A Large Ion Collider Experiment

ALICE Status Report

Predrag Buncic

ALICE

Data taking progress

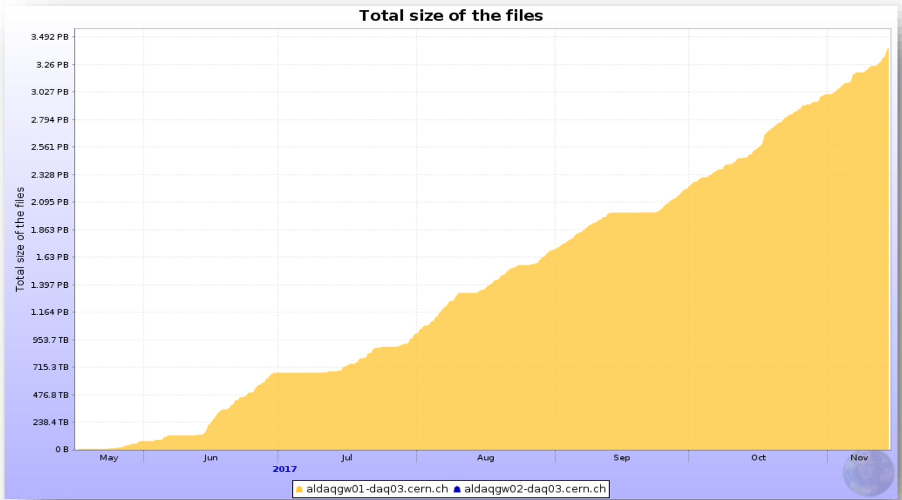

- Very smooth data taking, 4 PB registered and replicated at T1s
- Improved HLT compression commissioned
  - Raw data compression factor increased from 5.2 to 7.2

PERIOD	COLLISION	MAG FIELD (T)	INT. RATE (kHz)	n. Runs (ITS,TPC, TRD)	INT7 triggers (M)	muon_calo	CPass0/ CPass1	Time (d)	Manual Calib	PPass	Time (d)	QA and run lists
LHC17c	pp 13 TeV	B=-0.5	1-45	14	18	SETUP		0.16			0.17	Ongoing
LHC17d	pp 13 TeV	B=0	12	8	13			0.11				
LHC17e	pp 13 TeV	B=-0.5	5 -200	9	16			0.07			0.13	
LHC17f	pp 13 TeV	B=+0.5	4-22	5	12			0.06			0.1	Ongoing
LHC17g	pp 13 TeV	B=-0.2	4-70	37	144			1			2.1	
LHC17h	pp 13 TeV	B=-0.5	100-200	123	173			5.5		Next	9.8	
LHC17i	pp 13 TeV	B=-0.5	190	65	67			2.4		Running	4.4	
LHC17j	pp 13 TeV	B=+0.5	10-50	12	42			0.4			0.45	
LHC17k	pp 13 TeV	B=-0.5	190	140	150			6.7			11.8	
LHC17l	pp 13 TeV	B=-0.5	190	143	89			5.1		Next		
LHC17m	pp 13 TeV	B=-0.5	190	109	111	Running		4.5		Running		
LHC17n	XeXe 5.4 TeV	B=+0.2		2	1.9			0.25			0.16	
LHC17o	pp 13 TeV	B=-0.5	190	177	126	Running	Running	6.8				

- 2017 data reconstruction going equally well

Continuing work on data reduction

- Work on extra ~2.5% reduction of TPC raw data

Currently used option

New option
Qmax>5 cut at clusterization level, no effect on dE/dX

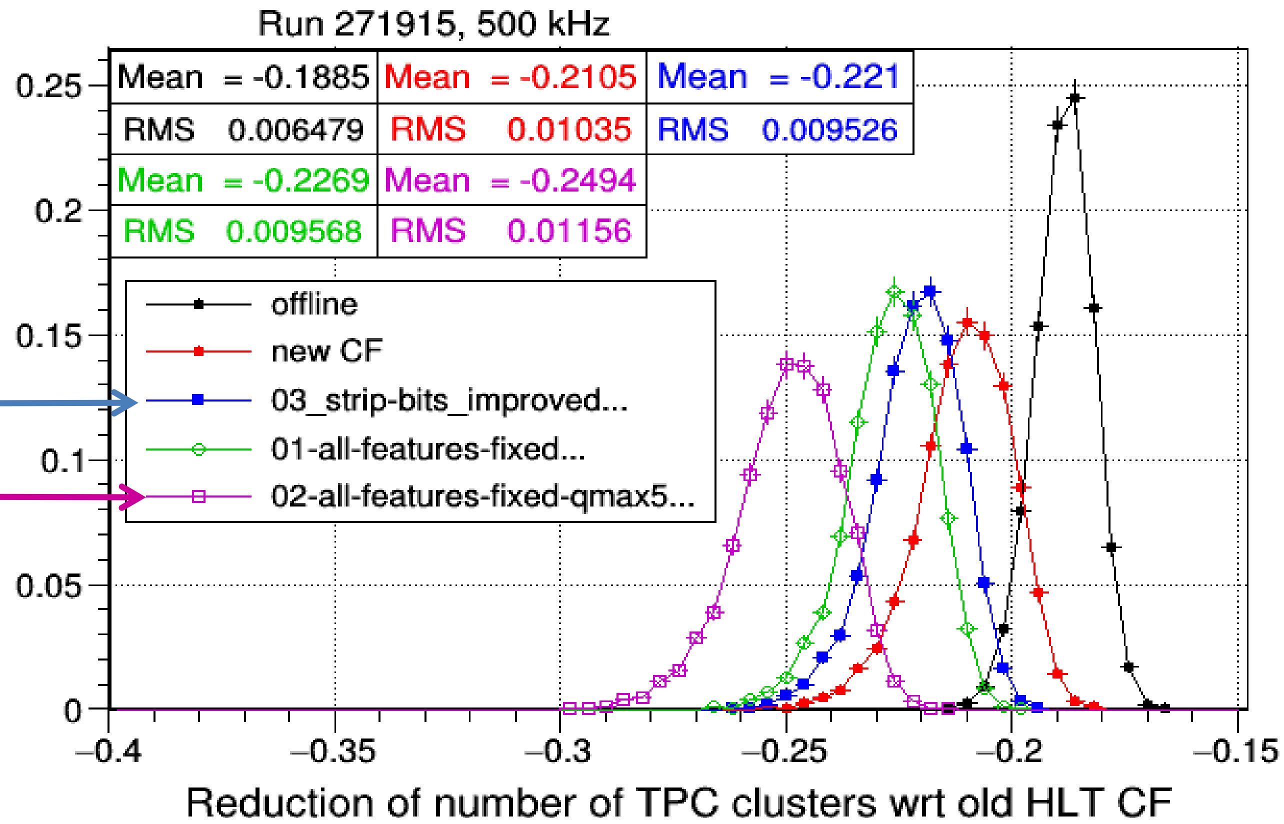

- Reduction of ESD/AOD size - optimization of V0 format and elimination of V0s not contributing to analyses in PbPb @ high IR

