



ALICE computing resources

Processing plans and requirements

Stefano Piano, Irakli Chakaberia






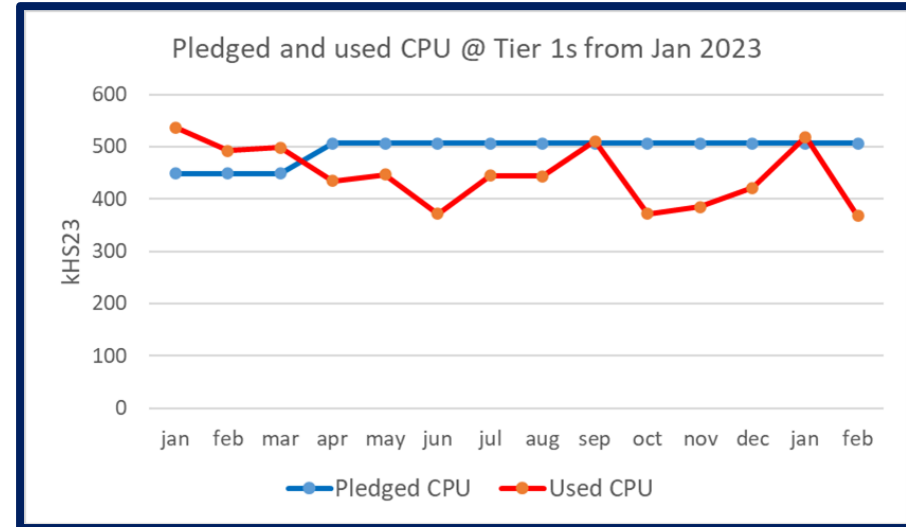
Resource usage

April 2023 – March 2024

CPU utilization





- Full utilization of the resources at T0, 95% T2s
- Used/Pledged 85% and Pledge / C-RSG 90% @ T1s
- Opportunistic CPU usage at the T0 and LBL_HPCS, Japan, Wigner and EPN (240 kHS23 only CPU, with 2.5 GPU speedup factor => 360 kHS23)



- Lack of reconstruction and Monte Carlo jobs at T1s:
 - 2022 Pb-Pb data taking postponed to 2023
 - Reconstruction of pp data exclusively at T0 and EPN (skimming and memory constraint)
- Energy-saving measures at some sites during the year



CPU breakdown by job types




- High activity for raw calibration and reconstruction of Run 3 runs
- Steady analysis activity for conferences and publications both on Run 2 and Run 3 data
- Lower MC share, but MC productions are ramping up at full scale


- No significant MC productions in October and November
- Substantial improvement mid-December to mid-January
- MC productions anchored to Run 3 data are gradually increasing pressure on the queue
- Crucial for overall efficiency (@T1 50->70%) and as backfill

T0 and T1 Efficiency

Tier-0: CPU efficiency




All Tier-1s: CPU efficiency




- T0 efficiency is slightly higher than the T1 due to the ongoing reconstruction
- Low MC job load highlights the poor efficiency of the analysis jobs resulting in an overall lower efficiency
- The analysis job efficiency improvement is crucial for better use of the CPU resources
- Run 2 analysis jobs are expected to gradually disappear, larger share of Run 3 analysis jobs should increase the efficiency
- The efficiency should be considered along with the CPU-time and dataset-size parameters before the job submission
- Our goal is to reach high 80% at all Tiers, **ongoing efforts to improve efficiency in all Run 3 workloads**

DISK and TAPE utilization




- The deployment of the pledged disk for 2024: 100% at T0, 102% at T1s (88% of C-RSG), 110% at T2s
- Current disk use 92% of pledged capacity at T0, 86% at T1s and 74% at T2s
- RU+JINR pledged 4.5 PB at T1 and 3.0 PB at T2s for 2023, while for 2024 only 3.0 PB at T2s



