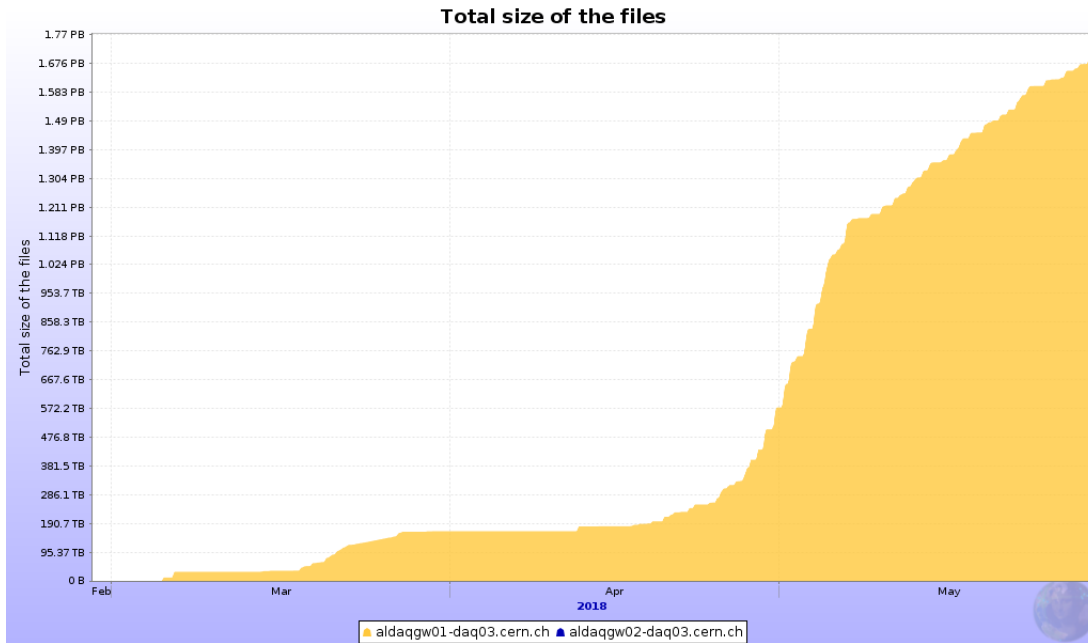




ALICE Status Report

Predrag Buncic

Data taking progress




- Smooth data taking, 1.6 PB registered and replicated at T1s
- Using Ar-CO₂ gas mixture in the TPC
 - Improvement in operation stability at higher rates, especially relevant for Pb-Pb data taking
 - Software correction methods are tested and used in production
 - About 10% increase in data volume, compensated by the improved HLT compression (now factor 8.3)

Description	Status	Events
LHC period LHC18c - Full production pass 1, CENT trigger, ALIROOT-7786	Running	0
LHC period LHC18e - Full production pass 1, TPC scan 2, ALIROOT-7824	Completed	6,721,305
LHC period LHC18b - Full production pass 1, TPC scan 2, ALIROOT-7824	Completed	6,225,345
LHC period LHC18c - Full production pass 1, FAST trigger, ALIROOT-7786	Running	0
LHC period LHC18b - Full production pass 1, TPC scan, ALIROOT-7779	Completed	7,573,142
LHC period LHC18b - Full production pass 1	Running	299,438,416
		319,958,208

- 2018 pp data reconstruction going very well


Data volume projection

Projection based on pp evt size of 1.7 MB (@ V0and = 150 kHz) and 430 Hz TPC read-out




Barrel event size in nominal running conditions

Our projections are based on a pp evt size of 1.7 MB (@ V0and = 150 kHz)



CPU UTILIZATION



- Good utilization of opportunistic CPU resources in particular at CERN
- CPU efficiency remains constant at ~83%
- Due to efficient use of available as well as the resources beyond the pledge CPU is not the problem but the disk space is considering our PbPb run at the end of this year

