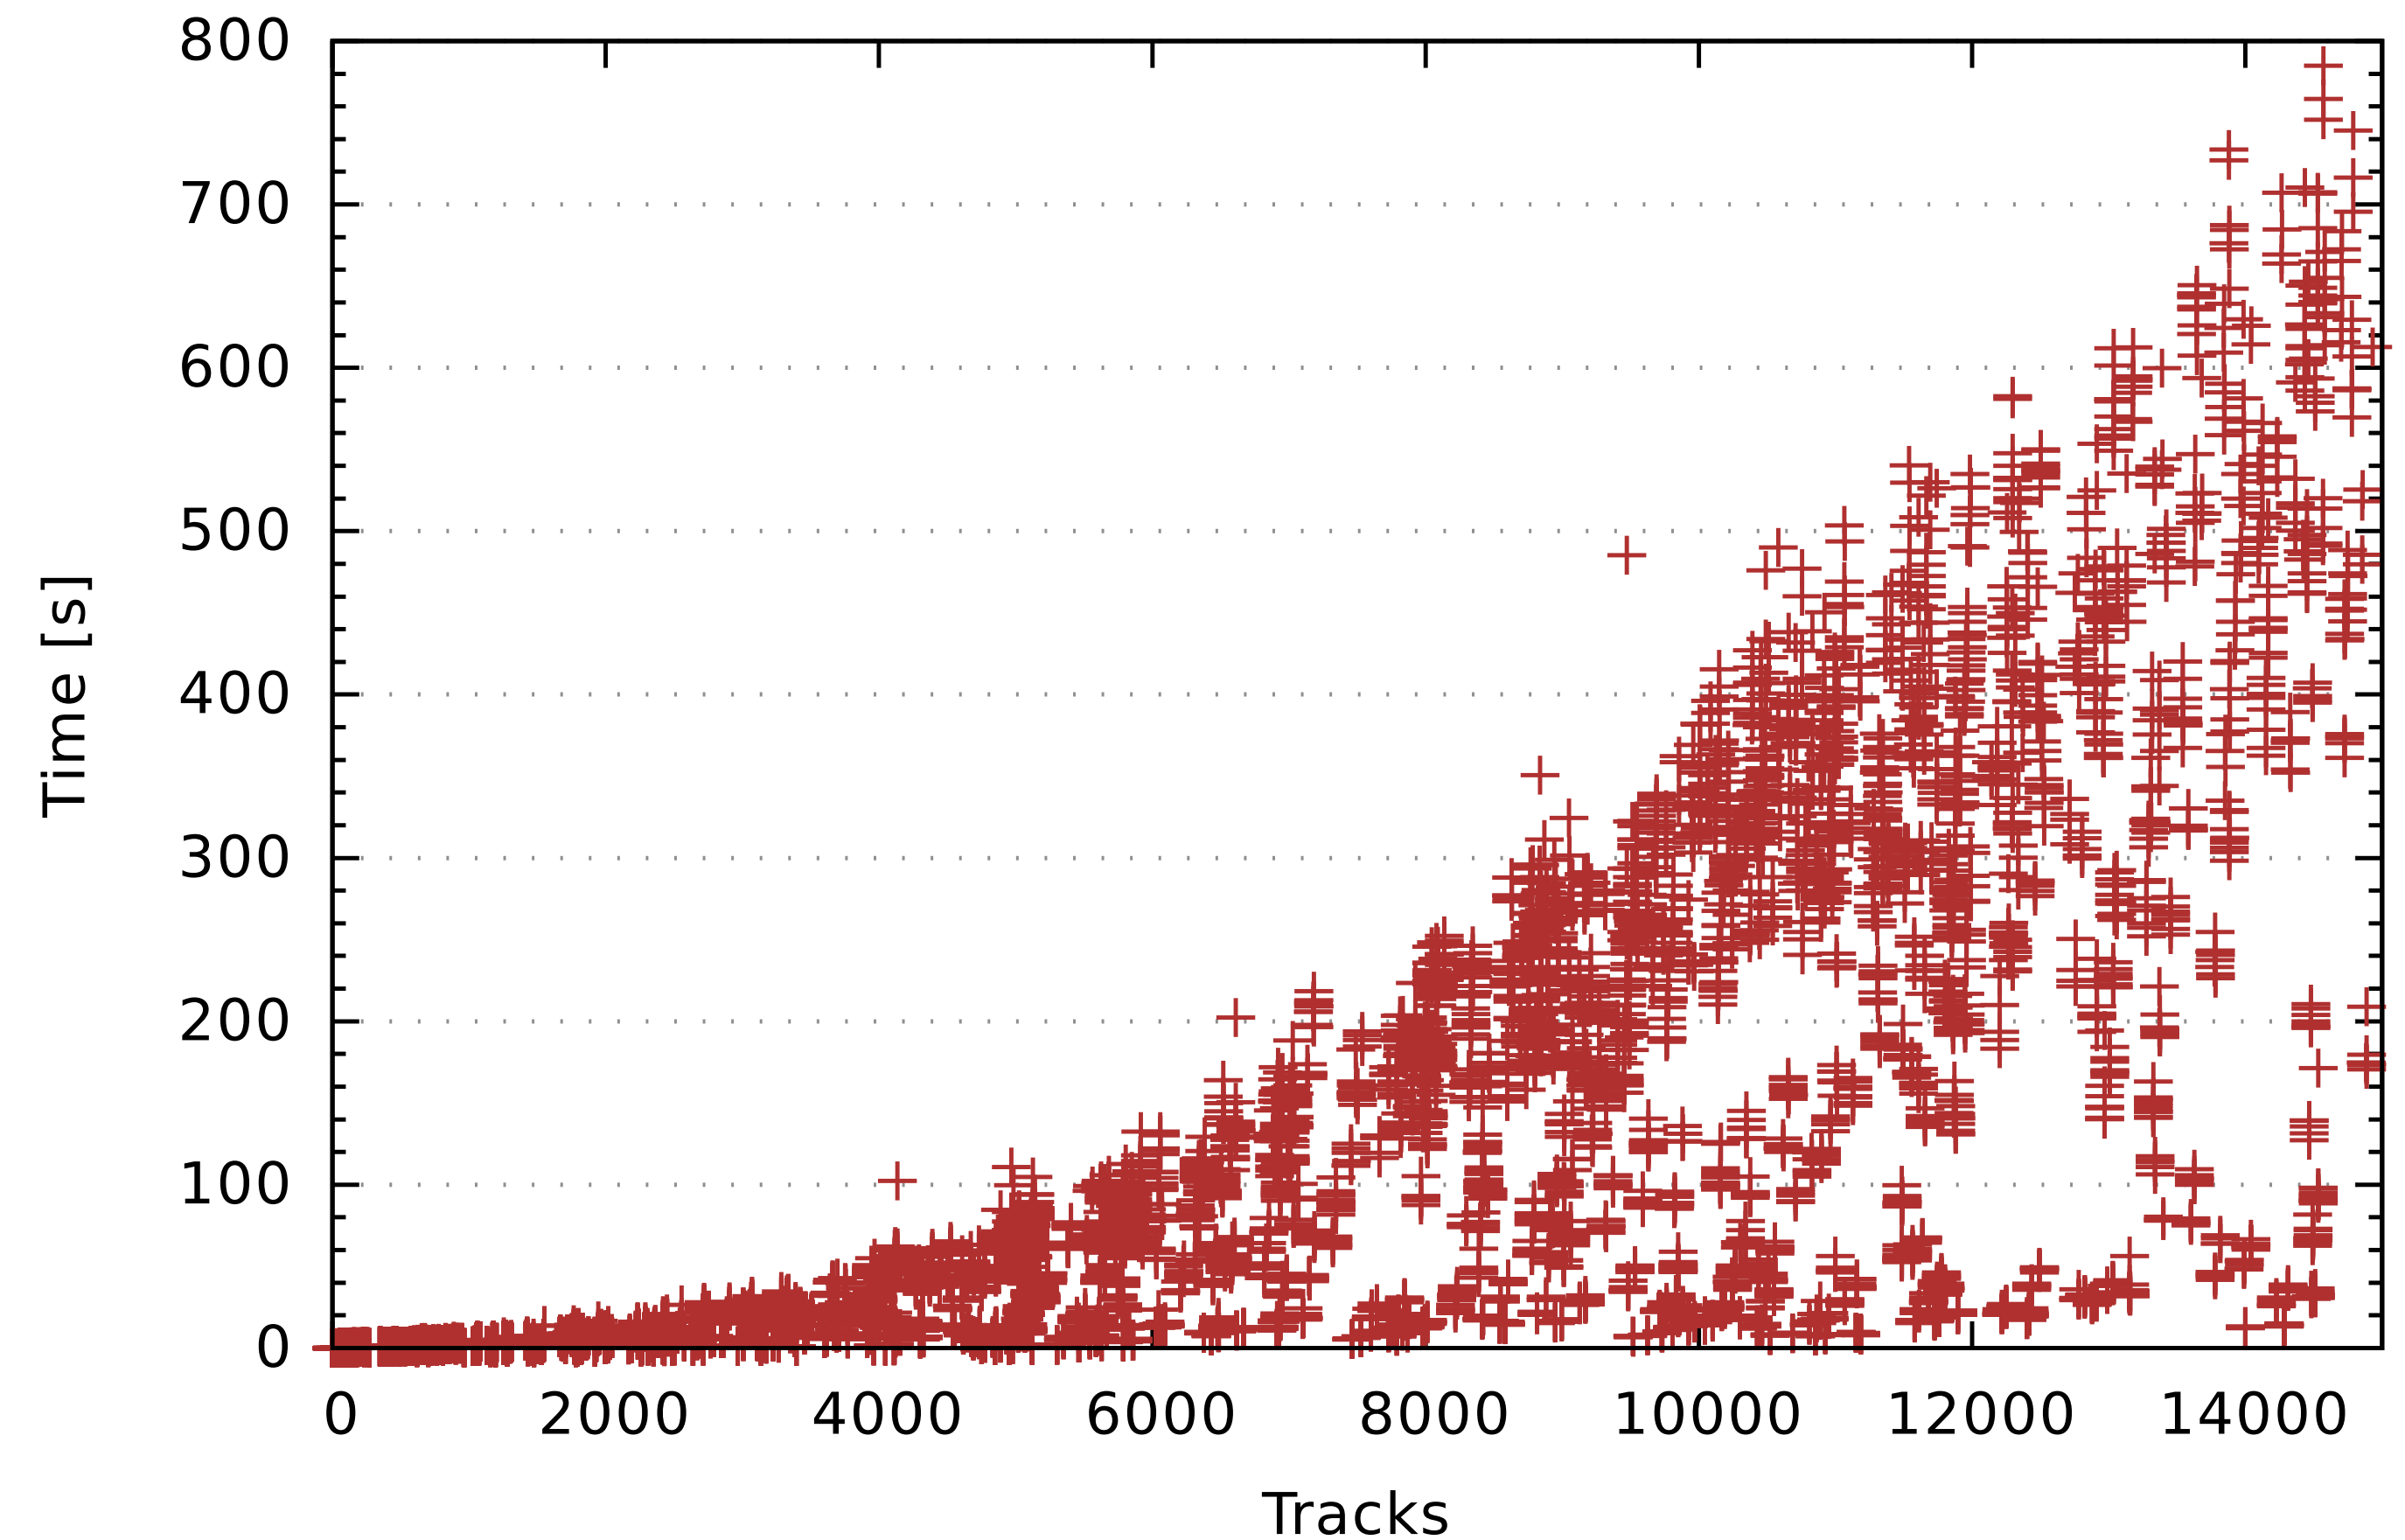
- Light blue bands:
  - event acquisition time for events that finish in a certain calibration period.