2017-2018 data taking scenario

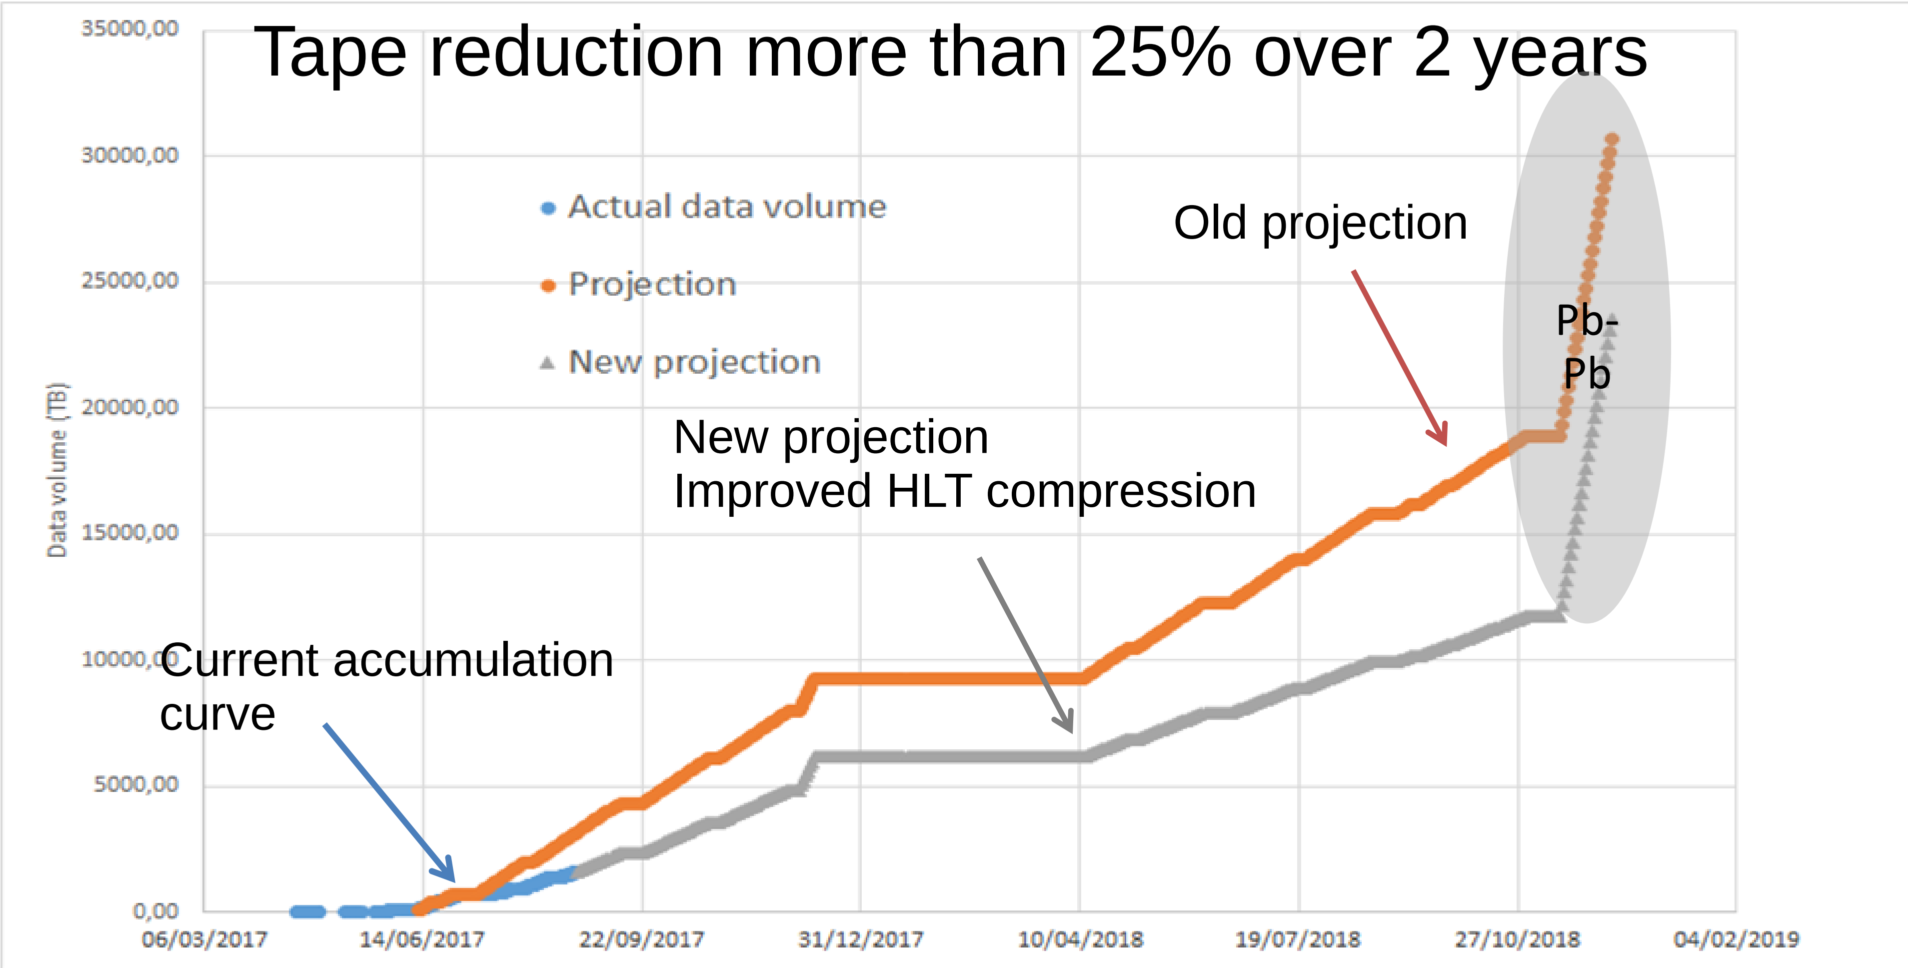

Continuing work on reducing the impact of PbPb simulation on CPU requirements

Local merging (WORKING):

All done in the same job

Generate 1 background event (only simulation)

Generate N signal events

and merge to the same background

- No need to write SDigits to AliEn (save disk)
- No I/O via network
- No need to change LPM scheme
- Background event is saved, may be used
- Reuse factor cannot be too large

Global merging:

Create a pool of background events

SDigits to be reused for several productions

- Reuse factor can be large
- Can be used by more PWGs
- No need to change LPM scheme
- Large I/O via network
- Complex workflow on LPM
- AliRoot/AliPhysics matching
  between background and signal