Large disk cleanup already done in 2017

- 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

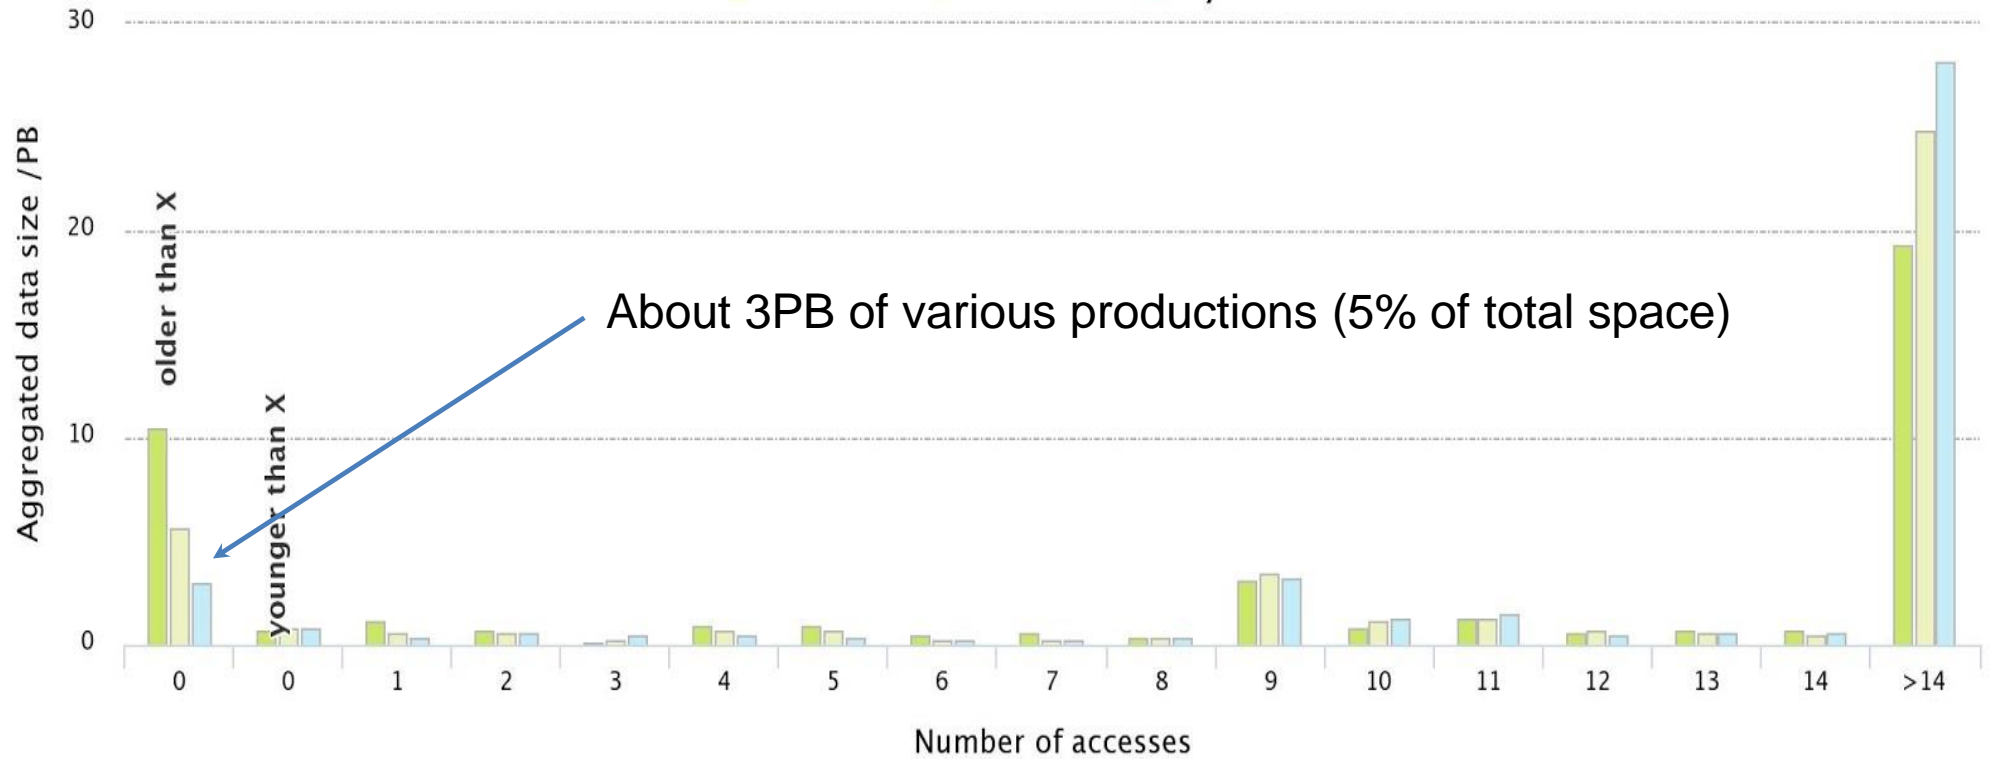




Chart generation startdate: July 2017

Conclusion of C-RSG excercies

ALICE		2017			2018		2019		
		CRSG recomm.	Pledged	Used	CRSG recomm.	Pledged	Request	2019 req. /2018 CRSR	C-RSG recomm.
CPU	Tier-0	292	292	389	350	350	430	123%	430
	Tier-1	256	235.5	295	307	279.5	365	119%	365
	Tier-2	366	279.6	299	312.9	312.9	376	120%	376
	HLT	n/a	n/a	26	n/a	n/a	n/a	n/a	0
	Total	914	807.1	1010	969.9	942.4	1171	121%	1171
	Others			39					
Disk	Tier-0	22.4	22.4	19.3	26.2	26.2	34.3	131%	34.3
	Tier-1	25.4	21.8	18.245	30.5	30.4	37.9	124%	37.9
	Tier-2	31.4	22.7	20.06	29	29	33.9	117%	33.9
	Total	79.2	66.9	57.6	85.7	85.6	106.1	124%	106.1
Tape	Tier-0	36.9	36.9	29.7	49.1	49.1	44.2	90%	44.2
	Tier-1	30.9	30.6	22.3	40.9	42.2	37.7	92%	37.7
	Total	67.8	67.5	52	90	91.3	81.9	91%	81.9

- After two months of intense discussion C-RSG accepted our arguments and approved our computing resource request for 2019 (and confirmed 2018 status at the level of existing pledges)

Resource evolution (I.Bird)



