

LHCC Referee Meeting

21/02/2017

ALICE Status Report

Predrag Buncic

CERN

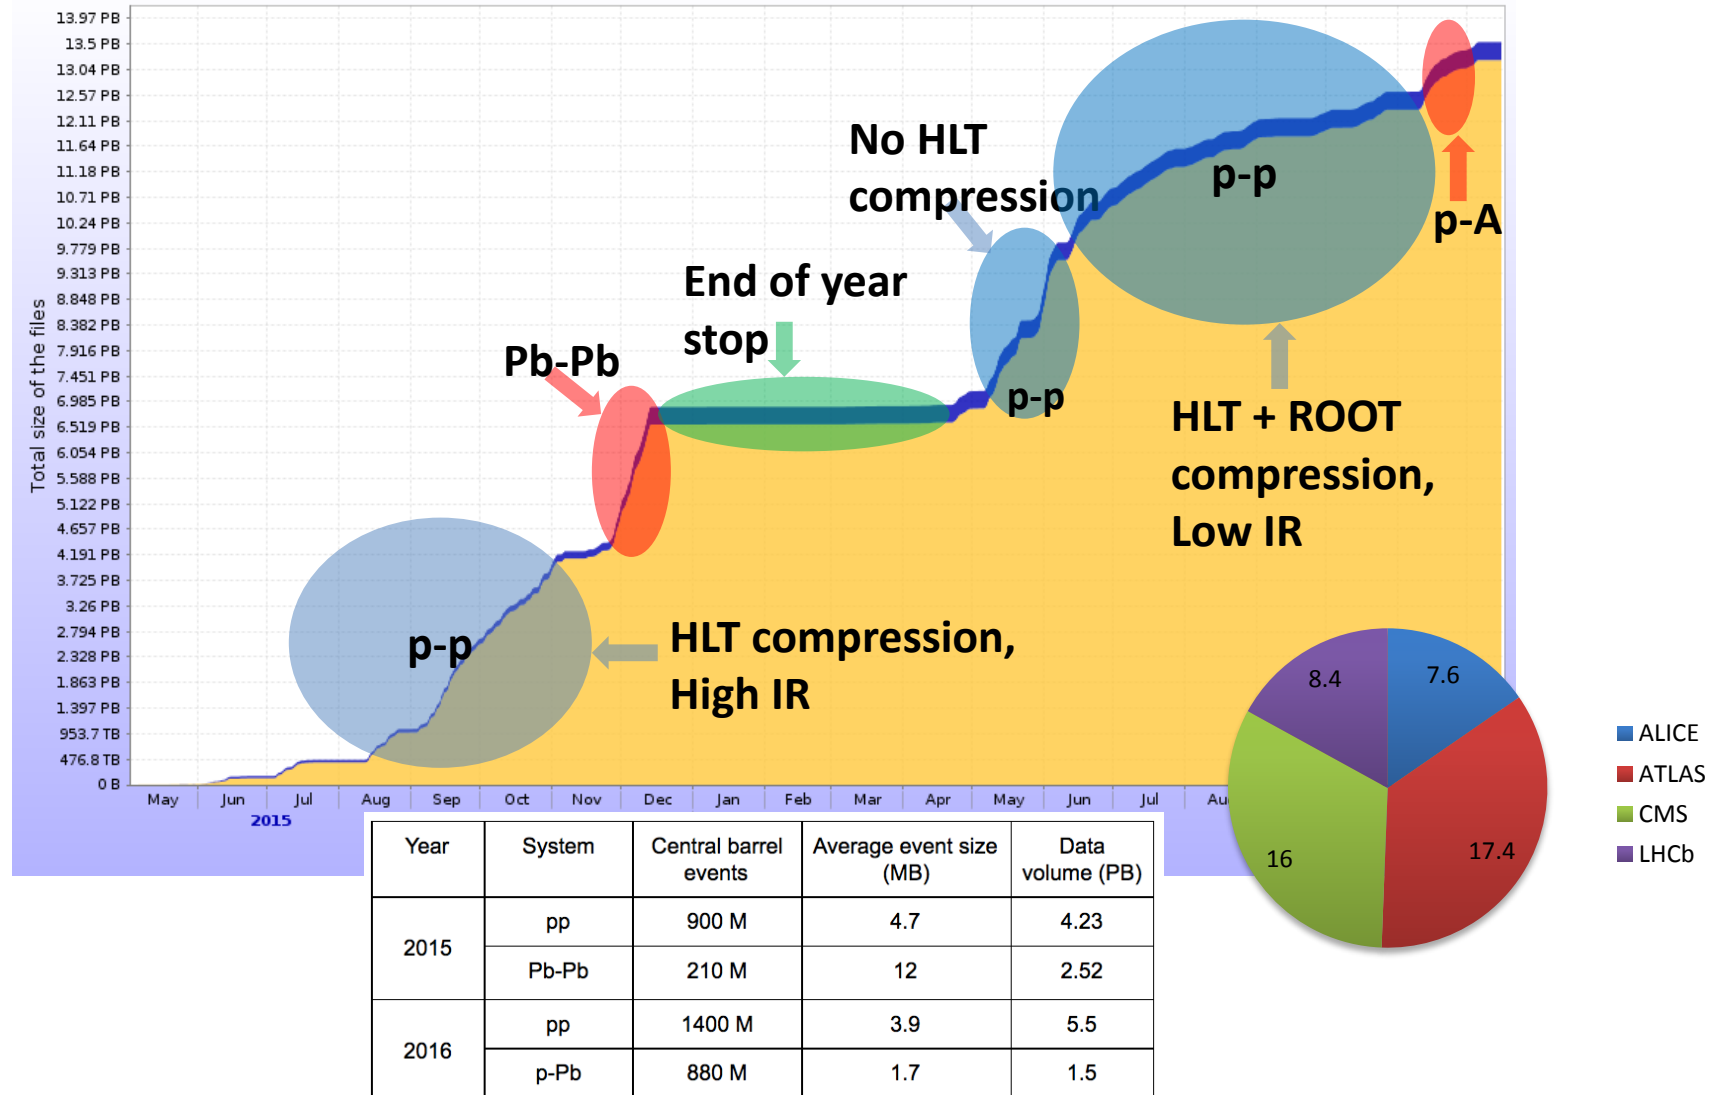


ALICE

Run 2 data taking objectives

- For Pb-Pb collisions:
 - Reach the target of 1 nb^{-1} integrated luminosity in Pb-Pb for rare triggers.
 - Increase the statistics of the unbiased data sample, including minimum bias and centrality triggered events.
- For pp collisions:
 - Collect a reference rare triggers sample with an integrated luminosity of 40 pb^{-1} , which is equivalent to the 1 nb^{-1} sample in Pb-Pb collisions.
 - Enlarge the statistics of the unbiased data sample, including minimum bias collisions at top energy.
 - Collect a reference sample of 10^9 events at the reference energy of 5.02 TeV
- For p-Pb collisions:
 - Enlarge the existing data sample, in particular the unbiased events sample at 5.02 TeV.


Total size of the files




- Readout electronics upgrade (RCU2) with goal to increase the readout rate by a factor of 2
 - Originally planned for 2015, final version installed from 2016

Measures taken to reduce data volume

- Improving HLT compression
 - ALICE does not store raw TPC data on tape
 - We store clusters
 - Standard online TPC cluster compression: RAW data reduction by factor ~ 4.3 to ~ 5.5
 - Differential Huffman compression algorithm developed in 2016 resulting in an additional 10-20% reduction of the TPC event size.