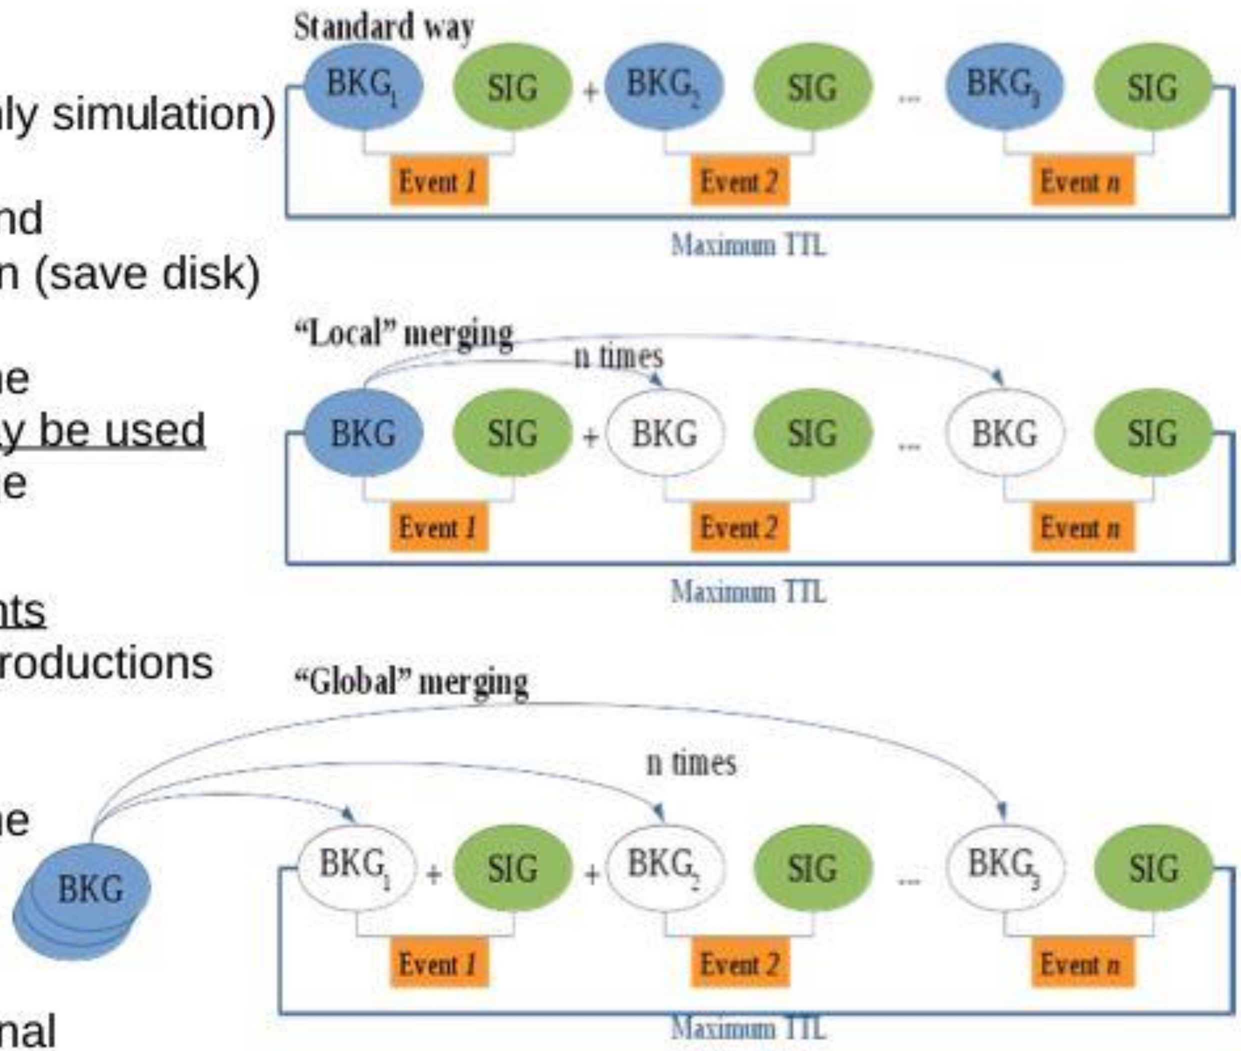

Disk cleanup initiated

- RAW and MC not accessed more than 1 year
  - 3PB deleted
  - 4PB reduced ESD replicas
- User space quotas
  - 0.5PB recuperated
- Removal of intermediate files
  - Unmerged (per chunk) QA and AOD – a lot of files, but limited gain in disk space (few PB)

What is left?

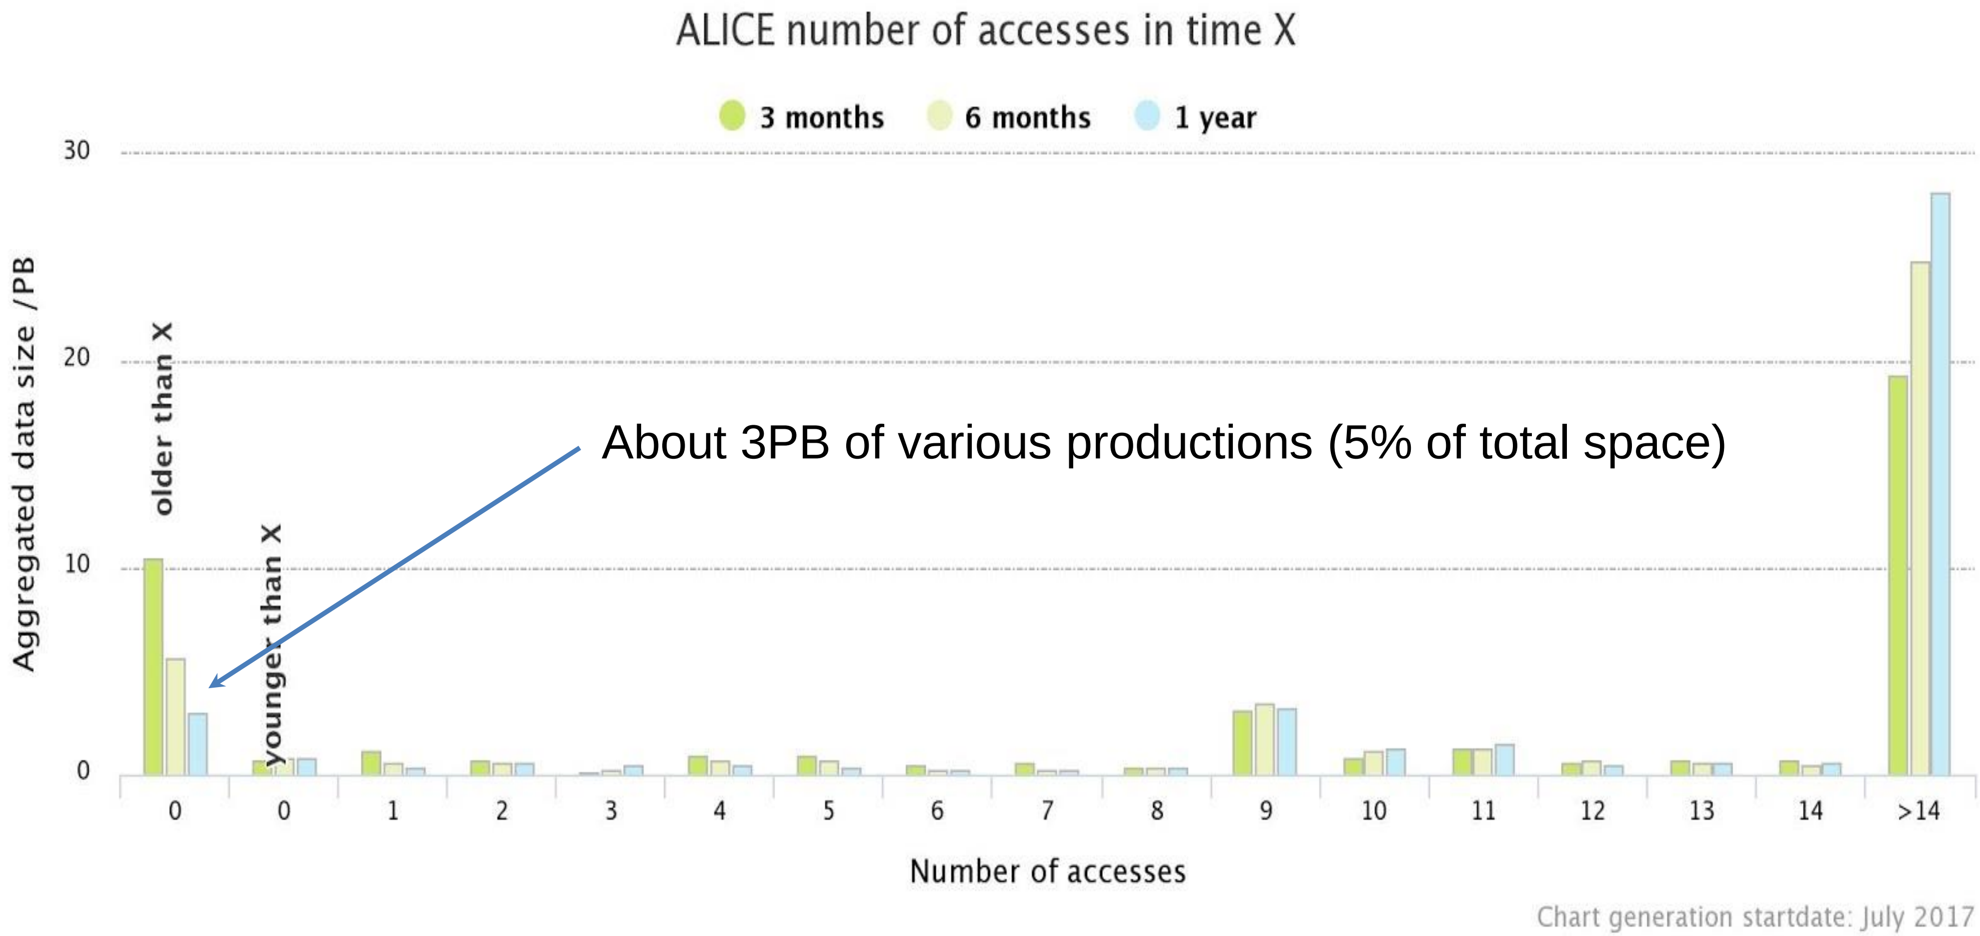
ALICE number of accesses in time X
3 months
6 months
1 year
Aggregated data size /PB
30
20
10
0
older than X
younger than X
About 3PB of various productions (5% of total space)
0
0
1
2
3
4
5
6
7
8
9
10
11
12
13
14
>14
Number of accesses
Chart generation startdate: July 2017