- 2010-2018 – pledges
- 2019, 2020 assume same as 2018
- 2021 assume 1.5 x 2018

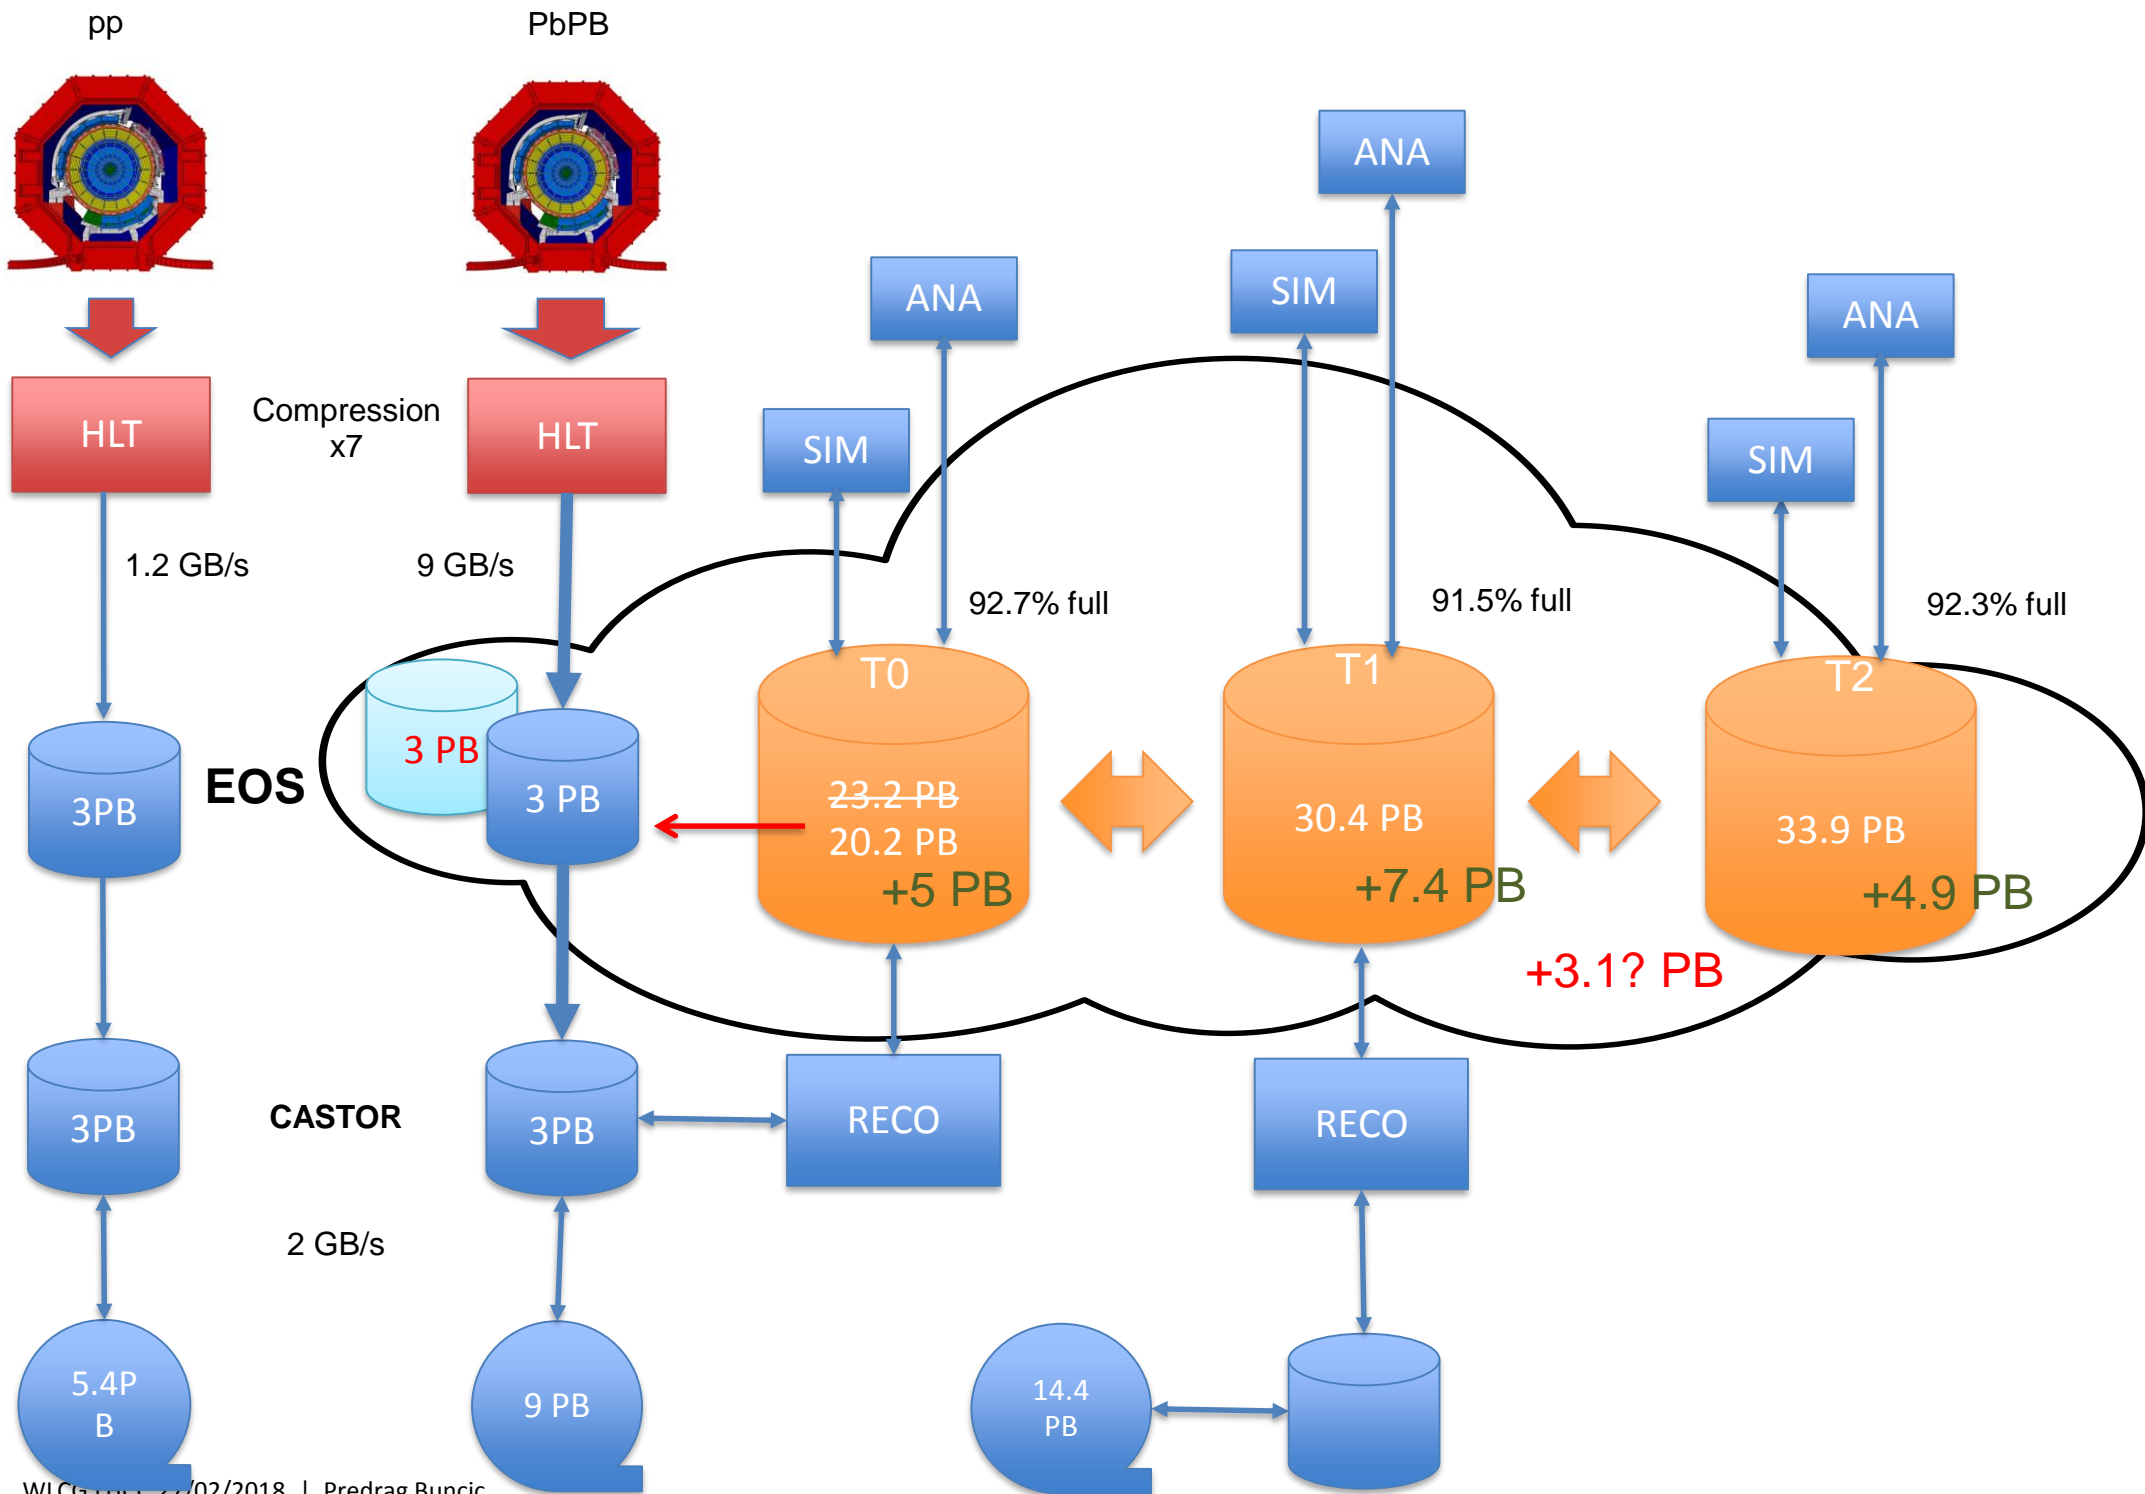
- Due to budget constraints, consolidation works foreseen in the CERN Computing Centre and a wish to maximize overall resources available to experiments at the start of Run 3 in 2021, CERN/IT has made their preliminary planning assuming that no new equipment purchase will take place during LS2.
- 0 PB disk increase in 2019 was the starting point for one month long discussion with CERN management

