Computing resources requirements 2013-2017

Wednesday, March 6, 13

This note presents an estimate of the ALICE computing resources needed in 2013-2014 during LS1 and a tentative estimate of the resources needed after LS1 and until LS2, i.e. for the period 2015-2017. It is based on our best knowledge of the LHC beam schedule and luminosity, the computing model procedures and the data taking and processing strategies.

Data processing

During LS1 (2013-2014)

At the start of LS1, the quantity of physics data recorded by ALICE since the beginning of the LHC operation is summarized in Table 1. These data are a mix of minimum bias events and rare events triggered by single and two muons, calorimeter, TRD, diffractive, ultra-peripheral collision events and high multiplicity events. The mix was varied from the beginning to the end of the data-taking period, giving an increased relative weight to rare events.

Table 1: Amount of physics events collected by ALICE for p-p, Pb-Pb and p-Pb collisions since 2010.

System	Number of events
p-p	2.5×10^{9}
Pb-Pb	5.0×10^{8}
p-Pb	1.8×10^{8}

The raw data and calibration data have been stored on the Tier0 taping system (10.2 PB) and raw events passing the physics selection have been replicated once in the Tier1s taping system (4.6 PB). All the raw data have been processed at least in one pass and the corresponding Monte-Carlo data have been produced, generating altogether disk space occupancy of about 18 PB (obsolete data have been permanently deleted).

During LS1 we will reprocess the entire set of data collected taking advantage of the best available calibration parameters and the optimal tuning of the reconstruction parameters.

During the same period a major upgrade of the whole offline environment for reconstruction, analysis and simulation is foreseen to improve the software quality and performance. In addition new developments resulting from the R&D program, including parallelization, vectorization, GPU algorithms, new algorithms, will be implemented. The new environment will become available in 2014. A partial reprocessing of the data will then be performed.

From LS1 to LS2 (2015-2017)

During LS1 ALICE will install additional detectors: the complete TRD covering from the now 60% to 100% azimuthal coverage and DCAL the electromagnetic calorimeter facing in azimuth the existing one. After LS1 p-p and Pb-Pb collisions will be provided at a new energy. For these reasons ALICE plans to collect in one year the same amount of p-p data as the one collected up to now. During the three years before LS2, ALICE plans to collect the missing data to reach the integrated luminosity of 1 nb^{-1} of PbPb collisions originally approved. To do so ALICE will take advantage of the four times higher luminosity anticipated for Pb-Pb collisions and acquire of the order the 10^{9} minimum bias data, the rest of the data being rare events triggered by MUON and EMCAL. This represents every year four times the amount of data collected in 2012.

We assume that LHC will operate in 2015, 2016 and 2017 as in a standard year, i.e., 5.2×10^{6} effective seconds of p-p collisions and 0.7×10^{6} effective seconds of Pb-Pb collisions, assuming an overall efficiency of 38%. We anticipate collecting 2.5×10^{9} p-p events during the first year and 2×10^{9} p-p events during the two following years, and 5.2×10^{9} Pb-Pb events in 3 years with a trigger mix assumed to be the same as in 2012 for pp and as in 2011 for Pb-Pb. Considering the higher beam energy and luminosity we have assumed an increase in the track multiplicity of 25%. We assumed that LS2 would start end of 2017.

Computing model parameters

The computing model parameters (processing power and data size) have been taken as the average values extracted from the 2012 data processing of p-p and Pb-Pb data. For the resources needed after LS1, estimates are based on the same CPU power for reconstruction and raw event size augmented by 25% to take into account the increase of the track multiplicity due to the higher beams energy and increased pile up. The computing powers needed to process one event are reported in Table 2. Note the differences with the parameters value presented in October 2012:

1. p-p reconstruction time has increased due to the more complex events due to increased pileup and the relative weight of rare triggered events with larger tracks multiplicity than minimum bias events.

2. Pb-Pb reconstruction time has substantially decreased because of large contribution of MUON triggered events that do not include TPC data and thus are much less costly to reconstruct.
3. The Monte-Carlo processing time has decreased for both p-p and Pb-Pb thanks to a more systematic use of signal embedding in existing minimum bias events.

