

A Large Ion Collider Experiment



Computing Resources for ALICE experiment

The 8th Asian Tier Center Forum
TIFR, India

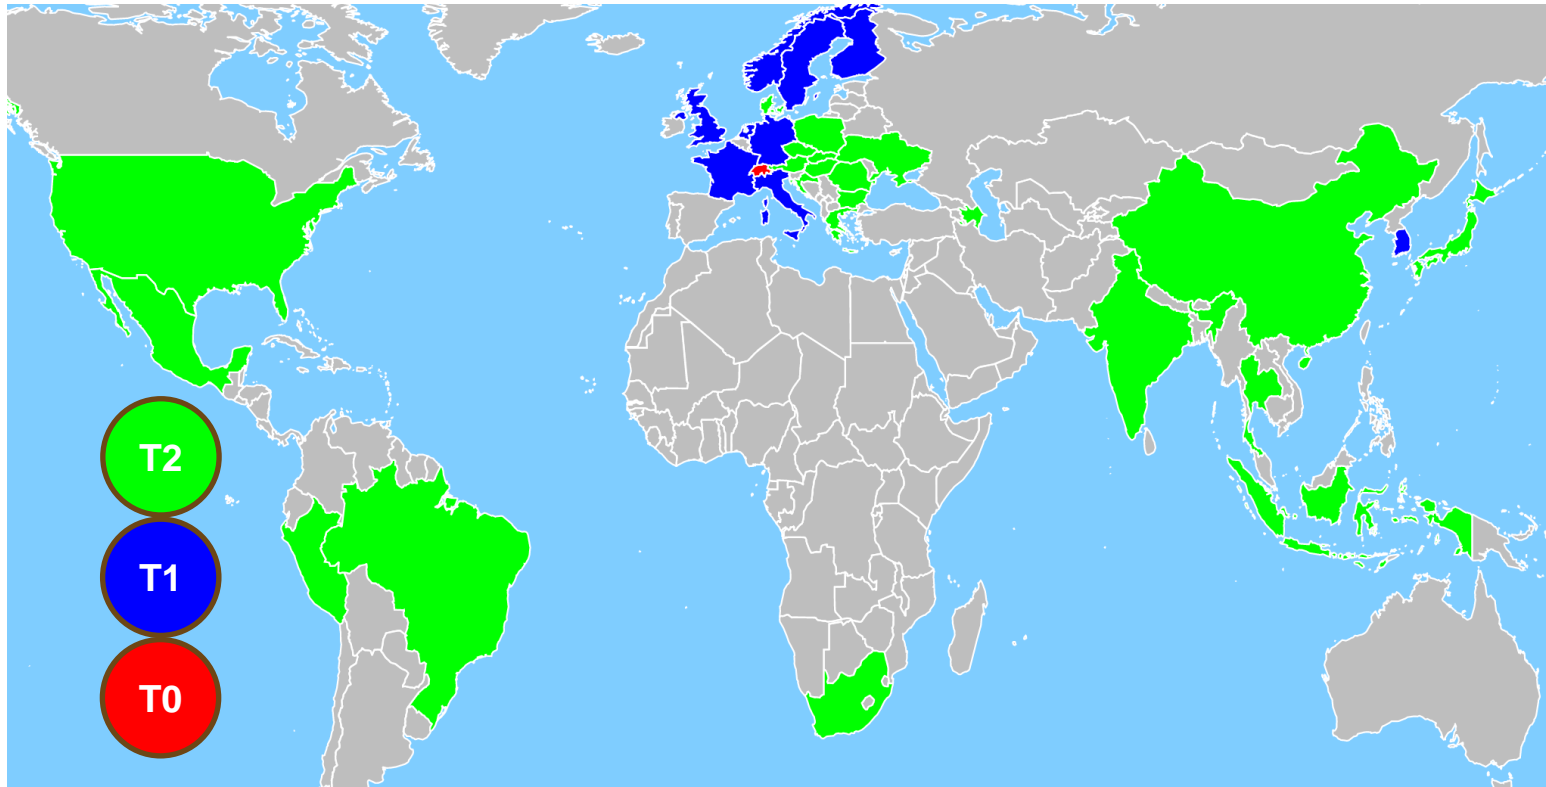
Irakli Chakaberia
Stefano Piano