- Pledged tape aligns with our T0 request, while 4.6 PB deficit at T1, where RU T1 has not pledged
- At T1 49 PB available, at T0 36 PB but 40 PB used to temporarily park 2023 pp at T0. The plan is to remove them.
- Archiving of 2023 Pb-Pb is completed at T0, T1s

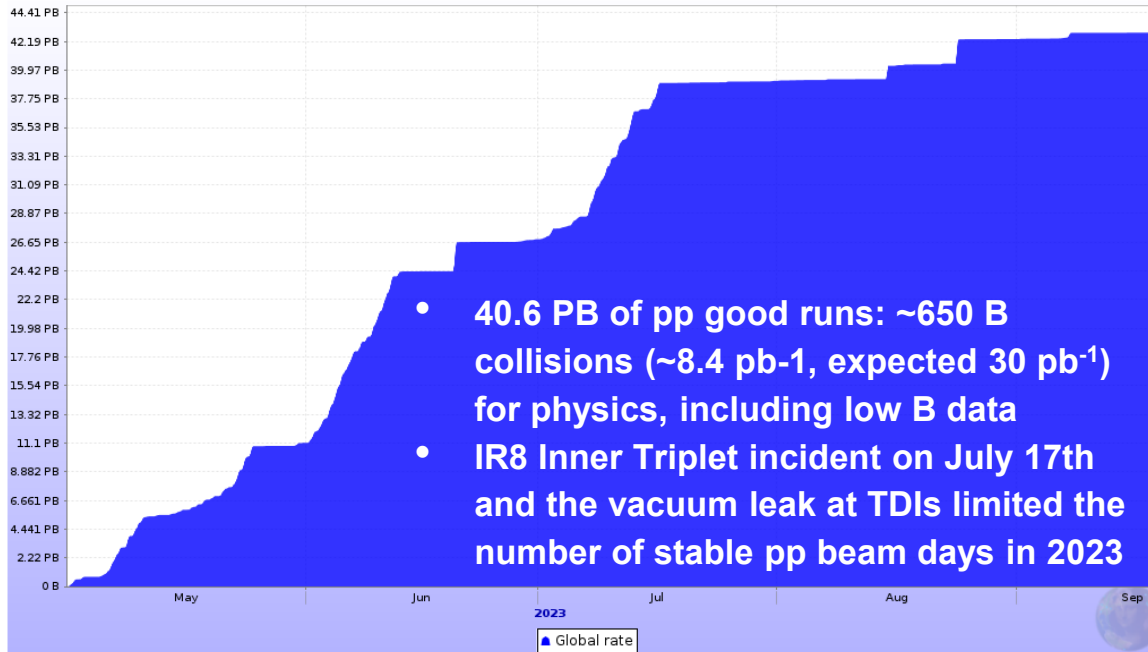


Status of data taking, calibrations and data processing



2023 pp data accumulation and processing

Global experiment data accumulation




Plan to skim the entire 2023 pp data sample at a 3% level:

- cpass (10%) and apass1 (100%) completed before HI but affected by calibration issue
- Three additional passes were required to achieve calibration quality sufficient for skimming
- Skimming is currently in progress and is expected to be completed by the end of the month


⇒ Once the skimming is validated, removing 2023 pp from disk buffer and T0 tape

2023 Pb-Pb data taking @ 5.36 TeV

Global experiment data accumulation



Impressive capability of the apparatus to collect data: recorded **4 PB in 24 hours (3 fills of 5h each + 3h of turnaround time)**




Average data collection rate in the last 10 days ~2.5 PB/day

Latest productions from HI period

- Asynchronous reconstruction pass of **2022 skimmed pp data**, including 1 pb^{-1} MinBias
- 4 asynchronous passes of **2023 pp data** in preparation for skimming
- Asynchronous passes of **2023 Pb-Pb** over the “golden” sample
 - 2 cycles ran, to allow reconstruction fixes and improvements
 - Extraction of ITS-TPC residuals for TPC distortion corrections for the sample affected by A11 IROC issue
 - Started the third pass for physics results

