ALICE Computing Resource

Stefano Piano

Outline

- Resources usage in 2022
- Status of data taking and processing
- ALICE 2023 Resource Requests
- Impact of the war in Ukraine on ALICE computing resources
- Pledged resources for 2023
- Estimates for 2024
- LHC energy consumption reduction

Resources usage in 2022

CPU utilization and breakdown by job types

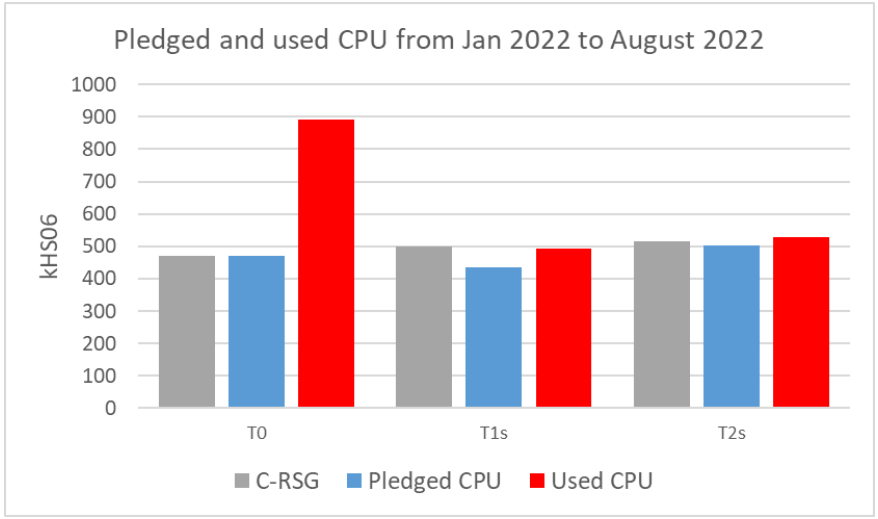

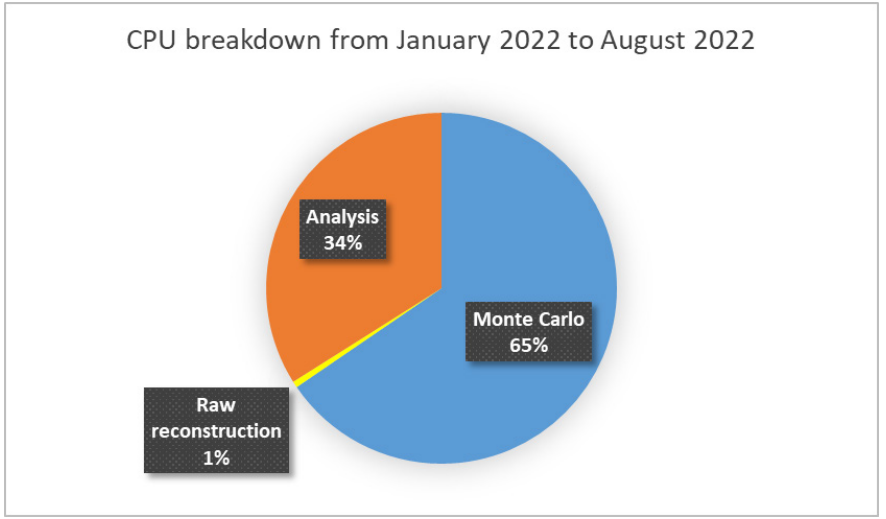

- Full utilization of the available resources with significant amount of opportunistic CPU usage at the T0, where Used/Pledged CPU ~1.9
- CPU capacity covers the pledges @ T0, T1s and T2s but pledge / C-RSG 90% @ T1 in 2022

- MC simulations in preparation for Run 3
- High analysis activity for summer conferences and publications and AO2D conversion campaign
- Raw reconstruction of the first Run 3 pp runs and cosmics data

DISK and TAPE utilization

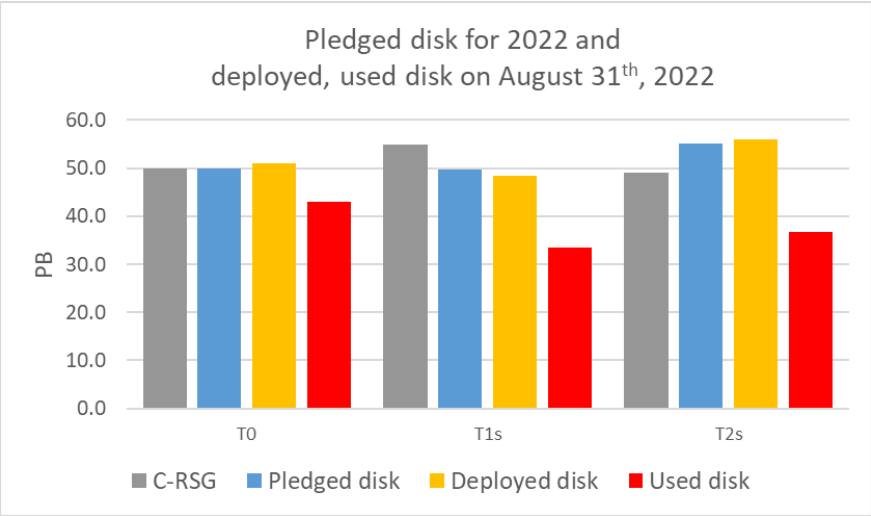

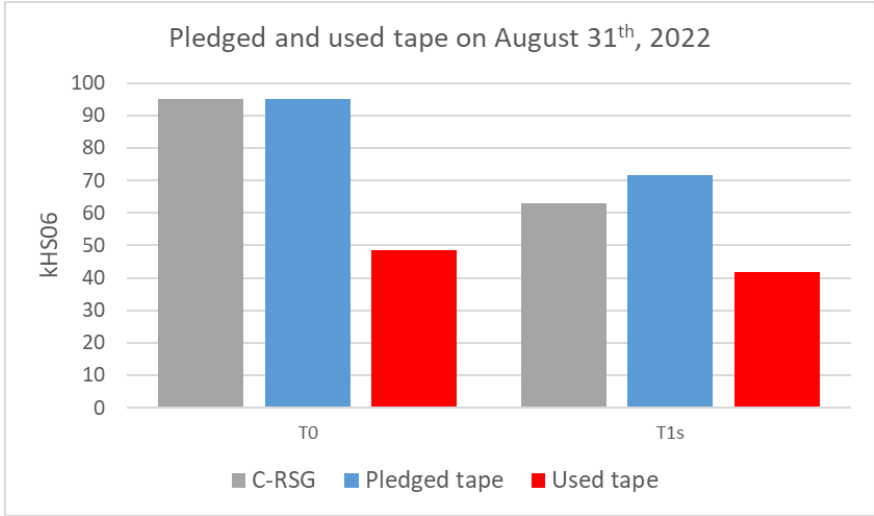

- The deployment of the pledged disk for 2022:
  - 100% at T0 and T2s, 98% at T1s
- Overall deployed / 2022 C-RSG ~100%:
  - Deficit at T1s, surplus at T2s
- Current disk use ~75% of deployed capacity

