



# ALICE Status Report

Update from November LHCC referees meeting


Stefano Piano



# Status of data processing and resources utilization

# Computing: status of data processing

- 2018 Pb-Pb Pass2 production started at the beginning of December 2020
  - ⇒ followed by general purpose MC
    - Fine calibration and code improvements obtained for the Pb-Pb 2018 RAW Pass2
- 2020 processing schedule: 2015 Pb-Pb data, all pp data from Run2, 2016 p-Pb data
  - uniform processing quality and software features for the entire Run 2 data
- 2020 schedule parameters:
  - CPU with two months LS2 extension
  - Disk: data removal campaign



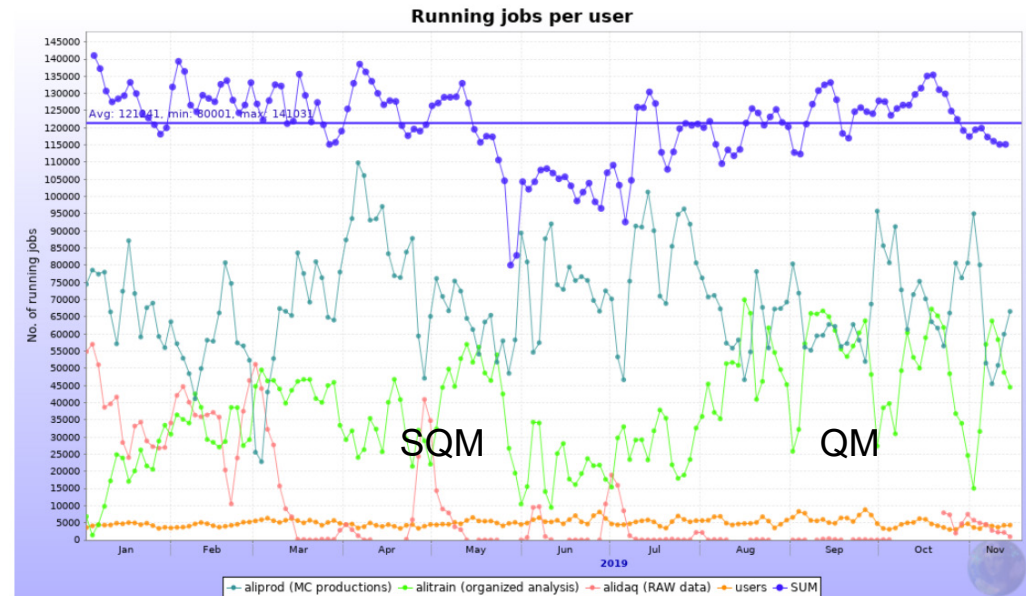
# CPU and disk usage by job type

- Increased amount of CPU use for analysis
  - Summer and autumn conference preparations (SQM/QM)
  - Full Run 2 data samples

Resource	Sim	Reco	Data Analysis
CPU	56%	7%	37%(*)
Disk	54%	39%	7%

Percentage of the CPU and disk storage  
from April 2019 to January 2020

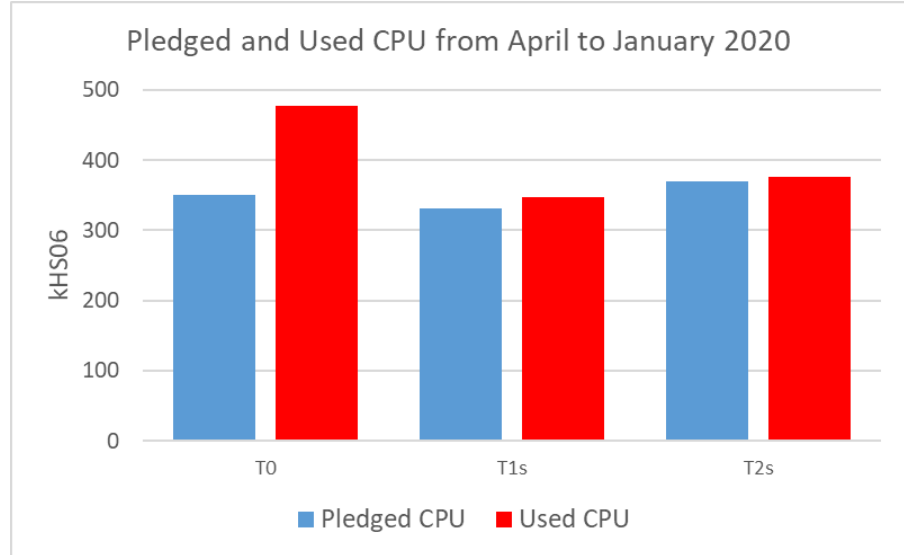
(\*) Organized Analysis (train): 35%  
Individual analysis at ~2% level