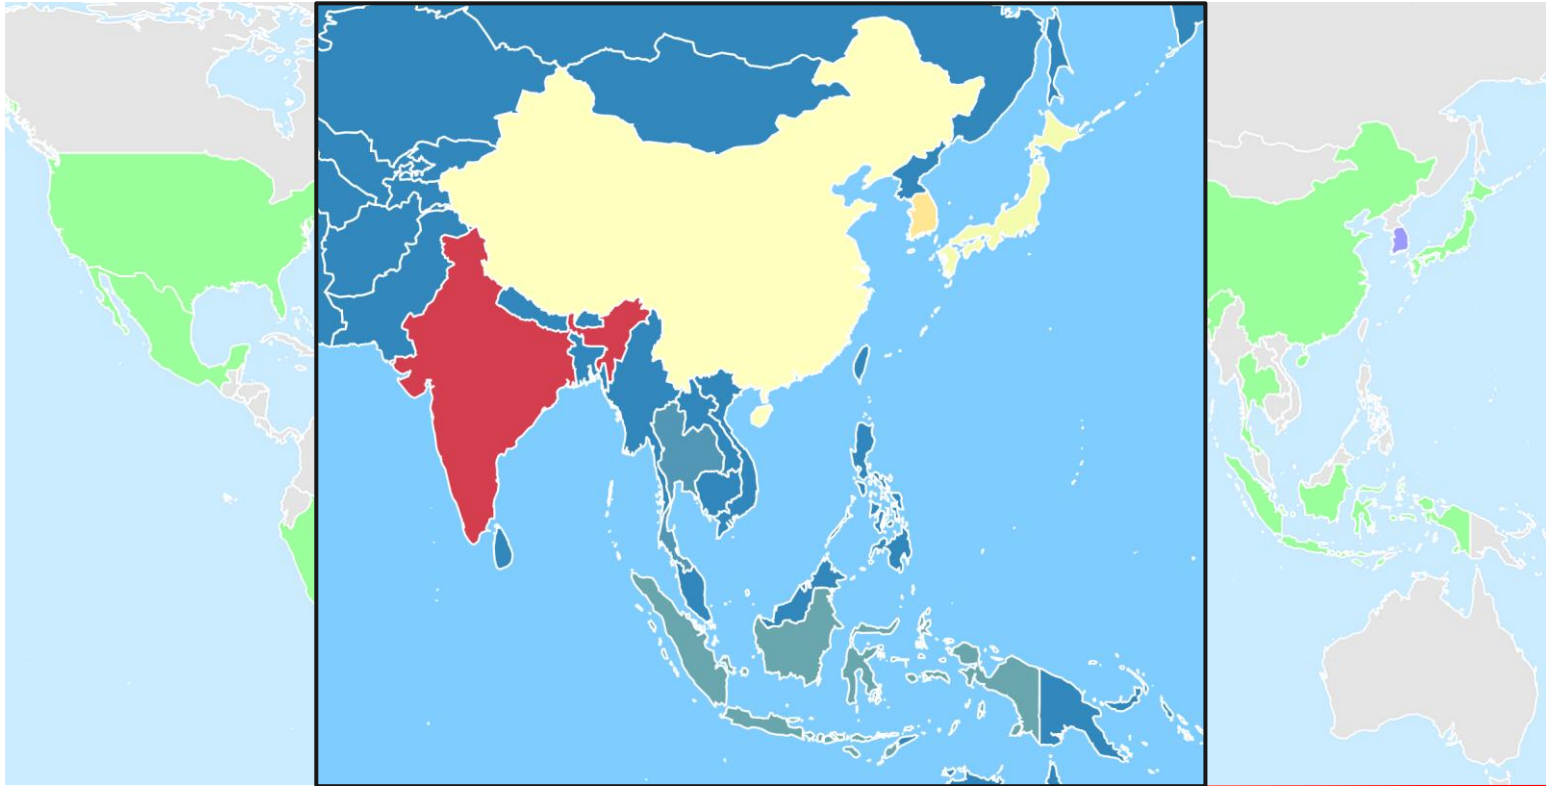


Setup

ALICE Computing Grid Map

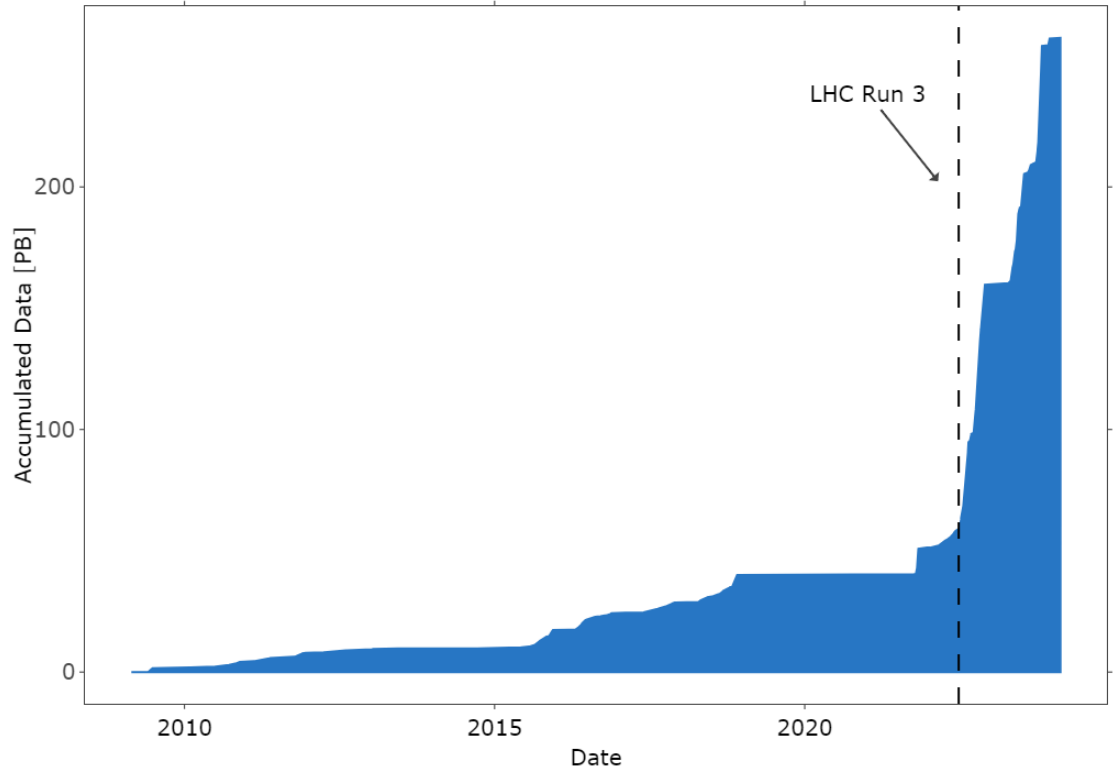


ALICE Representation Map – ATCF Section



ALICE in Run3

- ALICE in Run3 is essentially a new detector with a new readout and data processing chain
- Challenges related to computing resource, both storage and computing power, are substantial and require full participation for the success of the experiment