Raw production from 2023-10-01 to 2024-04-10

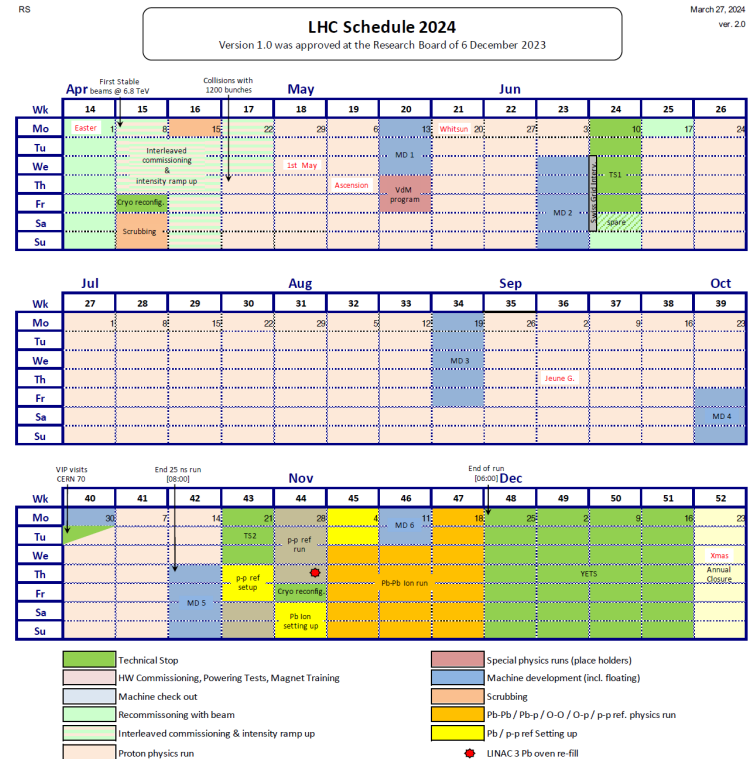


Processing plan and expected resource usage for 2024




2024 and 2025 LHC Schedule

- Overall 2024/2025 schedule was revised in December 2023, considering the shorter pp period in 2023
- Heavy Ions allocation for 2024/2025 runs:
 - 17 (16) days of Pb-Pb in 2024 (2025)
 - 6 (2) days of pp ref run just before PbPb in 2024 (2025)
 - p-O (3 days) and O-O (1 days) scheduled in 2025
- Recently, the CERN directorate has extended the 2024 pp run by 4 weeks and shortened the 2025 run by the same duration, maintaining overall beam allocation:
 - 171 (136) days of pp @ 650 kHz INEL including ramp up in 2024 (2025)
 - Totaling $\sim 97 \text{ pb}^{-1}$ of pp full field @ 650 kHz IR to be skimmed at 3% level
 - **3 pb^{-1} of pp low field in 2025**



Pledges and expected resource usage in 2024




- Pledged computing resources without Russia and JINR contribution for 2024 are sufficient for processing the collected data in 2023 and for collecting the anticipated data in 2024
- Lower amount of HI data collected in 2023 ($\sim 1.96 \text{ nb}^{-1}$) and anticipated in 2024 (exp. max 1.9 nb^{-1}) compared to the earlier projections (3.25 nb^{-1} each year)

