Appendix I. Detailed Resource Needs and Evolution

In this Appendix we describe the anticipated needs of computing and storage resources over the 3 years of the second LHC run. There is a set of common assumptions on which these requirement estimates are based. These are explained in the first section, followed by detailed explanations of the resource needs for each experiment. This represents a snapshot of the situation in Autumn 2013 and will evolve over time with experience and changing physics requirements.

1.1 General Assumptions

1.1.1 LHC Running time

It is assumed that the live time of the LHC and experiments will follow a similar pattern to that experienced during the 2010-2013 run; during the first year (2015) the LHC must be commissioned and ramp up in availability; 2016 is a nominal year of running, and 2017, being a year before a long shutdown and presumably when the LHC is running at peak performance may again be available for longer. These assumptions are listed in **Error! Reference source not found.** together with assumptions on efficiency and availability, again based on experience in Run 1.

Table 1: Assumptions for LHC pp running

	2015	2016	2017
LHC start date	1/05/2015	01/04/2016	01/04/2017
LHC end date	31/10/2015	31/10/2016	15/12/2017
LHC run days	183	213	258
Fraction of days for physics	0.60	0.70	0.80
LHC efficiency	0.32	0.40	0.40
Approx. running seconds	$3.0\ 10^6$	$5.0\ 10^6$	$7.0\ 10^6$

Assuming typical Pb-Pb or p-Pb running periods in each year as experienced in Run 1, gives the summary shown in **Error! Reference source not found.** for assumed running times during each year in Run 2.

Table 2:Assumed LHC live time (million seconds/year)

Year	pp ($x10^6$) sec	A-A ($x\ 10^6$) sec	Total ($x\ 10^6$) sec
2015	3	0.7	3.7
2016	5	0.7	5.7
2017	7	0.7	7.7
Total	15	2.1	17.1

1.1.2 Assumptions of pileup

ATLAS and CMS have in 2012 presented a complete set of arguments that the 50-ns mode of LHC operations at high pileup would cause several issues, not least a very substantial increase in the computing resources required, thus the assumption in this document is that there will be no extended physics data taking at 50-ns and high pileup values and that LHC will quickly move to 25-ns bunch spacing giving more moderate values of pileup albeit at the risk of somewhat reduced luminosity. Consequently, in 2015 LHC is assumed to achieve stable operation at the average pileup $\mu \approx 25$ for the luminosity of $10^{34}\text{cm}^{-2}\text{s}^{-1}$ at 25-ns bunch spacing. In 2016-2018, the luminosity according to the current LHC scenario could rise up to $1.5 \times 10^{34}\text{cm}^{-2}\text{s}^{-1}$ at 25-ns corresponding to an average pileup $\mu \approx 40$, with an corresponding increase in reconstruction (and pileup simulation) CPU times and event sizes.

For ATLAS and CMS the reconstruction times are highly dependent on the levels of pileup. Based on what is currently understood the likely LHC running conditions (25 ns bunch spacing, anticipated luminosities), the following assumptions are made for the average pileup anticipated in each year (**Error! Reference source not found.**).

Table 3: Assumptions for pileup in ATLAS and CMS

	2015	2016	2017
Average pileup	25	40	40

1.1.3 Efficiency for the use of resources

Since the start of the WLCG the resource requests have had some assumptions regarding the efficiency of being able to use resources. Based on empirical observation of the actual efficiencies during the first 3 years of LHC running, and following discussion and agreement with the Computing Resources Scrutiny Group (C-RSG), the values shown below in **Error! Reference source not found.** are assumed by all experiments. Note that following recent RRB meetings the efficiency of use of disk has been set to 1.0 to avoid confusion, while we have introduced a tape efficiency to account for the observation that there is some inefficiency due to various factors (somewhat site dependent), including how often repacking of tapes is done (recovering deleted tape file space), and the effects of not always being able to fill tapes.

Table 4: Efficiencies assumed for the use of resources

	CPU		Disk	Tape
	Organised	Analysis		
Efficiency	85%	70%	100%	85%

1.1.4 Resource Needs

The level of resource requests is driven directly by the LHC live times and the event trigger rates, for a given efficiency of being able to use those resources. Clearly the desire is to maximise the physics output balancing what is likely to be available in terms of resources.

In deriving the estimates for the resource needs here, we are guided by the assumptions that the technology advances anticipated follow the conclusions outlined in the

Appendix on technology expectations. There it was concluded that for a fixed cost, we could anticipate an annual growth in performance of approximately 20% for CPU, 15% for disk, and around 15% for tape.

Table 5: Estimated annual resource growth for fixed cost

	CPU	Disk	Tape
Effective annual growth for fixed cost	20%	15%	15%

1.2 ALICE

The goal for Run 2 is to reach the integrated luminosity of 1nb^{-1} of Pb-Pb collisions for which the ALICE scientific program was originally approved. Targeting this objective for Run 2 will allow the experiment to extend the reach of several measurements crucial for the understanding of the basic properties of the QGP and consolidate preliminary observations from Run 1 data.

The running scenario as presented in this document has been reported to the LHCC in June 2013. The objectives are as follows:

- For Pb-Pb collisions:
  - Reach the target of 1 nb^{-1} integrated luminosity in PbPb for rare triggers
  - Increase the statistics of the unbiased data sample, including minimum bias (MB), centrality triggered events.
- For pp collisions:
  - Collect a reference rare triggers sample with an integrated luminosity comparable to the one of the 1nb^{-1} sample in Pb-Pb collisions.
  - Enlarge the statistics of the unbiased data sample, including MB and high multiplicity triggered events.
- For p-Pb collisions:
  - Enlarge the existing data sample, in particular the unbiased events sample (the collision energy is still under discussion).

To reach these objectives ALICE will exploit the approximately 4 fold increase in instant luminosity for Pb-Pb collisions and will benefit from the consolidation of the readout electronics of TPC and TRD allowing for an increase in the readout rate by a factor of 2.

The increased data rate in the consolidated system will also increase the demands on the High Level Trigger system. The current architecture of the HLT system is expected to be scalable to the higher event rates. The performance of the General Purpose GPUs (GPGPU)-based TPC tracking algorithm has been demonstrated during Run 1 to meet the requirements of Run 2. The HLT will thus rely on the continued use of GPGPUs, which reduces the number of nodes in the farm. This will have the effect of doubling the event rate and the data throughput of the entire dataflow including the migration of data to the computing centre.

During LS1 ALICE will upgrade the existing detectors and install additional detectors: the TRD azimuthal coverage will increase from 60% to 100% and a second electromagnetic calorimeter (DCAL) diagonally opposite the existing EMCAL will be installed. This will increase the raw data output size by about 10%.

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 and the additional data volume from the TRD and DCAL upgrades. The computing power needed to process one event is reported in Table 6. The value for Pb-Pb and p-Pb reconstruction and MC has been increased compared to the values reported in April 2013 as all the events used for the present estimation include TPC data.

The data sizes at the various stages of processing are reported in Table 7.

A factor of 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 6: Processing power in kHEPSpec seconds per event

	Reconstruction	Analysis train	End user analysis	Monte Carlo
pp	0.22	0.17	0.01	1.37
PbPb	3.75	2.49	0.17	46.30
pPb	0.71	1.13	0.09	5.98

Table 7: Data sizes in MB/event

	Raw	ESD&AOD	Monte-Carlo
pp	1.05	0.16	0.37
PbPb	7.50	1.55	21.09
pPb	1.63	0.32	1.73

During LS1 ALICE is reprocessing the entire set of collected RAW data taking advantage of the best available calibration 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 directed towards the Upgrade program, including parallelization, vectorization, GPU algorithms, and new algorithms, will be implemented. The parts of the new environment will gradually become available after 2014. A partial reprocessing of the data will then be performed.

1.2.1 From LS1 to LS2 (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 important that the resources requested in Tier 0 are covered in order to process the first reconstruction pass of heavy-ion data promptly within 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 for processing the data (reconstruction and analysis) and producing 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 replicas 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. A deficit in pledged disk in Tier1s plus Tier2s could be recovered with an increase of the disk in Tier0. It is important that the sum of our disk requirements in Tier01, Tier1 and Tier2 are fulfilled.

Resources required in 2013-2017 are listed in Table 8(CPU), Table 9(Tape), Table 10(Disk); only resources for Run 2 have been updated.

Table 8: CPU requirements for 2015-2017[1]

	CPU (kHEPSPEC06)			
	Tier0	CAF	Tier1s	Tier2s
2015	130	45.0	120	200
2016	170	45.0	160	240
2017	200	45.0	210	270

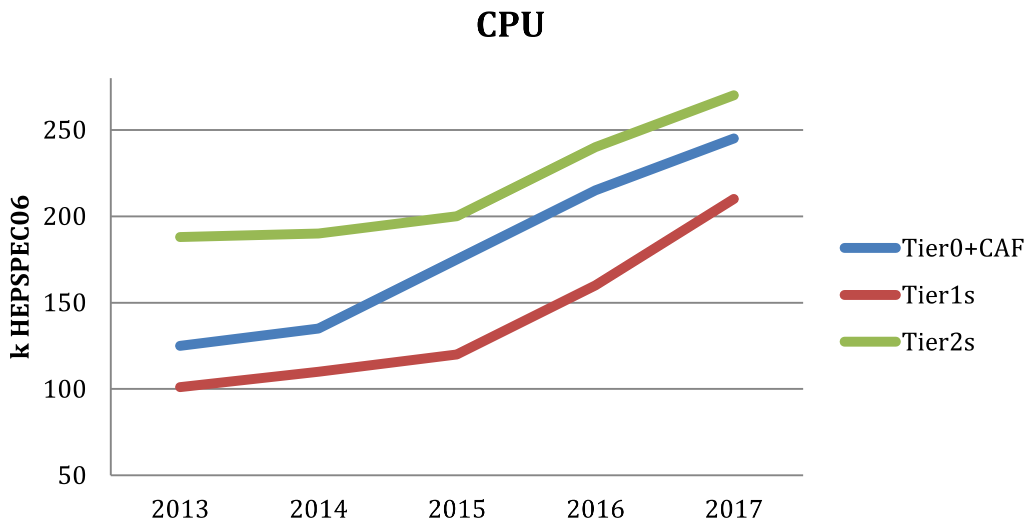

Figure 1: CPU requirement profile. Resources for a given year can be installed during the second semester.