- Reducing IR to 100 kHz
 - Benefiting from LHC performance





- Aiming to reach our data taking goals in more favorable conditions for detector
 - Reduced space charge (distortions) and pileup (event size)

Data processing status

- The TPC track distortion correction algorithms have been finalized and fully validated with both pp and 2015 Pb-Pb data at various interaction rates.
- This has allowed us to start RAW data processing of the
 - 2015 Pb-Pb
 - 2016 p-Pb period
 - longest pp data taking periods, both from 2015 and from 2016
 - complete 2015 and 2016 data for di-muon spectrometer
- The associated Monte-Carlo simulation, anchored to the conditions data and distortion corrections from the RAW data calibration cycles was completed in the beginning of 2017
 - in time for Quark Matter
- Full processing of remaining 2015 and 2016 data is ongoing

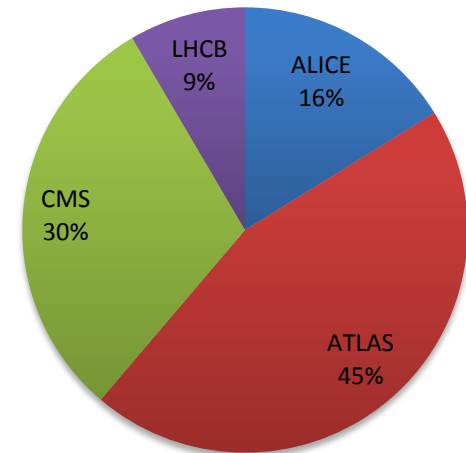

Grid usage

Running jobs per user




- High grid usage, good utilization of opportunistic resources
 - Average 76K parallel jobs, record 112K jobs
 - ALICE HLT cluster provides 4K jobs (5% of total)
- In spite of increased reconstruction load, stable shares of various workflows

ALICE resource usage and shares

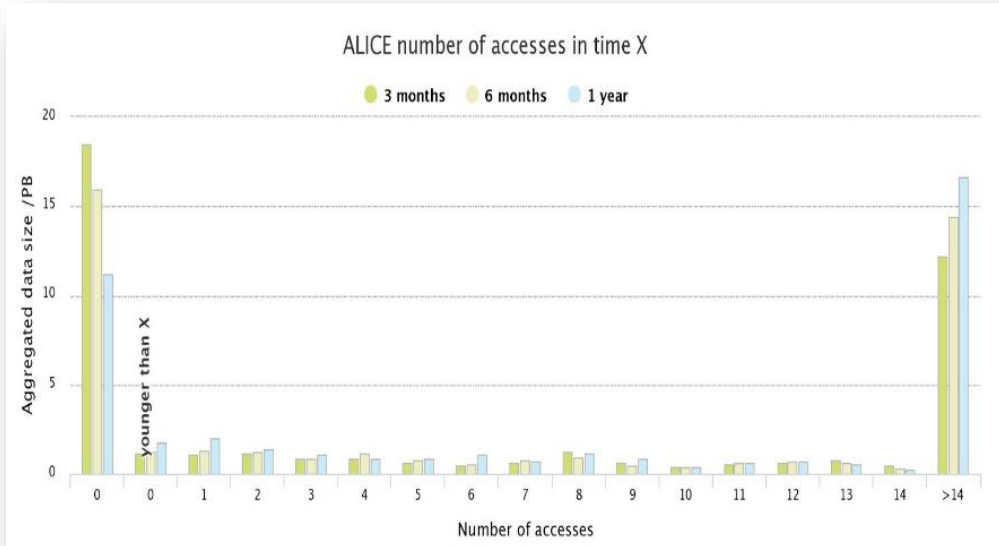


CPU delivered



Tape Used


Popularity and cleanup



- Each raw event is stored only once
- Old inactive MC productions removed
- All reconstruction output + associated MC
 - ESDs are stored as single replica
 - 2 replicas for AODs

Data occupying T0/T1/T2 Storage					
	Event Size [MB]	# of copies on disk		# of versions	# of copies on tape
		minimal	typical		
RAW	3 (pp) 11 (Pb-Pb)			1	2 (one at T0 + one at one of the T1s)
ESD	10 to 30% of RAW, depending on type of collision system and luminosity	1	2	1-3	
AOD	10 to 15% of RAW, depending on type of collision system and luminosity	1	2.6	1-4 per ESD version	
MC ESD	0.37 (pp) 2.7 (Pb-Pb)	1	2	1	
MC AOD	30% of MC ESD	1	2.6	2	