Agreement with CERN management

- Due to budget constraints, consolidation works foreseen in the CERN Computing Centre and a wish to maximize overall resources available to experiments at the start of Run 3 in 2021, CERN/IT has made their preliminary planning assuming that no new equipment purchase will take place during LS2.
- In order to partially implement ALICE 2019 request, CERN will pledge 5 PB in 2019 which should be sufficient to meet the foreseen raw data processing needs.
 - Any further disk needs at CERN should be met by the archiving more data to tape.
 - As a consequence, some of the ongoing analysis may have to be postponed or even canceled.

2018 data taking

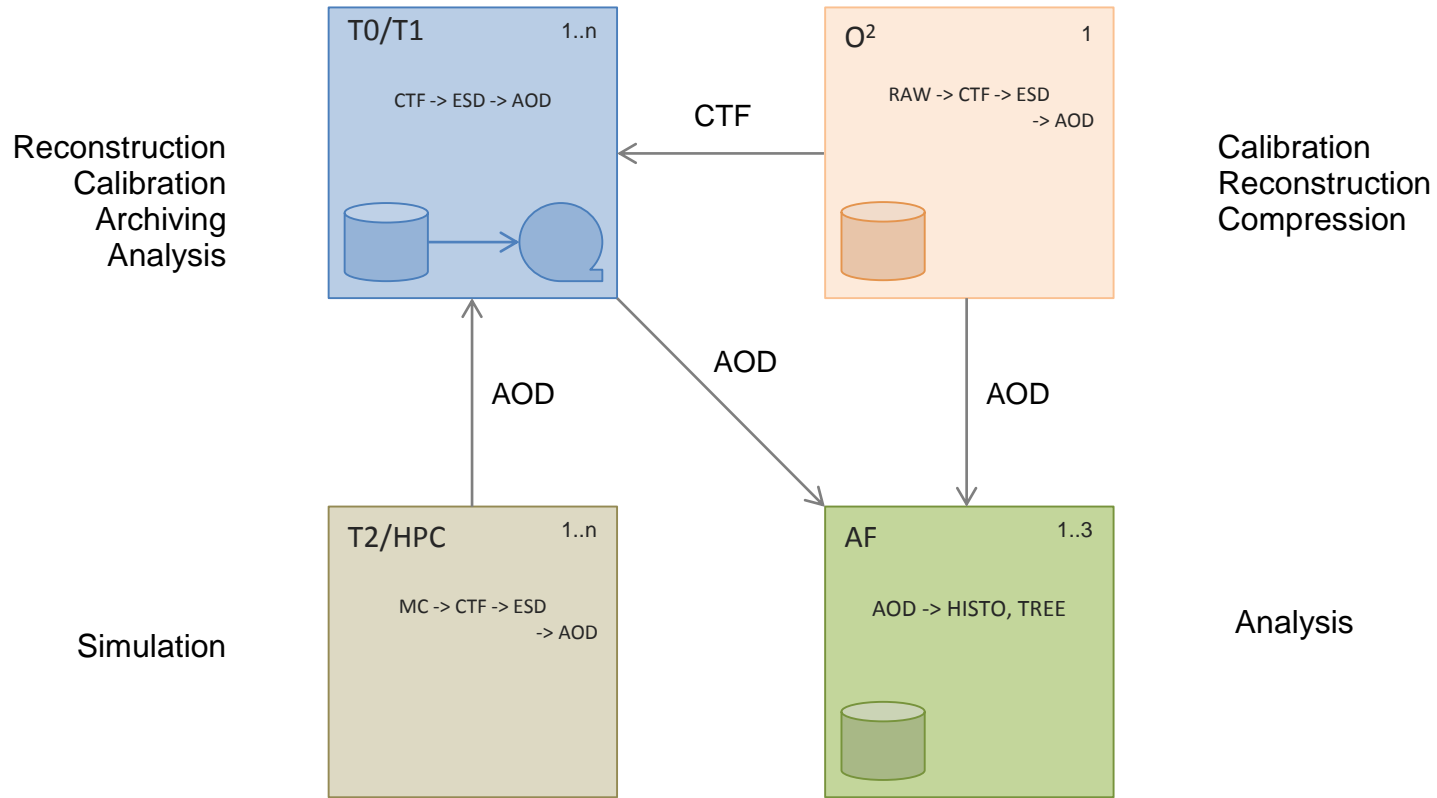
- As of today, CERN has fully implemented its disk pledge for ALICE for a total of 26.2 PB
 - In addition, CERN has provided 3 PB of temporary disk space for the duration of data taking.
- Given the expected data rate of up to 9 GB/s, the capability of tape system at CERN to archive data to tape at 2-3 GB/s and a total Pb-Pb raw data volume 9 PB, we estimate that an additional 3 PB of disk will be needed to provide a sufficient disk buffer to assure safe data taking.
 - This additional disk space will be liberated by archiving to tape some of the least used analysis data.



Measures to address shortage of disk

- We have to start archiving ESDs/AODs to tape
 - This will require
 - Cooperation and understanding our ALICE physics community that some datasets and analysis will have to be put temporary offline
 - Correction (increase) of our tape request
 - Developments and adjustments of our data management software
 - Increase in manpower dedicated to data management
- According to computing resource sharing procedure as defined by the ALICE Constitution, the CERN pledge reduced by 3.1 PB in 2019 will automatically increase obligations of other FAs in proportion to their M&O share.
 - We will be forced to reduce usage of CERN resources to reconstruction
 - Simulation will have to be run and output stored elsewhere on the Grid
 - That should be discussed with the representatives of FAs in preparation for October RRB.