CPU UTILIZATION IN 2017

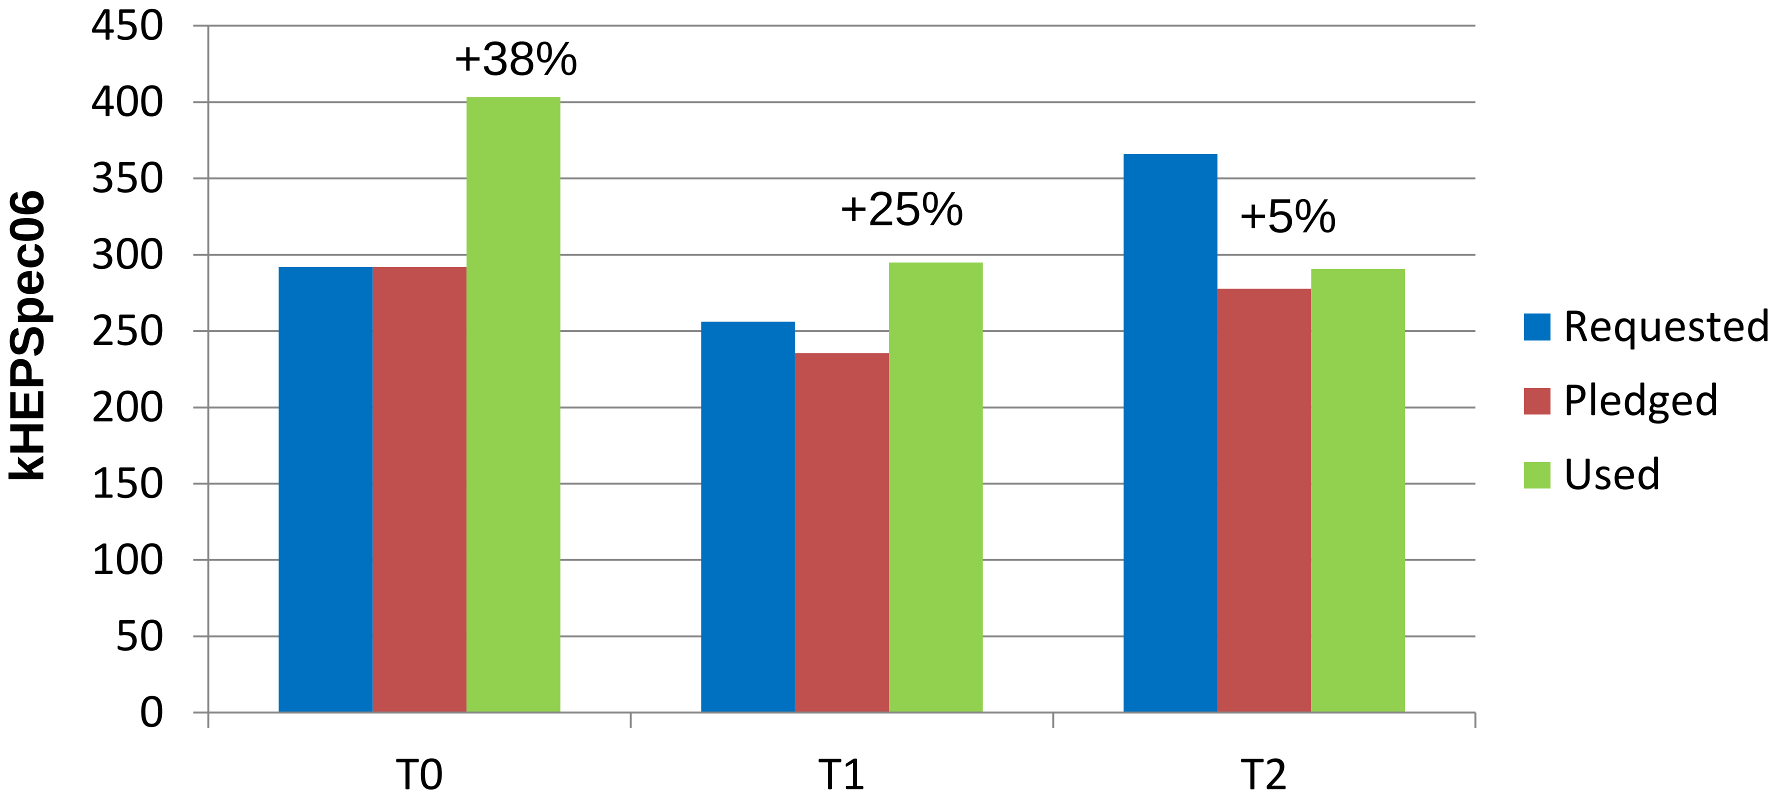

- Good utilization of opportunistic CPU resources
  - allowing for ahead of schedule processing of 2015/2016 raw data reconstruction and MC backlog
- CPU efficiency remains constant at ~83%

