



# Grid resources use

**ALICE Offline week**

30 March 2017

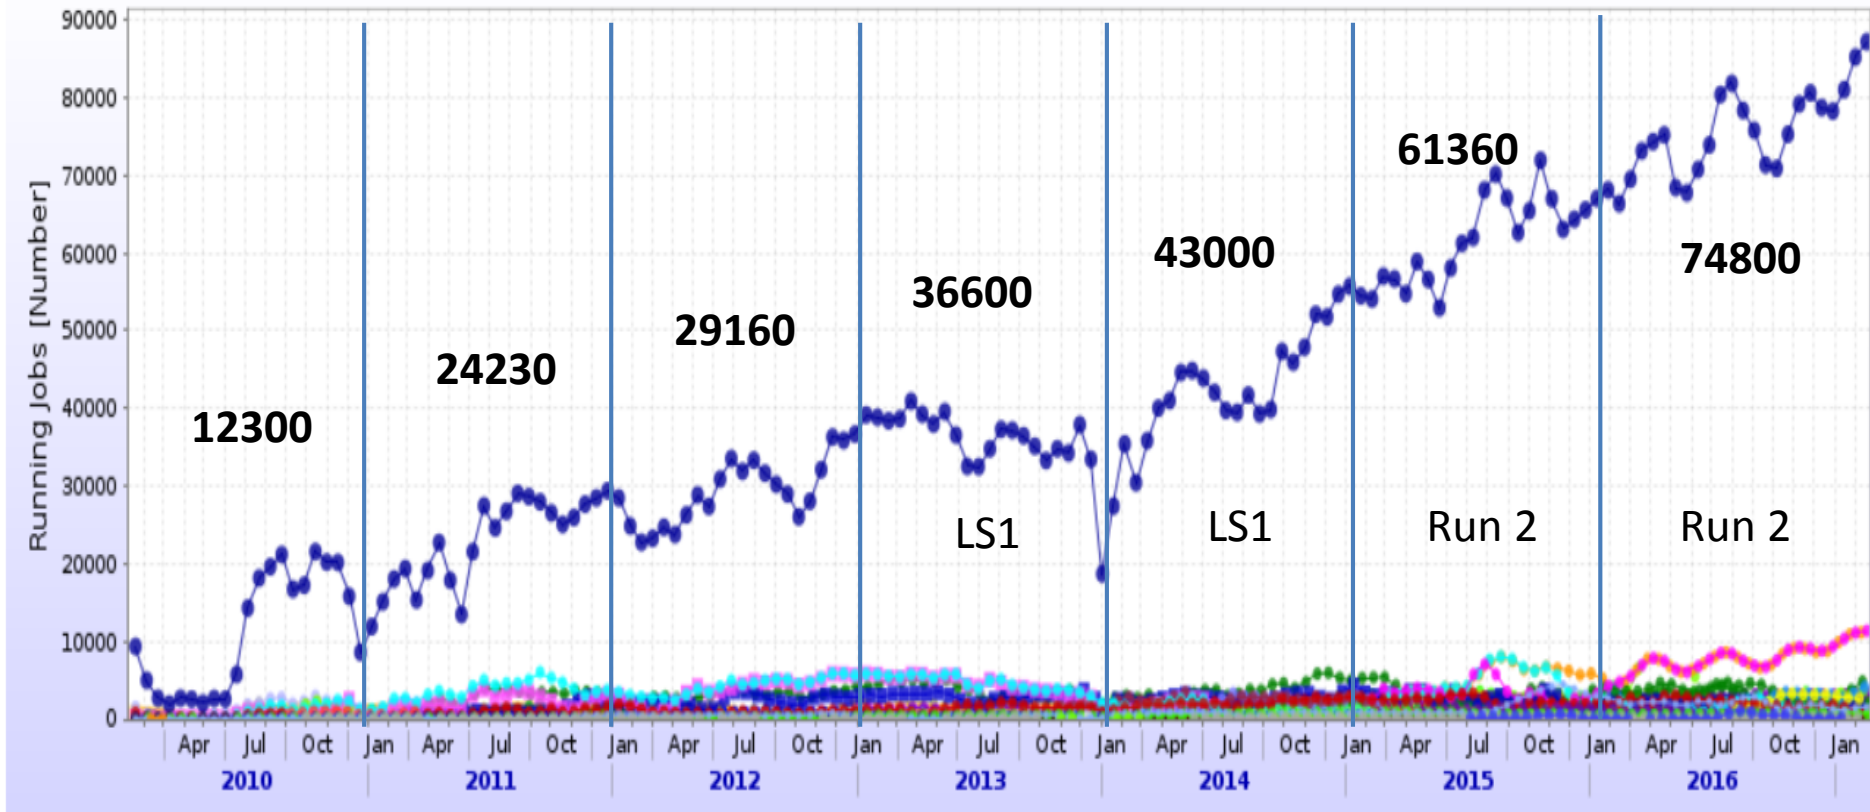
Latchezar Betev

# Job records are boring



# CPU resources evolution is exciting

Running Jobs



Year on year increase

+97%

+20%

+26%

+17%

**+43%**

+21%

The canonical 20% year-on-year increase is still there for CPU

# Disk resources evolution

- Pledged

Sites	2015 [PB]	2016 [PB]	Delta
T0	14.5	16.8	+16%
T1s	15.8	18.9	+19%
T2s	16.8	18.5	+10%