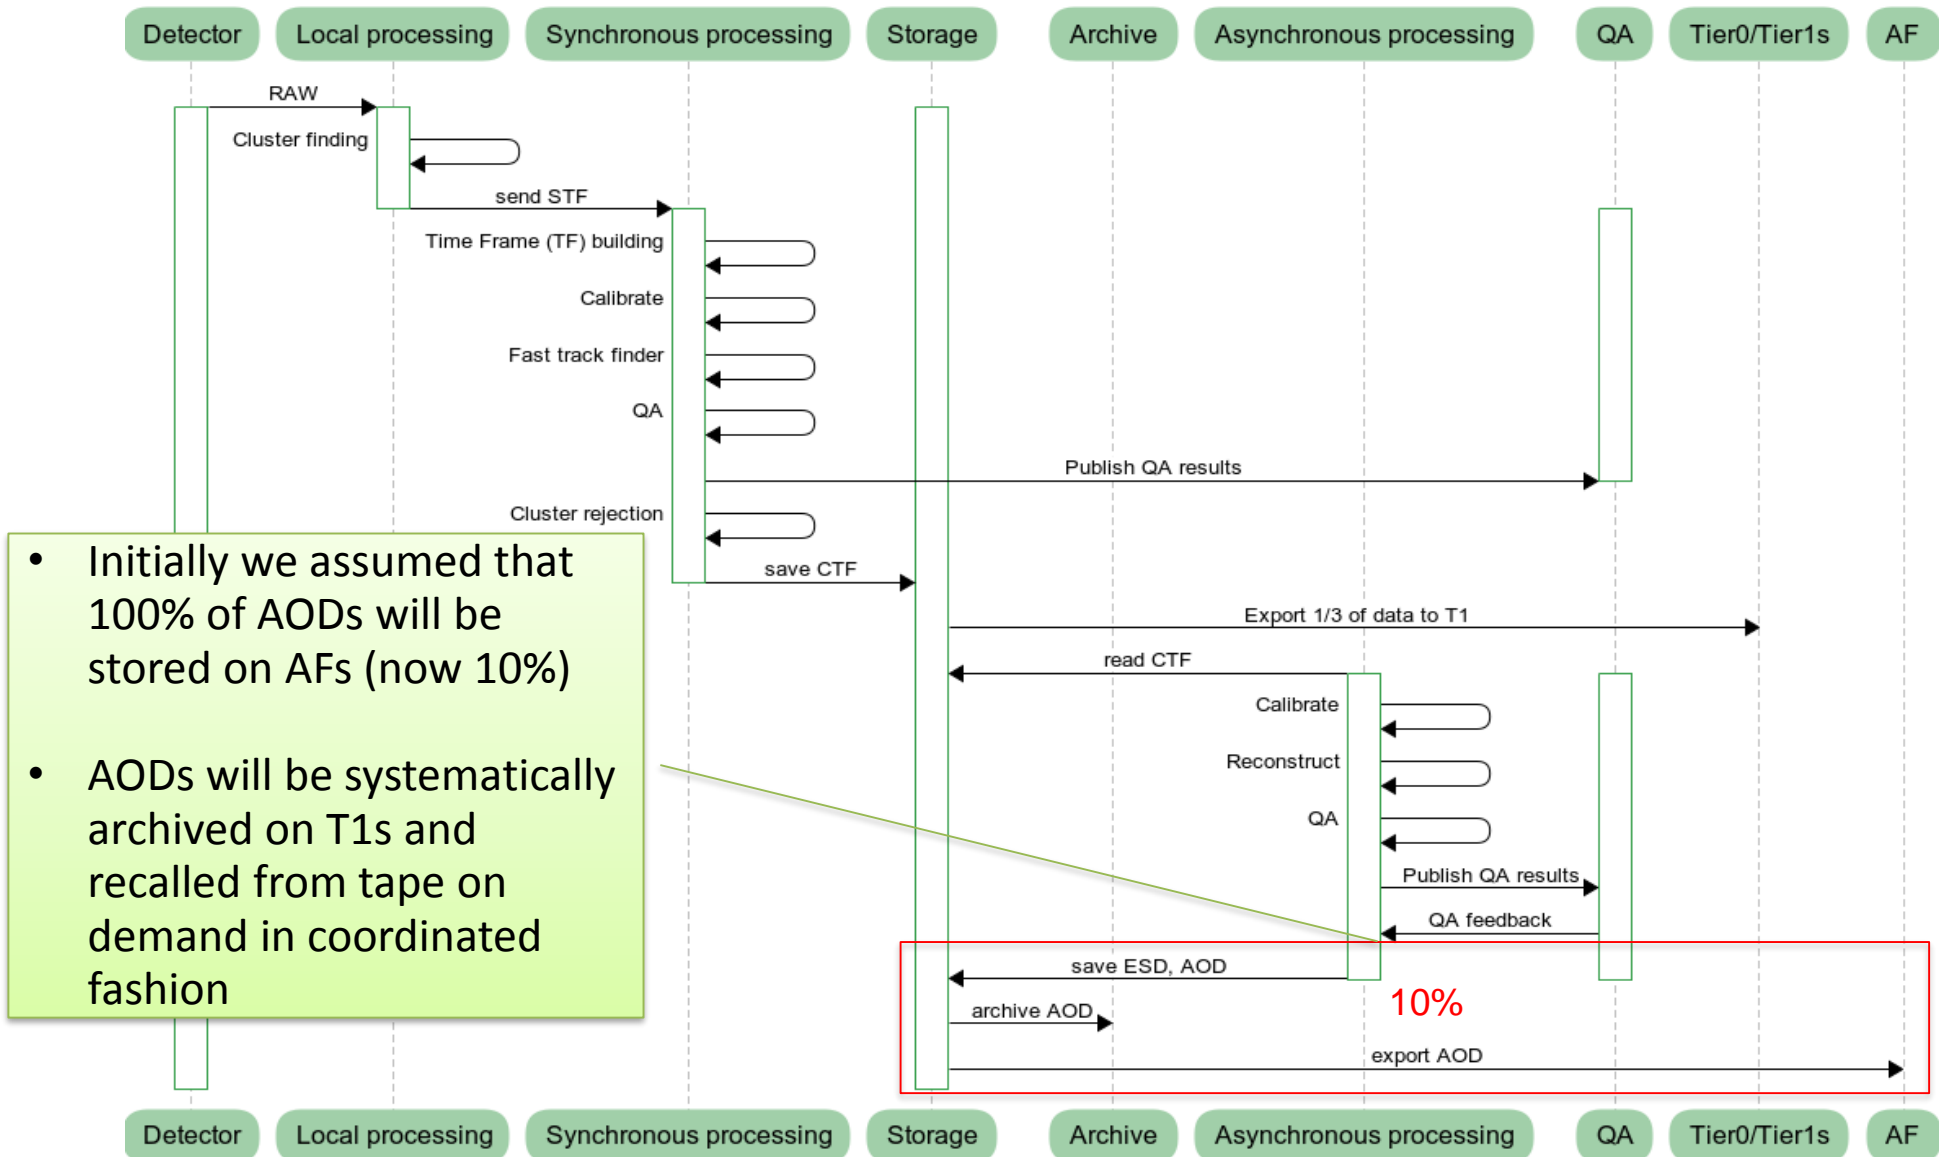
Outlook for Run 3



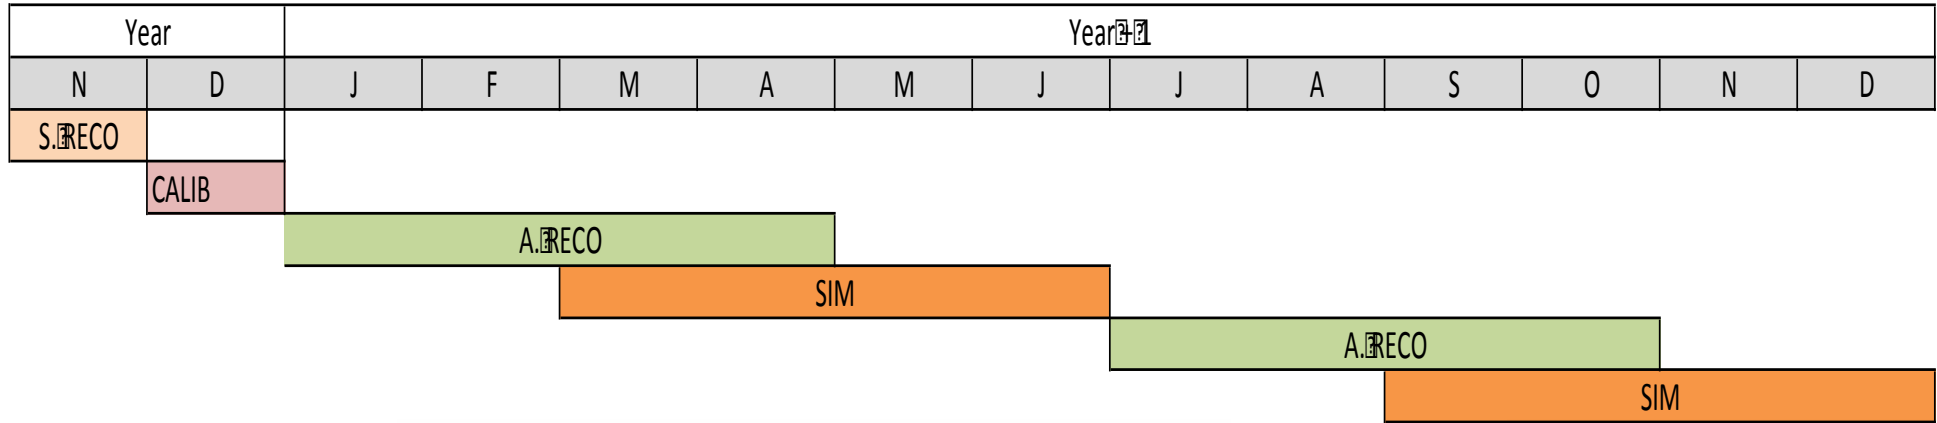
Sharing parameters

	S.RECO	A.RECO	SIM	ANA	CTF	AOD
O2	100%	60%				
T0		15%			60%	100%
T1		25%			40%	
T2			100%			
AF				100%		100%

Run 3 Processing sequence



Data Processing Timeline (PbPb)



Data Parameters		Pb+Pb	p+p	p+Pb	# of versions	Tape Copy		Pb+Pb	p+p	p+Pb
Raw data size (kb)	CTF	1600	30	110	1	1		1'600.00	30.00	110.00
	ESD	10%	100%	100%	1	0		160.00	30.00	110.00
Derived data size as % of CTF	AOD	50%	100%	50%	1	1		800.00	30.00	55.00
	HISTO	0.15%	0.15%	0.15%	1	0		2.40	0.05	0.17
	MC	150%	150%	150%	1	0		2'400.00	45.00	165.00
	MCAOD	150%	150%	150%	1	1		2'400.00	45.00	165.00

Estimated CPU/Disk needs								
S.RECO		A.RECO		SIM		ANA		
HS06s/ev.	MB/ev.	HS06s/ev.	MB/ev.	HS06s/ev.	MB/ev.	HS06s/ev.	MB/ev.	
pp	0.3	0.030	11.4	0.030	950.0	0.045	0.10	0.0000
pPb	1.3	0.110	45.5	0.055	1600.0	0.165	0.10	0.0002
PbPb	26.3	1.600	910.0	0.800	24000.0	2.400	0.90	0.0024

Estimates based on new and improved GPU based reconstruction

Conservative estimates based on the current performance of the Offline reconstruction

Data processing sharing between O2 and Grid

CPU and storage required for processing one year data sample

	Synchronous Reconstruction			Asynchronous Reconstruction			Simulation				Analysis	
	days	CPU (kHS06)	Disk (PB)	days	CPU (kHS06)	Disk (PB)	days	% of RAW	CPU (kHS06)	Disk (PB)	CPU (kHS06)	Disk (PB)
pp	14	16	1.6	30	119	0.8	120	7.5%	187	0.09	34	0.001
pPb	14	263	24.0	30	1911	6.0	270	5.0%	373	0.90	132	0.019
PbPb	24	605	37.8	120	1884	17.2	120	1.0%	497	0.52	226	0.052