# Resources utilization: CPU




- Full utilization of available CPU:  
>120 k jobs average
- Stable resources delivery at all tiers
- Significant amount of opportunistic CPU usage at the T0
- The level of opportunistic usage at T1s and T2s has recovered to 5% and 2%

→ expected after the period of extensive software updates



# Resource utilization: DISK

- Disk capacity is sufficient to proceed with all planned production and analysis activities (expected additional output until April  $\sim 4.4$ PB  $\Rightarrow$  occupancy 95%)
- Keep monitoring and removing the infrequently used data sets:
  - The amount of unpopular data remains at about 2% of the total used disk space
- On January 31<sup>st</sup> used/pledged: 99% at T0, 89% at T1s, 84% at T2s



# Resource utilization: TAPE

- Run 2 computing model:
  - Tape mainly used as a RAW data store
  - Occupied volume unchanged with respect to the previous period
  - Tape requirements for 2020 are equal to 2019 (0% increase)
- New computing model for Run 3:
  - All CTFs and the corresponding AODs will be archived to tape:
    - 2/3 at the T0 (+36 PB for Pb-Pb CTF's in 2021)
    - 1/3 at the T1s (+18 PB for Pb-Pb CTF's in 2021)




# Resource requirements for 2020 and for 2021



# Resource requirements for 2020

- Discussed and approved by RRB in October:
  - ALICE CPU and TAPE requirements will remain constant
  - Disk resources are required to grow by 16% at T1s and by 15% at T2s