Computing resource requests for 2025


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%



Status of data taking, calibrations and data processing



2024 pp data taking and data accumulation

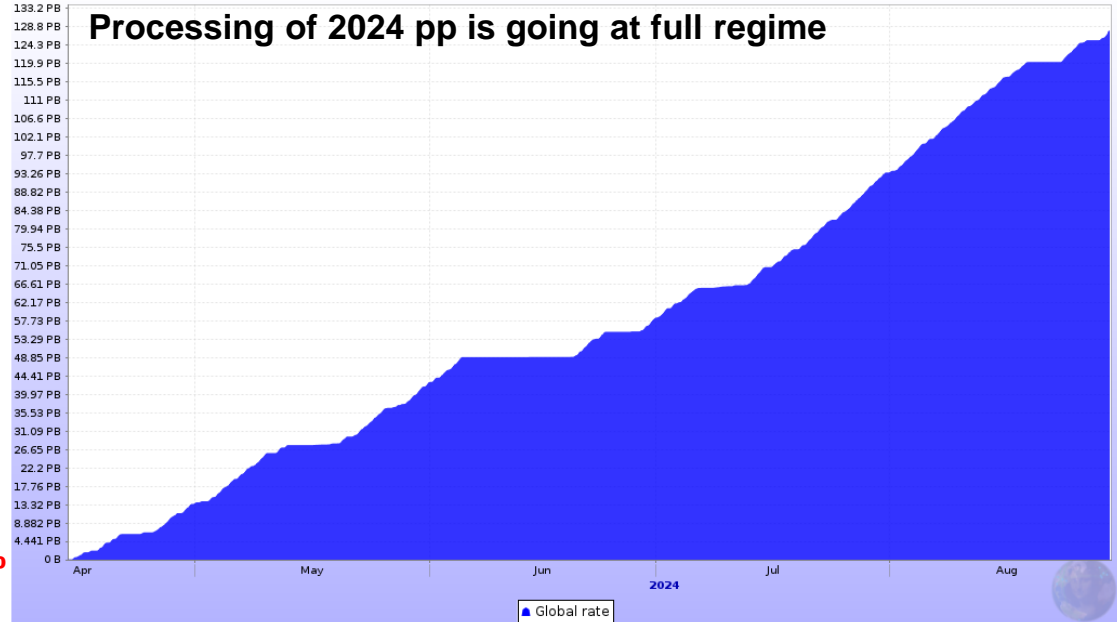


- About 36 pb⁻¹ collected so far
- r Efficiency systematically above 90%

2024 data accumulation:

- ~140 PB of pp runs (CTF) collected in 2024


Processing of 2024 pp is going at full regime






Resource usage status

DISK and TAPE utilization

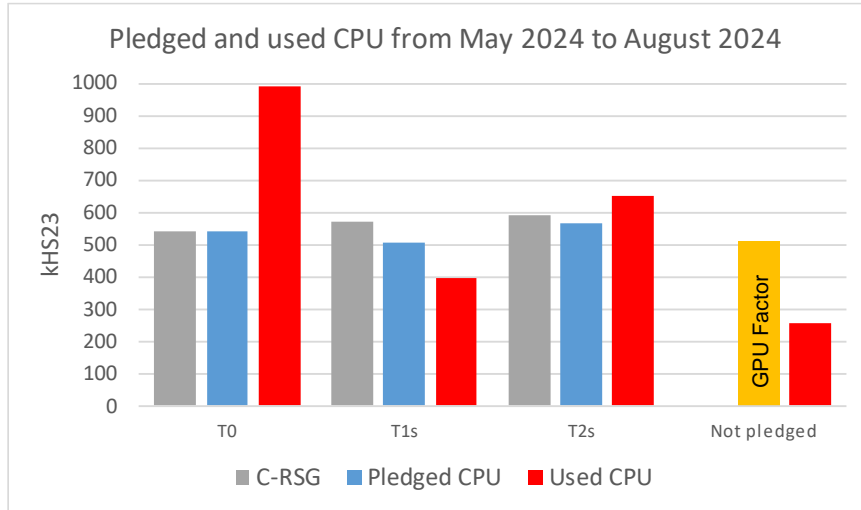


- Current disk use 93% of deployed capacity at T0, 91% at T1s and 88% at T2s
- ALICE operates very close to the 95% of storage capacity limit with continuous rebalancing and data removal activities




- 40 PB used to temporarily park 2023 pp at T0 removed. At T0 69 PB available, at T1 47 PB, **totaling 116 PB (requested +60 PB in 2025)**
- Tapes are sufficient for archiving the anticipated data in 2024 and 2025