- Archival of Run 3 commissioning data ongoing
- Pledges / C-RSG: 100% at T0 and 114% at T1s (8.8 PB surplus in 2022)
- **2022 expected usage in line with the requested resources as for CPU and disk**

Data accumulation on O2 Disk Buffer

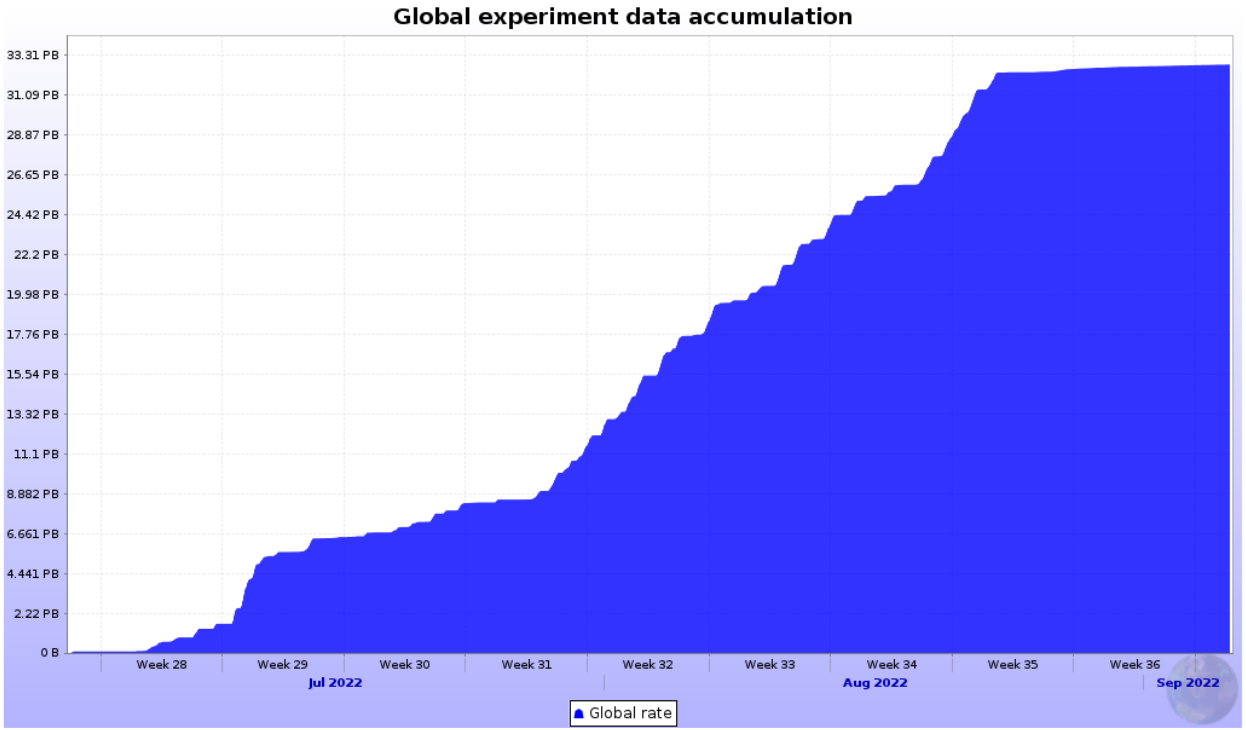

Used disk buffer to date: 47 PB
Buffer size: 91.5 PB

Since July 1st collected 32 PB:
80% of CTFs and 20% of rawTFs

About 15 PB of data on O2 disk buffer was collected before July

Archival and deletion of commissioning data ongoing

pp @ 500 kHz CTF will be removed once skimmed with the event selections

Status of data taking and processing

2022 data taking

- Collisions at injection energy: 6 stable beam fills in May and in June
  - Collected successfully about 116M pp collisions @ 900 GeV
- Data taking at top energy started the 5th of July:
  - Priority to commissioning of the apparatus and to HI preparation:
    - Low-rate magnet scans with different polarities of the solenoid
    - Interaction rate ramp-up from 50 kHz to 500 kHz + 1 MHz tests
    - High-rate tests at 2/3/4/5 MHz
    - Beam rate scan from 1 MHz to 0.6 kHz
  - Collected 278B collisions of pp @ 500 kHz:
    - Calibrations ongoing
    - Quality of data will be fully assessed once the calibrations are completed and included
    - To be processed with the first pass of asynchronous reconstruction
    - Followed by the event selections and the CTF skimming

Data and simulation processing

Period	Energy	B field	N. runs	N. runs processed
LHC22c,d,e	900 GeV	B-,B-,B+	47	42
LHC22f	13.6 TeV	B+	31	12
LHC22g	13.6 TeV	B+,low	2	2
LHC22h	13.6 TeV	B0	10	10
LHC22i	13.6 TeV	B-,low	3	2
LHC22j	13.6 TeV	B-	25	17
LHC22m	13.6 TeV	B-	174	4

About 75% of the runs collected before August already processed with first pass of async reconstruction for the commissioning

Large production will start in the next days on EPN farm (CPU+GPU) for the physics selections

In addition, TPC scan runs processed with residual extraction and pass 5 of Oct 2021 pilot beam data with updated alignment

- Extensive Run 3 simulation:
  - pp @ 13.6 TeV productions for the selection studies:
    - Completed the General Purpose MB and the signal embedding simulations
  - Production of MB pp @ 900 GeV anchored to pilot beam 2022
  - Ongoing Pb-Pb targets 2.6×10^8 events for the validation against 2018 Pb-Pb results
- Large ESD→AO2D conversion campaign:
  - All the Run 1 and Run 2 raw data periods + MC general purpose