Table 9: Tape requirements for 2015-2017[2]

	Tape (PB)	
	Tier0	Tier1
2015	16.2	10.2
2016	21.6	15.6
2017	25.7	19.7

---

[1] The 2015-2017 values have been updated with respect to the values presented to the RRB in April 2013. Values for 2013-2014 remain unchanged.

[2] Same as 1

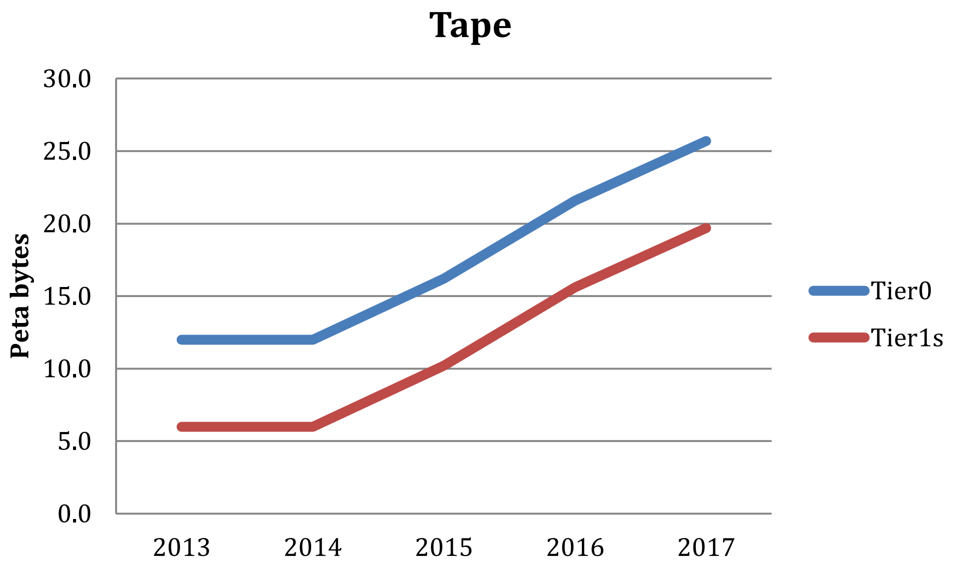

Figure 2: Tape requirement profile. Resources for a given year can be installed at the beginning of the following year.

Table 10: Disk requirements for 2013-2017[3]

	Disk (PB)			
	Tier0	CAF	Tier1s1)	Tier2s
2015	11.2	0.34	15.4	22.1
2016	13.4	0.44	18.6	26.8
2017	14.7	0.54	21.8	31.4

1) Excluding the 2.35 PB of disk buffer in front of the taping system

---

[3] Same as 1

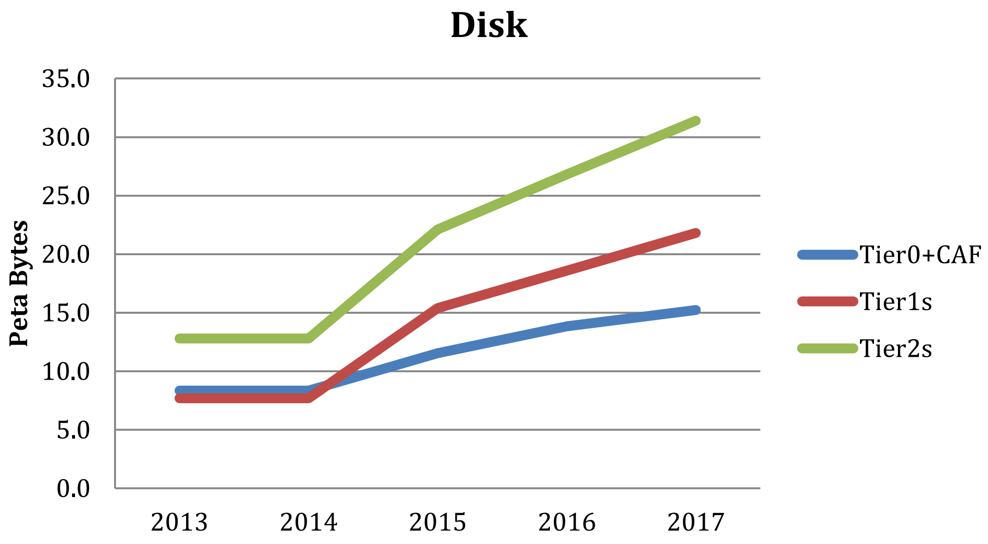

Figure 3: Disk requirement profile. Disks can be installed any time during a given year or during the previous year

1.3 ATLAS

Here the most important items for resource estimation, already described in previous chapters, are summarized:

- *Yearly reprocessing cycles:* In Run-2 there is one envisaged full data re-processing from RAW per year on average. In addition, on average two AOD2AOD re-processings of data and Monte-Carlo per year are expected, corresponding to a yearly effective multiplication factor of 2.7 for disk occupancy with respect to the AOD data volume from one (re-)processing. This is based on the assumption that, with the dynamic cleanup present, eventually only the state-of-the-art versions plus the version made at Tier-0 for each period will remain on disk (1 full (re-)processing and the factor 1.7 for the sum of partial AOD2AOD re-processings during the year).
- *Parked Data:* No delayed streams are envisaged for Run-2, which could be re-evaluated after the first year or two of the data taking.
- *Yearly MC Campaigns:* There will be at most two (full GEANT 4) simulation campaigns per year, one at the beginning of the LHC run with limited statistics and a subsequent full campaign with accurate LHC and detector parameters. This corresponds to an effective multiplication factor of 1.2 w.r.t. the total stored HITS volume from one campaign. The number of Monte-Carlo re-processings is identical to the real data (one full re-processing and on average two AOD2AOD re-processings per year, amounting to the effective multiplication factor for AOD of 2.7).
- *Analysis data:* With the introduction of the new analysis model improvements, the total volume of the group analysis (real and simulated) data is estimated to occupy twice the total volume of one AOD replica.
- *Placement of older data:* With the introduction of aggressive dynamic data deletion of pre-placed primary replicas, the total AOD disk volume at the end of a Run-2 data-taking year is supposed to represent the total volume from one re-processing for real and simulated samples, which are retained in the following year in one primary replica at Tier-2. In addition, the group analysis data volume projected to be kept in 2016 is 25% of the total group data volume produced in 2015. Likewise, in 2017, 50% of the 2016 group data volume will be kept (assuming the running conditions remain constant in 2016 and 2017).
- ***Opportunistic resources:*** In Run-1 ATLAS has benefitted from the availability of substantial beyond-pledge and opportunistic CPU resources. These additional resources proved extremely valuable, allowing ATLAS to pursue an even richer and more precise set of physics results than would otherwise have been possible in the same time frame. Our resource planning is based upon the physics programme that can be accomplished within achievable pledged resources, corresponding to a 'flat' spending budget, while ATLAS acknowledge that if our centres and funding agencies will continue to provide ATLAS with the invaluable resources beyond those pledged this will allow us to accomplish an optimal research programme and physics productivity.

Substantially reduced CPU consumption due to the improvements to ATLAS software, outlines in section ***Error! Reference source not found.*** *and described in detail in section* ***Error! Reference source not found.*** *is assumed in the resource requirements presented below.*

1.3.1 Summary tables of requirements for 2015 – 2018

The tables below summarize the ATLAS resource requirements for Run-2. The requirements are shown to conform to the expected 'flat budget' of cost, which is described by the scaling (modified Moore's law) to resource increase factor of 1.2/year for CPU and 1.15/year for disk and 1.15/year for tape with an uncertainty on the order of 10%. The square brackets [] for the 2015 request show the resource requirements of the ATLAS March 2013 resource request to the Computing Resources Review Board. The values for 2012-2014 shown in the diagrams represent the existing resources and/or validated pledges.

Table 11: Input parameters for ATLAS resource calculations.

