



ALICE Status Report

Update from June LHCC referees meeting

Stefano Piano




Status of data processing and resources utilization

Status of data processing

- p-p data
 - Pass 2, Pass 3 RAW of various p-p periods ongoing
- 2018 Pb-Pb data:
 - Pass2 production to be started before the end of the 2019
- Increased amount of CPU used for analysis (20% to 33%):
 - Full Run2 data samples
 - Analysis efforts for the 2019 conferences: (SQM – June and QM – November)

Resources utilization: CPU


- Continuous high utilization
- Stable resources delivery at all tiers,
- Migration from SL6 to CentOS7 and new gateways at all sites
 - Largely completed, ~10% still to go
- Slightly longer downtimes => less opportunistic resources utilization



Resources utilization: DISK

- Available disk capacity is sufficient to proceed with all planned production and analysis activities
- Keep monitoring and removing the infrequently used data sets:
 - The amount of unpopular data remains at about 2% of the total used disk space

Pledged, deployed and used disk up to 31 July 2019



Resources utilization: TAPE

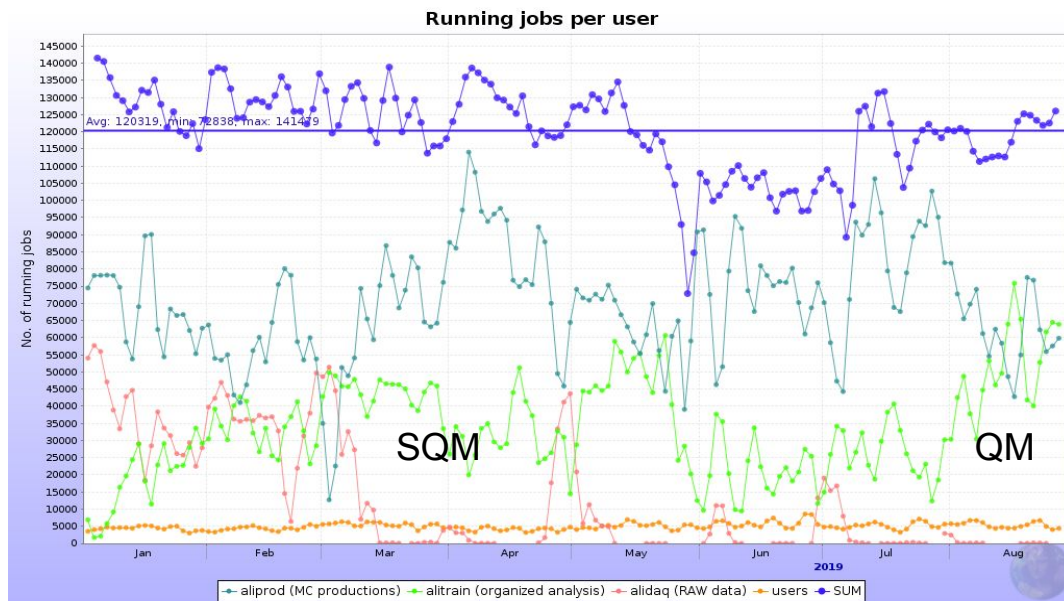
- No change since end of 2018
- Requirements for 2020 are adjusted to actual use
 - T0 - 41.4PB (94% of the pledged)
 - T1s - 35.8 PB (87% of the pledged)
 - Mainly due to the different than planned trigger composition in 2018 Pb-Pb period
- Uniform usage distribution across T1s
- Plan to archive to tape unpopular datasets still needed as physics reference

CPU and disk usage by job type

- Increased amount of CPU use for analysis
 - Summer and autumn conference preparations (SQM/QM)

Resource	Sim	Reco	Data Analysis
CPU	61%	6%	33%(*)
Disk	58%	36%	6%

(*) Organized Analysis (train):
98% of Analysis Jobs
Individual analysis at ~2% level