ALICE 2023 Resource Requests

Processing plan for 2022 and 2023

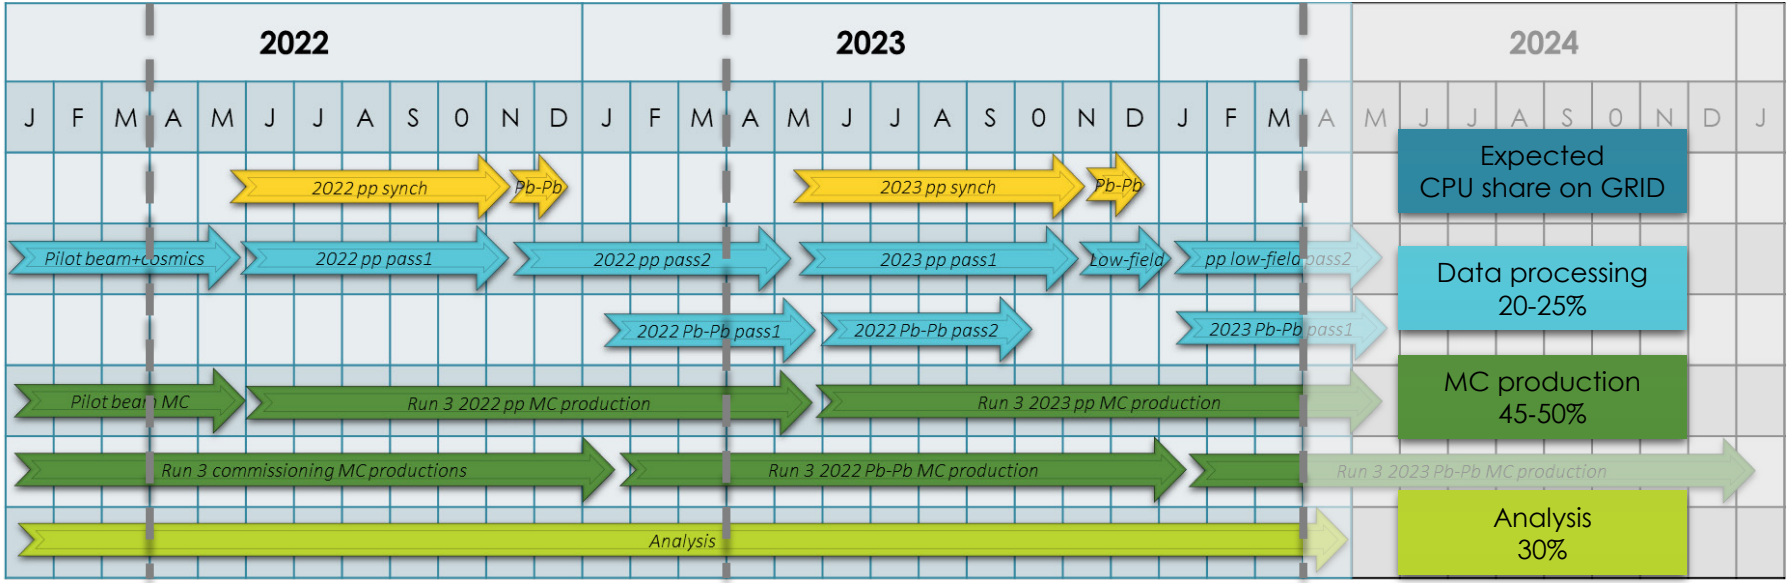

- 2022 CPU and disk requests: commissioning + 2022 pp data + 65% of 1st 2022 Pb-Pb async pass
- 2023 CPU and disk requests: 35% (70%) of 1st 2022 (2023) Pb-Pb async pass + 100% of 2nd 2022 Pb-Pb async pass + 2023 pp full and low field data

2023 Resource Requests

ALICE		2022		2023			
		C-RSG	Pledge	Req.	C-RSG	C-RSG 2023 / C-RSG 2022	C-RSG 2023 / Pledge 2022
CPU [kHS06]	Tier-0	471	471	541	541	115%	115%
	Tier-1	498	448	572	572	115%	128%
	Tier-2	515	517	592	592	115%	115%
	ALICE Total	**1484**	**1436**	**1705**	**1705**	**115%**	**119%**
Disk [PB]	Tier-0	50.0	50.0	58.5	58.5	117%	117%
	Tier-1	55.0	49.7	63.5	63.5	115%	128%
	Tier-2	49.0	55.2	57.5	57.5	117%	104%
	ALICE Total	**154.0**	**154.9**	**179.5**	**179.5**	**117%**	**116%**
Tape [PB]	Tier-0	95.0	95.0	131.0	131.0	138%	138%
	Tier-1	63.0	71.8	82.0	82.0	130%	114%
	ALICE Total	**158.0**	**166.8**	**213.0**	**213.0**	**135%**	**128%**

- Resource requirements for 2023 are unchanged since the previous report
- The C-RSG endorsed the ALICE 2023 resource requests in April RRB
- Standard growth for CPU (+15%) and disk (+17%) in 2023 compatible with flat budget, step for tape (+55 PB), where strategy B have been considered
- More CPU and disk growth needed to reduce the deficit at Tier 1 (+28%)

Impact of the war in Ukraine on ALICE computing resources

Pledges for 2023 and compensation for RU

- FAs entered their 2023 pledge input into CRIC by September 16th
- RU sites and JINR continue to provide resources to ALICE
- But ALICE cannot rely on the RU sites for the long-term retention of data:
  - T1 tapes will not be used
  - Disk is used for replicas only
  - CPU is used until network evolution allows
- The resources planned to be deployed in Tier 1 RU needed to be replaced in the other Tier 1 sites:
  - The expected 14.5 % contribution of RU to be distributed among Tier 1 sites
  - Considered the Tier 1 share (with and) without RU in agreement with WLCG
  - Priority for the experiment: tape, disk and CPU
  - The cost of energy (and the hardware market trend) is a major concern for all FAs
  - All FAs have offered to help ALICE in case of shortfall