LHC and data taking parameters		**2012 pp** **actual**	**2015 pp** **μ=25 @ 25 ns**	**2016 pp** **μ=40 @ 25 ns**	**2017 pp** **μ=40 @ 25 ns**
Rate [Hz]	Hz	400 + 150 (delayed)	1000	1000	1000
Time [sec]	MSeconds	6.6	3.0	5.0	7.0
Real data	B Events	3.0 + 0.9 (delayed)	3.0	5.0	7.0
Full Simulation	B Events	2.6 (8 TeV) + 0.8 (7 TeV)	2	2	2
Fast Simulation	B Events	1.9 (8TeV) + 1 (7 TeV)	5	5	5
Simulated Data					
Event sizes					
Real RAW	MB	0.8	0.8	1	1
Real ESD	MB	2.4	2.5	2.7	2.7
Real AOD	MB	0.24	0.25	0.35	0.35
Sim HITS	MB	0.9	1	1	1
Sim ESD	MB	3.3	3.5	3.7	3.7
Sim AOD	MB	0.4	0.4	0.55	0.55
Sim RDO	MB	3.3	3.5	3.7	3.7
CPU times per event					
Full sim	HS06 sec	3100	3500	3500	3500
Fast sim	HS06 sec	260	300	300	300
Real recon	HS06 sec	190	190	250	250
Sim recon	HS06 sec	770	500	600	600
AOD2AOD data	HS06 sec	0	19	25	25
AOD2AOD sim	HS06 sec	0	50	60	60
Group analysis	HS06 sec	40	2	3	3
User analysis	HS06 sec	0.4	0.4	0.4	0.4

Table 12: ATLAS Tier 1 CPU

Tier-1 CPU ***(kHS06)***	***2015***	***2016***	***2017***
Re-processing	38	30	43
Simulation production	154	89	102
Simulation reconstruction	194	245	280
Group (+user) activities	76	187	267
Total	***462 [478]***	***552***	***691***

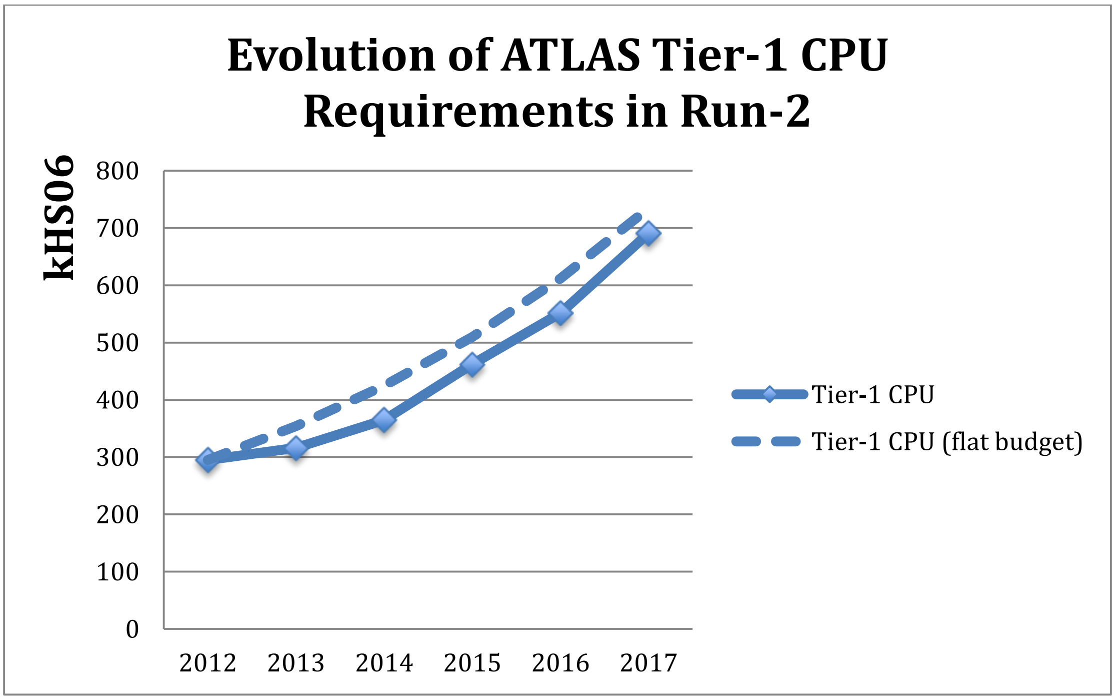

Figure 4: Evolution of ATLAS Tier 1 CPU

Table 13: ATLAS Tier 1 disk

Tier-1 Disk (PB)	***2015***	***2016***	***2017***
Current RAW data	2.4	5.0	7.0
Real ESD+AOD+DPD data	5.6	7.9	11.1
Simulated RAW+ESD+AOD+DPD data	9.2	11.4	11.4
Calibration and alignment outputs	0.3	0.3	0.3
Group data	7.5	8.0	10.4
User data (scratch)	2.0	2.0	2.0
Cosmics	0.2	0.2	0.2
Processing and I/O buffers	3.0	3.0	3.0
Dynamic data buffers (30%)	9.0	10.9	12.6
Total	***39 [47]***	***49***	***58***

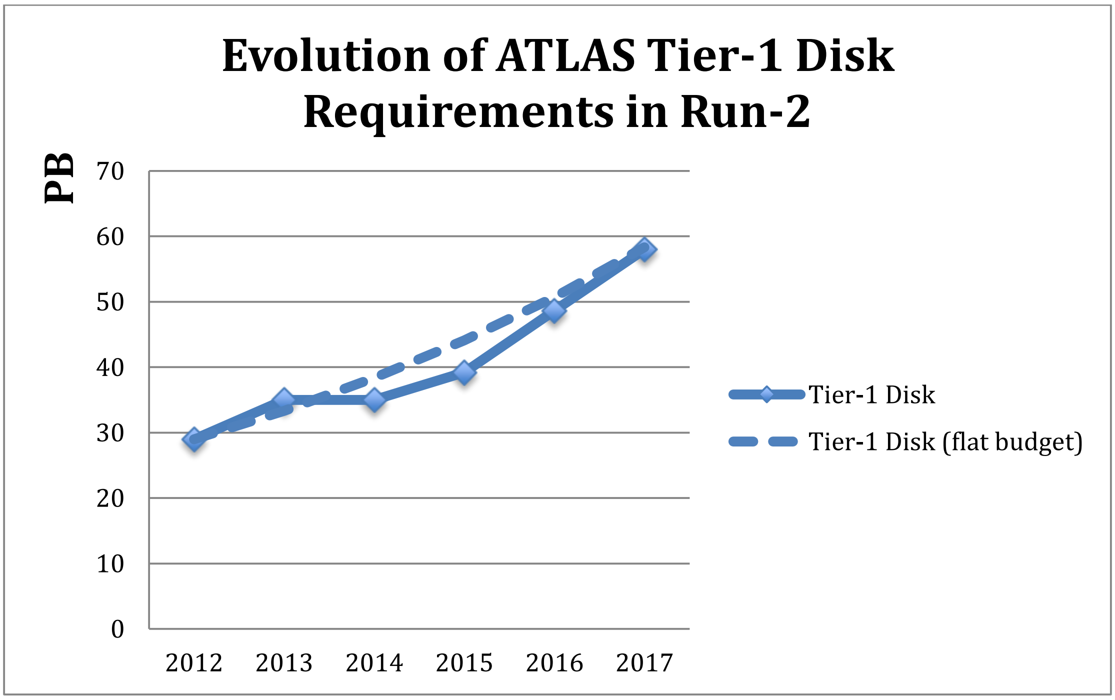

Figure 5: Evolution of ATLAS Tier 1 Disk

Table 14: ATLAS Tier 1 tape

Tier-1 Tape (PB) Cumulative	***2015***	***2016***	***2017***
Real RAW+AOD+DPD data	17	26	39
Cosmics and other data	4	4	4
Group + User	7	8	9
Simulated HITS+AOD data	37	46	56
Total	***65 [74]***	***84***	***108***

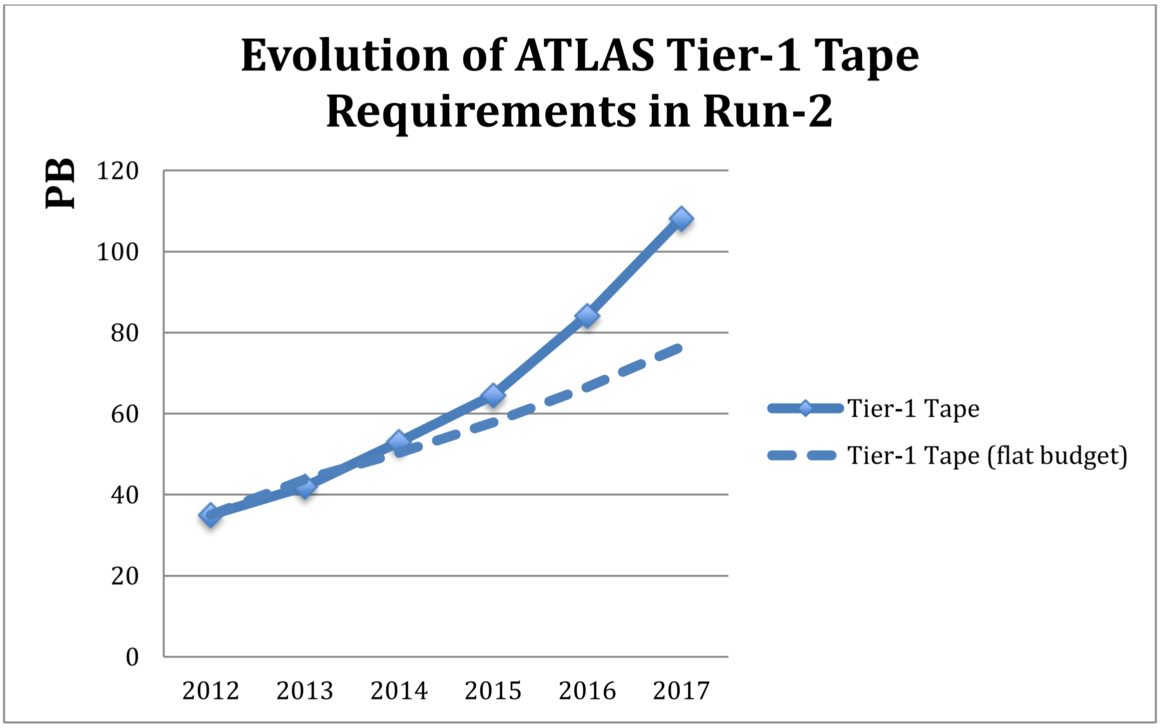

Figure 6: Evolution of ATLAS Tier 1 Tape

Table 15: ATLAS Tier 2 CPU

Tier-2 CPU (kHS06)	***2015***	***2016***	***2017***
Re-processing	20	33	47
Simulation production	338	347	396
Simulation reconstruction	77	61	70
Group + User activities	96	166	219
Total	***530 [522]***	***608***	***732***