ALICE

2015 data Pass 2 completed

PERIOD	COLLISION	MAG FIELD (T)	INT. RATE (kHz)	n. Runs (ITS,TPC, TRD)	INT7 triggers (M)	muon_calo	CPass0/ CPass1	Time (d)	Manual Calib	PPass	Time (d)	QA and run lists
LHC15f	pp 13 TeV	Both polarities	3-25	83	143	Pass2				Pass2		
LHC15g	pp 13 TeV	B=+0.5	1-380	13	39	Pass2		2.1			1.2	
LHC15h	pp 13 TeV	B=+0.5	20-500	50	159	Pass2		6.8		Pass2	3.1	
LHC15i	pp 13 TeV	B=+0.5	30-300	112	162	Pass2		7.9		Pass2	7.1	
LHC15j	pp 13 TeV	B=+0.5	300-400	140	64	Pass2		3.2		Paused	3	
LHC15k	pp 13 TeV	B=-0.5	1-350	21	43	Pass2		0.7			1.3	
LHC15l	pp 13 TeV	B=+0.5	400	124	31	Pass2		13.3		Pass1	7.2	
LHC15m	pp 13 TeV	B=+0.5	2-6	0	0	Pass2						
LHC15n	pp 5 TeV	B=+0.5	20-50	27	133	Pass2				Pass4		
LHC15o	Pb-Pb 5 TeV	B=+0.5										

- Completed processing of 2015 and 2016 data with the latest software

ALICE

MC productions anchored to RAW data

GENERAL PURPOSE PRODUCTIONS							
ANCHOR PERIOD	**PASS**	**GENERATOR / DESCRIPTION**	**PRODUCTION NAME**	**JIRA TICKET**	**COLLISION SYSTEM**	**ALIROOT**	**EVENTS/ SAMPLING**
HC17n	pass1	HIJING	***LHC17j7***	ALIROOT-7531	Xe-Xe 5.44 TeV	v5-09-18	1M
HC17k							
HC17j	pass1	PYTHIA8	***LHC17h11***	ALIROOT-7470	pp 13 TeV	v5-09-14	25%
HC17i							
HC17h							
HC17g	pass1	PYTHIA8	***LHC17h3***	ALIROOT-7430	pp 13 TeV	v5-09-14	100%
HC17c,e,f	pass1	PYTHIA8	***LHC17h1***	ALIROOT-7426	pp 13 TeV	v5-09-12	25%

PWG PRODUCTIONS BEING PROCESSED							
PWG	**GENERATOR / DESCRIPTION**	**PRODUCTION NAME**	**JIRA TICKET**	**COLLISION SYSTEM**	**ANCHOR PERIOD**	**ALIROOT**	**EVENTS/ SAMPLING**
DQ	PYTHIA+HF->ee (low B)		ALIROOT-7422	pp 13 TeV	LHC17g		
GA	PYTHIA GJ+JJ triggered	***LHC17g5**** *(4 cycles)*	ALIROOT-7267	pp 8 TeV	LHC12c-i	v5-08-13zc-cookdedx	
GA	PYTHIA JJ	***LHC17f8**** *(11 cycles)*	ALIROOT-7268	pp 13 TeV	LHC16d,e,f, g,h,i,j,k,l,o,p	v5-09-02, v5-09-12, v5-08-22, v5-08-13q-p7-cookdedx	
GA	PYTHIA GJ+JJ triggered	***LHC17i3**** *(6 cycles)*	ALIROOT-7269	pp 13 TeV	LHC16d,e,f, g,h,i,j,k,l,o,p	v5-09-12	
GA	PHOS single particle		ALIROOT-7391	pp 13 TeV	LHC16g-o		
GA	DPMJET+JJ	***LHC17g8**** *(4 cycles)*	ALIROOT-7270	p-Pb 5+8 TeV	LHC16q,r,s,t	v5-08-13zc-cookdedx	
GA	DPMJET GJ+JJ triggered	***LHC17g6**** *(9 cycles)*	ALIROOT-7271	p-Pb 5+8 TeV	LHC16r,s LHC13d,e,f	v5-08-13zc-cookdedx	
GA	HIJING+JetJet	***LHC16h2* (9 cycles)***	ALIROOT-6822	Pb-Pb 5 TeV	LHC15o	v5-08-13-q-p5	60M
HF	EPOS + HF-jets	***LHC17h6***	ALIROOT-7404	p-Pb 5 TeV	LHC16q,t	v5-08-13zc-cookdedx	
HF	PYTHIA with cc, bb	***LHC17h8****	ALIROOT-7443	pp 13 TeV	LHC16d,e,g, h,j,o,p	v5-08-22, v5-09-02	
HF	PYTHIA + Lc		ALIROOT-7487	pp 13 TeV	LHC16d,e,g, h,j,o,p		
HF	PYTHIA + Lc		ALIROOT-7487	p-Pb 5 TeV	LHC16q,t		
LF	PYTHIA+Strange		ALIROOT-7507	pp 13 TeV	LHC16o,p		
LF	HIJING+nuclei	***LHC16h7a,b,c (_g4)***	ALIROOT-6825	Pb-Pb 5 TeV	LHC15o	v5-08-13o	900k
LF	HIJING+phi,K*		ALIROOT-7529	Pb-Pb 5 TeV	LHC15o		500k
MM	PYTHIA8 + ITSRecP	***LHC17h10a***	ALIROOT-7461	pp 13 TeV	LHC15f	v5-07-05	1M
MM	PYTHIA8 + ITSRecP	***LHC17h10b***	ALIROOT-7461	pp 13 TeV	LHC16l	v5-08-13q-p7-cookdedx	1M
MM	HIJING	***LHC17j6***	ALIROOT-7542	XeXe 5.44 TeV	LHC17n	v5-09-17	300k
MM	EPS-LHC	***LHC17j6b***	ALIROOT-7542	XeXe 5.44 TeV	LHC17n	v5-09-17	300k
UD	MC for diffractive	***LHC17h7a,b***	ALIROOT-7432	pp 13 TeV	LHC17j	v5-09-13	
UD	Central diffraction with DRgen		ALIROOT-7402	pp 13 TeV	LHC16*		

- Raw data processing is complemented by the general purpose MC productions
  - Number of generated events is driven by the analysis requirements
    - Fixed number of events
    - Percent of raw data events
- In addition we run the special purpose MCs for a given analysis