Expectations for 2017 and 2018



- During pp data taking mode will be set to limit the TPC readout rate to 400 Hz
 - The goal is to reach the statistics objective set for Run 2 in all trigger categories as well as at the reference energy of 5.02 TeV
 - The total amount of data recorded will be 17.5 PB
- During the Pb-Pb run in 2018, assuming the HLT compression of a factor of 6 we anticipate a total readout rate of 10 GB/s
 - The total amount of data recorded will be 12 PB

Summary of mitigation measures

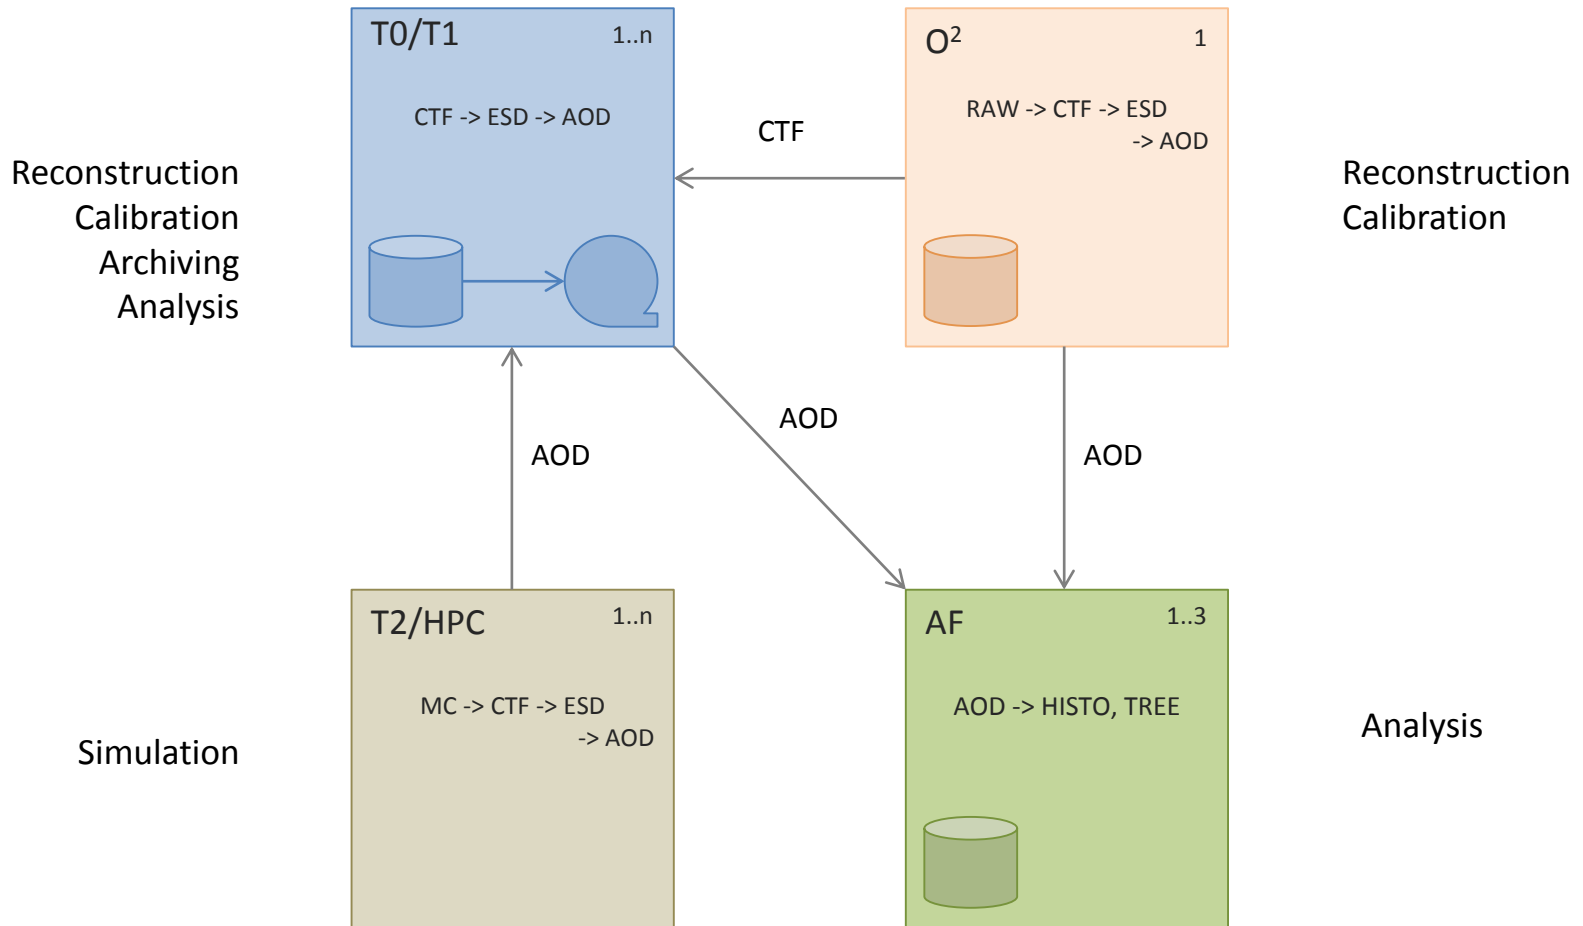
Possible improvements	Gain	Feasibility	Impact on analysis
Simulation - reusing the underlying event for Pb-Pb simulations	% of 70% of total CPU budget	Manpower needed for improvement of the digitization largely overlaps with reconstruction experts who are now fully engaged on upgrade activities	Acceptable (if fully validated and CPU neutral)
Simulation – improve digitization performance			
Simulation – use of G4 transport		Transport with G4 till factor 1.6 slower than G4	
Reconstruction	% of 15% of total CPU budget	Effort directed mostly towards quality improvements. Goals achieved with reasonable (5%) performance penalty.	Positive but comes with performance hit
Analysis	% of 15% of total CPU budget	Moving away from ESD analysis and reducing ESD replicas to one	Positive but not universally applicable