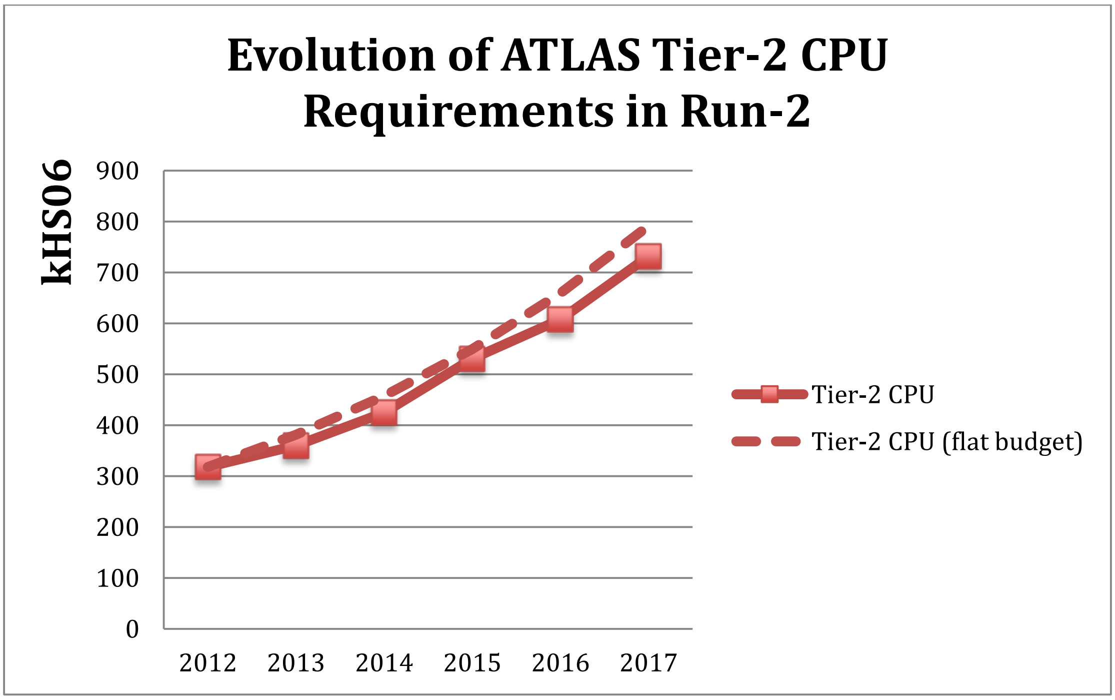

Figure 7: Evolution of ATLAS Tier 2 CPU

Table 16: ATLAS Tier 2 disk

Tier-2 Disk ***(PB)***	***2015***	***2016***	***2017***
Real AOD+DPD data	4.1	6.3	10.6
Simulated HITS+RDO+ESD+AOD	10.6	16.6	21.6
Calibration and alignment outputs	0.2	0.2	0.2
Group data	20.4	29.3	41.6
User data (scratch)	4.0	4.0	4.0
Processing and I/O buffers	3.0	3.0	3.0
Dynamic data buffers (30%)	12.7	15.3	16.8
Total	***55 [65]***	***75***	***98***

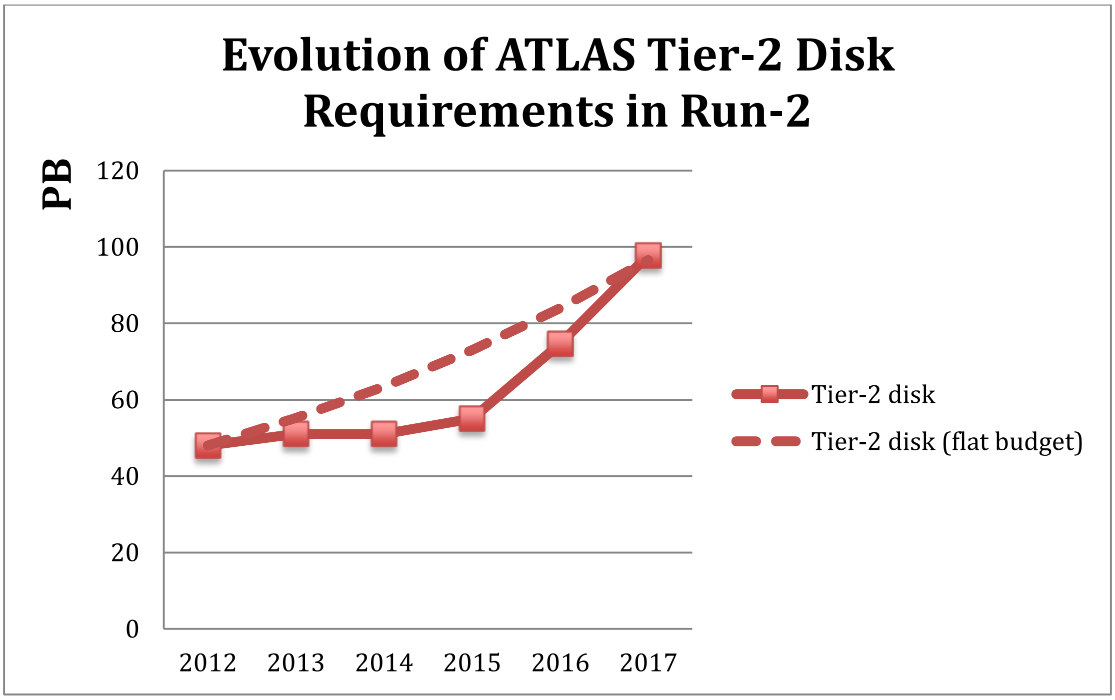

Figure 8: Evolution of ATLAS Tier 2 Disk

Table 17: ATLAS CERN CPU

CERN CPU (kHS06)	***2015***	***2016***	***2017***
CERN CPU Total	***205 [240]***	***257***	***273***
Tier-0 subtotal	***156***	***199***	***199***
T0: Full reconstruction	133	175	175
T0: Partial processing and validation	12	12	12
T0:Merging and monitoring	4	5	5
T0: Automatic calibration	5	5	5
T0: Servers	2	2	2
CAF subtotal	***49***	***58***	***73***
CAF: Partial reconstruction, debugging and monitoring	13	18	18
CAF: Non-automatic calibrations	4	4	4
CAF: Group activities	15	19	27
CAF: User activities	5	6	13
CAF: Servers	12	12	12

