

LHCb computing requests for 2025

Lucio Anderlini
INFN Firenze

2024-09-04



Istituto Nazionale di Fisica Nucleare
SEZIONE DI FIRENZE

LHCb data taking in 2024

After the difficulties of last year, LHCb started the 2024 data taking at full steam.


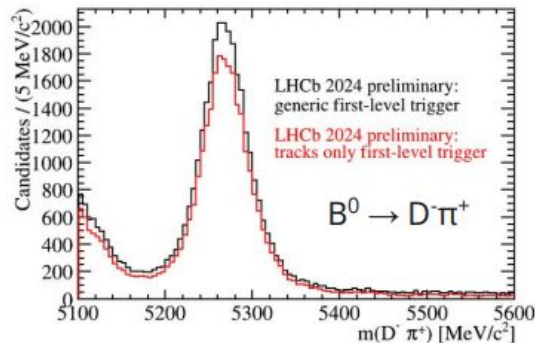

In 2024, we have collected **6.23 fb^{-1}** (3.23 fb^{-1} in Run1 + 7.09 fb^{-1} in Run2)

Improved trigger efficiency at lower momentum thanks to the full-software trigger.

More data


Stress on the computing resources


LHCb Integrated Recorded Luminosity in pp by years 2010-2024

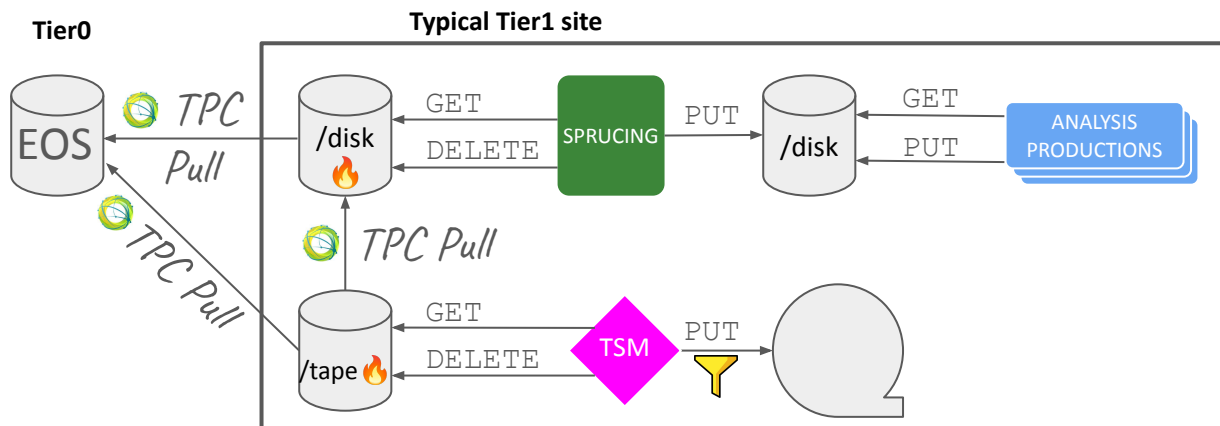
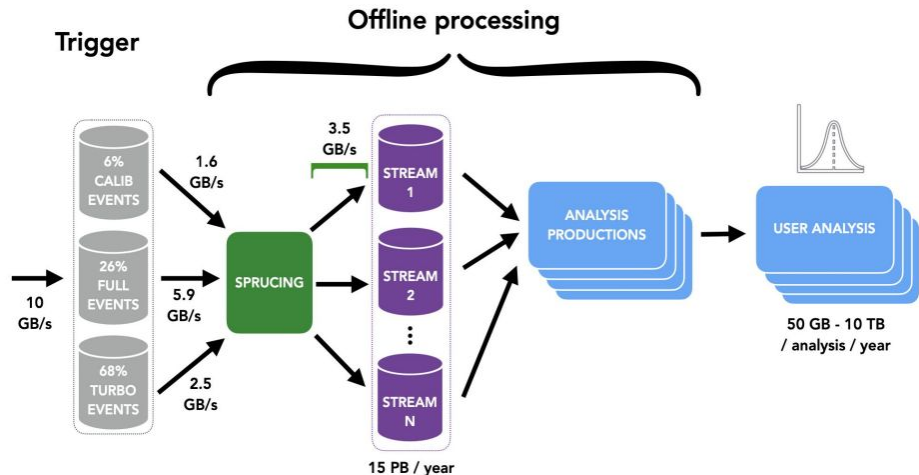


Data processing model

The data from the HLT (RAW) is streamed at **10 GB/s** to the Tier1 sites, where it is processed and archived (**2nd copy**).