Summary of mitigation measures



Possible improvements	Gain	Feasibility	Impact on analysis
Removing noise clusters in TPC	10% of raw data	In progress. Depends on the QA result after change of TPC gas	Acceptable (conditionally)
Adapt data taking scheme	10% of raw data	Possible if machine and background conditions remain favorable.	Acceptable (conditionally)
Reduce raw data replication	50 % of tapes	Generally possible but carries increased risk of data loss.	Not acceptable
Reduce number of ESD/AOD replicas	n/a	Not feasible, already at minimum	Not acceptable
Partial parking/delayed replication and processing of 2018 pp data	9 PB of tapes at T1s	Possible but this could have an impact on ALICE competitiveness with the other LHC experiments.	Acceptable

Backup

Run 3 Computing Model



Changing Disk/CPU ratio on T2s



- Expecting Grid resources (CPU, storage) to grow at 20% per year rate
 - Large number of disk will be used by Run 1 and Run 2 data
- Since T2s will be used almost exclusively for simulation jobs (no input) and resulting AODs will be exported to T1s/AFs, we expect to significantly lower needs for storage on T2s

Deletion/parking policy



- With the exception of raw data (CTF) and derived analysis data (AOD), all other intermediate data created at various processing stages is transient (removed after a given processing step) or temporary (with limited lifetime)
 - CTF and AODs are archived to tape
- Given the limited size of the disk buffers in O2 and Tier 1s, all CTF data collected in the previous year, will have to be removed before new data taking period starts.
- All data not finally processed during this period will remain parked on tapes until the next opportunity for re-processing arises: LS3