Impact on Tier 1 and 2 share

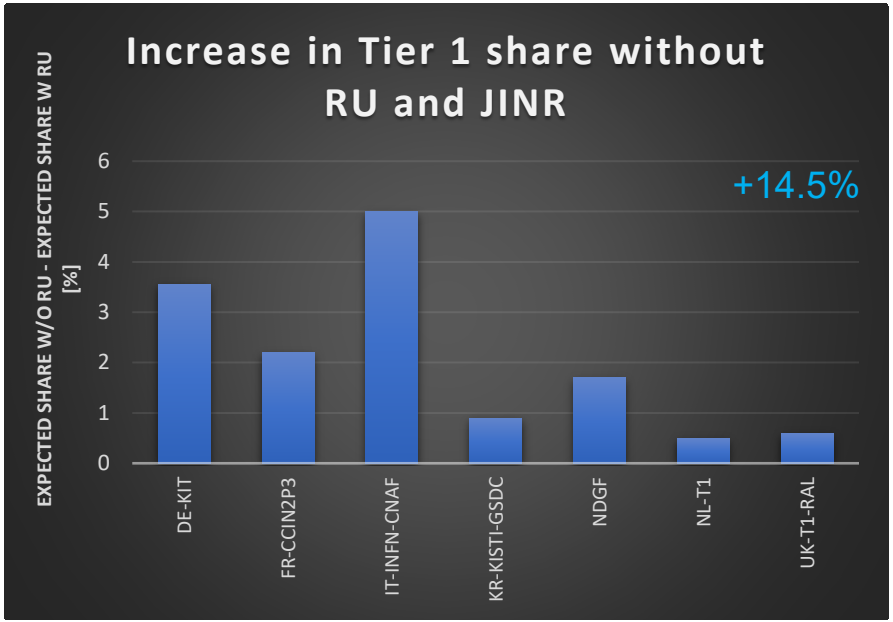

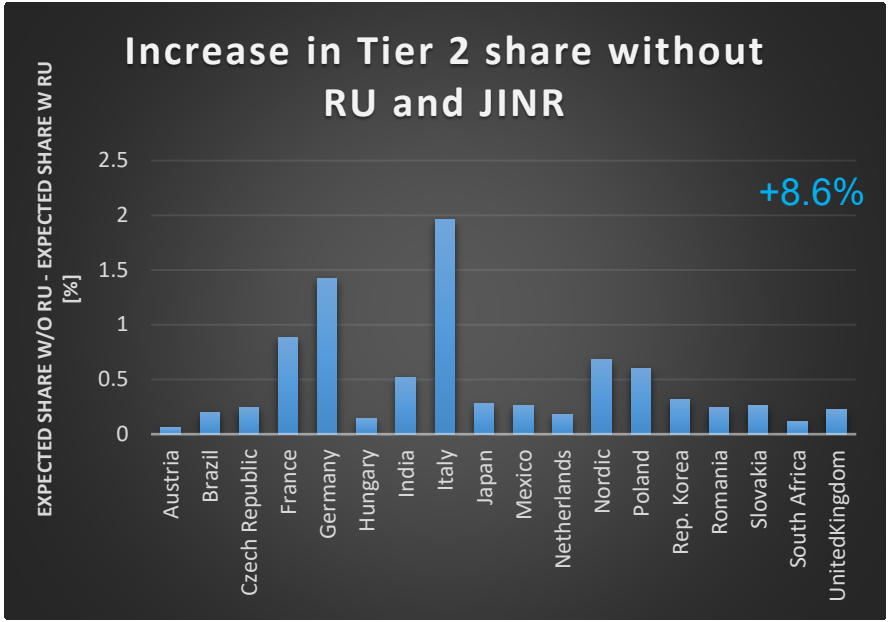

- Larger impact in countries that provide Tier 1, expected contribution proportional to the share of M&O-A members
- Tier 2 share of 8.6% to be divided among all the countries
- Priority for 2023: tape, disk, CPU

Pledged resources for 2023

Pledges for 2023

Pledged resources sufficient to support ALICE data taking in 2023

- Tape inline with our request at Tier-0 and surplus at Tier-1 (+5.7 PB) which perfectly compensates the tape pledged by RU
  - Commitment of all FA's to support data taking
  - 12 PB KISTI Tapeless Archive Storage

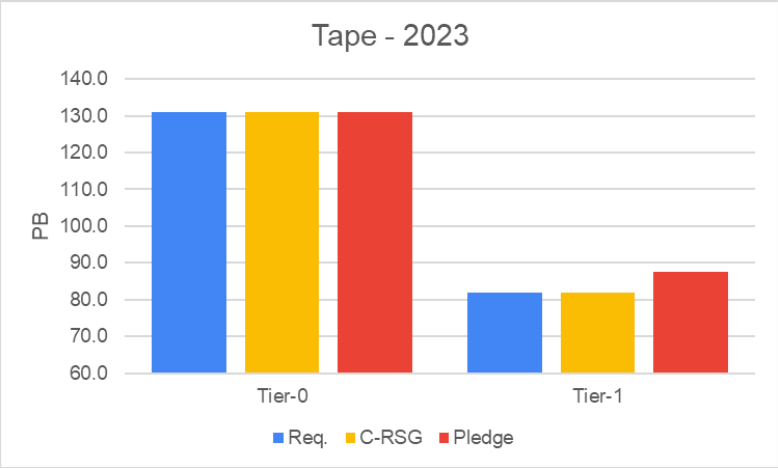

Pledges for 2023

Pledged resources sufficient to support ALICE data taking in 2023

- Tape inline with our request at Tier-0 and surplus at Tier-1 (+5.7 PB) which perfectly compensates the tape pledged by RU
  - Commitment of all FA's to support data taking
  - 12 PB KISTI Tapeless Archive Storage

But they are not enough to support ALICE data processing in 2023

- 3.7 PB deficit in the pledged disk
  - Tier 0 in line with our request
  - Tier 1 under pledged disk (-5.7 PB)
  - Tier 2 surplus of 2.2 PB

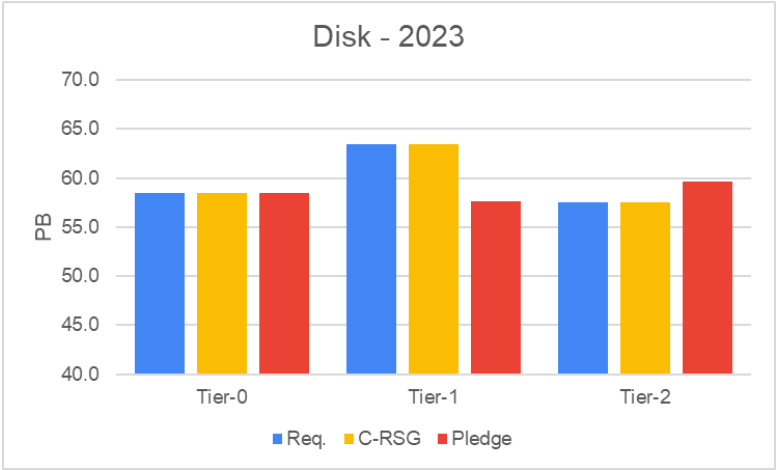

Pledges for 2023

Pledged resources sufficient to support ALICE data taking in 2023

- Tape inline with our request at Tier-0 and surplus at Tier-1 (+5.7 PB) which perfectly compensates the tape pledged by RU
  - Commitment of all FA's to support data taking
  - 12 PB KISTI Tapeless Archive Storage

But they are not enough to support ALICE data processing in 2023

- 3.7 PB deficit in the pledged disk
  - Tier 0 in line with our request
  - Tier 1 under pledged disk (-5.9 PB)
  - Tier 2 surplus of 2.2 PB
- Slight deficit of pledged CPU (-5%)
  - Mainly due to Tier-1 under pledges (-66 kHS06)