The expected **bottleneck**  in the data processing is the bandwidth to write to tape.

/disk and /tape volumes are subject to a **high IOPS** including deletions .



Computing the CNAF share

Average throughput at a given Tier1.

Throughput from HLT2: **10 GB/s**

Fraction of Run1-2 data processed in that Tier 1


$$\langle T_{T1} \rangle > T_{HLT2} f_{T1}$$

[[wlcg docs](#)]

Site	Throughput GB/s
CNAF	1.72
IN2P3	1.25
KIT	2.23
NCBJ	1.32
NLT1	1.07
RRC-KI	0.25
PIC	0.2
RAL	2.96

The CNAF **bottleneck** was exercised in multiple tape challenges during the last few years, showing little contingency.

Tier-1 migration to Tecnopolo (→ new tape libraries) expected to solve.



LHCb@CNAF – Summer 2024: operations seen with FTS

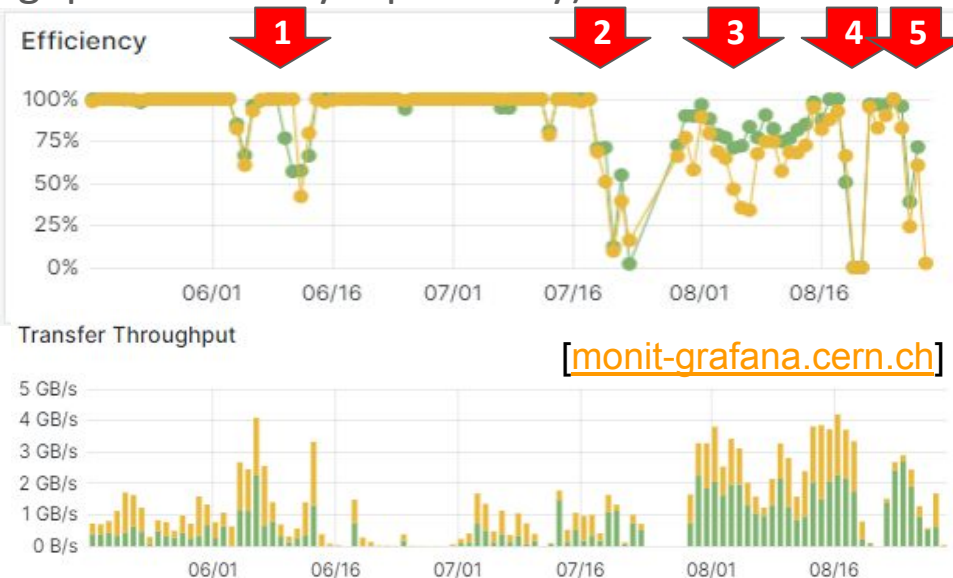


1. Operations instabilities due to the delay in the provisioning of disk pledge
2. Migration to the new disks (disk hardware installed at technopole)
3. Unsustainable load of sprucing jobs running at full steam
4. Tape buffer exhausted (for migration throughput limited by tape library)
5. Tape buffer exhausted (for migration throughput limited by GPFS).



Downtime since August 28th

Analysis jobs on other sites deprioritized.
Consequences of a prolonged downtime
are being assessed.





	max	avg
— CERN → CNAF DISK	2.70 GB/s	657 MB/s
— CNAF DISK → TAPE	2.35 GB/s	599 MB/s



Other high-throughput Tier-1s

[\[monit-grafana.cern.ch\]](https://monit-grafana.cern.ch)