- Pledged and Installed in 2016

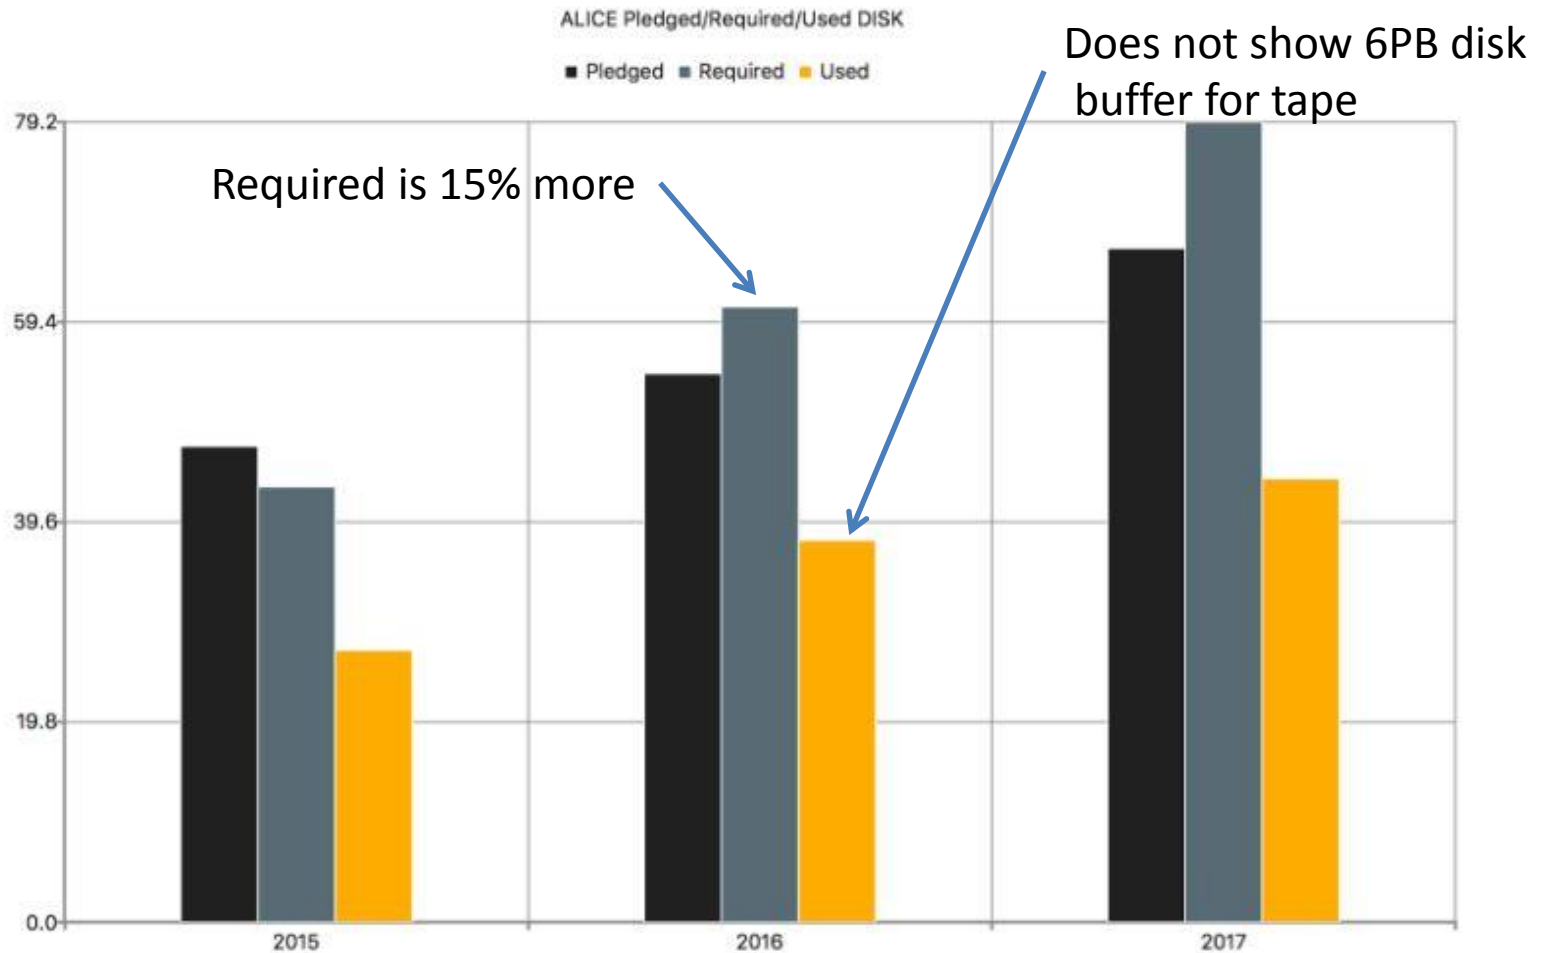
Sites	Pledged [PB]	Installed [PB]	Delta
T0	16.8	16.6	0%
T1s	18.9	18.7	0%
T2s	18.5	14.5	-22%

# Resources evolution 2015-2016

- Real 20% CPU increase
  - Standard increase per year
  - Matches the site pledges
- Pledged 15% per year storage increase
  - Standard increase per year
  - However, the **real** increase is just **6%**
    - Number of T2s are not fulfilling their pledges