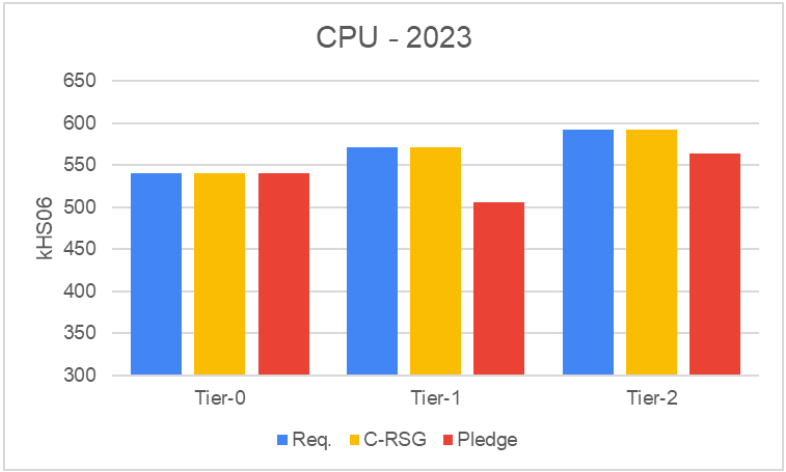

Pledges trend

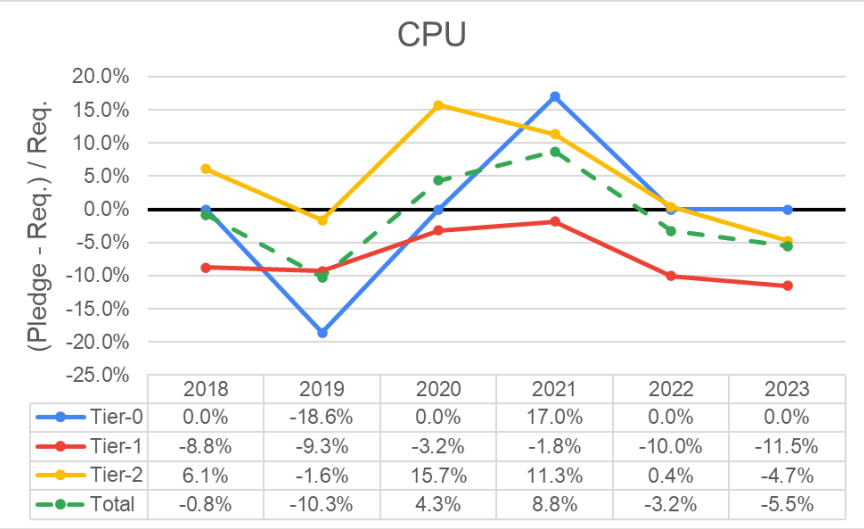

	2018	2019	2020	2021	2022	2023
Tier-0	0.0%	-18.6%	0.0%	17.0%	0.0%	0.0%
Tier-1	-8.8%	-9.3%	-3.2%	-1.8%	-10.0%	-11.5%
Tier-2	6.1%	-1.6%	15.7%	11.3%	0.4%	-4.7%
Total	-0.8%	-10.3%	4.3%	8.8%	-3.2%	-5.5%

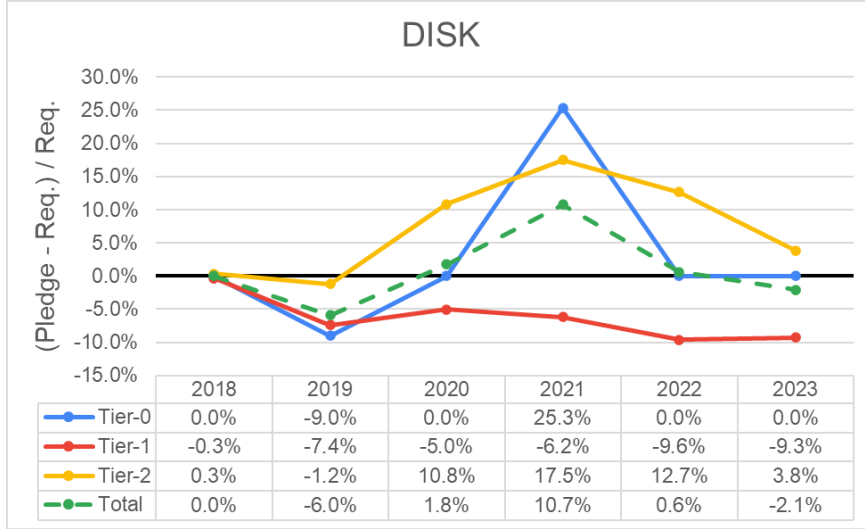

	2018	2019	2020	2021	2022	2023
Tier-0	0.0%	-9.0%	0.0%	25.3%	0.0%	0.0%
Tier-1	-0.3%	-7.4%	-5.0%	-6.2%	-9.6%	-9.3%
Tier-2	0.3%	-1.2%	10.8%	17.5%	12.7%	3.8%
Total	0.0%	-6.0%	1.8%	10.7%	0.6%	-2.1%

- Difficult to bridge the CPU and disk resources gap at T1s:
  - Some FAs prefer to balance their contribution between T1s and T2s
  - FAs with fixed budget prioritize tape for data taking
  - Larger budget share needs to be reserved for the increasing electricity costs

Estimates for 2024

Baseline scenario for 2024

2023 Running Conditions for Computing estimates including contingency

- ATLAS/CMS luminosity: <100/fb
- ATLAS/CMS average pile-up: <50 (peak 52)
- LHCb luminosity: <15/fb
- ALICE luminosity (pp): <100/pb
- Running time pp: 6×10^6 seconds
- Running time ions (PbPb): 1.2×10^6 seconds

Same conditions can be assumed for now in 2024 and 2025 for long term projections, except there will be pPb instead of PbPb run in 2024

- **Assumed for 2024 the same running condition of 2023 with p-Pb collisions during the HI period and short oxygen-beam run**

- **Considered as upperlimit:**
  - **140 days of pp in 2024:**
    - pp collisions at 500kHz
  - **Short O-O and p-O run:**
    - 1 nb^{-1} and 5 nb^{-1}, respectively
  - **28 days of HI in 2024:**
    - **28 HI days dedicated to p-Pb collisions (strategy B aggressive)**

pp period in 2024

- 500 kHz IR:
  - Assuming 50%x 90% efficiency:
  - 35 pb^{-1}
  - 2.72 10^{12} events
  - 1.2% selection:
    - 0.7 PB CTF
    - 0.25 PB AOD
    - Needed ~ 510 kHS06 (~200 CPU + GPU EPNs)
- 1 MHz:
  - x2 processing capacity to perform online selections wrt 500kHz
  - Crucial to assess the impact of GPU and number of available EPNs for selection (also for 2022)