Redistribution of CNAF share



Status of the Chinese Tier-1 site

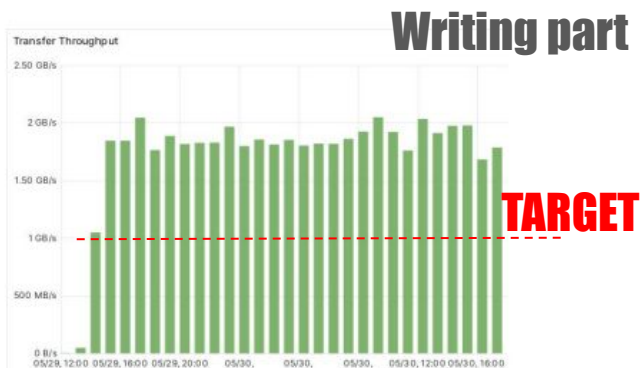
Data Challenge successfully completed in May 2024.
Target rate (1 GB/s) exceeded with very low failure rate.

WLCG Overview Board confirmed IHEP as Tier-1 on June 27th

Work in progress: adding IHEP to Dirac production configuration.



Institute of High Energy Physics
Chinese Academy of Sciences



CERN Disk → IHEP Disk



IHEP Disk → IHEP Tape



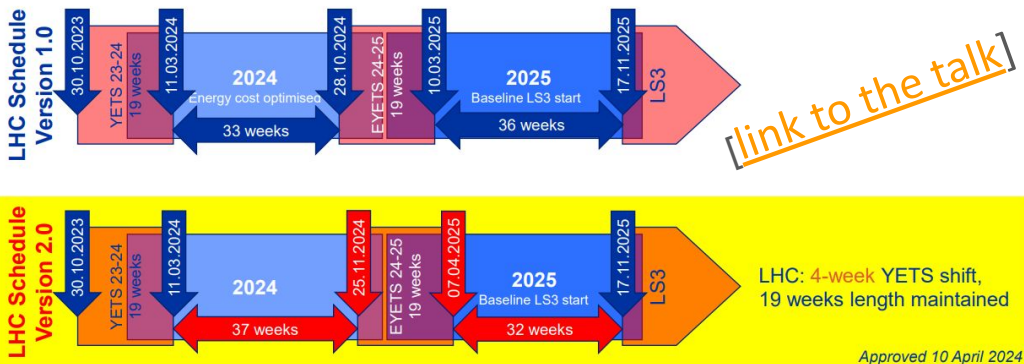
IHEP Tape → IHEP Disk

Anticipated data taking

LHC modified is schedule to postpone the technical stop by 4 weeks.

This sets additional strains on 2024 storage resources for all experiments including LHCb.

Changes made to the 2024 and 2025 LHC Schedule



[\[link to the talk\]](#)

The additional physics time of the 2024 run goes to proton physics
The **integrated luminosity target for 2024** was updated from 90 fb^{-1} to **110 fb^{-1}**



03.06.2024

R. Steerenberg | LHCP - 12th Large Hadron Collider Physics Conference

5

Recomputing the 2024 requests, **LHCb would need an additional pledge of 20 PB of disk (2 PB at CNAF*)** and 10 PB of tape (**covered by INFN the overpledge 2024***).

Requests are being discussed with WLCG for all experiments.

*) Assuming CNAF will resume operations quickly




Computing Requests for 2025

(including 4 weeks of data-taking anticipated to 2024)