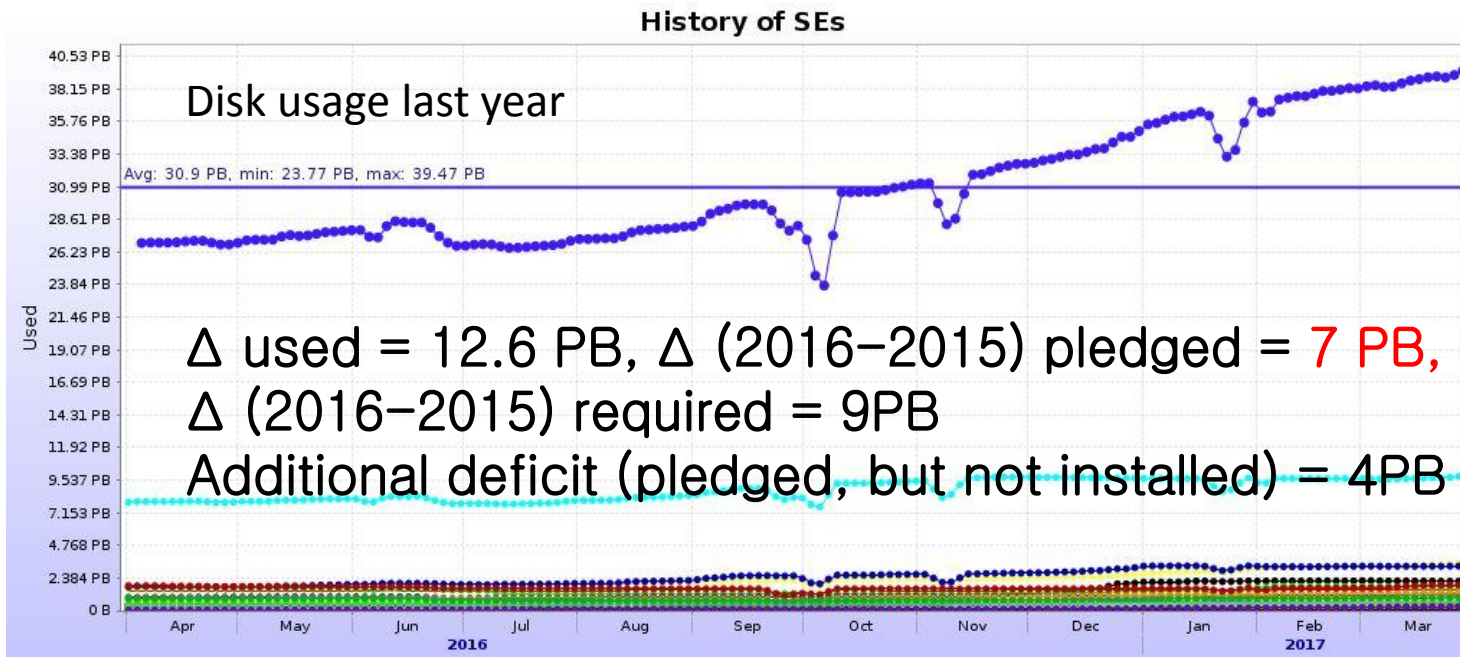
# Things become worse...

- ... If we consider the requirements



# Real situation with storage today

- Installed disk, excluding buffer – 48 PB (should be 56 PB if requirements fulfilled)
- Used disk, excluding buffer – 40PB (83%)