CPU utilization and breakdown by job types



- Full utilization of the available resources at T2s with opportunistic CPU usage at the T0 and LBL, Japan, Wigner and EPN (250 kHS23 only CPU, with 2.5 GPU speedup factor => 510 kHS23)
- At T1s: Used/Pledged 80% and Pledge / C-RSG 86% in 2024



- Healthy share of MC, RECO and Analysis activity
- New data format and framework (hyperloop) based analysis activity share is steadily increasing



Computing resource estimates for 2026 and for LS3

Two scenarios for 2026 LHC run



No data taking

- Until October:
 - Finish the processing of the data collected in 2025 (Pb-Pb, pp ref run, full field pp, low field pp)
- From October to December:
 - Recall 2023 Pb-Pb from tape, it will only begin once the O2 disk buffer is cleared
- From January onwards:
 - 2023 Pb-Pb reprocessing (expected 70%)

Full year of data taking

- Until October:
 - Finish the processing of the data collected in 2025 (Pb-Pb, pp ref run, full field pp, low field pp)
 - 2026 pp data acquisition and the related processing for skimming
- From November onwards:
 - 2026 HI calibration and processing, **only 30% of the first pass**, the remainder processed during LS3
 - 2026 pp pass2 and reprocessing of the previous year's pp data

CPU needs for 2026



If there is no data taking, all the computing capacity (CPU+GPU) of the EPN farm is available, and there is no need for further growth except for +5% at Tier 2 to meet MC productions

If there is data taking, the computing capacity (CPU+GPU) of the EPN farm is only partially available for offline use; an overall growth of +6% is needed.

First estimate of the resources needed during LS

- First year:
 - Low growth of CPU at T2s, flat at T0 and T1s, where EPN will be fully utilized for reconstructions
 - Disk space increase +15% to complete the reprocessing of the previous year HI data
 - Tape request of ~20 PB is needed to archive the legacy data
 - Corresponding disk space will be freed in the following year
- Following years:
 - CPU: same as first year but last year of the LS, when the EPN will be upgraded. The impact on CPU will depend on how long the EPN is unavailable
 - Tape: archiving campaign of legacy data will continue throughout the entire LS3 (~20 PB/year)
 - Disk: archiving legacy data will reduce the requests at T0 and T1 sites, but additional disk space will be needed at T2 sites for MC productions covering the entire Run 3 statistics

Computing resource estimates for 2026

ALICE		2024			2025		2026		
		Req.	C-RSG	Pledge	Req.	C-RSG	Est. No DT / C-RSG 2025	Est. DT / C- RSG 2025	
CPU [kHS23]	Tier-0	600	600	600	680	680	100%	104%	
	Tier-1	630	630	540	690	690	100%	104%	
	Tier-2	650	650	641	730	730	105%	111%	
	Total	1880	1880	1782	2100	2100	102%	107%	
Disk [PB]	Tier-0	67.5	67.5	67.5	78.0	78.0	115%	115%	
	Tier-1	71.5	71.5	61.9	79.0	79.0	115%	116%	
	Tier-2	66.5	66.5	69.8	77.0	77.0	116%	117%	
	Total	205.5	205.5	199.2	234.0	234.0	115%	116%	
Tape [PB]	Tier-0	A	181.0	181.0	181.0	220.0	220.0	104%	120%
		B						114%	
	Tier-1	A	107.0	107.0	102.4	123.0	123.0	107%	121%
		B						114%	
	Total	A	288.0	288.0	283.0	343.0	343.0	105%	121%
		B						114%	

Summary

- Currently deployed resources are being used with dynamic rebalancing where necessary
- ALICE and LHC operate with very high efficiency
 - This results in a lot of high-quality data
 - ... but in a lot of pressure on the required resources
- In particular, this puts high pressure on T0 disk
- Disk usage in every tier is very close to the operating capacity of 95%
- With the loss of russian resources every FA needs to step up to support smooth ongoing operations of the experiment
- Computing resource estimates for 2026, both running scenarios, are ready and submitted to the C-RSC