

# Nouvelles WLCG et LCG-France

D. Bouvet  
L. Duflot







# Plan

- ◆ Pledges 2026 et perspectives pour 2027
- ◆ Interface aux HPC en France
- ◆ WLCG tech news
- ◆







# Part du CC dans les pledges



Nouvel



# Part des T2 FR dans les pledges



Nouvel



# ALICE preliminary requests for 2027

ALICE		2025		2026		2027		
		RRB approved	Pledged	Request	2026 req. / 2025 RRB	RRB approved	Prelim Request	2027 req. / 2026 RRB
CPU	Tier0	680	680	710	104%	710	720	101%
	Tier1	690	596	720	104%	720	730	101%
	Tier2	730	750	810	111%	810	850	105%
	Total	2100	2026	2240	107%	2240	2300	103%
	HLT							
Disk	Tier0	78.0	78.0	90.0	115%	90.0	96.0	107%
	Tier1	79.0	69.1	90.0	114%	90.0	100.0	111%
	Tier2	77.0	82.5	89.0	116%	89.0	101.0	113%
	Total	234.0	229.6	269.0	115%	269.0	297.0	110%
	Tape	220.0	220.0	292.0	133%	292.0	306.5	105%
Tape	Tier1	123.0	117.4	155.0	126%	155.0	169.5	109%
	Total	343.0	337.4	447.0	130%	447.0	476.0	106%

## 2027/2026 requested resource increases:

- CPU: 3% (+60 kHS23)
- Disk: 10% (+28 PB)
- Tape: 6% (+29 PB)

## Plans for 2027

- Extensive Pb-Pb data (re)processing
  - Entire 2023 dataset
  - 50% of the 2024 dataset
  - Second pass of entire 2026 dataset
- Extensive MC production

## Missing details in the report to fully understand the resource requests

- More information provided on request



# ATLAS preliminary requests for 2027

ATLAS		2025		2026		2027		
		RRB approved	Pledged	Request	2026 req. / 2025 RRB	RRB approved	Prelim Request	2027 req. / 2026 RRB
CPU	Tier0	1100	1100	1265	115%	1265	1285	102%
	Tier1	1635	1639	1802	110%	1802	1802	100%
	Tier2	1998	2297	2202	110%	2202	2202	100%
	<b>Total</b>	<b>4733</b>	<b>5036</b>	<b>5269</b>	<b>111%</b>	<b>5269</b>	<b>5289</b>	<b>100%</b>
	HLT			438			1000	
Disk	Others						800	
	Tier0	56.0	56.0	65.0	116%	65.0	68.0	105%
	Tier1	186.0	186.7	199.0	107%	199.0	214.0	108%
	Tier2	227.0	218.9	243.0	107%	243.0	262.0	108%
	<b>Total</b>	<b>469.0</b>	<b>461.6</b>	<b>507.0</b>	<b>108%</b>	<b>507.0</b>	<b>544.0</b>	<b>107%</b>
Tape	Tier0	258.0	258.0	302.0	117%	302.0	312.0	103%
	Tier1	561.0	567.6	692.0	123%	692.0	741.0	107%
	<b>Total</b>	<b>819.0</b>	<b>825.6</b>	<b>994.0</b>	<b>121%</b>	<b>994.0</b>	<b>1053.0</b>	<b>106%</b>

2027/2026 requested resource increases:

- CPU: 0.4% (+20 kHS23)
- Disk: 7% (+37 PB)
- Tape: 6% (+59 PB)

## Plans for 2027

- **No major reconstruction campaigns** of the Run 3 data are foreseen
  - Derivation of new combined performance recommendations (calibrations, etc) takes too long
- Significant MC production for Run 3 analyses

**Opportunistic CPU resources**, beyond HLT farm, considered in the planning!

- About 11% of the total requirements

**Opportunistic tape resources** provided by two Tier2 sites (DESY, North-East)

- Used for additional MC dataset replicas

C-RSG congratulates ATLAS for **efficient use of disk resources**

- Dynamic replication of popular datasets, low volume of cold data on disk, efficient disk clean-up procedures

10



## CMS preliminary requests for 2027

CMS		2025		2026		2027		
		RRB approved	Pledged	Request	2026 req. / 2025 RRB	RRB approved	Prelim Request	2027 req. / 2026 RRB
CPU	Tier0	1180	1180	1350	114%	1350	1350	100%
	Tier1	1100	1166	1200	109%	1200	1200	100%
	Tier2	1900	1830	2000	105%	2000	2000	100%
	Total	<b>4180</b>	<b>4176</b>	<b>4550</b>	<b>109%</b>	<b>4550</b>	<b>4550</b>	<b>100%</b>
	HLT Others			350			860	
Disk	Tier0	70.0	70.0	81.0	116%	81.0	81.0	100%
	Tier1	142.0	133.8	164.0	115%	164.0	180.0	110%
	Tier2	175.0	159.6	198.0	113%	198.0	217.0	110%
	Total	<b>387.0</b>	<b>363.4</b>	<b>443.0</b>	<b>114%</b>	<b>443.0</b>	<b>478.0</b>	<b>108%</b>
Tape	Tier0	442.0	442.0	515.0	117%	515.0	530.0	103%
	Tier1	445.0	411.5	540.0	121%	540.0	610.0	113%
	Total	<b>887.0</b>	<b>853.5</b>	<b>1055.0</b>	<b>119%</b>	<b>1055.0</b>	<b>1140.0</b>	<b>108%</b>