Evolution of analysis

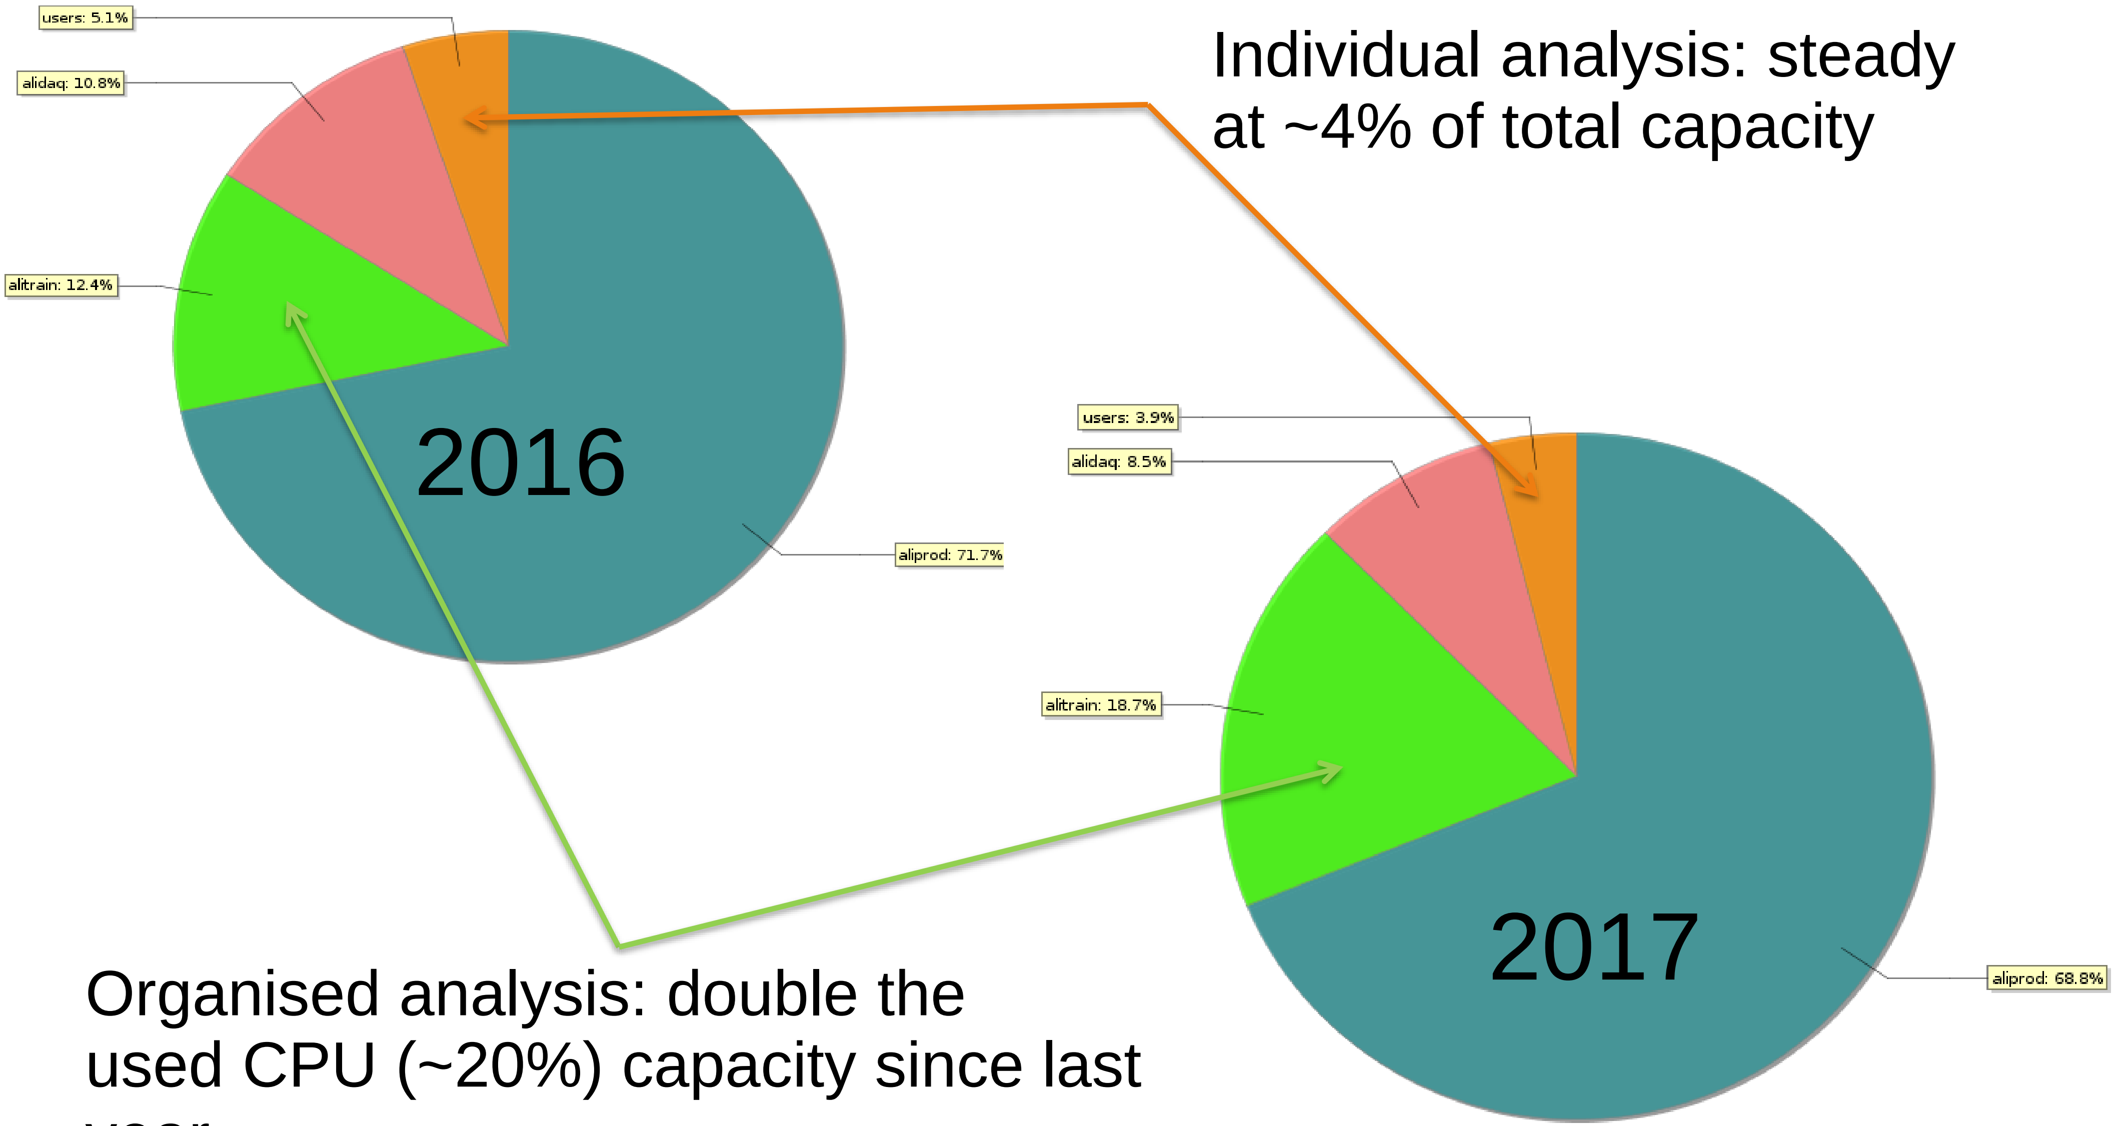

Individual analysis: steady at ~4% of total capacity

Organised analysis: double the used CPU (~20%) capacity since last year

Grid performance

- Good overall performance, 30% capacity increase

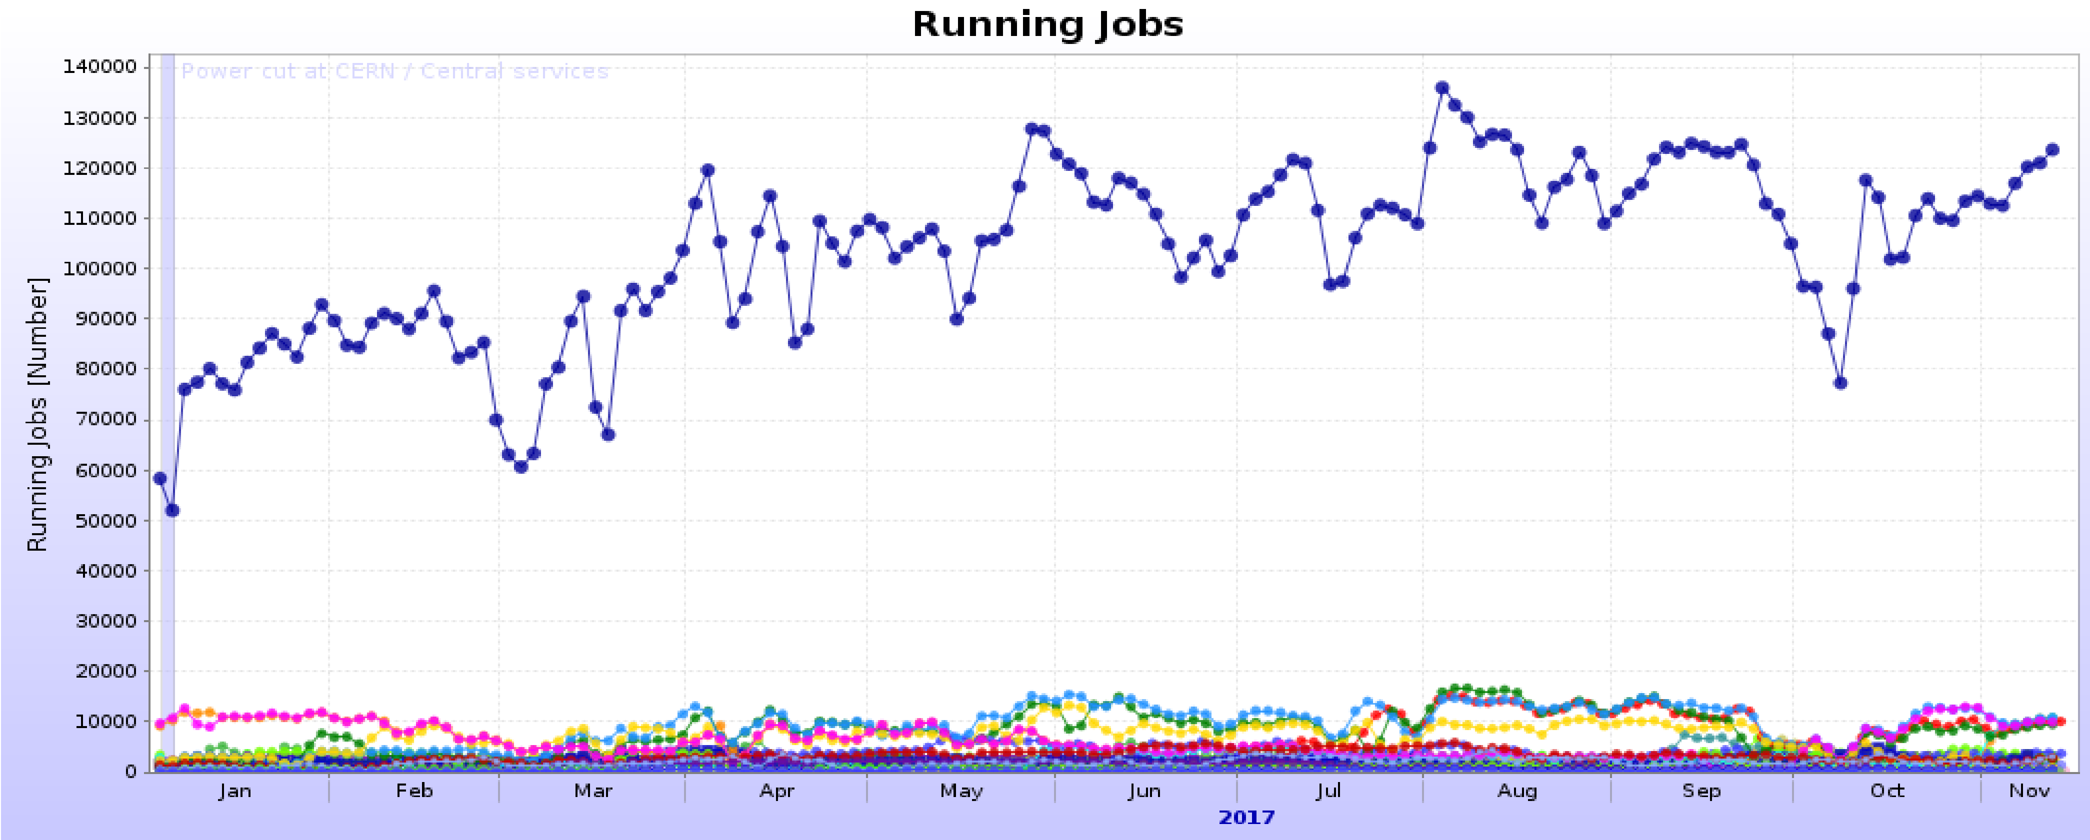

CNAF accident: Consequences for ALICE

- Capacity loss (temporary)
  - 4PB of disk (90% used, 7% of total)
  - 4500 CPU cores (4% of total, 20% of T1 capacity)
  - 6.2PB of RAW data replicas (20% of replicated data)
- Data loss
  - No RAW data loss, one more copy exist at CERN
  - Disk – max 10% (MC) to 15% (RAW) of single-replica ESDs, perhaps some single-replica user files

CNAF accident: Consequences for ALICE

- Performance loss
  - CNAF is an ***excellent*** T1, used for RAW data replication and reconstruction and has very good efficiency for analysis
  - The single copy ESDs are distributed ~uniformly across all computing centres => ESD-based analysis may lose ~10% statistics
  - No evidence yet of detrimental effect on organized analysis
- In the coming weeks we will evaluate the full impact of the incident on ALICE

Final resource request 2018/19

Resource	Site	2018 request	2019	Growth
CPU (kHS06)	T0	350.0	496.0	0.42
	T1	307.0	465.0	0.51
	T2	389.0	589.0	0.51
Disk (PB)	T0	26.2	30.7	0.17
	T1	30.5	35.8	0.17
	T2	35.2	42.0	0.19
Tape (PB)	T0	49.1	49.1	0.0
	T1	40.9	40.9	0.0

- Reducing expected readout rate from 2 to 1.25 kHz for PbPb data taking resulted downscaling 2019 request
- Reduction in tape request is due to HLT improvements and less split clusters in TPC

Table 9: Resources requested for 2018-2019.

Resource	Site	2018 request	2019	Growth