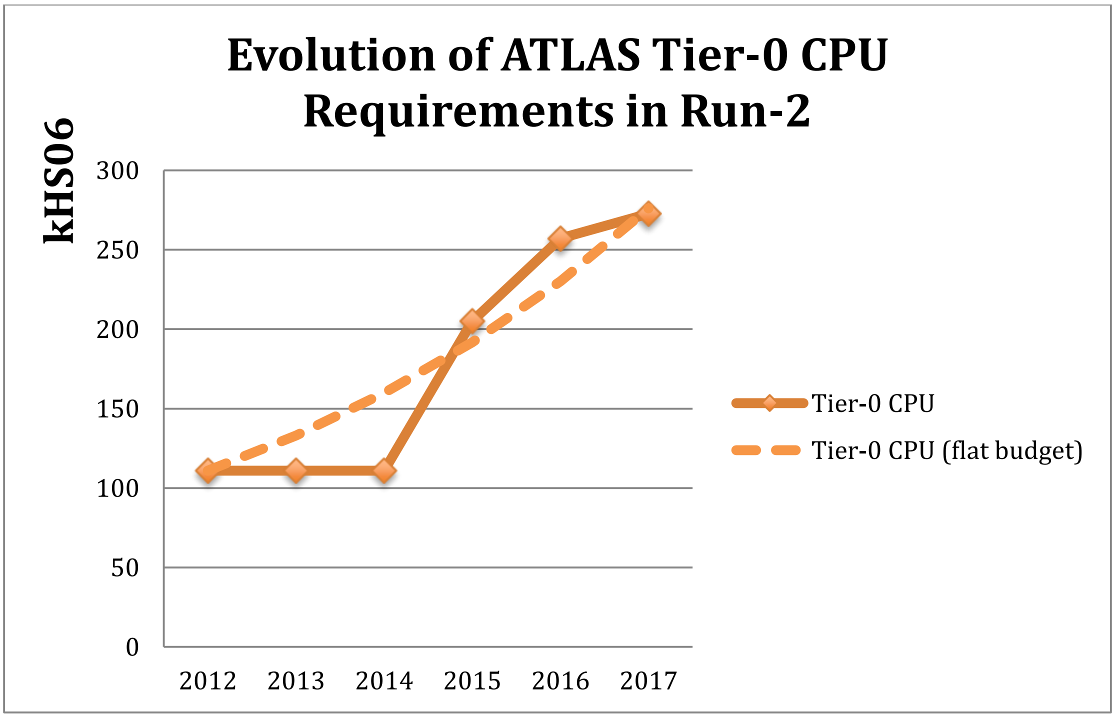

Figure 9: Evolution of ATLAS Tier 0 CPU

Table 18: ATLAS CERN disk

CERN Disk (PB)	***2015***	***2016***	***2017***
CERN Disk Total	***14.1 [15.3]***	***17.0***	***19.1***
Tier-0 Disk Subtotal	***3.40***	***3.40***	***3.40***
Buffer for RAW and processed data	3.00	3.00	3.00
Buffers for merging	0.30	0.30	0.30
Tape buffer	0.10	0.10	0.10
CAF Total	***10.7***	***13.6***	***15.7***
CAF: Calibration and alignment	0.5	0.5	0.5
CAF: Derived detector data	2.0	2.8	3.9
CAF: Derived simulated data	6.7	8.8	8.8
CAF: Group data	1.0	1.0	2.0
CAF: User data	0.5	0.5	0.5

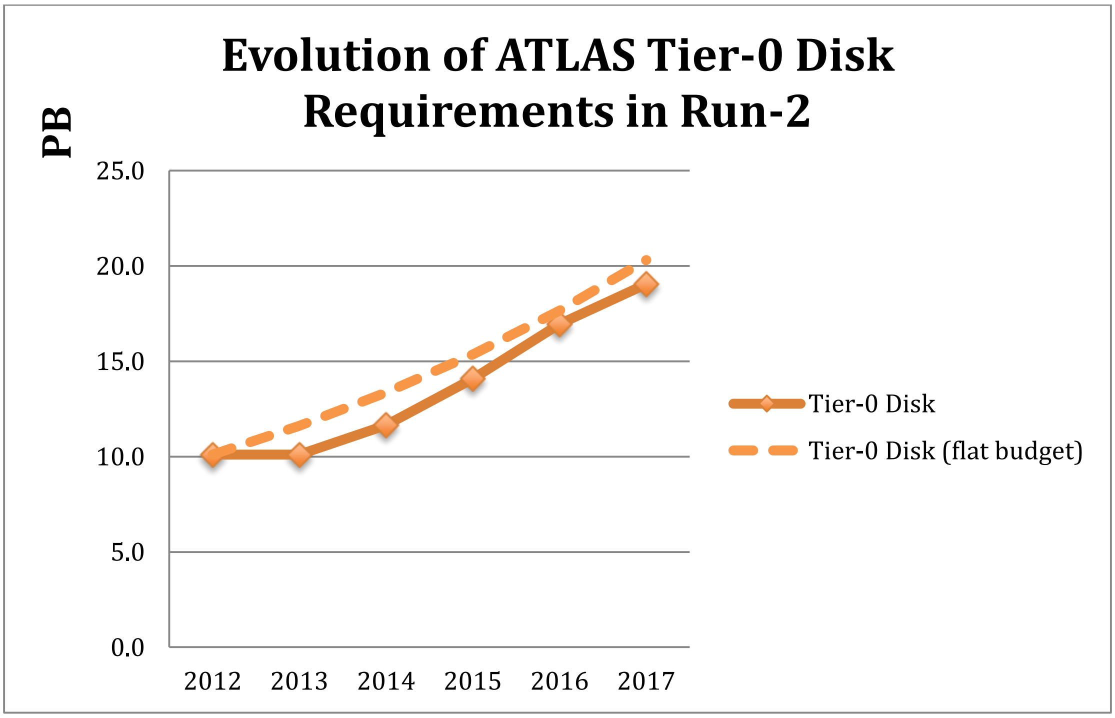

Figure 10: Evolution of ATLAS Tier 0 Disk

Table 19: ATLAS CERN tape

CERN Tape (PB) Cumulative	*2015*	*2016*	*2017*
Total	***33 [35]***	***42***	***54***

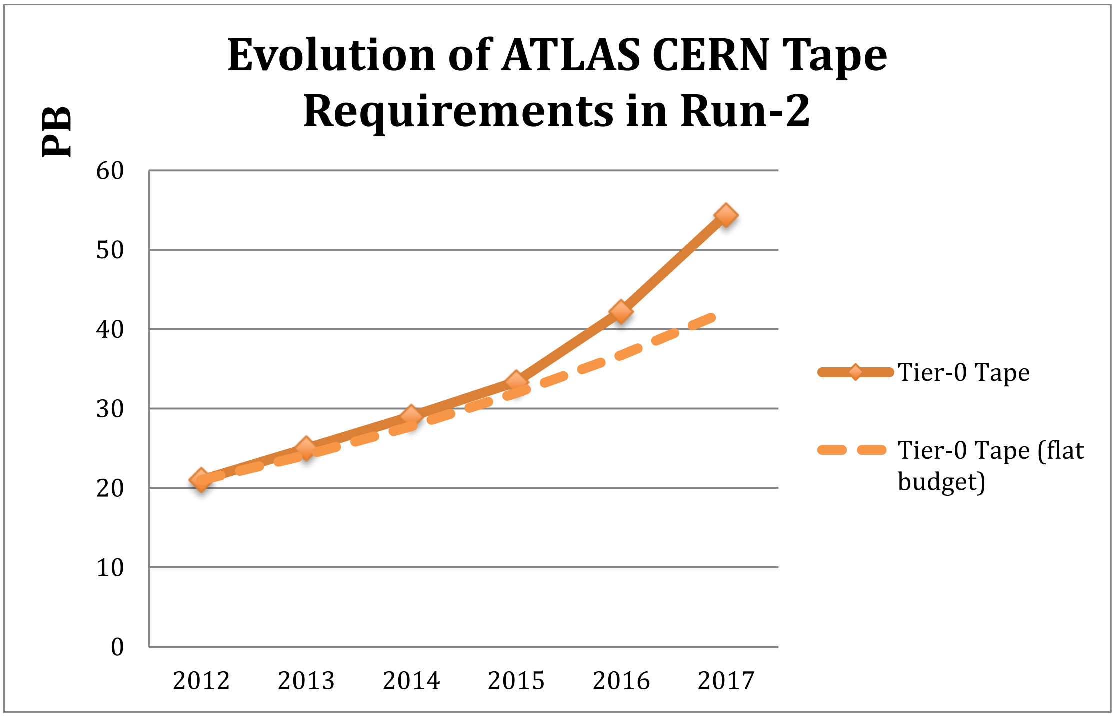

Figure 11: Evolution of ATLAS Tier 0 Tape

1.4 CMS

In 2014 CMS asked for only a small increase in tape capacity to handle new processing and simulation samples. Any additional processing resources were provided by allocating the Tier 0 capacity and using the newly commissioned Higher Level Trigger farm (HLT).

Looking forward to 2015, LHC predicts 3Ms LHC run time in 2015 and a DAQ rate of 1 kHz on average. The increase is required in order to maintain the same physics capability with the foreseen increase in instantaneous luminosity and energy (the latter accounting for a 15% increase in the rate). This will imply a 2.5 increase factor in computing needs with the higher trigger. While luminosity increase will increase the pileup and therefore the reconstruction time, which will imply a further increase factor of 2.5 in computing processing needs. The move to running at 25 ns, together with the increased luminosity, will increase the out of time pileup. Early studies indicated that this could result in another factor of two, but in the planning it is assumed that the problem can be solved with code improvements.

In order to address the situation CMS will seek to gain from operation efficiency and access to opportunistic resources, described in the rest of the document. The following optimisations in the computing model have a significant impact on the needs for computing resources:

- Access to the HLT during the shutdown period and whenever else is possible. Without this change the Tier-1 sites would need to be 25% larger than requested;
- Reducing to 1 re-reconstruction campaign during the year for data and simulation. This reduces the original Tier-1 activities by 50%;
- Deploying half the prompt reconstruction to Tier-1s during the running period. Without this the Tier-0 request would have been approximately 40% larger;
- Optimizing event formats.

If CMS had maintained the data model and workflows used in Run 1 and not made these optimisations, it would have faced an increase in the computing work to be done of a factor of 6 (or 12 with the out-of-time pileup effect). However, the optimisations and improvements described here result in a less than factor of 2 increase in the required processing resources.

1.4.1 HLT Resources

The use of the HLT during the shutdown period increases the overall Tier-1 resources by 25%. CMS has commissioned a cloud interface to allow rapid deployment of a high fraction of the number of HLT systems for use in offline processing.

1.4.2 Yearly re-processing cycles

In Run 2 there is one envisaged full data re-processing from RAW per year on average at the end of the running year when the maximum capacity can be used from the HLT farm. In addition targeted reprocessing passes of individual primary datasets are expected throughout the year, but in total add to a fraction of the full reprocessing pass.

1.4.3 Parked Data

No delayed streams are envisaged for the Run 2, which could be re-evaluated after the first year or two of the data taking.