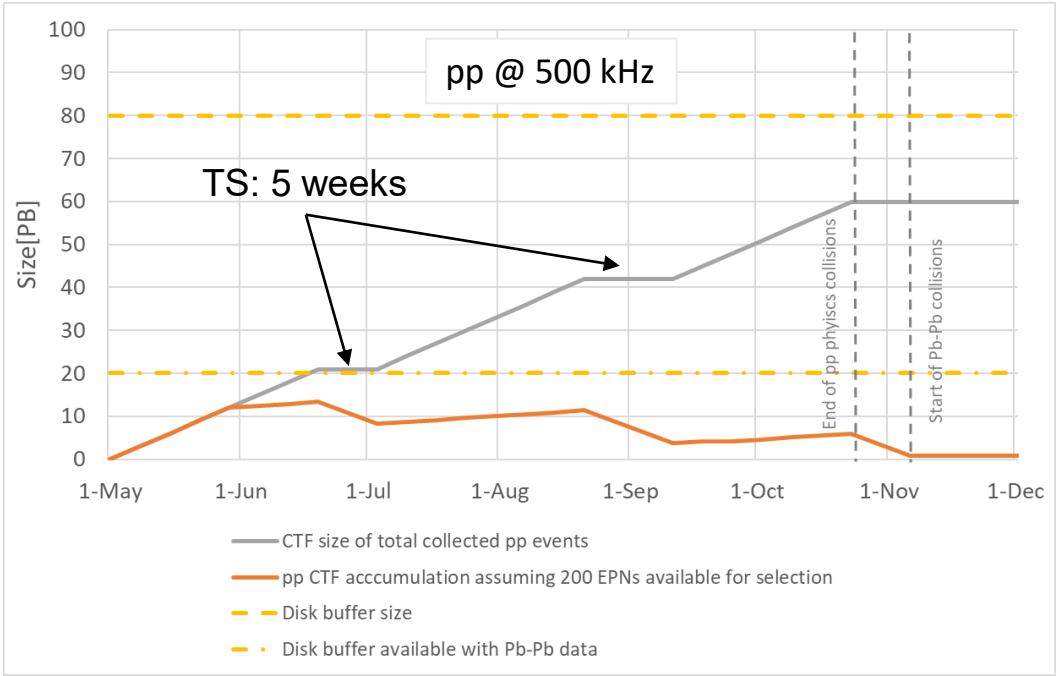

p-Pb period in 2024

Assumed IR leveled at 1 MHz
(1MHz / 2.12 barn = 4.7 10^{29} cm^{-2} s^{-1}):

- 310 nb^{-1} ~545 kHz average IR
- 6.57 10^{11} collected events
- 95% efficiency: 6.24 10^{11} events (x27 Pb-Pb)

Estimates from Run 2 data (multiplicity scaling and AOD conversion):

- CTF:
  - 76.1 kB/event
  - Totaling 47.5 PB (same as Pb-Pb)
- AOD:
  - 17 kB/event
  - 10.6 PB per pass per copy (x2 Pb-Pb)
- MC:
  - 2% of total events
  - 1500 kHS06 per year

Filling scheme	$\mathcal{L}_{tot}$ IP1/5	$\mathcal{L}_{tot}$ IP2	$\mathcal{L}_{tot}$ IP8
1240b_1240_1200_0	677 [687]	306 [313]	0 [0]
1240b_1144_1144_239	634 [629]	309 [316]	45 [53]
1240b_1088_1088_398	605 [596]	308 [317]	73 [87]
1240b_1032_1032_557	583 [563]	311 [319]	103 [121]
1240b_976_976_716	558 [531]	312 [312]	135 [154]
733b_733_702_468	415 [416]	287 [294]	86 [89]

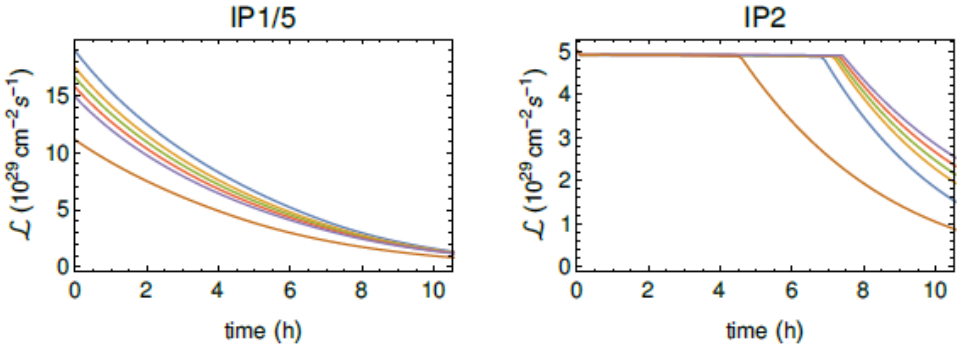

Computing resource estimates for 2024

ALICE		2022		2023				2024		
		C-RSG	Pledge	C-RSG	RU (incl. JINR)	C-RSG 2023 / C-RSG 2022	C-RSG 2023 / Pledge 2022	Est.	Est. 2024 / C-RSG 2023	Est. 2024 / (C-RSG - RU) 2023
CPU [kHS06]	Tier-0	471	471	541	0	115%	115%	622	115%	115%
	Tier-1	498	448	572	83	115%	128%	655	115%	134%
	Tier-2	515	517	592	50	115%	115%	683	115%	126%
	Total	**1484**	**1436**	**1705**	**133**	**115%**	**119%**	**1960**	**115%**	**125%**
Disk [PB]	Tier-0	50.0	50.0	58.5	0.0	117%	117%	67.5	115%	115%
	Tier-1	55.0	49.7	63.5	9.2	115%	128%	71.5	113%	132%
	Tier-2	49.0	55.2	57.5	4.9	117%	104%	66.5	116%	126%
	Total	**154.0**	**154.9**	**179.5**	**14.1**	**117%**	**116%**	**205.5**	**114%**	**124%**
Tape [PB]	Tier-0	95.0	95.0	131.0	0.0	138%	138%	167.0	127%	127%
	Tier-1	63.0	71.8	82.0	11.9	130%	114%	102.0	124%	145%
	Total	**158.0**	**166.8**	**213.0**	**11.9**	**135%**	**128%**	**269.0**	**126%**	**134%**

- First estimates for 2024 discussed with C-RSG in preparation of RRB
- Standard growth for CPU (+15%) and disk (+14%) in 2024 compatible with flat budget, step for tape (+56 PB)
- Without the expected RU+JINR contribution, the estimated growth exceeds the flat budget
- Resources of RU+JINR contribution will be missing unless they are not compensated by the other FAs.