  - ~2% take longer than 5 minutes.
  - Early events do not pollute the calibration period - end up in a different time bin, which can be used offline.
- Rate: 31Hz.
- Each calibration period contains 98% of processed events.
- Delay < 7 minutes.

Online Calibration Processing Rate and Latency in the High Level Trigger Online Compute Farm



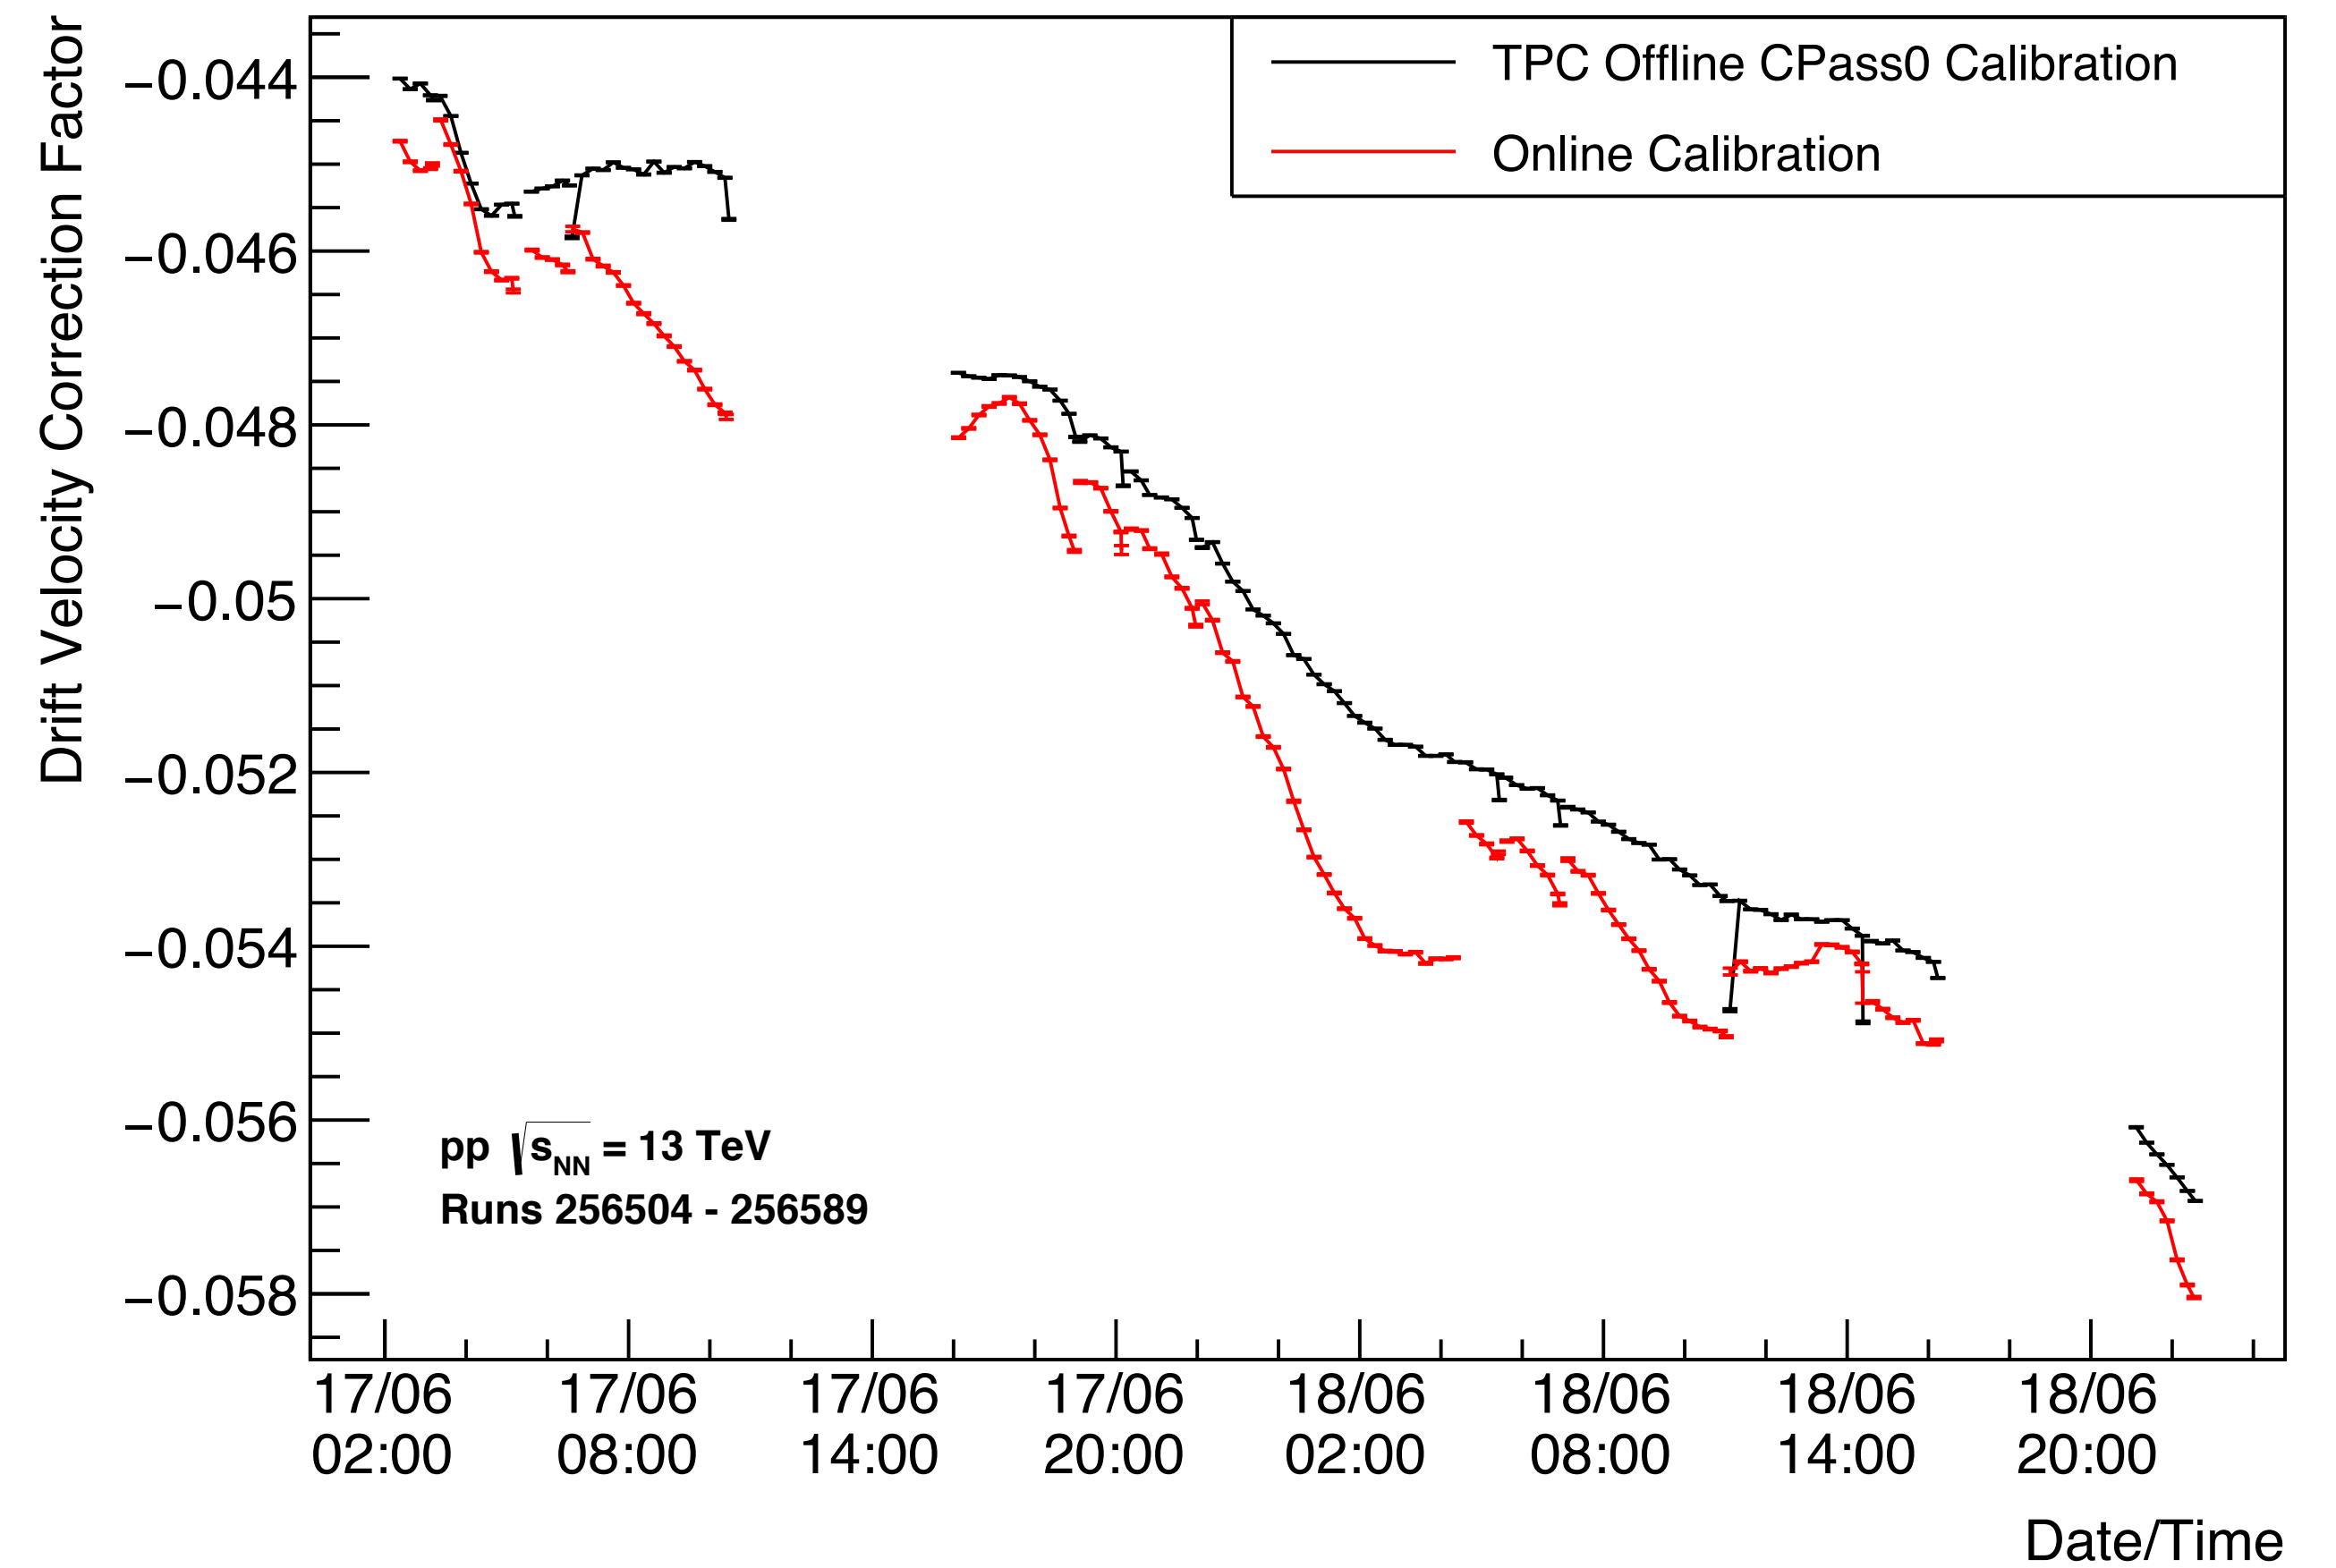
# Processing time

- Calibration is based on track matching