CPU (kHS06)	T0	350.0	430.0	0.23
	T1	307.0	375.0	0.22
	T2	398.0	475.0	0.19
Disk (PB)	T0	26.2	30.7	0.17
	T1	30.5	35.8	0.17
	T2	35.1	39.7	0.13
Tape (PB)	T0	49.1	49.1	0.0
	T1	40.9	40.9	0.0

- The disk request for 2018 was reduced to account for recently execute cleanup campaign that removed 7PB of rarely used datasets

Status of 2018 pledges

Tier	Pledge Type	ALICE	Required	Balance
Tier 0	CPU (HEP-SPEC06)	350000	350000	0.00%
Tier 0	Disk (Tbytes)	26200	26200	0.00%
Tier 0	Tape (Tbytes)	49100	49100	0.00%
Tier 1	CPU (HEP-SPEC06)	279549	307000	-9.00%
Tier 1	Disk (Tbytes)	30359	30500	0.00%
Tier 1	Tape (Tbytes)	42183	40900	3.00%
Tier 2	CPU (HEP-SPEC06)	312924	398000	-21.00%
Tier 2	Disk (Tbytes)	28950	35100	-18.00%

- With correction of our request for 2018 we are now in reasonably good situation with pledges for 2018

Outcome of C-RSG process

- In spite of reduction of our resource requests and many iterations with C-RSG we could not find a common ground
  - our 2018 request was approved at the current level of pledges on T0/T1
  - C-RSG questioned our requests based on perceived change of a ration of MC/RAW events that we generate
- We were referred to the LHCC to justify the MC/RAW ratio
  - 2018 resource request for T2 resources is put on hold
  - 2019 request was not reviewed

Clarification on ALICE Simulation request

- In the document presented to C-RSG we so far quoted general purpose simulation requirements and treated separately special MC productions
- In order to be fully transparent and to simplify the calculations as per the C-RSG request, we exposed the average combined general purpose and dedicated MC to raw event ratio.
  - 1 .0 -> 1.4 (pp) and 0.18 -> 0.3 (PbPb).