### 2027/2026 requested resource increases:

- CPU: 0%
- Disk: 8% (+35 PB)
- Tape: 8% (+85 PB)

### Plans for 2027

- The increase in requested disk and tape storage is mainly driven by legacy **processing of Run 3 data and the corresponding MC simulation**
- CMS plans to delete a significant amount of data from tape
  - Requires coordination with sites

C-RSG commends the **significant reduction of storage requirements** through an optimized disk data placement and replication strategy

- AOD production for MC only on demand
- Reduce initial number of miniAOD replicas from 2 to 1



## CMS recommendations

1. The C-RSG recommends that CMS strengthen its data access monitoring capabilities to **enable more precise identification of cold data on disk for removal**. In particular, CMS is encouraged to collaborate with the XrootD development team to improve the monitoring of remote data access.
2. The C-RSG further recommends that CMS assess the effort required to **enhance different aspects of data management**, in order to prioritize activities that will deliver the greatest impact. In particular, the committee highlights the importance of adopting practices already in use by other LHC experiments, such as dynamic dataset replication based on access popularity and maintaining minimal disk placement of infrequently accessed data.
3. The C-RSG commends the recent **changes in data disk placement policies**, specifically the reduction from two to one initial copy of miniAOD datasets and the production of AOD data format for MC only on demand, which have resulted in a substantial reduction in storage needs. The C-RSG requests that the experiment **evaluate the resulting gains** with respect to the requested disk and tape resources for **2025 and 2026**, and that it subsequently reassess its storage requirements for 2027 in light of these policies.

# LHCb preliminary requests for 2027

LHCb		2025		2026		2027		
		RRB approved	Pledged	Request	2026 req. / 2025 RRB	RRB approved	Prelim Request	2027 req. / 2026 RRB
CPU	Tier0	283	283	344	122%	344	408	119%
	Tier1	928	849	1127	121%	1127	1337	119%
	Tier2	518	535	629	121%	629	747	119%
	Total	1729	1667	2100	121%	2100	2492	119%
	HLT			600			2700	
	Others			100			100	
Disk	Tier0	54.9	54.9	70.9	129%	70.9	73.3	103%
	Tier1	89.9	89.3	107.1	119%	107.1	108.0	101%
	Tier2	17.4	15.2	20.7	119%	20.7	20.9	101%
	Total	162.2	159.4	198.7	123%	198.7	202.2	102%
Tape	Tier0	170.4	170.4	202.2	119%	202.2	211.9	105%
	Tier1	194.8	181.2	233.7	120%	233.7	249.0	107%
	Total	365.2	351.6	435.9	119%	435.9	460.9	106%

## 2027/2026 requested resource increases:

- CPU: 19% (+392 kHS23)
- Disk: 2% (+3.5 PB)
- Tape: 6% (+25 PB)

## Plans for 2027

- Full reprocessing of Run 3 dataset
- Extensive MC production

Small fraction (2%) of opportunistic CPU resources, beyond HLT farm, considered in the planning

**Factor of two increase in full simulation CPU work per event relative to previous estimations!**

- Previous calculations not based on real measurements
- Impact in 2027 mitigated by availability of HLT farm and Tier0 resources for offline use
- C-RSG strongly recommends to **invest effort** in speeding up the full simulation and developing mitigation strategies

## LHCb recommendations

1. The CPU work required to simulate a single event in the full simulation has doubled compared to the number used in previous years to estimate the simulation CPU needs. The C-RSG requests that LHCb **quantify the impact of this increase on the CPU resources already approved for 2025 and 2026**, and clarify whether the allocated resources remain sufficient to meet the experiment's simulation goals.
2. The C-RSG asks LHCb to discuss **software optimization plans to reduce the CPU needed to fully simulate an event**.
3. The C-RSG asks LHCb to detail the actions being taken to **mitigate the CPU increase in the full simulation** without relying on additional pledged or opportunistic resources, such as the adoption of alternative simulation approaches, or adjustments in the event production strategy.
4. The C-RSG asks LHCb to provide **monitoring data on the ratio of full- to fast-simulated events** used by end-user analysis jobs in future usage reports.
5. The C-RSG encourages LHCb to include in future reports activities aimed at **optimizing the use of pledged resources**. Areas of focus may include more efficient CPU utilization, reduction of generated event sizes, enhancement of memory-efficient multi-threading support, and exploration of emerging architectures, such as ARM CPUs and GPUs.