Resource requirements for 2020 and estimates for 2021

Resource Requirements for 2020

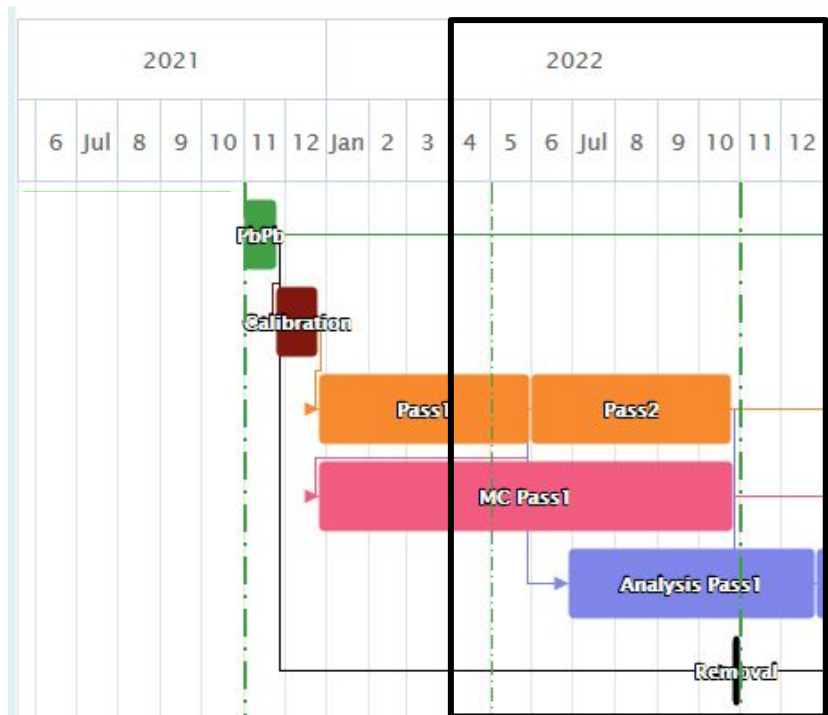
ALICE		2019			2020		
		Req.	C-RSG	Pledge	Req.	2020/2019 CRSG	2020/2019 pledge
CPU	Tier-0	430	430	350	350	-18.6%	0.0%
	Tier-1	365	365	331	365	0.0%	10.3%
	Tier-2	376	376	368	376	0.0%	2.2%
	Total	1171	1171	1049	1091	-6.8%	4.0%
Detector	Tier-0	34.3	34.3	31.2	31.2	-9.0%	0.0%
	Tier-1	5.1	5.1	44.0	44.0	16.1%	25.4%
	Tier-2	3.4	3.4	39.0	39.0	15.0%	16.7%
	Total	9.7	9.7	114.2	114.2	7.6%	14.5%
Tape	Tier-0	44.2	44.2	44.2	44.2	0.0%	0.0%
	Tier-1	37.7	37.7	41.1	37.7	0.0%	-8.3%
	Total	81.9	81.9	85.3	81.9	0.0%	-4.0%

Endorsed by CRSG in April 2019
No changes with respect to previous report

Running scenario for pp period in 2021

- Expect to run most of the year for commissioning
 - New detectors calibration and performance
 - For (new) physics
- Target run plan for pp in 2021 defined:
 - Run event selection algorithms inspecting the full 1 MHz pp interaction rate
 - Strong filter (1‰) to select interesting events during the processing stage on EPN only
 - $2.7 \cdot 10^{10}$ ($6.7 \cdot 10^{10}$) events in the baseline (upper limit) scenario
- Discussion and development of on-line event selection started
- **The total resources needed for the processing of selected events at T0 and T1s are less than 2‰ (5‰) for the baseline (upper limit) scenario**, compared to the resources needed for the processing of the Pb-Pb data
- ***Essentially no increase of computing resources at T0 and T1s wrt the previous estimates***

Estimates of requirements for the 2021 Pb-Pb run



- The two asynchronous passes executed in the 10 months following the Pb-Pb data taking with O2, T0 and T1s resources:
 - Data volume: ~50PB CTFs, ~20 PB AODs
 - Computing power: ~640 kHS06 for 10 months
- AODs will be persistent on disk and archived on tape at the T0/T1s.

Estimated MC requirements

Results from in-depth evaluation of the Physics Working Groups input:

- For Pb-Pb estimated 2.0% of the number of physics events (0.5% in TDR)
 - The necessary computing power for 2 passes in 5+5 months becomes too high
- Mitigation for the computing resources
 - Reduce the MC passes to 1, i.e. about 510 kHS06 for 10 months, corresponding to our original 2021 CPU request for all T2s
 - Assisted by a considerable reduction of the CPU activities on the Run2 data
- For pp estimated MC requirements in MB equivalent signal events: 75% of RAW data events: 150 kHS06 (baseline scenario) for 1 pass in 150 days