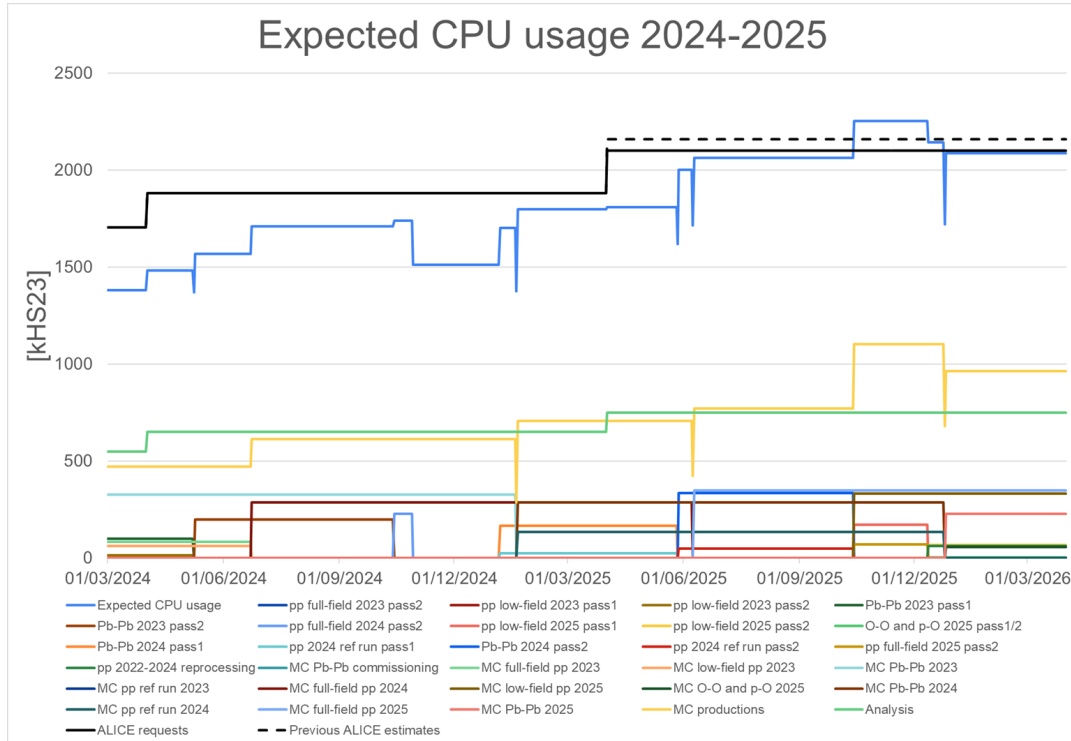
Russian and JINR computing resources

- **Impact on computing resources of the council decision not to renew the ICA after expiration on 30 Nov 2024 with Russia and to terminate all MoUs with Russian Institutes at the same date:**
 - The Russian Tier-1 (except for JINR) and Tier-2 sites will stop being part of WLCG
 - Russia T1 did not pledge resources for 2024
 - The 13% Tier 1 RU share is to be distributed among the 7 other Tier 1 sites, while the 7% Tier 2 RU share is to be distributed among all 23 countries
 - A site not part of WLCG can still be used opportunistically by the experiments
 - However, doing so the spirit of CERN Council's resolution would be contradicted even if there is no official recommendation from CERN
 - WLCG MB recommended not to use Russian sites even opportunistically after 30 Nov 2024
- **No decision yet on JINR (ICA expires 27 January 2025):**
 - We can continue to use their resources until the council deliberates
 - These resources amount to ~1.5 PB of disk and ~15 kHS23 of CPU
 - JINR T2 pledges are included in the RU pledges (total: 3 PB disk and 35 kHS23 CPU)

Computing resource requirements for 2025



CPU needs for 2024 - 2025



- The blue line is the minimum CPU capacity needed to process all the planned productions on GRID, while the solid (dashed) black line shows the ALICE requests (previous estimates)
- The requests for 2025 revised downward due to the lower amount of data expected to be collected compared to the earlier projections

Disk and tape needs for 2025

- **Disk:** Pb-Pb AO2D and MC AO2D data, the average event size turns out to be 380 kB/event and 2260 kB/event, instead of 220 kB/event and 1500 kB/event from the first asynch pass, pp unchanged
- **Tape:** considered the compression strategy A (lossless) in 2024 and 2025:
 - Not yet optimal calibrations still do not allow the adoption of more aggressive compression
- The increase in disk and tape due to much larger event sizes that initially estimation is a point of concern raised during the meeting with C-RSG