LHCb Requests to the WLCG for 2025

LHCb-PUB-2024-002

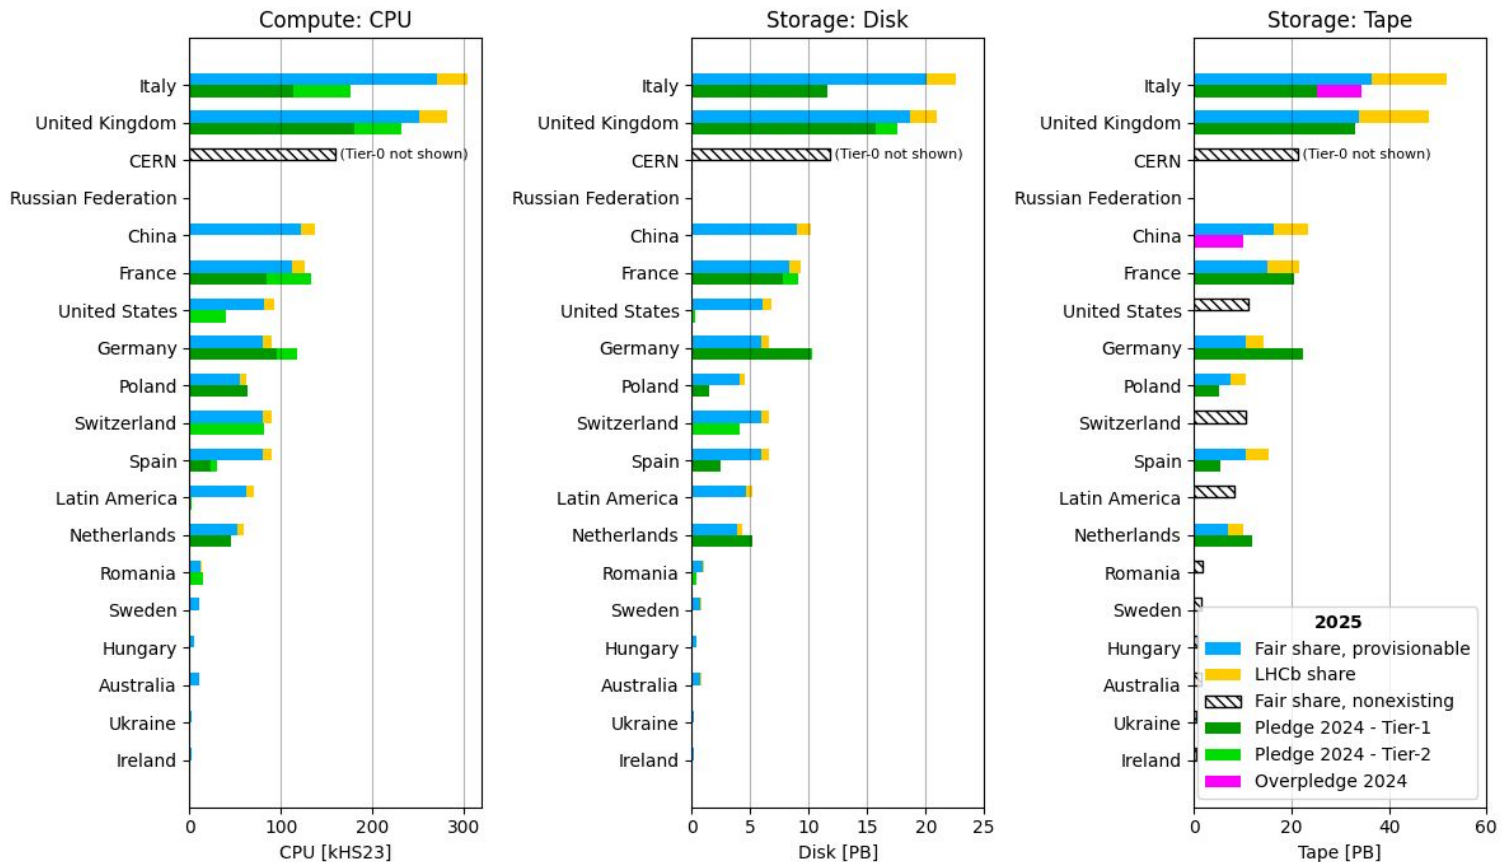
LHCb		2024					2025	
		Request	Pledge	Pledge/req	2024 req./ 2023 CRSG	2024 req. / 2023 pledge	Request	2025 req. / 2024 CRSG
WLCG CPU [kHS23]	Tier-0	123	174	141%	57%	57%	283	162%
	Tier-1	404	542	134%	57%	68%	928	162%
	Tier-2	224	394	176%	57%	52%	518	162%
	HLT	50			100%	100%	50	100%
	Sum	801	1110		59%	62%	1779	160%
Others		50			100%	100%	50	100%
Total		851	1,110	130%	60%	63%	1,829	157%
Disk [PB]	Tier-0	35.4	30.6	86%	117%	117%	54.9	180%
	Tier-1	50.8	53.0	104%	84%	93%	89.9	147%
	Tier-2	9.8	9.4	96%	84%	124%	17.4	147%
	Total	96.0	93.0	97%	94%	103%	162.2	157%
Tape [PB]	Tier-0	97	117	121%	107%	107%	170.4	146%
	Tier-1	118	125	106%	75%	88%	194.8	146%
	Total	215.0	242.2	113%	87%	96%	365.2	146%