The data sizes at the various stages of the processing are reported in Table 3. A factor 4 for raw data compression has been considered. Replication of the reconstructed data is now limited to two instances instead of three as adopted in the previous years.

Table 2: *Processing power in kHEPSpec seconds per event*

kHEP06×s/event		Reconstruction	Analysis train	End user analysis	Monte Carlo
pp	2012	0.17	0.13	0.01	1.01
PbPb	2012	0.71	1.99	0.14	24.69
pPb	2012	0.38	0.90	0.07	3.19
pp	2017	0.22	0.17	0.01	1.37
PbPb	2017	0.89	2.49	0.17	30.86

Table 3: *Data sizes in MB/event*

MB/event		Raw	ESD&AOD	Monte-Carlo
pp	2012	0.84	0.13	0.26
PbPb	2012	6.00	1.24	16.87
pPb	2012	1.30	0.26	1.38
pp	2017	1.05	0.16	0.37
PbPb	2017	7.50	1.55	21.09

Computing resources requirements

2013-2014

The requirements have been evaluated based on our original computing model for data processing. The reprocessing of the Pb-Pb collisions data set drives the CPU requirements in 2013-2014. We request to process these data in no more than 5 months to provide sufficient time to fully analyse the data in time for the major conference in the field (QM2014) scheduled beginning of 2014. This re-processing completed with the corresponding Monte-Carlo production (about

35% of p-p real data, 30% of Pb-Pb data and 100% of p-Pb data) and associated data analysis sets the amount of CPU required in 2013-2014 (Table 4). The distribution of the CPU share in tier categories assumes that all reconstruction tasks are performed in Tier0 and Tier1s and that Monte-Carlo production and analysis are shared among Tier2s (50%) and Tier0+Tier1s. The resulting distribution is 27% in Tier0 and Tier1s and 46% in Tier2s.

All together the 2013-2014 productions will generate 9.2 PB of data (where 21% is for reconstructed data, 9% for analysis-generated data and 69% for Monte-Carlo data) to be available on disk, including a replication factor of 2. The share of each tier category is 30% in Tier0, 30% in Tier1s and 39% in Tier2s. To cope with the pledges for 2013, we will reduce the disk occupation of existing data by limiting the number of instances to one. As of today the amount of data on disk is 22 PB, which will shrink to 11 PB after removal of the replicas. There is no request to increase the tape capacity, on the contrary we are ready to reduce our requirements and trade tape for disk. If possible this strategy would smoothen the strong disk storage increase required in 2015.

2015-2017

The time profile of the required resources assumes that the heavy-ion runs are scheduled toward the end of the year. Within this scenario the resources required for a given year can be installed during the second semester. It is imported that the resources requested in Tier0 are covered allowing us to process the first reconstruction pass of heavy-ion data promptly in 4 months. The share in Tier1s and Tier2s can be further adjusted depending on the pledges, however the sum of the requested resources in Tier1s and Tier2s is essential to allow us processing the data (reconstruction and analysis) and produce the associated Monte-Carlo data within the year following the heavy-ion run. The disk usage has been estimated in a way to store on disk one reconstruction pass with two replica of all data collected between 2015 and 2017 plus a fraction of the associated Monte-Carlo data limited to keep the amount of requested disk storage at a “reasonable” level. New disk storage can be installed any time during a given year and also during the preceding year. Any new disk can be quickly used and will help to process more efficiently analysis tasks.

Resources required in 2013-2017 are listed in Table 4-Table 6.

Table 4: *CPU requirements for 2013-2017*[1]

CPU (kHEPSPEC06)	Tier0	CAF	Tier1s	Tier2s
2013	91	35.0	101	188
2014	100	35.0	110	190
2015	145	45.0	110	200
2016	150	45.0	130	220
2017	150	45.0	160	240