Expected growth of CPU and disk space in 2022-2024

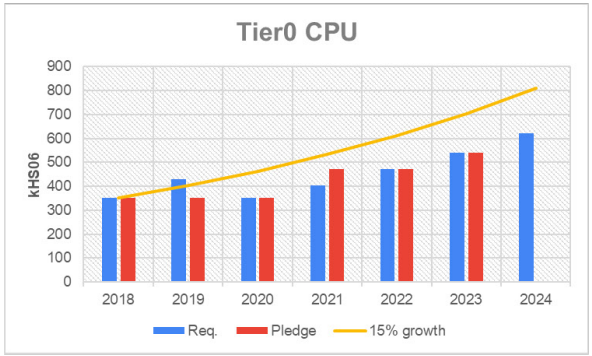

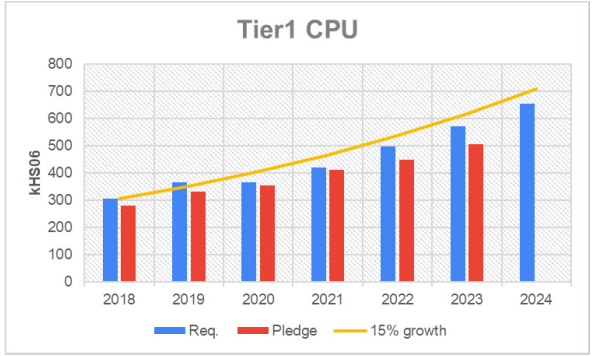

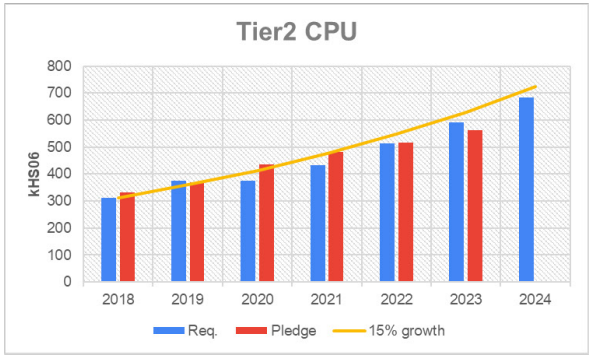

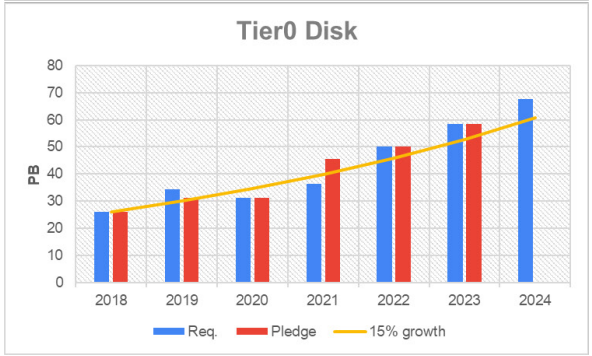

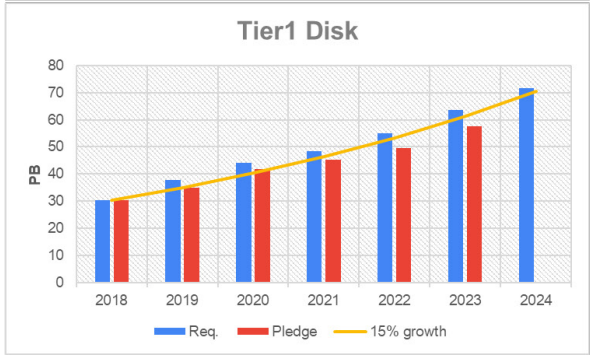

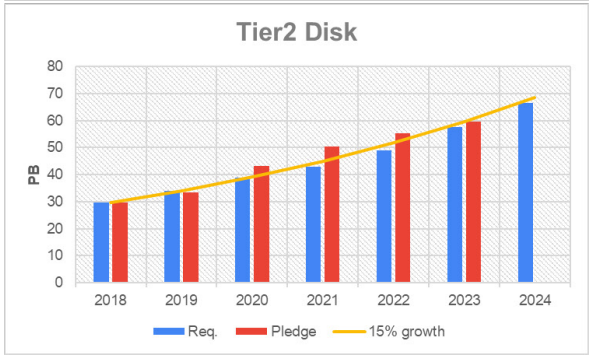

------ 15% growth: annually compounded rate from 2018 along LS2

Expected growth of TAPE space in 2022-2024

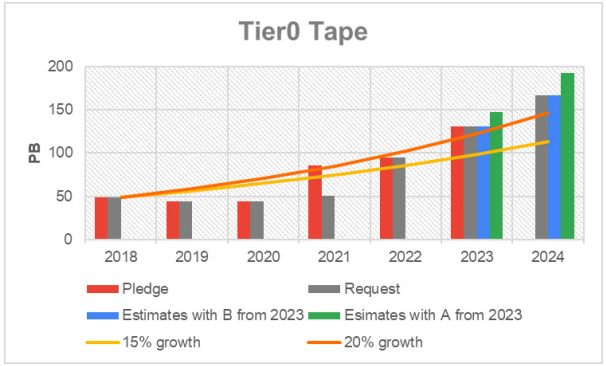

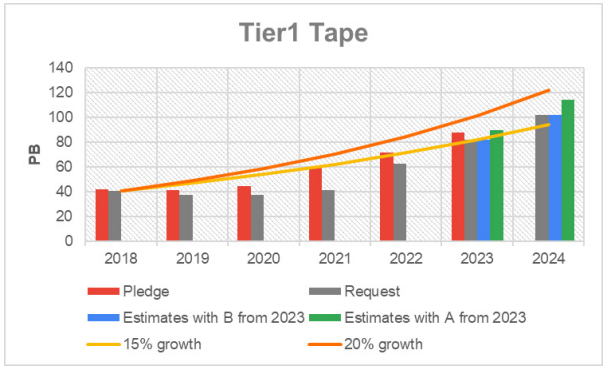

- TAPE in 2022 (strategy A):
  - Requested 158 PB
  - Pledged 167 PB
- 79 PB needed in 2022:
  - 22.0 PB for commissioning + pp
  - 57.0 PB for Pb-Pb
- Flexibility of FA's to delay the deployment:
  - commissioning in February 2022
  - Pb-Pb in October 2022
- TAPE in 2023:
  - + 44 (+66) PB for Pb-Pb with Strategy B (A)
  - + 11 PB for low and high field pp
- TAPE in 2024:
  - + 48 PB for p-Pb with Strategy B
  - + 2 PB for pp and +6 PB for Run 2 archival campaign
- Crucial to adopt strategy B