## Overall findings and recommendations summary

- The large-scale data reconstruction campaigns planned by the experiments for 2027 will require sustained high **read rates** from tape
  - At present, tape throughput is not formally included among the computing requirements
  - Suggest discussion between experiments and WLCG about **I/O rate requirements**
- Computing resource requirements are expected to remain largely stable during the shutdown period, but a substantial increase is foreseen at the start of Run 4 for CMS and ATLAS
  - **Long term projections of computing requirements** are essential for effective coordination with funding agencies
  - Computing CDR/TDR reports for HL-LHC from CMS/ATLAS expected during 2026

# Interface aux HPC en France


- ◆ Dans le cadre de NUMPEX il y a une volonté de proposer une interopérabilité avec les autres DC et de faciliter l'accès
  - ◆ Proof of Concept avec proposition concrète de service d'interface ==>
  - ◆ Intérêt de LCG-France :
    - ◆ Adriano pour CMS (en premier AAI)
    - ◆ ATLAS? Faisabilité d'un SE+ARC-CE comme site grille devant ce service?



# Proposition POC

De l'extérieur on peut écrire sur des disques temporaires et soumettre des jobs avec des paramètres prédéfinis pour les applications enregistrées

Le service d'interface a des disques temporaires, un catalogue d'applications avec les paramètres d'exécution et un workflow manager



Le centre HPC est dans zone de confiance avec le service d'interface, il peut copier les données et tourner les jobs en mode "pull" donc sans être ouvert au monde  
Les mainteneurs sont responsables vis-à-vis du centre HPC de la sécurité et adéquation du code

Les "mainteneurs" proposent des codes validés HPC, installent sur les machines et publient



# DOMA: prévisions CMS pour DC27



## Network Connectivity Overview

Run-4		T0 Export		Re-Reco		Re-Mini		MC (incl. PU)	
All in Gbs		egress	ingress	egress	ingress	egress	ingress	egress	ingress
CERN	1550			380	0	0	0	380	20
KIT		130		30	40	10	0	5	40
PIC		60		20	20	5	0	5	25
IN2P3		130		30	40	10	0	5	40
CNAF		130		30	40	10	0	5	40
RAL		90		20	30	5	0	5	40
Serbia		60		20	20	5	0	5	25
NCBI		60		20	20	5	0	5	25
JINR		100		30	30	5	0	5	40
FNAL		510		130	170	35	5	380	45
Summed T1s		1270		330	410	90	5	420	320
Summed T2s		270		260	660	20	90	130	485

Connectivity is given in Gbs including a factor 4 for over provisioning

- Divide by 32 to get throughput in GB/s (no overprovisioning)

Some values are rounded down to 0

Certain flows are missing

- Can be guessed by extrapolating from Run-3?

## ◆ Overview Board

- ◆ T1 polonais: déjà en place pour ALICE, tout va bien pour CMS, approuvé
- ◆ Serbie et Roumanie toujours sur le chemin pour devenir T1 (CMS et ALICE respectivement)
- ◆ Task force on capacity planning (Sabine pour FR)
  - ◆ Pas mal de questions intéressantes (pledges, pledging/accounting infra, ressources opportunistes, SLA plutôt que Tier, GPU, ...) → veut-on se faire une opinion comme LCG-France?



- ◆ Accounting
  - ◆ WLCG demande d'adopter HS23 pour la publication.
  - ◆ Rapporter le nom du benchmark dans l'accounting : instructions pour déclarer le nom du benchmark <https://twiki.cern.ch/twiki/bin/view/LCG/AccountingClientHEPscoreDeployment>
    - ◆ le site doit s'assurer que le nom du benchmark est exact.
  - ◆ WLCG se dirige vers l'adoption à long terme de AUDITOR comme système générique d'accounting au niveau des sites.
- ◆ SHA-1 dans les CA
  - ◆ SHA-1 désactivé dans EL9.
  - ◆ Quelques CA utilisent SHA-1 pour les root certificates et les CRL.
  - ◆ CA root certificate :
    - ◆ pb seulement avec XrootD < v5.7.0.
  - ◆ CRL :
    - ◆ pas réellement un pb que les CRL soient absentes mais qq fonctionnalités de sécurité pourraient manquer dans les services.
      - ◆ CRL doivent passer à SHA-2 dès que possible.
  - ◆ Date limite de prise en charge de SHA-1 dans WLCG : fin 2026.
    - ◆ laisse 1 an pour passer à une nouvelle CA ou en créer une nouvelle.