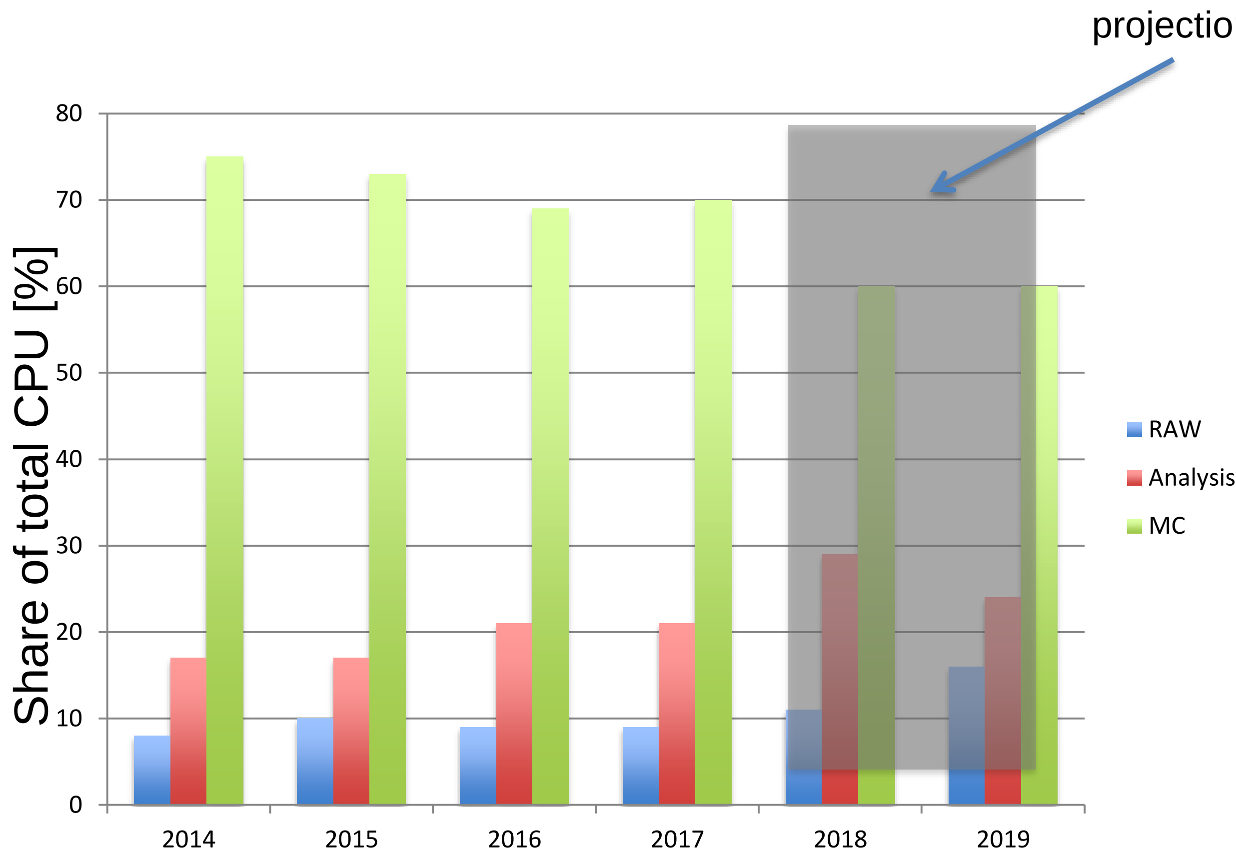

- This did not increase the number of MC events in our projections for 2018 and 2019 resulting in no change to disk and CPU requirements.
- The projected share of MC in total CPU budget is shrinking because of increased impact of analysis and CPU intensive reconstruction of 2018 PbPb data

C-RSG process

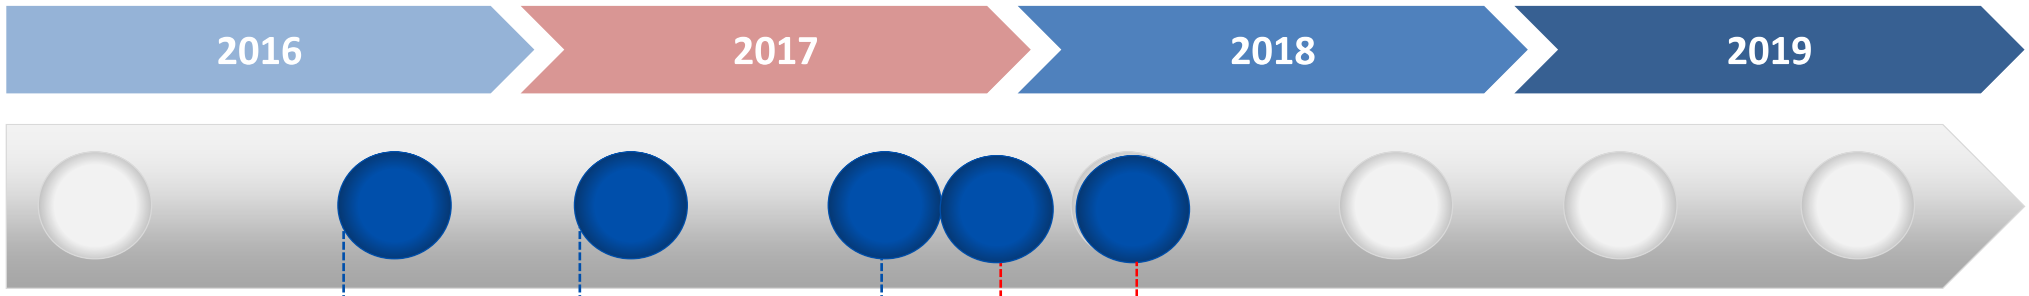

April RRB
Agreement on 2018 request
Preview of 2019 request
Pledges installed

October RRB
Presentation of 2018 request
Confirmation of 2017 pledges in REBUS

October RRB
Presentation of 2019 request
Confirmation of 2018 pledges in REBUS

- Discussion with LHCC on MC/raw data ration
- April RRB
  - Confirmation (?) of our 2019 request
  - Preview of 2020 request

Summary

- 2015/2016 RAW data reconstruction completed, no backlog in 2017 reconstructions
- Reduction of RAW data volume by 30% in p-p due to new TPC gas and improved HLT compression algorithm
- Removed 7.5PB of unpopular data from disk
- We continue to work on several fronts to reduce resource needs
  - Tape (compression)
  - Disk (reduction of number of replicas, removal of data, reduction of ESD/AOD size)
  - CPU (MC to MC embedding for PbPb MC
- We hope to streamline the future interactions with C-RSG