# A lot of red

SE Name	AliEn name	Size	Used	Free	Usage ▲
44. RRC-KI - SE	ALICE::RRC-KI::SE	243 TB	241.7 TB	1.346 TB	99.45%
25. ITEP - SE	ALICE::ITEP::SE	20.01 TB	19.82 TB	187.7 GB	99.08%
45. RRC_KI_T1 - EOS	ALICE::RRC_KI_T1::EOS	2.267 PB	2.239 PB	28 TB	98.79%
6. CCIN2P3 - SE	ALICE::CCIN2P3::SE	1.972 PB	1.939 PB	33.11 TB	98.36%
14. FZK - SE	ALICE::FZK::SE	4.5 PB	4.409 PB	92.77 TB	97.99%
11. CNAF - SE	ALICE::CNAF::SE	3.405 PB	3.329 PB	77.64 TB	97.77%
42. Prague - SE	ALICE::Prague::SE	1.591 PB	1.555 PB	36.74 TB	97.74%
38. NIPNE - EOS	ALICE::NIPNE::EOS	185.2 TB	180.8 TB	4.409 TB	97.62%
34. Legnaro - SE	ALICE::Legnaro::SE	1.128 PB	1.095 PB	33.78 TB	97.08%
37. NIHAM - FILE	ALICE::NIHAM::FILE	829.2 TB	798 TB	31.27 TB	96.23%
16. GRIF_IPNO - SE	ALICE::GRIF_IPNO::SE	669.3 TB	641 TB	28.27 TB	95.78%
39. ORNL - EOS	ALICE::ORNL::EOS	956.4 TB	912.7 TB	43.76 TB	95.43%
46. SaoPaulo - SE	ALICE::SaoPaulo::SE	203.9 TB	194 TB	9.932 TB	95.13%
12. CyberSar_Cagliari - SE	ALICE::CyberSar_Cagliari::SE	70.87 TB	66.63 TB	4.24 TB	94.02%
29. KISTI_GSDC - SE2	ALICE::KISTI_GSDC::SE2	1.446 PB	1.355 PB	93.91 TB	93.66%
2. Birmingham - SE	ALICE::Birmingham::SE	418.4 TB	389.2 TB	29.2 TB	93.02%
15. Grenoble - SE	ALICE::Grenoble::SE	319.8 TB	291.2 TB	28.56 TB	91.07%
24. ISS - FILE	ALICE::ISS::FILE	421.5 TB	383.8 TB	37.7 TB	91.05%
5. Catania - SE	ALICE::Catania::SE	1.1 PB	1.001 PB	101.6 TB	90.98%
22. IPNL - SE	ALICE::IPNL::SE	37.3 TB	33.8 TB	3.5 TB	90.61%
10. Clermont - SE	ALICE::Clermont::SE	298.3 TB	270.1 TB	28.19 TB	90.55%
52. Subatech - EOS	ALICE::Subatech::EOS	1.339 PB	1.21 PB	131.8 TB	90.39%
21. IHEP - SE	ALICE::IHEP::SE	270.7 TB	242.2 TB	28.52 TB	89.47%
27. JINR - SE	ALICE::JINR::SE	289 TB	255.2 TB	33.72 TB	88.33%
51. Strasbourg_IRES - SE	ALICE::Strasbourg_IRES::SE	65.86 TB	57.66 TB	8.203 TB	87.55%
8. CERN - EOS	ALICE::CERN::EOS	11.55 PB	9.918 PB	1.632 PB	85.87%



# What happened

- Furious pace of reconstruction of 2015 + 2016 RAW data (especially heavy) and MC
  - This will continue
- No more quiet and “easy” cleanup possible
  - Dark data (in SE, but not catalogue) < 1%, checked few weeks ago
  - Some CPass0/CPass1 to be cleaned, but little to be gained, mostly catalogue space
  - Users occupy ~10% of the SEs, some a lot over (300GB) soft quota – introduce strict quotas and recuperate the space

# What more can be done

- Next round of production cleanups
  - Between 5-10 PB on disk not accessed in last year
  - Can this be released for replica reduction?
- Urgency of this task is becoming high!
- Some respite will come when the 2017 resources are installed (April and beyond)
  - This will be temporary, as the pledges are below the requirements