between the TPC and the inner tracker (ITS).
- Due to combinatorics ~quadratic dependence of processing time to event size.
- First optimisation: don't process large events.
- The calibration code could also be optimised.



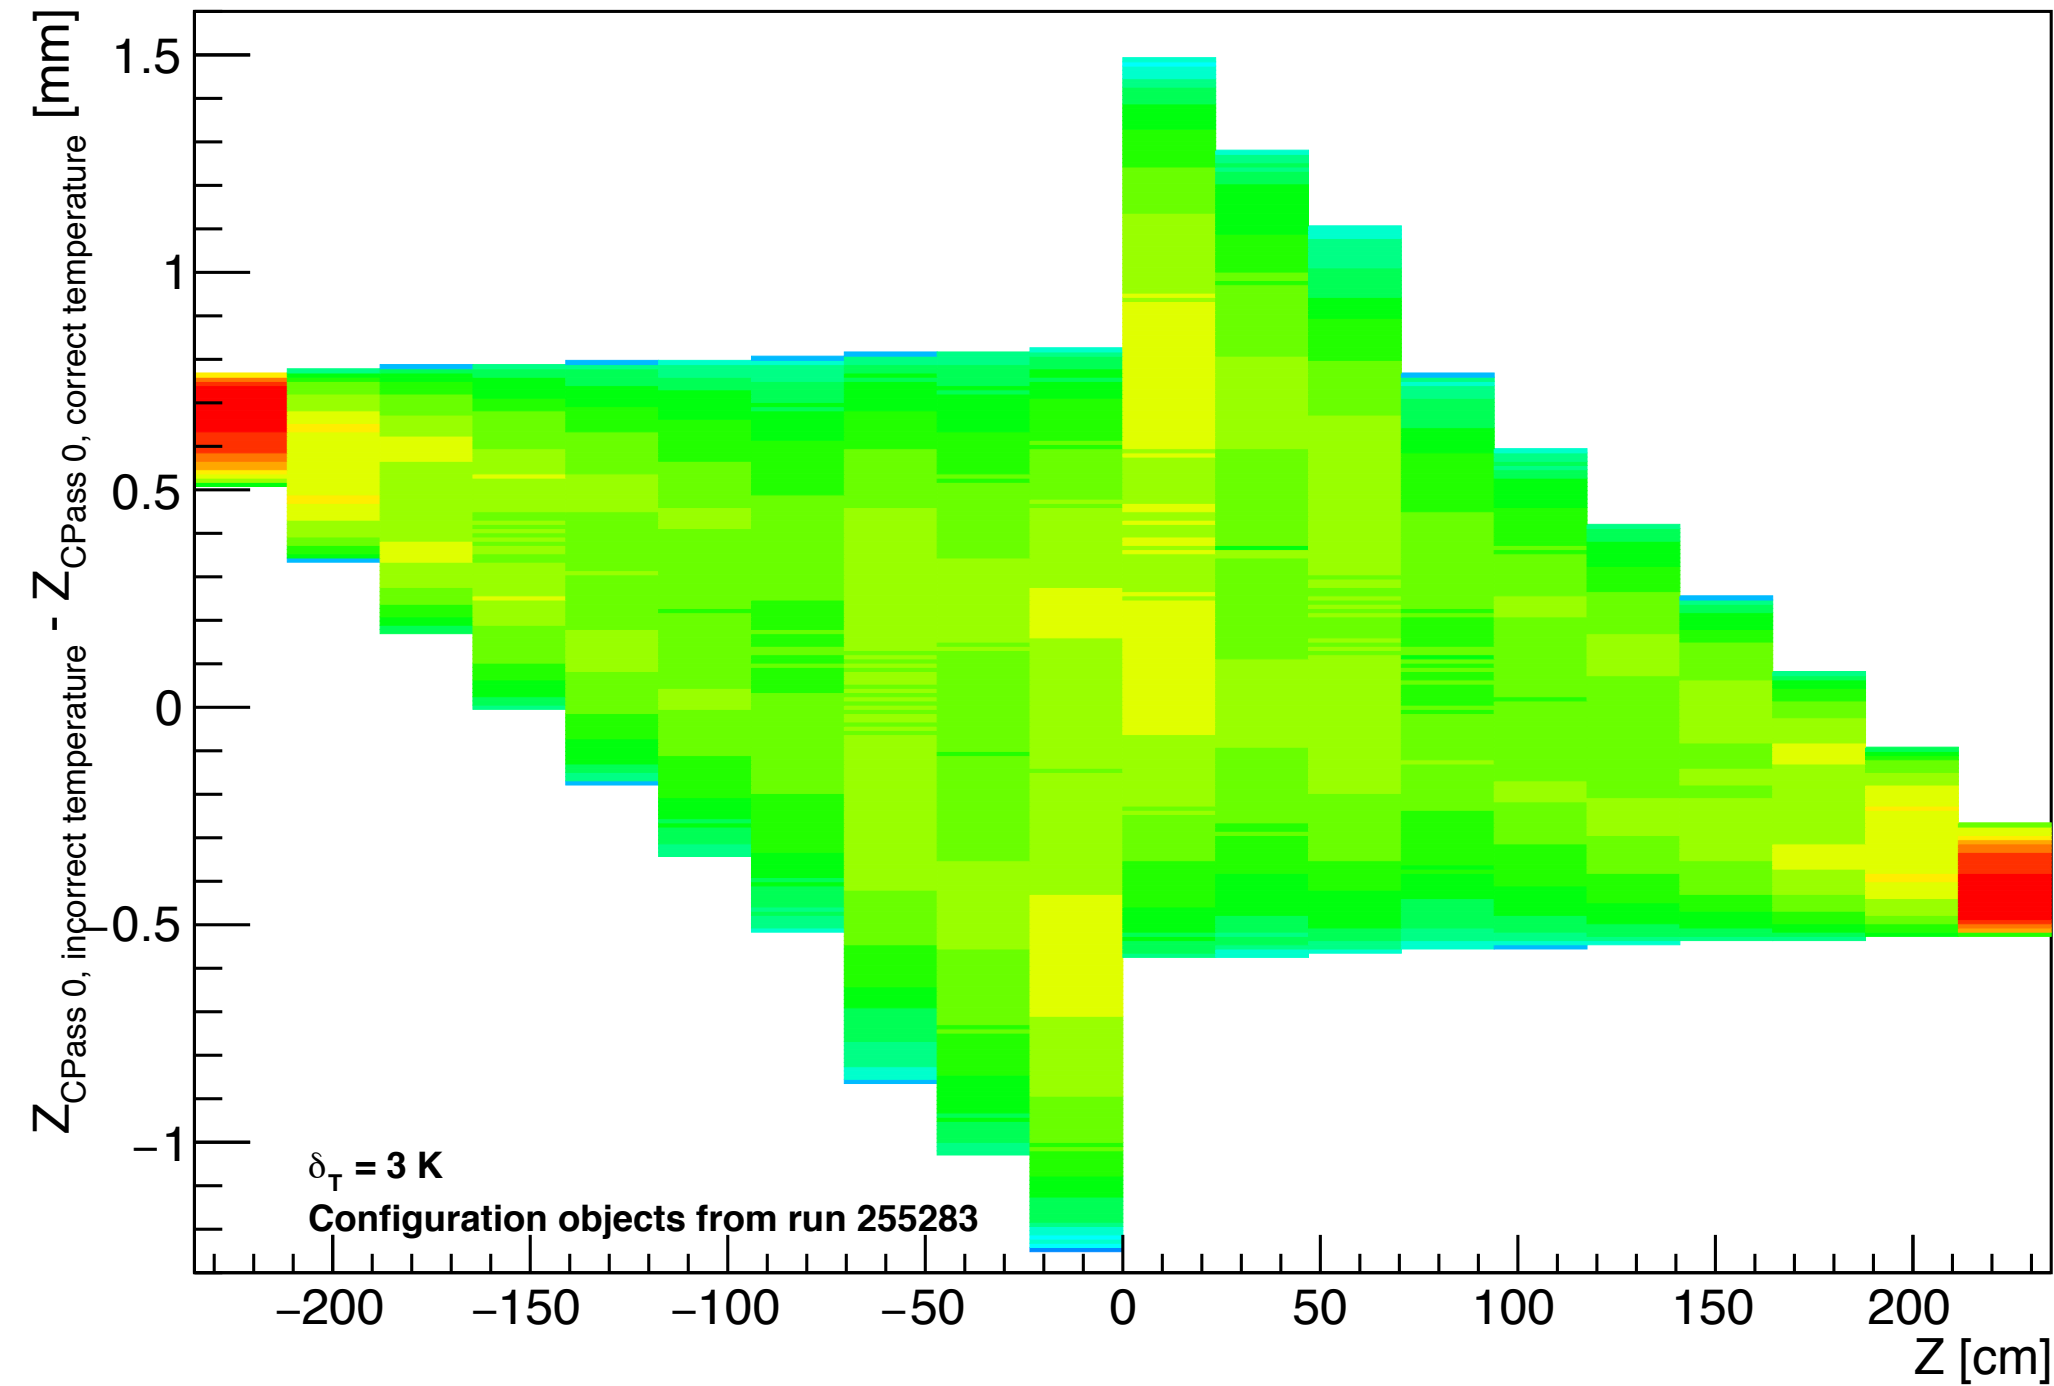
# Calibration results

- Compared to the offline drift velocity calibration over the course of many runs:
- The correction factor close, follows the trend as function of time.