# Resource requirements for the 2021 pp period

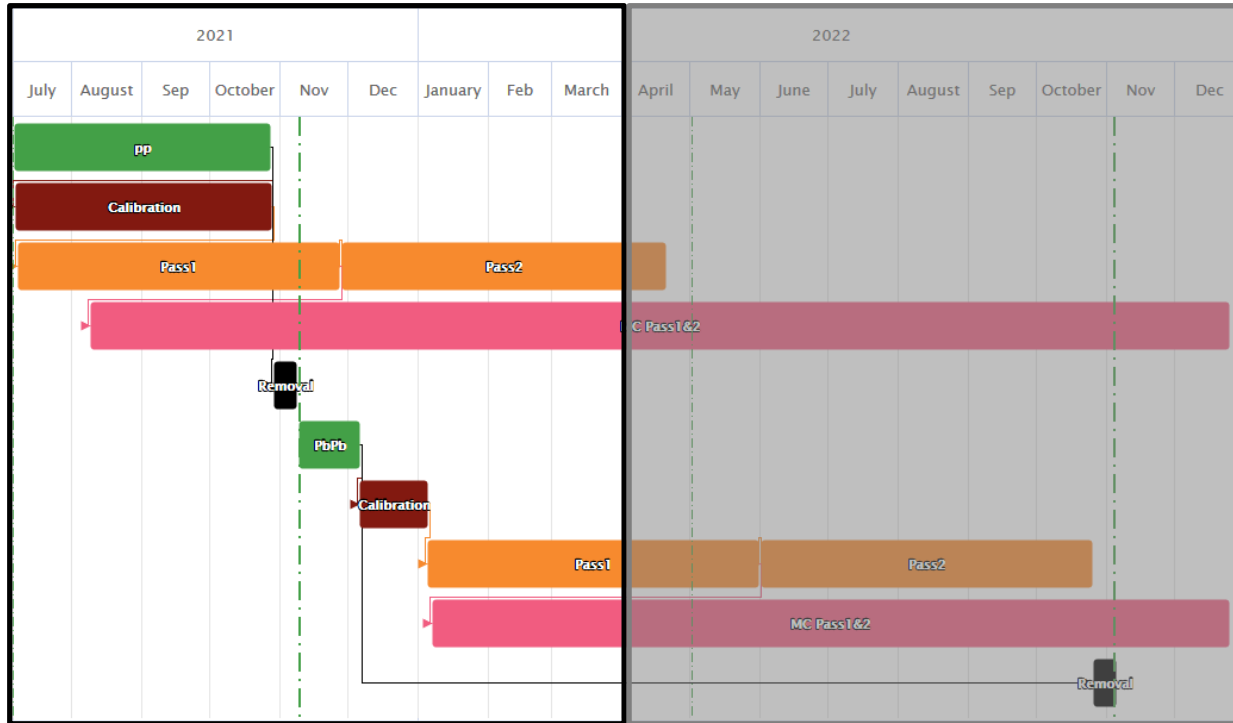
- ALICE plans to run during the full pp period in 2021
  - Commission and gain experience with the upgraded detector
  - Keep a fraction of about 1‰ of partially compressed and uncompressed TFs and the corresponding CTFs for reference (+3 PB archived to tape)
  - Collect data samples that enable a complete validation of the calibration, processing and analysis chain (+3.5 PB archived to tape)
- LHC running conditions for the computing estimates for 2021:
  - Assumption for the proton run is to deliver up to  $45 \text{ pb}^{-1}$
  - ALICE interaction rate will be adjusted to about 1 MHz for most of the period
  - Two weeks of pp collisions at 5 MHz interaction rate:
    - Equivalent load of Pb-Pb collisions at 50 kHz interaction rate

# Resource requirements for the 2021 Pb-Pb period

- The assumption is to record collisions for 28 days with:
  - IR up to 50 kHz and an average rate of 23 kHz
  - Combined efficiency factor of 48%  
⇒  $3.5 \text{ nb}^{-1}$  of integrated luminosity
  - Considered as an upper limit for the 2021 HI period
- $2.7 \times 10^{10}$  recorded events at 2.0 MB/event:
  - Total CTF volume of about 54 PB
  - Initially stored on the O2 disk buffer
  - 1/3 of CTFs will be exported, archived to tape and subsequently processed on T1s
  - 2/3 of CTFs are stored at the T0 and processed by O2 facility and T0
  - Total AOD output of ~20 PB (2 replicas of 5 PB for 2 reconstruction passes)

LHC parameters	Pb-Pb
Run time (days)	28
Run time (s)	$2.42 \times 10^6$
Efficiency factor	48%
Effective time (s)	$1.15 \times 10^6$
# of collisions	$2.66 \times 10^{10}$

# Resource requirements for the 2021 run



- **The two asynchronous passes** executed in the 10 months following the Pb-Pb data taking with O2, T0 and T1s resources:
  - Data volume: ~54PB CTFs, ~20 PB AODs
  - Computing power: ~430 kHS06 for 10 months
- AODs will be persistent on disk and archived on tape at the T0/T1s
- 1 MC pass: ~500 kHS06 for 11 months