# Reduction of reco output

Production	Description	Status	Run Range	Runs	Chunks	Size	Chunks	Size		
LHC16j_pass1	LHC period LHC16i - Full production pass 1	Running	255515 - 255618	21	168,070	282.5 TB	167,676	99%	45.46 TB	16%
LHC period LHC16f - Full production pass 1	LHC period LHC16f - Full production pass 1	Completed	253659 - 253978	24	265,854	449.8 TB	263,184	98%	14.25 TB	3%
LHC16t_pass1_CENT_wSDD	LHC period LHC16t - Full production pass 1, CENT trigger selection, with SDD, ALIROOT-7068	Completed	267161 - 267166	5	48,397	80.83 TB	48,379	99%	12.31 TB	15%
LHC16s_pass1_CENT_wSDD	LHC period LHC16s - Full production pass 1, CENT trigger selection, with SDD, ALIROOT-7068	Completed	266405 - 267110	25	58,950	93.87 TB	58,647	99%	36.13 TB	38%
LHC16t_pass1_CENT_woSDD	LHC period LHC16t - Full production pass 1, CENT trigger selection, without SDD, ALIROOT-7068	Completed	267161 - 267166	5	48,397	80.83 TB	48,379	99%	12.13 TB	15%
LHC16t_pass1_FAST	LHC period LHC16t - Full production pass 1, FAST trigger selection, without SDD, ALIROOT-7068	Completed	267161 - 267166	5	48,397	80.83 TB	48,378	99%	12.41 TB	15%
LHC16s_pass1_CENT_woSDD	LHC period LHC16s - Full production pass 1, CENT trigger selection, without SDD, ALIROOT-7068	Completed	266405 - 267110	25	58,950	93.87 TB	58,629	99%	35.97 TB	38%
LHC16s_pass1_FAST	LHC period LHC16s - Full production pass 1, FAST trigger selection, without SDD, ALIROOT-7068	Completed	266405 - 266998	2	3,551	5.431 TB	3,551	100%	849 GB	15%
LHC16r_pass1_CENT_woSDD	LHC period LHC16r - Full production pass 1, CENT trigger selection, without SDD, ALIROOT-7068	Completed	265589 - 266318	53	130,043	208.9 TB	128,710	98%	72.97 TB	35%
LHC16r_pass1_CENT_wSDD	LHC period LHC16r - Full production pass 1, CENT trigger selection, with SDD, ALIROOT-7068	Completed	265589 - 266318	53	130,043	208.9 TB	128,429	98%	73.15 TB	35%
LHC16r_pass1_FAST	LHC period LHC16r - Full production pass 1, FAST trigger selection, without SDD, ALIROOT-7068	Completed	265630 - 266318	9	17,048	26.67 TB	17,039	99%	4.17 TB	15%
LHC16q_pass1_CENT_woSDD	LHC period LHC16q - Full production pass 1, CENT trigger selection, without SDD, ALIROOT-7068	Completed	265305 - 265525	35	305,861	512.3 TB	304,725	99%	75.89 TB	14%
LHC16q_pass1_CENT_wSDD	LHC period LHC16q - Full production pass 1, CENT trigger selection, with SDD, ALIROOT-7068	Completed	265305 - 265525	35	305,861	512.3 TB	304,769	99%	77.11 TB	15%
LHC16q_pass1_FAST	LHC period LHC16q - Full production pass 1, FAST trigger selection, without SDD, ALIROOT-7068	Completed	265305 - 265525	35	305,861	512.3 TB	304,967	99%	77.83 TB	15%
LHC16k_pass1	LHC period LHC16k - Full production pass 1, ALIROOT-6947	Completed	256504 - 258537	213	696,219	1.134 PB	691,464	99%	250.8 TB	21%
LHC16l_pass1	LHC period LHC16l - Full production pass 1, ALIROOT-6853	Completed	258883 - 260014	89	199,244	329.2 TB	198,454	99%	105.3 TB	32%
<b>16 productions</b>					<b>2,790,746</b>		<b>4,531 PB</b>	<b>2,775,380</b>		<b>907.3 TB</b>


Reconstruction output size is ~20% of RAW

Is more aggressive pileup removal possible?

Relevant for upgrade, perhaps worth starting now

# Back to some positive news

- Analysis for QM'2016



# Performance

- The reconstruction/simulation/analysis drive for QM'2016 was by all accounts a success
- The infrastructure performed well throughout the year
- To be acknowledged
  - Costin and Miguel for the tireless central services support and improvements
  - Maarten for keeping a tab on all events on the Grid and not letting anything go wrong
  - Steady work of all T0/T1/T2experts which keep the infrastructure running

# Summary

- 2016 was another good year for the Grid computing
- Through the dedicated efforts of all experts, the infrastructure grew and delivered what was expected of it
- Keeping up with the demand for resources is still a struggle, ***especially storage***
  - This will be an issue for the observable future and we must find a way to avoid crisis situations