		2025									
		pp 2024	pp ref 2024	Pb-Pb 2024	pp 2025	pp ref 2025	Pb-Pb 2025	O-O and p-O 2025	pp low field 2025	Needed in 2025	Req 2025
ALICE	Tier-0	0.1	1.6	6.9	1.6	0.5	2.6	0.3	1.6	15.2	78.0
	Tier-1	0.2	2.0	8.8	2.5	0.6	3.2	0.3	1.8	19.4	79.0
	Tier-2	0.2	1.9	8.1	2.2	0.6	3.0	0.3	1.7	18.0	77.0
	Total	0.5	5.5	23.8	6.3	1.7	8.8	0.9	5.1	52.6	234.0
Disk [PB]	Tier-0	0.2	0.0	0.0	8.1	3.9	33.3	0.7	9.8	56.0	220.0
	Tier-1	0.1	0.0	0.0	7.0	2.0	16.7	0.4	4.9	31.1	123.0
	Total	0.3	0.0	0.0	15.1	5.9	50.0	1.1	14.7	87.1	343.0
Tape [PB]	Tier-0	0.2	0.0	0.0	8.1	3.9	33.3	0.7	9.8	56.0	220.0
	Tier-1	0.1	0.0	0.0	7.0	2.0	16.7	0.4	4.9	31.1	123.0
	Total	0.3	0.0	0.0	15.1	5.9	50.0	1.1	14.7	87.1	343.0

To be processed in 2025 

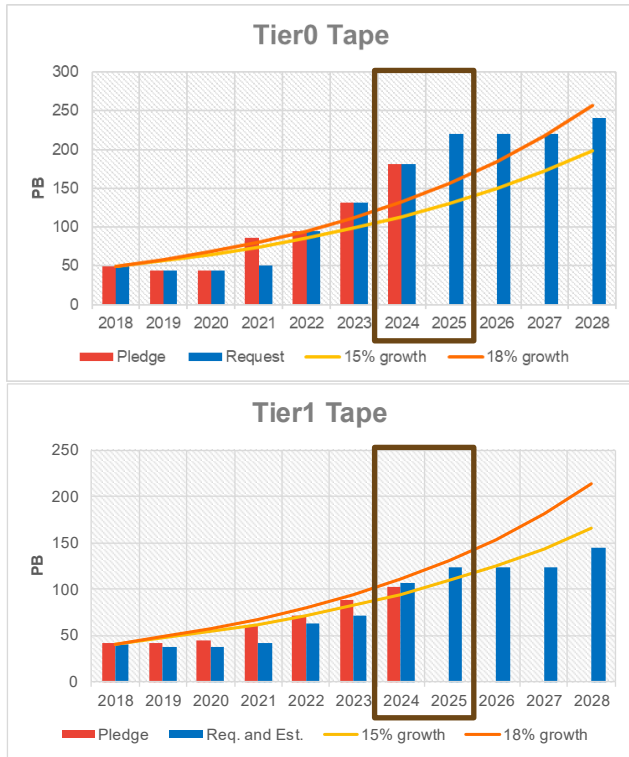
Computing resource requests for 2025

Resource requests for 2025 are being discussed with C-RSG in view of April RRB

ALICE		2024					2025				
		Req.	C-RSG	Pledge	RU + JINR pledge	Pledge - (RU+JINR)	Prev. Est.	Req.	Req. 2025 / C-RSG 2024	Req. 2025 / (Pledges - RU) 2024	
CPU [kHS23]	Tier-0	600	600	600		600	690	680	113%	113%	
	Tier-1	630	630	540	0	540	725	690	110%	128%	
	Tier-2	650	650	641	33	608	750	730	112%	120%	
	Total	1880	1880	1782	33	1748	2165	2100	112%	120%	
Disk [PB]	Tier-0	67.5	67.5	67.5		67.5	78.5	78.0	116%	116%	
	Tier-1	71.5	71.5	61.9	0.0	61.9	82.5	79.0	110%	128%	
	Tier-2	66.5	66.5	69.8	3.2	66.5	77.5	77.0	116%	116%	
	Total	205.5	205.5	199.2	3.2	195.9	238.5	234.0	114%	119%	
Tape [PB]	Tier-0	181.0	181.0	181.0		181.0	226.0	220.0	122%	122%	
	Tier-1	107.0	107.0	102.4	0.0	102.4	135.0	123.0	115%	120%	
	Total	288.0	288.0	283	0	283.4	361.0	343.0	119%	121%	