15% growth: annually compounded rate from 2018 along LS2

LHC energy consumption reduction

- CERN is drawing up a plan to reduce its energy consumption this winter:
  - YETS will start 2 weeks in advance (reduced 2022 LHC run) on Monday, November 28th
  - Expected an overall reduction of Run 3 running time ~20% both for pp and for HI
  - Updated Run 3 LHC schedule not yet available but no HI in 2022
  - Number of days dedicated to HI per year strongly impact on ALICE computing needs
  - Expect a reduction of our needs in 2022
- Impact of no HI in 2022:
  - We will collect ~7 weeks of pp that will requires more resources than planned for pp
  - But ... Less CPU resources for async processing (T0,T1s) and MC (T0,T1s,T2s)
  - And less AODs to be saved to disk and CTF to be archived on tape
  - The computing resource surplus in 2022 will be carried in 2023 (and 2024)
- New estimates for 2023 and 2024 computing needs as soon as the LHC schedule will be available

Impact of no HI in 2022 in ALICE computing needs

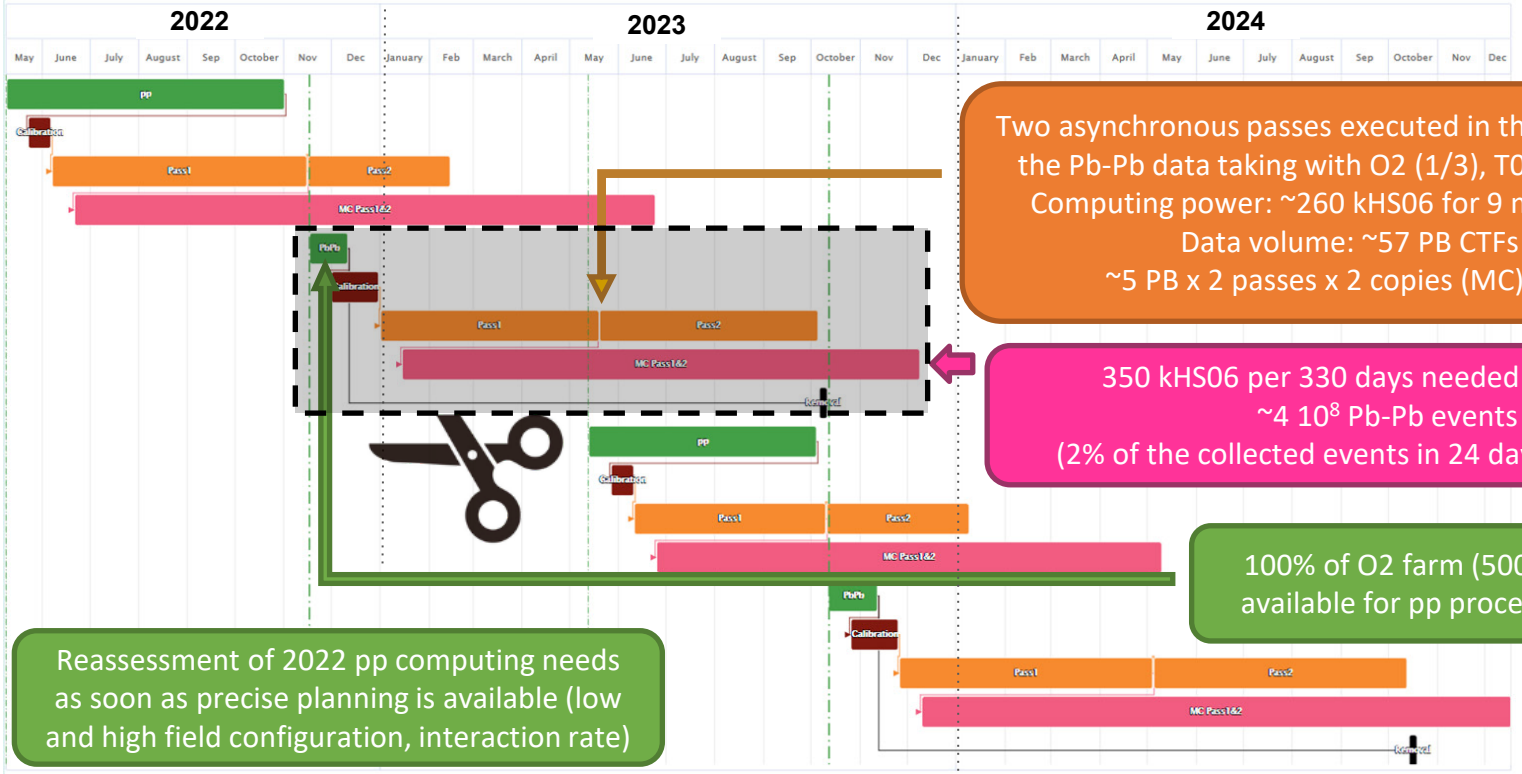

Two asynchronous passes executed in the 10 months following the Pb-Pb data taking with O2 (1/3), T0 (1/3) and T1s (1/3) :
Computing power: ~260 kHS06 for 9 months @ T0 and T1
Data volume: ~57 PB CTFs on Tape
~5 PB x 2 passes x 2 copies (MC) AODs on Disk

350 kHS06 per 330 days needed to simulate
~4 10^8 Pb-Pb events
(2% of the collected events in 24 days of HI period)

100% of O2 farm (500 CPUs + 2000 GPUs) available for pp processing until May 2023

Reassessment of 2022 pp computing needs as soon as precise planning is available (low and high field configuration, interaction rate)

Summary

- **Computing resource utilization:**
  - Full utilization of CPU resources
  - Disk and tape expected usage in line with the requested resources
  - O2 disk buffer: full capacity installed, and the performances exceed the expectations
  - Archival and deletion of commissioning data ongoing
- **Run 3 data and simulation processing :**
  - Good progress of asynchronous reconstruction of pp commissioning data
  - Precise timeline for the reconstruction and the selection of pp data for physics
  - Completed Run 3 pp MC productions for the selection studies
  - Ongoing MC productions anchored to Run 3 pp data and in preparation for Pb-Pb
- **Resource requests for 2023 and estimates for 2024:**
  - CPU and disk in 2023 and 2024 compatible with flat budget considering RU contribution
  - Without the expected RU+JINR contribution, the estimated growth exceeds the flat budget
  - Unavoidable step increase of tape but computing centres can deploy capacity gradually
  - CERN is considering different scenarios for energy reduction
  - Strong impact on ALICE computing resource needs for Run 3

Thank you for providing reliable resources to ALICE collaboration!