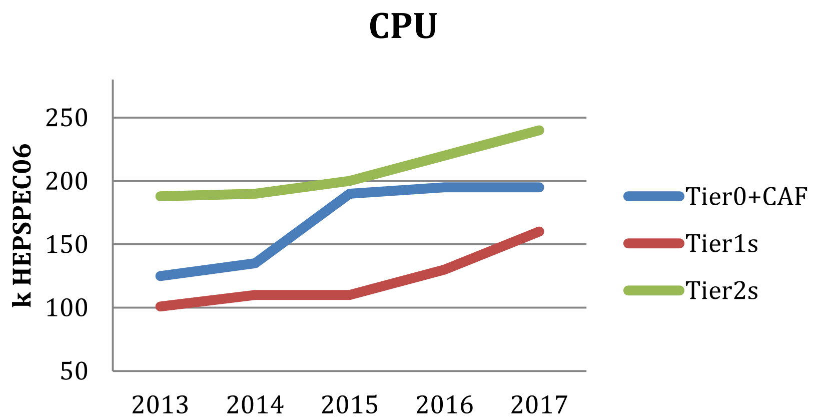

Figure 1: CPU requirement profile. Resources for a given year can be installed during the second semester.

Table 5: *Tape requirements for 2013-2017*

Tape (PB)	Tier0	Tier1
2013	12.0	6.0
2014	12.0	6.0
2015	27.0	21.0
2016	40.8	34.8
2017	54.6	48.6

[1] The 2013 values are not the ones approved by the RRB in October 2012 (see Table 1 of Usage document), but have been updated to reflect the presently pledged values as reported by REBUS (see Table 2 of Usage Document).

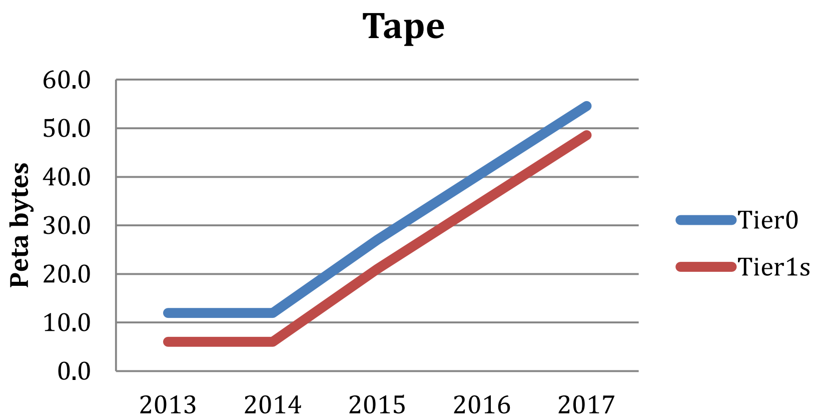

Figure 2: Tape requirement profile. Resources for a given year can be installed at the beginning of the following year.

Table 6: ***Disk requirements for 2013-2017*[2]**

Disk (PB)	Tier0	CAF	Tier1s	Tier2s
2013	8.1	0.24	7.7	12.8
2014	8.1	0.24	7.7	12.8
2015	10.5	0.34	11.2	16.1
2016	11.9	0.44	14.2	20.5
2017	15.3	0.54	17.3	24.8

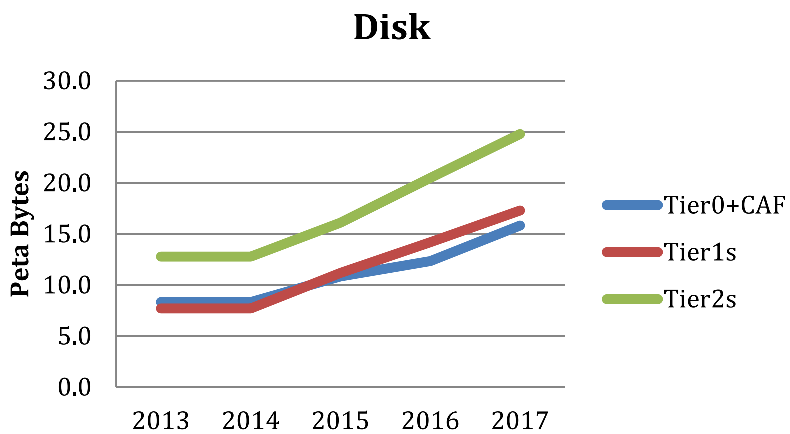

Figure 3: Disk requirement profile. Disks can be installed any time during a giving year or during the previous year.

[2] Same as footnote 1