Run 3 computing upgrade updates

O2 Simulation

- Development completed for all large and majority of small detectors
 - Entered stage of refinements and algorithm optimization
- Reduction of CPU and output in Embedding techniques methods
- Review of O2 TPC digitization code
 - Important gain for signal embedding:
 - Substantial improvement in CPU performance while keeping constant physics quality
 - Largely simplified GEM Amplification scheme
- Optimization of transport time (G3/G4) through cuts and geometry configuration
- Fast simulation
- Virtual MC supports simulation using several transport engines
- Successful tests on multicore Grid queues and SC at CINECA/Marconi


Status of O2 reconstruction tasks

		CPU	GPU	Comments
TPC	Tracking	done	done	
	dE/dX	done	done	Calibration not applied yet
	Compression	done	done	Depends on Common entropy coding library (in progress)
	Clustering	In progress		
ITS	Tracking finding/fitting	done	done	Extra passes for secondaries to be tuned
	ITS-TPC matching	done	Q3/2019	Afterburner for deep secondaries being finalized
	Compression	done	-	Depends on Common entropy coding library (in progress)
TRD	Matching to ITS-TPC	Done (in "HLT")	done	To be adapted for O2 TRD once its simulation works
TOF	Matching to ITS-TPC	done	Q2/2019	
EMCAL	Data decoding, clustering	done	-	Clustering (analysis stage) will be added in September
PHOS	Data decoding	Q2/2019	-	Foreseen in September
CPV	Data decoding	Q2/2019	-	Foreseen in September
MUON	MCH clustering, tracking	Q4/2019	-	
	MID tracking	done		
MFT	Tracking (standalone)	done	-	
	Matching to MCH	Depends on MCH		
FIT	T0+ reconstruction	done	-	
	V0+ reconstruction	ongoing		
	FDD reconstruction	done		Reconstruction from digits
ZDC	ZDC reconstruction			Code-sprint foreseen in September
HMPID	Clustering, matching	Q2/2019	-	Clusterization started

- Development almost completed for all large and majority of small detectors
- The work is on track and we are confident will be finished in time
- Sufficient amount of code to be tested during the vertical test exercise



CR0



- 2.2MW total IT power
- All containers and cooling units installed and operational
- Optical fibers to CR1 installed
- Network between CR0 and CERN data center installed by CERN-IT:
 - Currently 4 links of 100 GbE
 - Final installation - 20 links of 100 GbE over two fiber pairs for redundancy

Vertical slice test

- InfiniBand purchase delayed due to a dispute about CERN general terms and conditions
 - Expecting delivery in coming weeks
 - Installation will start immediately as soon as the InfiniBand components are delivered
- EPN servers:
 - 50 servers from Run-2 HLT will be equipped with the new InfiniBand HDR100 HCAs (100 Gb/s)
- Gateway and infrastructure nodes:
 - EPN Login node (access from GPN only)
 - InfiniBand to Ethernet gateway (access to Storage – EOS)

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
- Preparations ongoing for Vertical Slice Test of the entire data acquisition, simulation and processing chain



Resources requirements for 2019-2020 and estimates for 2021

ALICE		2019			2020			2021	
		Req.	C-RSG	Pledge	Req.	2020/2019 CRSG	2020/2019 pledge	Req.	2021/2020 Req
CPU	Tier-0	430	430	350	350	-18.6%	0.0%	471	34.6%
	Tier-1	365	365	331	365	0.0%	10.3%	498	36.4%
	Tier-2	376	376	370	376	0.0%	1.6%	515	37.0%
	Total	1171	1171	1051	1091	-6.8%	3.8%	1484	36.0%
Disk	Tier-0	34.3	34.3	31.2	31.2	-9.0%	0.0%	45.5	45.8%
	Tier-1	37.9	37.9	35.1	44	16.1%	25.4%	53.3	21.1%
	Tier-2	33.9	33.9	33.5	39	15.0%	16.4%	44.8	14.9%
	Total	106.1	106.1	99.8	114.2	7.6%	14.4%	143.6	25.7%
Tape	Tier-0	44.2	44.2	44.2	44.2	0.0%	0.0%	80.0	81.0%
	Tier-1	37.7	37.7	41.1	37.7	0.0%	-8.3%	55.0	45.9%
	Total	81.9	81.9	85.3	81.9	0.0%	-4.0%	135.0	64.8%

MC Projections

- Embedding Signal into Pb-Pb events:
 - Run2 experience: systematics limit reached by Heavy Flavor productions:
 - 2×10^8 MB equ. events
 - about 9 other channels on the $\sim 2 \times 10^7$ level
- For Run3 assumed 10 channels with 2×10^8 MB equ. events:
 - 2×10^9 events (nominally 2% of data)
- Refinements expected from further experience with 2018 Pb-Pb analysis and simulation challenges