# Resource Requirements for 2020 and for 2021

ALICE		2019			2020				2021	
		Req.	C-RSG	Pledge	Req.	C-RSG	Pledge	2020/ 2019 pledge	Req.	2021/ 2020 Req
CPU	Tier-0	430	430	350	350	350	350	0.0%	471	34.6%
	Tier-1	365	365	331	365	331	353	10.3%	498	36.4%
	Tier-2	376	376	370	376	376	410	1.6%	515	37.0%
	Total	1171	1171	1051	1091	1057	1113	3.8%	1484	36.0%
Disk	Tier-0	34.3	34.3	31.2	31.2	31.2	31.2	0.0%	45.5	45.8%
	Tier-1	37.9	37.9	35.1	44	44	41.8	25.4%	53.3	21.1%
	Tier-2	33.9	33.9	33.5	39	39	41.0	16.4%	44.8	14.9%
	Total	106.1	106.1	99.8	114.2	114.2	114.0	14.4%	143.6	25.7%
Tape	Tier-0	44.2	44.2	44.2	44.2	44.2	44.2	0.0%	86.0	94.6%
	Tier-1	37.7	37.7	41.1	37.7	37.7	44.4	-8.3%	57.0	51.2%
	Total	81.9	81.9	85.3	81.9	81.9	88.6	-4.0%	143.0	74.6%



# Run 3 computing upgrade updates

# GPU Reconstruction and Benchmarks

- **Time performance benchmarks**
  - Input needed for CPU's and GPU's EPN specification
  - Benchmark tests ongoing on NVIDIA and AMD GPUs
  - Computing time performance optimization ongoing, in particular for AMD
- **Majority of Barrel Tracking reconstruction is GPU ready**
  - All known issues fixed, GPU output matches CPU output
  - Latest element implemented is the **TPC Cluster Finder**:
    - Estimate of  $\leq 500$  GPU turns out to be valid
    - Final number of GPUs still under investigation
- **GPU Memory Requirements**
  - 14 GB for 128 LHC orbits (11.4 ms)
  - Fits on 16 GB GPU
  - TDR: timeframe in range of 11.4-22.7 ms

# Preparation for Mock Data Challenges

- **Move from tailored (pragmatic) individual workflows like ...**
  - Run2 AOD conversion -> Run3 analysis
  - TPC/ITS Run3 simulation -> reconstruction benchmarks
- **... to full (realistic) Run3 workflows**
  - Simulation (including MC-to-MC embedding) -> Reconstruction -> Analysis
  - Analysis challenge ready in July 2020 (Milestone PDP.7)
- **Local integration tests have started**
  - Fixing now a few loose ends
- **In parallel: simulation performance optimization**
  - Transport
  - Digitization



# Focus on Calibration

- **Definition of the role of calibration in the overall data processing flow**
  - Calibration data is collected from different sources: FLP, EPN, DCS, LHC
  - ... and plays a different role for different reconstruction stages: synchronous, asynchronous
- A major item is the real-time conditions data distribution in the synchronous stage
  - Solved through multicasting, tested and working
- Aggregation of calibration data and storing into the Constants and Calibration DB
  - A new system to aggregate the data is being developed (similar to one used in RUN1/2)
- Integration of the CCDB into Data Processing Layer - ongoing according to plan