Current situation on CRIC (pledges 2024)



Requests 2025 (compared to pledges 2024)



LHCb Requests to the Italian Federation for 2025

	Pledge '24	Request '25	Increment	Δ [k€]
CPU Tier-1 [HepScore23]	113430	173801	+60371	+604
Disk Tier-1 [TB]	11561	20096	+8535	+853
Tape Tier-1 (RRB) [TB]	25261	36483	+11222	+112
Tape Tier-1 (overpledge) [TB]	9068	15432	+6364	+64
CPU Tier-2 [HepScore23]	62595	97014	+34419	+344

Requested pledges computed as

(pledge) = (requests@Tier1 + requests@Tier2) %_{FTE} with %_{FTE} = 18.7%

(%_{FTE} increased wrt. 2024, 15.5% → 18.7%, for: +15 INFN and -61 Russian members)

Tape (pledge+overpledge) is requested based on the FTE of federations provisioning tape: %_{FTE} = 26.7%

CERN-RRB-2024-046

2025	PhD eq.	
	total/	%
AUSTRALIA	4	0.7
BRAZIL	19	3.3
COLOMBIA	2	0.3
COSTA RICA	4	0.7
FRANCE	45	7.7
BMBF GERMANY	30	5.2
MPG, GERMANY	2	0.3
HUNGARY	2	0.3
IRELAND	1	0.2
INFN ITALY	109	18.7
NETHERLANDS	21	3.6
P. R. CHINA	49	8.4
POLAND	22	3.8
HHNIPNE ROMANIA	5	0.9
RUSSIA INST.*	0	0.0
RUSSIA UNI.*	0	0.0
SPAIN	32	5.5
SWEDEN	4	0.7
SWITZERLAND	32	5.5
UKRAINE**	1	0.2
UK	101	17.4
USA	33	5.7
CERN	64	11.0
	582	100.0

Assessing the effect of Russian exclusion and new Tier1s

Assuming same Russian FTE as in 2025

	Pledge '24	Request '25	Increment	Effect
CPU Tier-1 [HepScore23]	113430	157313	+43883	+16488
Disk Tier-1 [TB]	11561	18189	+6628	+1907
Tape Tier-1 (RRB) [TB]	25261	33022	+7761	+3461
Tape Tier-1 (overpledge) [TB]	9068	12155	+3087	+3277
CPU Tier-2 [HepScore23]	62595	87810	+25215	+9204

The exclusion of the Russian community from the LHCb Collaboration increased the Italian %_{FTE}

16.9% → 18.7%

the effect on the requests to the Italian community for 2025 is **515 k€**.

Assuming no new Tier1

	Pledge '24	Request '25	Increment	Effect
CPU Tier-1 [HepScore23]	113430	173801	+60371	00
Disk Tier-1 [TB]	11561	20096	+8535	00
Tape Tier-1 (RRB) [TB]	25261	36483	+11222	00
Tape Tier-1 (overpledge) [TB]	9068	26337	+17269	-10905
CPU Tier-2 [HepScore23]	62595	97014	+34419	00

Widening the number of Tier-1s, including Beijing and Warsaw reduced the need for an Italian overpledge by 11 PB.

Conclusion

The upgraded LHCb experiment is taking good data steadily, distributing them through its Tier-1s.

CNAF is the most important Tier-1 site for LHCb.

Its **migration to the Technopole** and the subsequent difficulties with disk provisioning are **heavily impacting LHCb operations, with back-pressure reaching the online farm.**

The updated LHC schedule with **one additional month of data** taking will add **further disk pressure. Ideally, 2 PB of the 2025 requests could be anticipated to late 2024.**

Less pressure on tape due to the higher priority given by funding agencies (*e.g.* INFN overpledge)

The exclusion of the Russian colleagues (10% of the Collaboration) caused a significant increase of the Italian share, compensated only partially by the **promotion of NCBJ and Beijing sites to Tier1s.**

Requirements on bandwidth to tape were computed as yearly averages with a fully available site.

To recover from the incident we would need **more bandwidth to tape**, which requires infrastructure.