1.4.4 Yearly MC Campaigns

A simulation campaign equivalent to 1.5 times the number of data events collected is budgeted for 2015, dropping to 1.3 times data in 2016, and 1 times data in 2017. The factor changes with the number of events collected, which is lowest in 2015, and with the expected activities and measurements. CMS expect to produce samples using fast simulation, which has been heavily used for upgrade samples. The difficulty of fast simulation has been to move it to a transient format. The simulation is fast enough that it is frequently better to reproduce it than to store is persistently.

1.4.5 Placement of data

With the introduction of the xrootd-based data federation and the use of monitoring of the access level through the popularity service, CMS has reduced the replication factor for data at the Tier-2s. This is reducing the slope of the disk increase in the Tier-2 disk planning,

1.4.6 Summary tables of requirements for 2015 – 2018

The tables below summarize the CMS resource requirements for Run 2. The requirements are shown to conform to the expected ‘flat budget’ of cost, which with is described by the scaling (modified Moore’s law) to resource increase factor of 1.2/year for CPU and 1.15/year for disk, with an uncertainty on the order of 10%.

The data taking parameters for the last running year and the next three running years are outlined in Table 20.

Table 20: Input parameters for CMS resource calculations.

LHC and data taking parameters		**2012 pp actual**	**2015 pp μ=25**	**2016 pp μ=40**	**2017 pp μ=40**
Rate [Hz]	Hz	400 + 600 (parked)	1000	1000	1000
Time [sec]	MSeconds	6.6	3.0	5.0	7.0
Real data	B Events	6B	3.0	5.0	7.0
Full Simulation	B Events	5	3.0	6	7
Simulated Data					
Event sizes					
RAW Data	MB	0.5	0.65	.95	0.95
RAW Data	MB	0.5	0.65	.95	0.95
RECO Data	MB	0.75	0.8	0.9	0.9
RECO Data	MB	0.75	0.8	0.9	0.9
AOD Data	MB	0.28	0.30	0.35	0.35
AOD Data	MB	0.28	0.30	0.35	0.35
RAW Sim	MB	1.5	1.5	1.5	1.5
RAW Sim	MB	1.5	1.5	1.5	1.5
RECO Sim	MB	0.80	0.85	0.95	0.95
AOD Sim	MB	0.3	0.35	0.40	0.40
CPU times per event					
Full Simulation	HS06 Sec	500	500	500	500
Fast sim	HS06 sec	50	50	50	50
Real recon	HS06 sec	300	525	920	920
SIM RECO	**HS06 sec**	**400**	675	1050	1050

The table and figure below show the expected evolution during Run 2, and the total changes since 2012 for the Tier-1 processing capacity. The red line in Figure 12 indicates how the resources would evolve with a flat funding profile. The Tier-1 processing is one of the few areas that CMS calculates needs to progress at larger than flat funding.

Table 21: CMS Tier 1 CPU

Tier-1 CPU ***(kHS06)***	***2015***	***2016***	***2017***
Re-processing	100	150	200
Simulation production	150	200	225
Simulation reconstruction	50	50	100
Group (+user) activities	0	0	0
Total	***300***	***400***	***525***

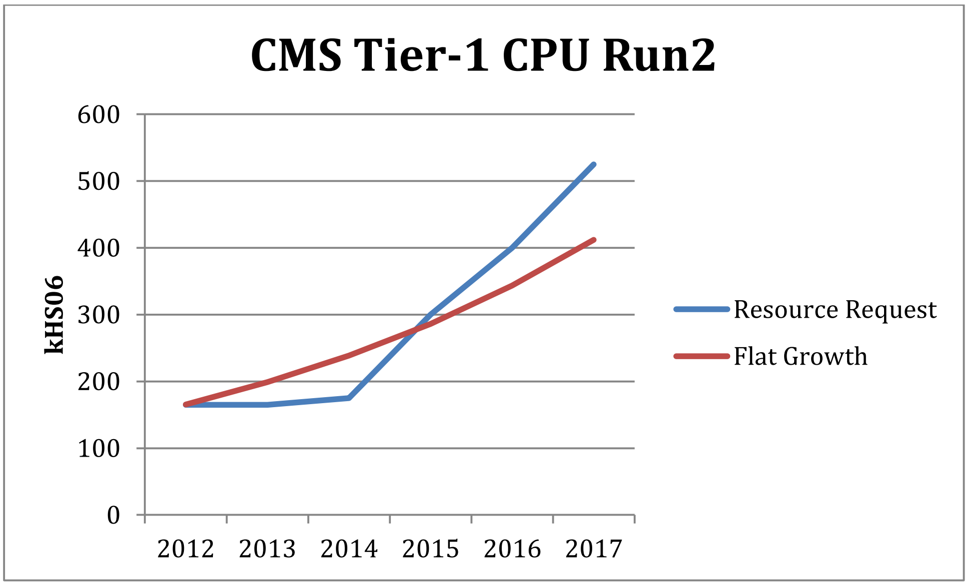

Figure 12: Evolution of CMS Tier 1 CPU

The table and figure below show the expected evolution during Run 2, and the total changes since 2012 for the Tier-1 disk storage capacity. The red line in Figure 13 indicates how the resources would evolve with a flat funding profile.

Table 22: CMS Tier 1 disk

Tier-1 Disk (PB)	***2015***	***2016***	***2017***
Current RAW data	2.0	3.0	4.0
Real RECO+AOD	9.0	11.0	14.0
Simulated RAW+RECO+AOD	8.0	11.0	14.0
Skimming data	3.0	4.0	5.0
User data (scratch)	2.0	2.0	2.0
Dynamic data buffers	2.5	4.0	6.0
Total	***27***	***35***	***45***

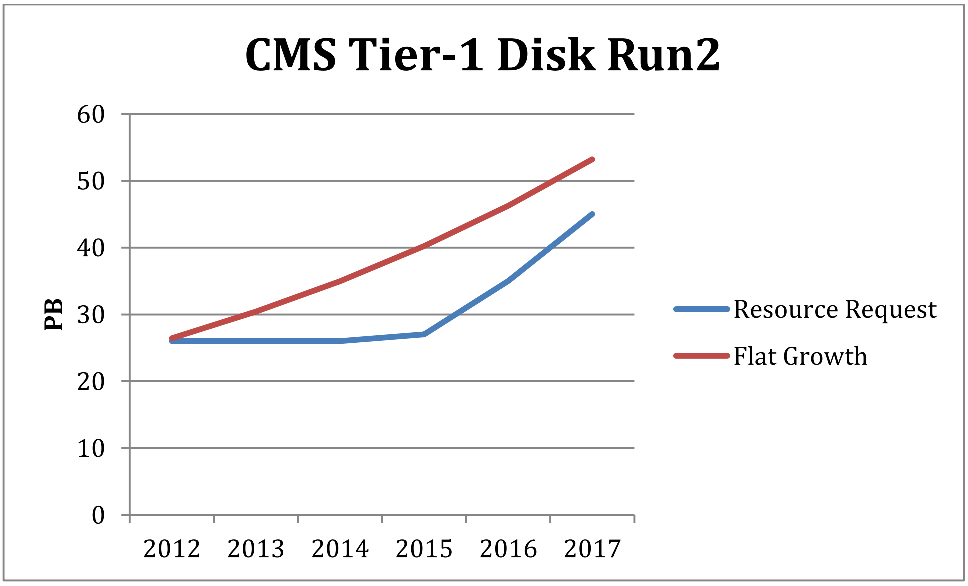

Figure 13: Evolution of CMS Tier 1 Disk

The table and figure below show the expected evolution during Run 2, and the total changes since 2012 for the Tier-1 tape storage capacity. The red line in Figure 14 indicates how the resources would evolve with a flat funding profile.

Tier-1 Tape (PB) Cumulative	***Run 1 (2014)***	***2015***	***2016***	***2017***
RAW Data	5	9	14	19
RAW Simulation	16	18	20	24
RECO Data and Simulation	19	22	26	34
AOD Data and Simulation	15.5	24.5	40.0	58.0
Total	***55***	***73.5***	***100***	***135***

Table 23: Evolution of CMS Tier 1 Tape

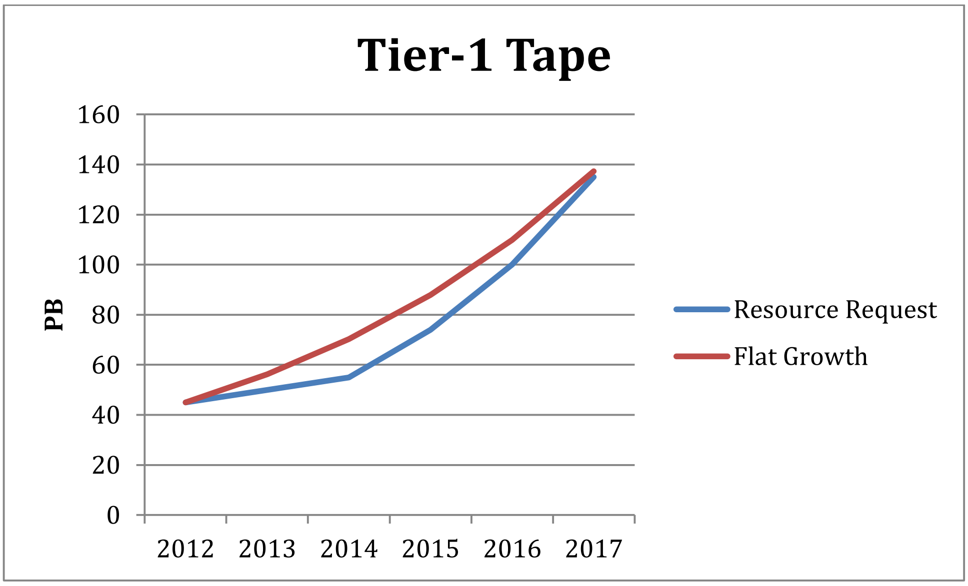

Figure 14: Evolution of CMS Tier 1 Tape

The table and figure below show the expected evolution during Run 2, and the total changes since 2012 for the Tier-2 processing capacity. The red line in Figure 15 indicates how the resources would evolve with a flat funding profile.