# Vertical slice test

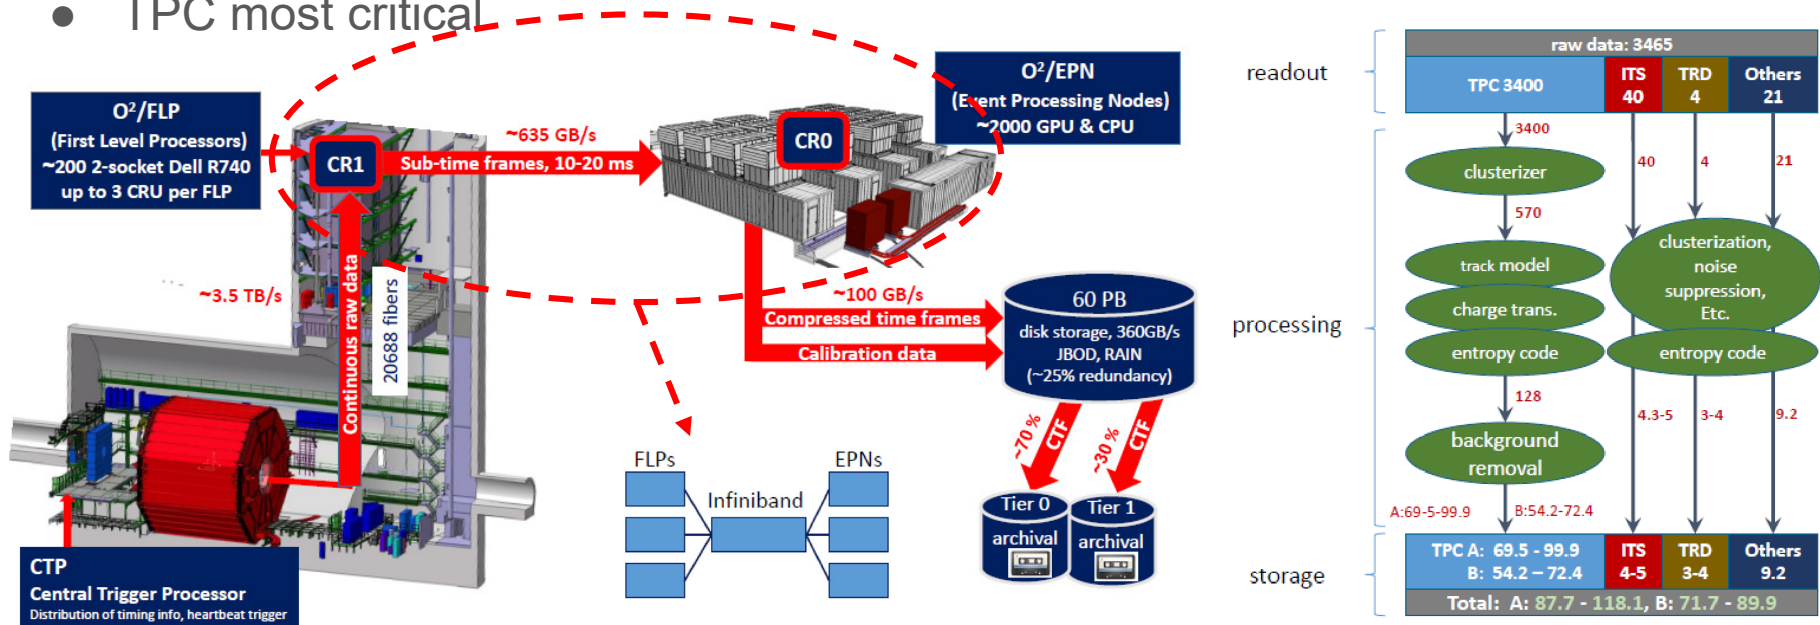
- Elements for the VS available in October 2019
- **VS started in November 2019**
  - IP traffic over the InfiniBand network
  - Storage tests (in connection with IT ST and IT CS)
  - Tests going on
  - Integration work taking momentum
- Grow into the final system (while supporting global commissioning)

# Summary

- Physics analysis and Pass2/3 of 2018 RAW data are ongoing
- Continuous full utilization of the computing resources
- ALICE is in the critical phase of the Run3 upgrade preparation
- Software algorithms and updated computing model allow to fit into the standard Grid resource growth
- Small step of requests in 2021 for tapes, CPU and disk compatible with standard resources growth over the entire 2018-2021 period
- Focus on computing upgrade activities - results on track with the expectations
- Vertical Slice Test and integration of the entire data acquisition, simulation and processing chain started on Q4 2019

# Status of O2 reconstruction tasks

- Up to 50 kHz of Pb-Pb continuous readout, data compression mandatory
- TPC most critical



[Shahoyan, R. CHEP19, Plenary](#)

# GPU Reconstruction Software Status

- GPU components for baseline scenario almost finished:
  - Material lookup during tracking not finished (not strictly needed for TPC).
  - TPC Track Merger still runs certain steps on the CPU, not critical.
  - Junk identification below 10 MeV/c missing (still searching for a good algorithm, affects compression ratio).
  - TPC entropy compression on GPU missing (not strictly needed, can run on CPU).
- Optimistic scenario for better GPU utilization in asynchronous reconstruction, fully work in progress.