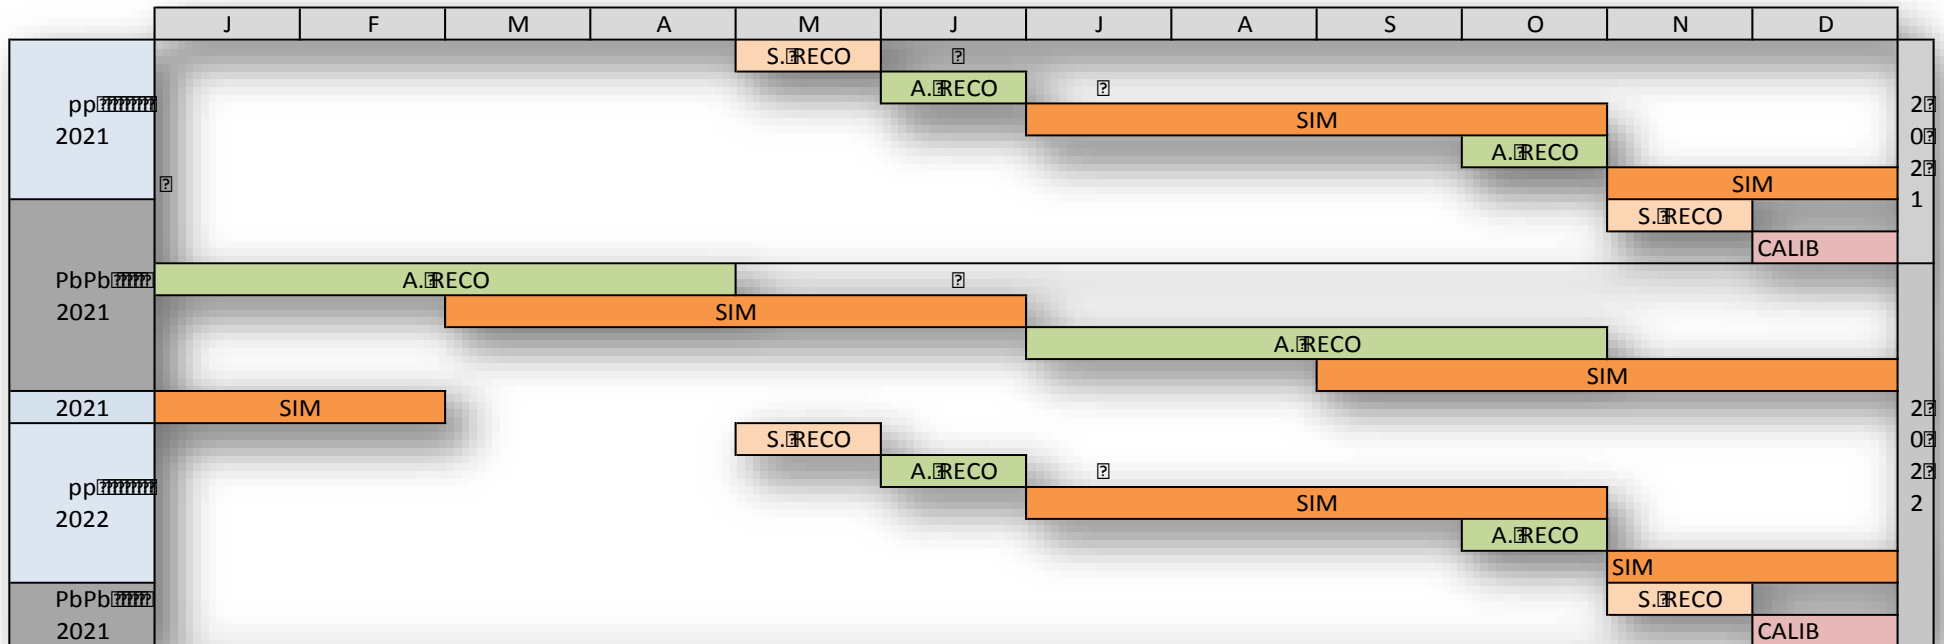


	O2		T0			T1			T2		AF	
	CPU (kHS06)	Disk (PB)	CPU (kHS06)	Disk (PB)	Tape (PB)	CPU (kHS06)	Disk (PB)	Tape (PB)	CPU (kHS06)	Disk (PB)	CPU (kHS06)	Disk (PB)
pp	119	2.3	18	0.6	2.2	30	0.6	0.6	187	0.1	34	0.9
pPb	860	33.5	287	4.5	28.4	478	7.5	7.5	373	0.9	132	6.9
PbPb	848	52.9	283	12.9	50.6	471	13.7	13.7	497	0.5	226	17.7
pp+PbPb	848	55.2	283	13.5	52.9	471	14.3	14.3	497	0.1	260	19

Mapping of the requirements to O2 and Grid Tiers

- Computing resource requirements for processing data collected in a typical Run 3 data taking year (pp + Pb run)
- O2 disk buffer has to be big enough to host CTF data corresponding to one year of data taking as well as transient and temporary files (ESD)

Estimating requirements for 2021/2022



“Standard” data taking year

(1) pp+PbPb	848	55.2	283	13.5	52.9	471	14.3	14.3	497	0.1	260	19
-------------	-----	------	-----	------	------	-----	------	------	-----	-----	-----	----

C-RSG	2020	O2		T0			T1			T2		AF	
		CPU (kHS06)	Disk (PB)	CPU (kHS06)	Disk (PB)	Tape (PB)	CPU (kHS06)	Disk (PB)	Tape (PB)	CPU (kHS06)	Disk (PB)	CPU (kHS06)	Disk (PB)
				20%	25%		20%	25%		20%	5%	20%	25%
				450.00	41.80	36.60	364.00	45.60	42.10	376.00	39.00	250.00	16.00
	2021	1710.00	60.00	540.00	52.25	89.50	436.80	57.00	56.45	451.20	40.95	300.00	20.00
	(2) 2022	1710.00	60.00	648.00	65.31	142.39	524.16	71.25	70.80	541.44	42.36	360.00	25.00
	(2) (1)	862.34	4.83	365.45	-0.43	0	46.41	-0.10	0	44.64	1.30	1.90	-0.48

- Assuming a nominal 20-25% growth over 2020 resources presented (but not yet approved) to C-RSG, we should be in position to process Run 3 data in 2021/22 (and beyond).

2021 vs 2018 comparison

	CPU	Disk	Tape	Year on year growth (%)			2021 vs 2018 (%)		
2018	967	86.09	78.81						
2019	1172	106.26	78.81	21%	23%	0%			
2020	1191	126.41	78.81	2%	19%	0%			
2021	1428	150.2	149.52	20%	19%	90%	48%	74%	90%
2022	1713.6	178.95	220.34	20%	19%	47%			

Summary

- 2018 data taking and processing is progressing well
- This year's C-RSG exercise was finally concluded with positive recommendations for ALICE
 - ALICE computing resources request was re-based to the level of pledges in 2018 and assumes approximately 20% growth of CPU and disk storage at all centers
 - The estimated resource needs in 2021 are consistent with flat budget guideline
- Unfortunately, CERN is not in position to fully implement the request in 2019
 - Mitigation plan was agreed to support 2018 PbPb data taking
 - Disk space will have to be liberated by moving EOD/AODs to tape and postponing some ongoing analysis
- ALICE Computing Model for Run 3 already assumes that AODs will be systematically archived to tape
 - Continuous quick turnaround analysis at AF(s) on 10% sample
 - Full dataset available on T1a on request and after approval