Table 24: CMS Tier 2 CPU

Tier-2 CPU (kHS06)	*2015*	*2016*	*2017*
Analysis	400	550	600
Simulation production	100	150	200
Total	***500***	***700***	***800***

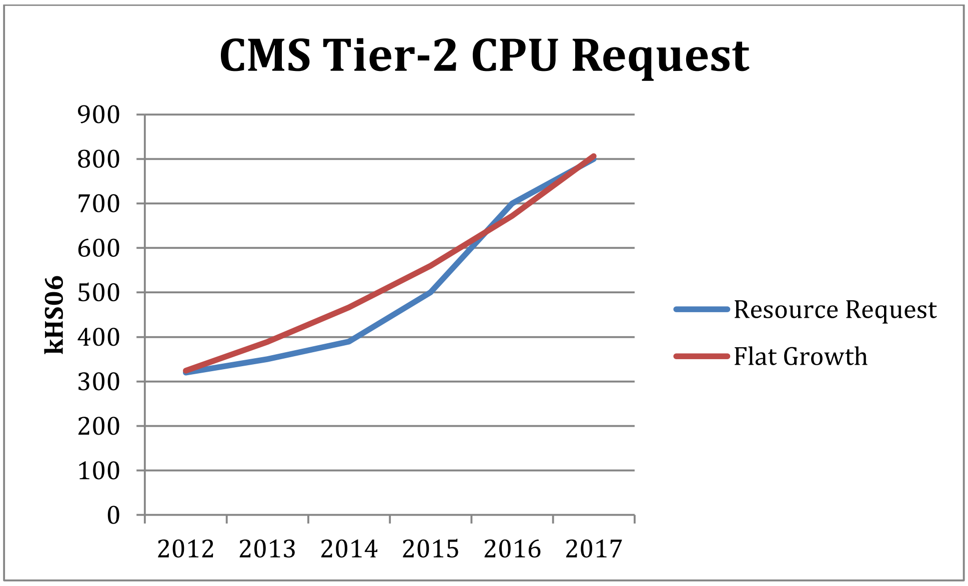

Figure 15: Evolution of CMS Tier 2 CPU

The table and figure below show the expected evolution during Run 2, and the total changes since 2012 for the Tier-2 storage capacity. The red line in Figure 16 indicates how the resources would evolve with a flat funding profile

Table 25: CMS Tier 2 disk

Tier-2 Disk (PB)	***2015***	***2016***	***2017***
Real RECO + AOD	9.0	12.0	15.0
Simulated RECO + AOD	12.0	15.0	17.0
Production Data	2.0	2.0	2.0
User data	8.4	11.0	14.0
Total	***31.4***	***40***	***48***

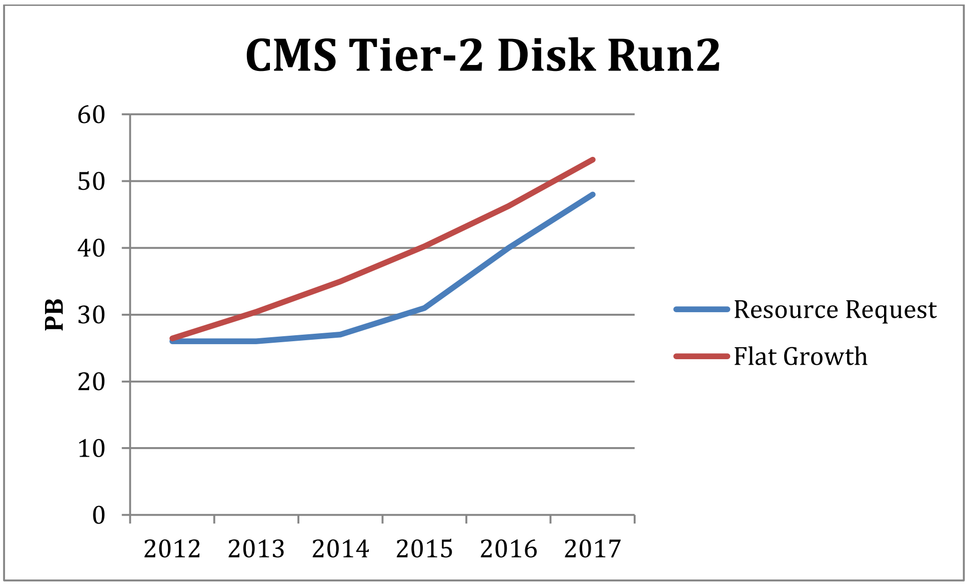

Figure 16: Evolution of CMS Tier 2 Disk

The table and figure show below the expected evolution during Run 2, and the total changes since 2012 for the Tier-0 processing capacity. The red line in Figure 17 indicates how the resources would evolve with a flat funding profile.

Table 26: CMS CERN CPU

CERN CPU (kHS06)	***Run-1 (2012)***	***2015***	***2016***	***2016***
CERN CPU Total	***135***	***271***	***315***	***365***
Tier-0 subtotal	***121***	***256***	***300***	***350***
T0: Full reconstruction	83	210	246	292
Express	12	17	21	21
T0:Repacking	8	8	10	12
T0: Automatic calibration	6	6	6	6
T0: Servers	12	15	17	19
CAF subtotal	***15***	***15***	***15***	***15***

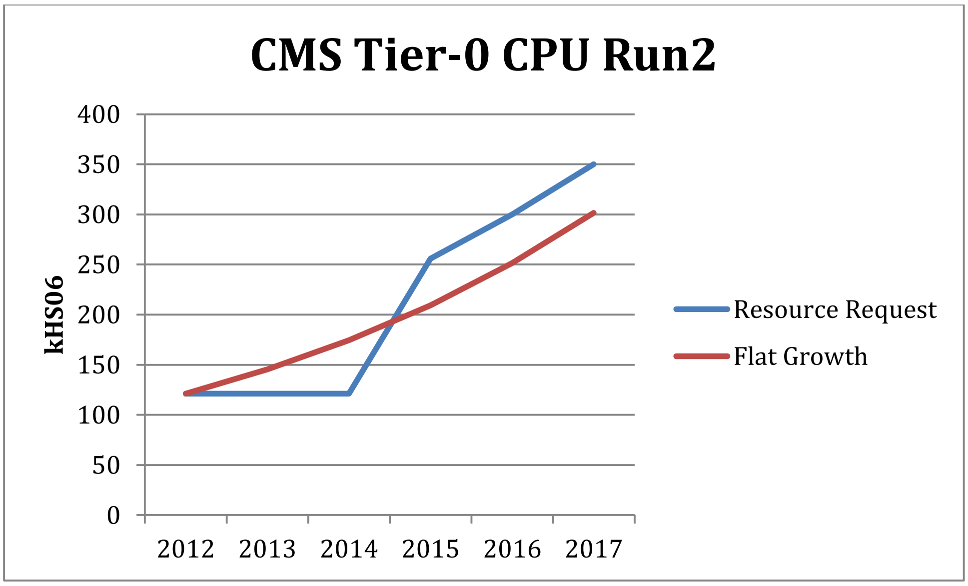

Figure 17: Evolution of CMS Tier 0 CPU

Table 27: CMS CERN disk

CERN Disk (PB)	*Run 1 (2014)*	*2015*	*2016*	*2017*
CAF and Analysis	*9.0*	*12*	*13*	*14*
Tier-0 and Data Distribution	*0.0*	*3.2*	*3.2*	*3.2*
Total	*9.0*	*15.2*	*16.2*	*17.2*

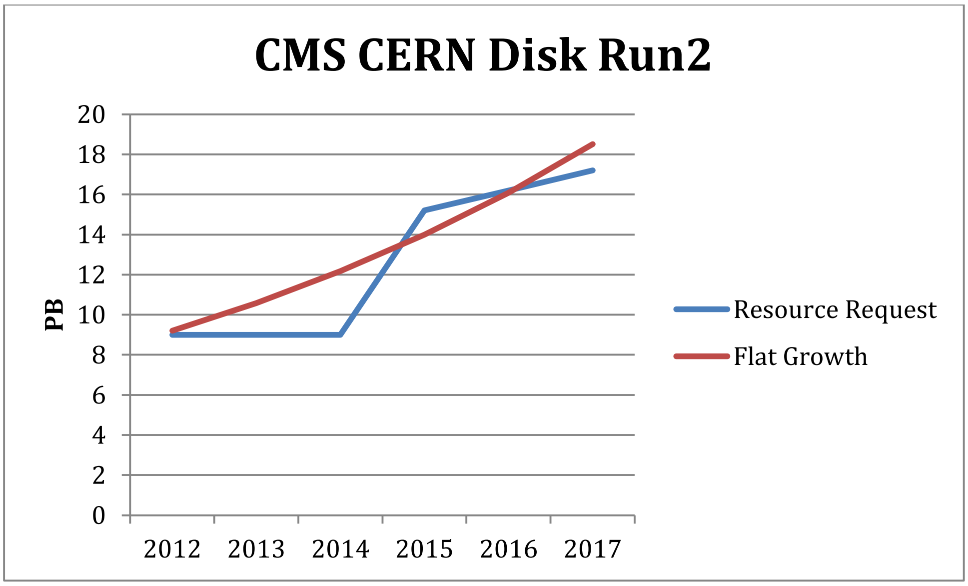

Figure 18: Evolution of CMS Tier 0 Disk

Table 28: CMS CERN tape

CERN Tape (PB) Cumulative	*Run 1 (2014)*	*2015*	*2016*	*2017*
Total	*26*	*31*	*38*	*50*