- Difference due to wrong values for other relevant quantities, i.e. temperature, gas composition. (Offline uses time evolution of sensor data, HLT not yet).
- Calibration compensates the wrong assumptions in the correction factor.

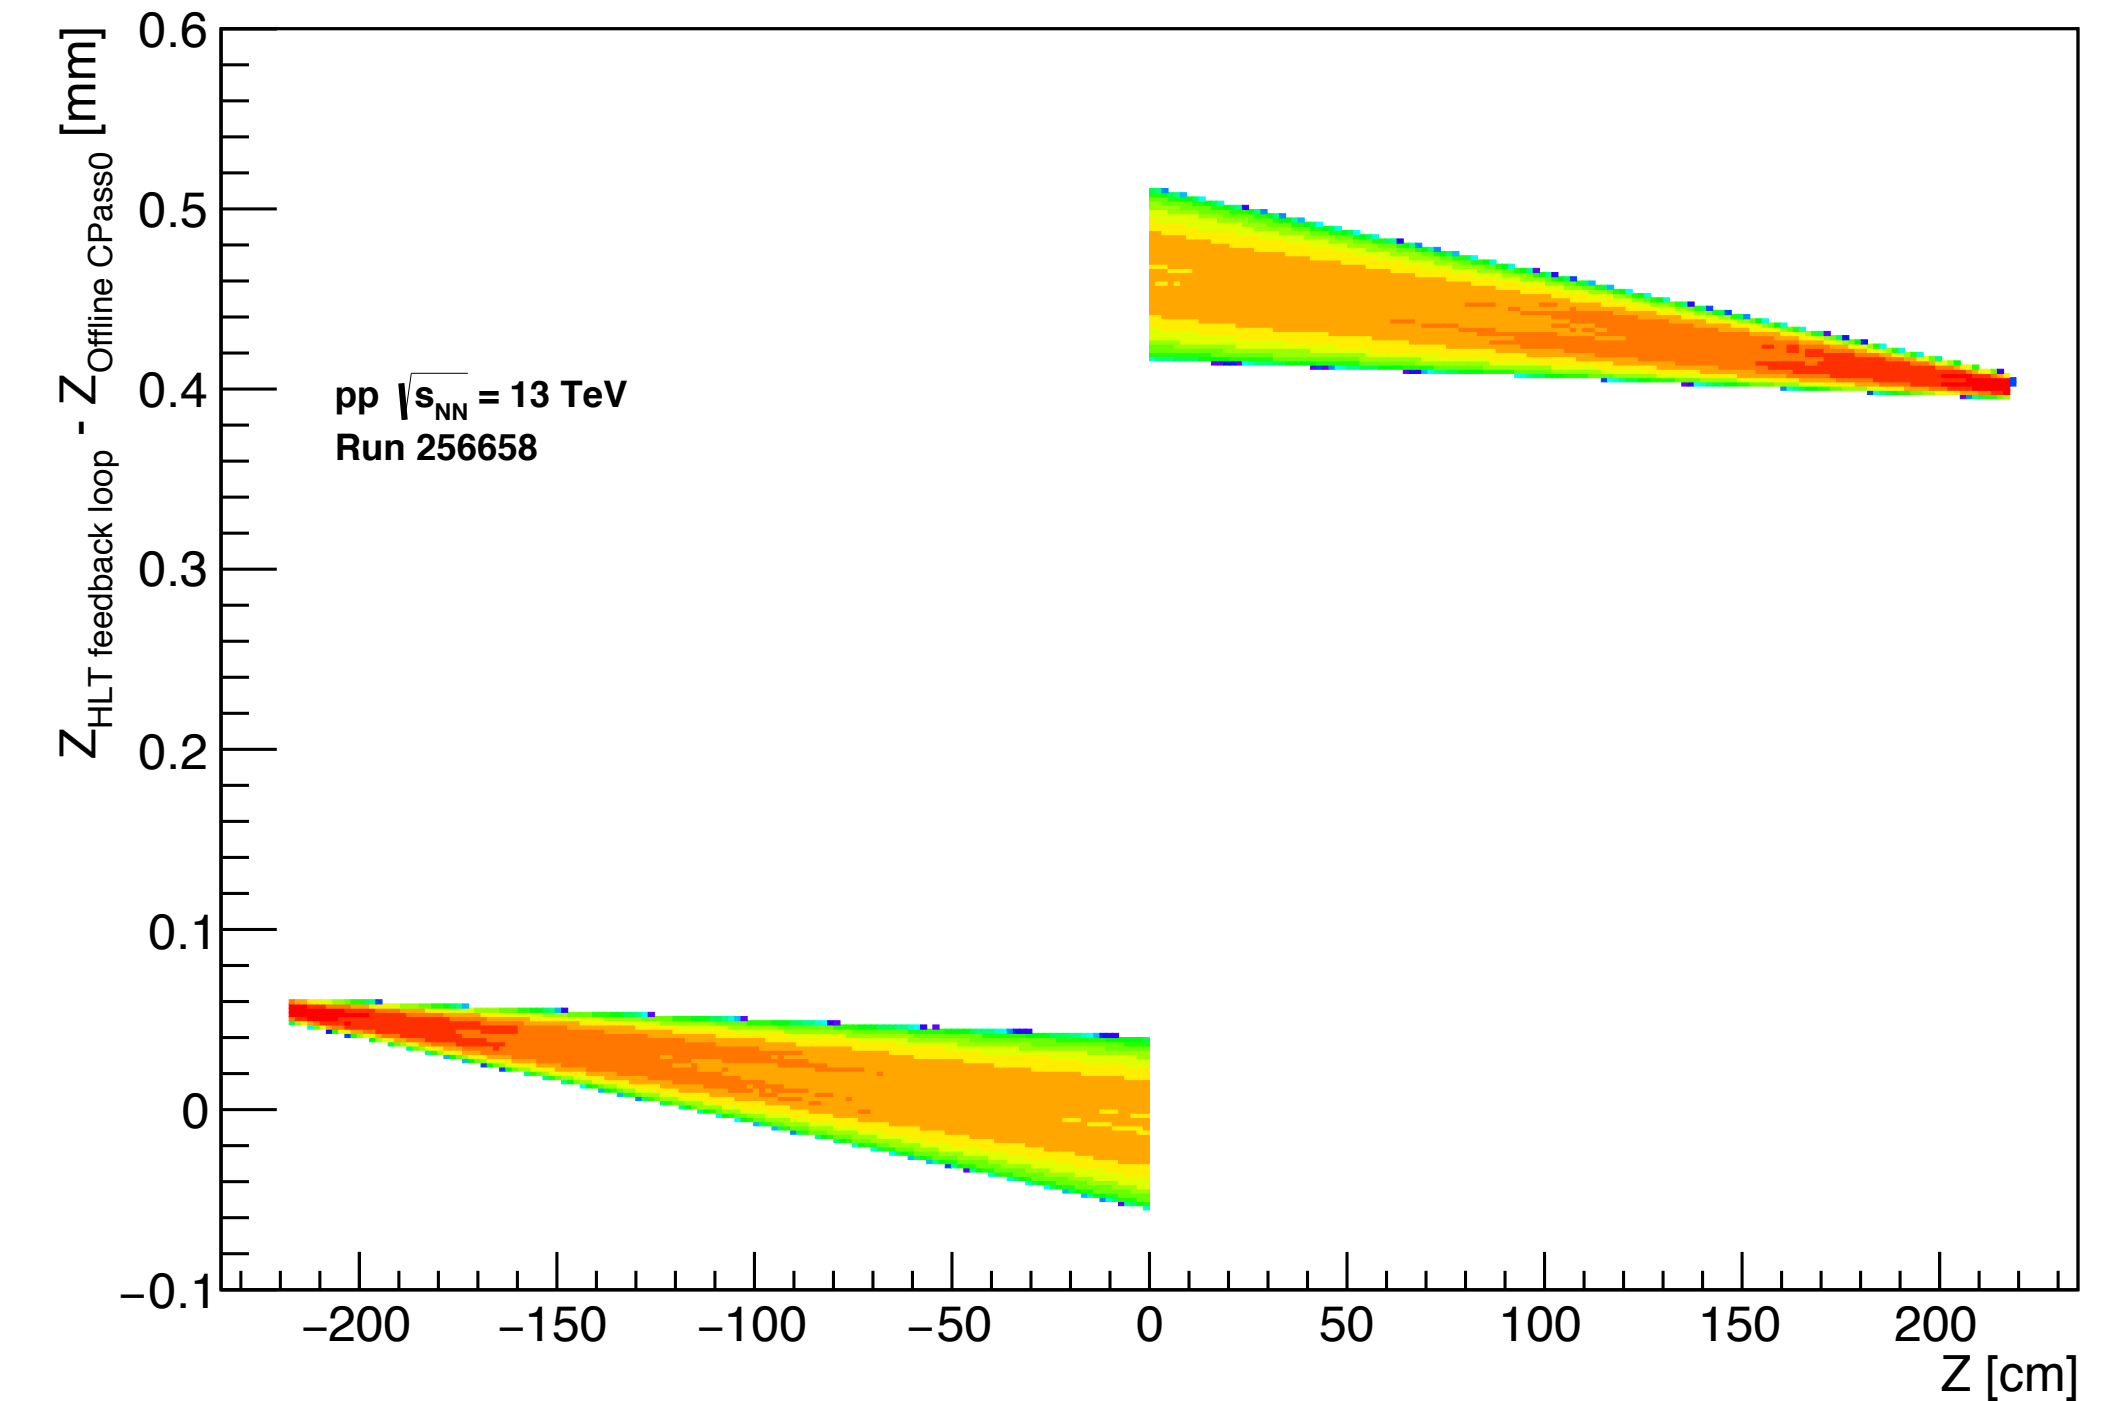


# Calibration results

- Cluster position difference between clusters used in online tracking and fully calibrated offline clusters.



- No online calibration.
- Static calibration object taken from a previous run.
- HLT and offline differ (order of centimetres).



- Online calibration with feedback loop.
- Difference  $< 0.5 \text{ mm}$ , within cluster resolution.

# Summary and afterthought

- Online calibration first deployed during the PbPb period in December 2015.
  - Time dependent TPC calibration.
  - 5 CPU cores per node (on 120 nodes)  $\rightarrow$  31Hz,  $\sim$ 80Hz after some initial optimisation.
  - $\sim$ 3000 events per interval processed in 5 minutes.
  - 140 seconds to distribute and prepare new transformation maps for online reconstruction.
- Calibration equivalent to offline within cluster resolution.
- Running ever since, load in pp  $< 3$  cores per node.
- More calibrations to come: new TPC distortion calibration (Run2 & Run3).
  - Mandatory in Run3.