- Standard growth for CPU (+10%) and disk (+14%) in 2024 compatible with flat budget
 - Wrt the pledges, the estimated growth exceeds the flat budget especially for the T1 sites, where the RU T1 contribution still needs to be compensated
- 20% step for tape:
 - Considered compression strategy A
 - If compression strategy B is adopted, the surplus tape will be utilized in 2026 for archiving Run 3 AO2Ds

ALICE tape needs compared to 15% annual increase



- Planned to store only one instance of HI raw data file (CTF) on tape to achieve a sustainable growth
- Two copies are required for pp: the risk of archiving only one copy of skimmed data is too high (x1000 MB/lumi).
- Adopting the more aggressive compression strategy B in 2024 or 2025 could result in a 33% improvement in compression
- The computing model plans to archive exclusively the CTF files and one copy of AO2Ds for long-term preservation, which can be postponed until 2028. This mitigates the required growth during Run 3. **Considering the LS3 period: <18% at T0 and <15% at T1**

CPU and disk resources needed during LS3

- A clearer picture of the computing estimates needed during LS will be available for the October RRB
- 26 PB of disk space is required solely to complete the processing of data collected in 2025
- Additionally, during LS, we plan to recall and reprocess all data collected in Run 3, along with related simulations → CPU and disk space requirements will heavily depend on the collected statistics
- The growth of CPU resources can be mitigated if the EPN farm (CPU+GPU) remains available during LS for asynchronous reconstruction (590 kHS23 + 2240 MI50 + 560 MI100)
- To be clarified with Tier 0 how much growth they will provide during the LS

Summary

- Ongoing
 - High activity for raw calibration and reconstruction of Run 3 runs
 - Run 3 simulations and analyses are ramping up to full scale
- Expectations for 2024
 - Enough resources are pledged to cover the missing contributions from the Russian institutions
- Plans for 2025
 - Resource needs have been revised based on the actual 2023 data amount and the new LHC running plan
 - Missing resources from the Russian institutions will need to be provided by the other funding agencies
- (MC)AO2D event size reduction, compression strategy B implementation, and job efficiency improvement are the critical aspects for the ALICE computing resources

2023 data taking

- pp data taking 13.6 TeV
 - Production for physics at 650 kHz INEL interaction rate and at different IRs
 - IR8 Inner Triplet incident on July 17th + two vacuum leaks on the bellows of two TDIS modules
 - ~650 B collisions (~8.4 pb⁻¹, expected 30 pb⁻¹) for physics, including low B data
 - Small sample at 900 GeV for recommissioning in September (~20 M collisions)
- Pb-Pb data taking 5.36 TeV:
 - During the initial fills, high background measured on ITS, not for physics (~6% of the Pb-Pb sample)
 - Data taking mostly at 25 kHz interaction rate with reduced number of bunches
 - To mitigate quenches and losses in the ramp
 - 2023: $L_{\text{int}} = 2.16 \text{ nb}^{-1}$ (1.96 nb⁻¹ delivered without beam background, expected 3.25 nb⁻¹)
 - In Run 1+2 1.5 nb⁻¹ was delivered
- pp reference run
 - Expected $L_{\text{int}} = 3 \text{ pb}^{-1}$ at 1 MHz INEL interaction rate, postponed to end of 2024
 - Only one fill with single bunches delivered: ~700 M collisions (~9 nb⁻¹) for physics