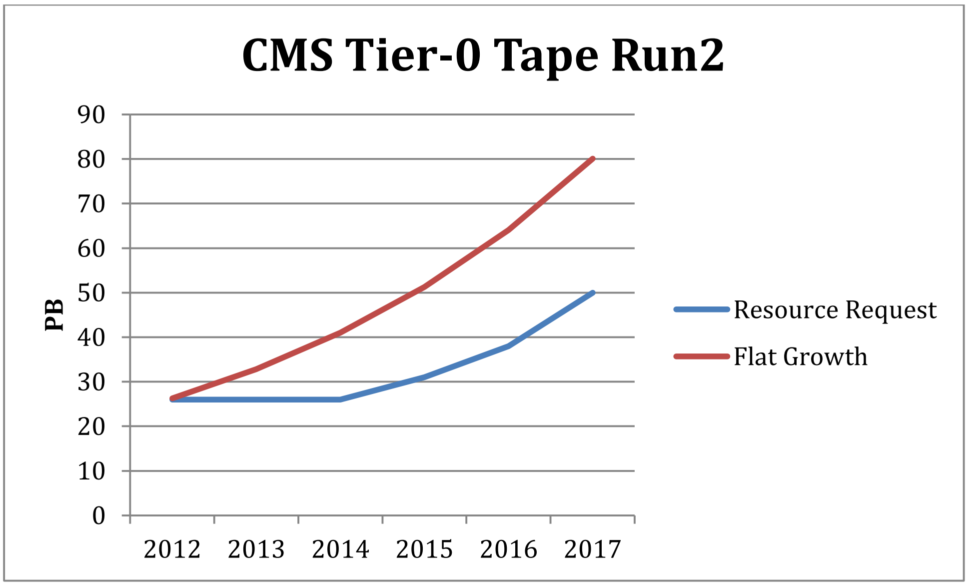

Figure 19: Evolution of CMS Tier 0 Tape

1.5 LHCb

The factors that determine the processing and storage requirements are described in Chapter 2. In summary:

- For 2015 and beyond the event sizes (and therefore processing times) are expected to remain roughly the same as in Run 1 due to the trade off between increased interaction complexity and 25 ns bunch crossing
- Event sizes are given in the relevant table in Chapter 2.
- The trigger rate will be 12.5 kHz (10 kHz to be reconstructed +2.5 kHz "Turbo"). The concept of data parking has been introduced for a fraction of the 10 kHz, but it is not expected to need to use it in 2015 or 2016, a decision for 2017 will be made based on available resources.
- It is assumed that the LHC will run with a bunch spacing of 25ns; this is an important parameter for the event size (and therefore computing resources requirements).
- Given currently available tape resources and the expected growth rate in tape requirements, the LHCb data preservation archives consist of a single tape copy, which makes these archives vulnerable to inevitable tape failures. This is clearly an area of concern but it is thought unrealistic to request the additional resources that would be required for a second archive copy.

Data operations

The detailed scheme for data processing/stripping/re-processing for each year of data in 2015-2017 is described in detail and given in the relevant Table in Chapter 2.

Simulation campaigns

For simulation LHCb's model differs from that of ATLAS or CMS, in that the bulk of the MC production is done **after** the year of data-taking, when beam and trigger conditions are known. Once the software is frozen, the simulation runs continuously for up to a year, using idle resources at all Tiers, as well as opportunistic (unpledged) resources.

In 2015 most analyses of 2011-2012 data are expected to be in an advanced state and to have satisfied most of their simulation needs. Simulation efforts are likely to concentrate on further simulations for the LHCb Upgrade studies, and on tuning the simulation to the observed 2015 data-taking conditions.

Since it is not planned to reprocess the 2015 data during the 2015-2016 winter shutdown, it is foreseen to start a massive MC production for analysis of 2015 data as soon as the 2015 run ends. LHCb expect to have satisfied approximately 50% of the simulation needs for the analysis of 2015 data before the restart of the LHC in spring 2016 (i.e. during the 2015 WLCG accounting period).

The simulation for 2016 is concentrated in the 2016-2017 winter shutdown and continues at lower priority throughout 2017, using any CPU resources not needed to process the real data. The simulation for 2017 is largely postponed to LS2.

1.5.1 CPU requirements

Table 29 presents, for the different activities, the CPU work estimates for 2015, 2016, 2017. Note that in this table there are no efficiency factors applied: these are resource requirements assuming 100% efficiency in using the available CPU. The last row shows the power averaged over the year required to provide this work, after applying the standard CPU efficiency factors of **Error! Reference source not found.**.

Table 29: Estimated CPU work needed for the LHCb activities

LHCb CPU Work in WLCG year (kHS06.years)	2015	2016	2017
Prompt Reconstruction	19	31	43
First pass Stripping	8	13	9
Full Restripping	8	20	9
Incremental Restripping	0	4	10
Simulation	134	153	198
User Analysis	17	17	17
Total Work (kHS06.years)	**185**	**238**	**286**
Efficiency corrected average power (kHS06)	**220**	**283**	**339**

The required resources are apportioned between the different Tiers taking into account the computing model constraints and also capacities that are already installed. This results in the requests shown in Table 30. The table also shows resources available to LHCb from sites that do not pledge resources through WLCG.

Table 30: CPU Power requested at the different Tiers

Power (kHS06)	**Request 2015**	**Forecast 2016**	**Forecast 2017**
Tier 0	44	53	63
Tier 1	123	148	177
Tier 2	52	62	74
Total WLCG	**219**	**263**	**315**
HLT farm	10	10	10
Yandex	10	10	10
Total non-WLCG	**20**	**20**	**20**

The request for 2015 has been sized to satisfy entirely with WLCG resources the requirement presented in the Table. This is partly because the contribution from the HLT farm is uncertain (the farm would in any case only be available during the winter shutdown, when many maintenance activities are also required) but also to allow a ramp up, within a constant budget, to the resources required in 2016 and 2017.

1.5.2 Storage requirements

Table 31 presents, for the different data classes, the forecast total disk space usage at the end of the years 2015-2017.

Table 31: Breakdown of estimated disk storage usage for different categories of LHCb data

LHCb Disk storage usage forecast (PB)	2015	2016	2017
Stripped Real Data	7.3	13.1	14.7
Simulated Data	8.2	8.8	12.0
User Data	0.9	1.0	1.1
MDST.DST	1.5	1.9	
FULL.DST	3.3		
RAW buffer	0.4	0.5	0.3
Other	0.2	0.2	0.2
Total	**21.7**	**25.4**	**28.2**

Table 32 shows, for the different data classes, the forecast total tape usage at the end of the years 2015-2017 when applying the models described in the previous sections. The numbers include the standard 85% tape efficiency correction described in **Error! Reference source not found.**, which is probably pessimistic for RAW data that is written sequentially to a dedicated tape class, and never deleted.

Table 32: Breakdown of estimated tape storage usage for different categories of LHCb data

LHCb Tape storage usage forecast (PB)	2015	2016	2017
Raw Data	12.6	21.7	34.5
FULL.DST	8.7	15.2	19.7
MDST.DST	1.8	5.2	7.7
Archive	8.6	11.5	14.7
Total	**31.7**	**53.7**	**76.6**

The disk and tape estimates shown in Table 31 and Table 32 are broken down into fractions to be provided by the different Tiers using the distribution policies described in LHCb-PUB-2013-002. These numbers are shown in Table 33 and Table 34.

As can be seen the increase in disk storage can be managed to fit inside a reasonable growth envelope by adjustments in the details of the processing strategy.

On the other hand, the growth in the tape storage requirement is more challenging but largely incompressible: in Table 32 one can see that the major part of the increase is due to RAW data that, if not recorded, is lost. “Parking” of some fraction of this raw data will only reduce by a corresponding fraction the growth rate of tape for FULL.DST (note that Table 32 already assumes parking of 50% of the RAW data in 2017).

Table 33:LHCb disk request for each Tier level. Note, that for countries hosting a Tier 1 it is left it up to the country to decide on the most effective policy for allocating the total Tier 1+Tier 2 disk pledge. For example the Tier 2 share could also be provided at the Tier 1

LHCb Disk (PB)	**2015 Request**	**2016 Forecast**	**2017 Forecast**
Tier0	6.7	8.3	9.5
Tier1	12.5	14.2	15.4
Tier2	2.5	2.9	3.3
Total	**21.7**	**25.5**	**28.3**

Table 34: LHCb Tape requests for each Tier

LHCb Tape (PB)	**2015 Request**	**2016 Forecast**	**2017 Forecast**
Tier0	10.4	15.9	21.6
Tier1	21.3	37.8	55.0
Total	**31.7**	**53.7**	**76.6**