Computing Resources Scrutiny Group Report

José M. Hernández (CIEMAT)

for the Computing Resources Scrutiny Group

October 27-28, 2025

CERN-RRB-2025-75

C-RSG membership

M Schulz (CERN)	ALICE, lead	R Aaij (Netherlands)	ATLAS, CMS
J Amundson (USA)	ALICE	T Mkrtchyan (Germany)	LHCb, lead
L del Debbio (UK)	ALICE	D Gingrich (Canada)	LHCb
J Kleist (Nordic C.)	ATLAS, CMS, lead	K Dziedziniewicz-Wójcik(CERN)	scient. secr.
E Fede (France)	ATLAS, CMS	J Hernández (Spain)	chair, LHCb
D Elia (Italy)	ATLAS, CMS		

- Katarzyna Dziedziniewicz-Wójcik takes over from A Valassi as Scientific Secretary
- Luigi del Debbio (University of Edinburgh) replaces Chris Alton (Swansea University) as the UK-proposed member of the C-RSG
- Thomas Hartmann (DESY) has been proposed as the successor of Tigran Mkrtchyan (DESY) by the German funding agency
- C-RSG invites the RRB to formally confirm their nomination
- C-RSG thanks the C-RSG Scientific Secretary, the collaboration computing representatives and CERN management for their support

Autumn 2025 Scrutiny Process

- C-RSG met with LHCC WLCG referees and chair in advance
  - Coordination to avoid duplication of effort and ensure that no critical aspects are overlooked
  - Agreed that the C-RSG will focus on monitoring parameters related to CPU and storage usage, while the LHCC will oversee parameters tied to physics and data policy, such as trigger rates, the ratio of simulated to recorded data events, and the scientific relevance of both data and simulation samples
- LHC Collaborations submitted report on
  - Computing activities and plans for 2027
  - Preliminary resource requests for 2027
- C-RSG responded with written questions
  - Met via Zoom with Collaboration computing representatives
  - Provided Collaborations with draft report for any eventual corrections of facts
- C-RSG made recommendations regarding requests and plans

Computing resource pledges for 2026

Balance [%] 2026 Computing resources: (Pledged - RRB_Aproved) / RRB_Approved

		ATLAS	CMS	ALICE	LHCb	Total
Tier0	CPU	0	0	0	0	0
	Disk	0	0	0	0	0
	Tape	0	0	0	0	0
Tier1	CPU	1	6	-13	-14	-3
	Disk	0	-6	-14	-2	-4
	Tape	0	-7	1	-8	-3
Tier2	CPU	8	-3	1	-3	2
	Disk	-1	-11	6	-16	-4

- Funding agencies were requested to submit pledged resource contributions for 2026 by the deadline of 15 September 2025.
- ~14.1 M HS23 CPU, ~1.4 EB disk, 2.9 EB tape
- **Under-pledges at Tier1 and Tier2 levels** are a cause of concern

Balance of computing resource pledges

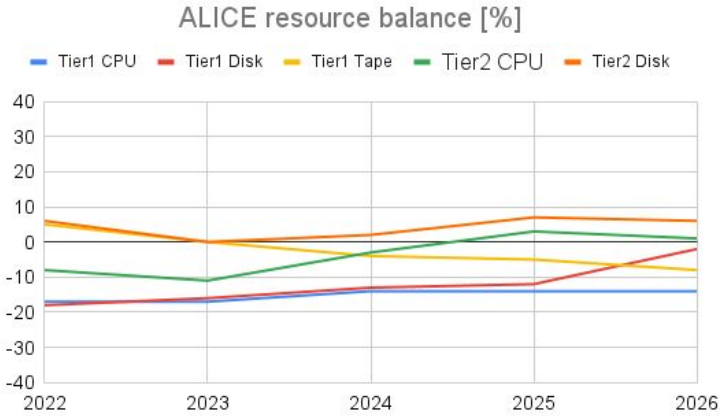

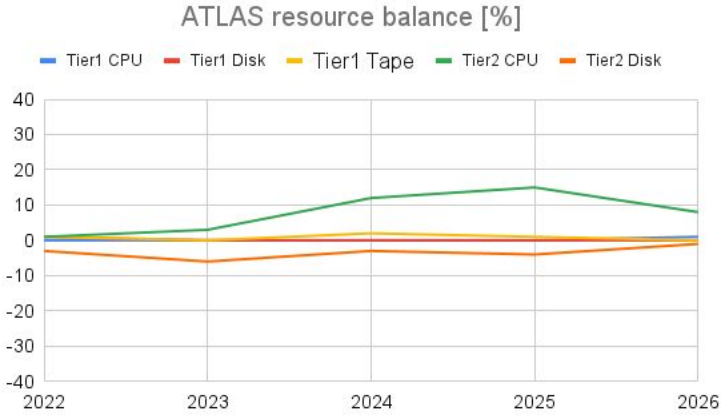

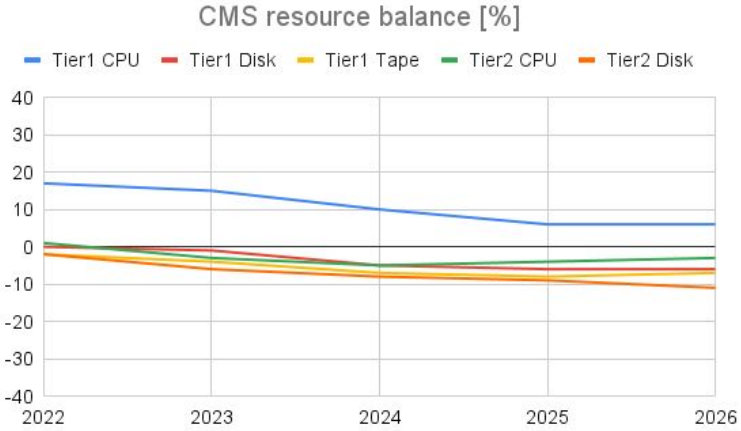

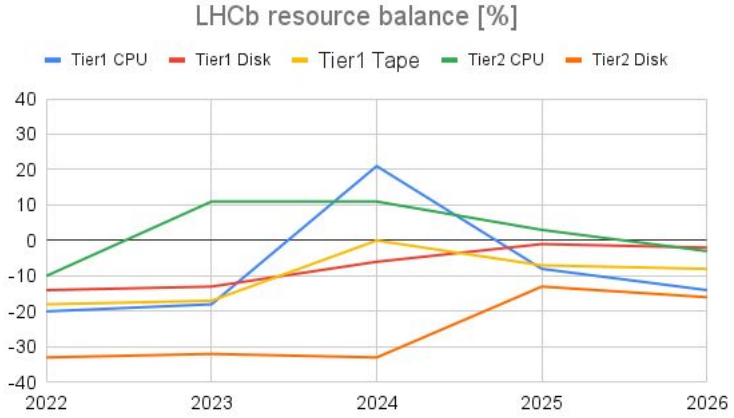

Background for 2027 computing resource requests

- In the Autumn scrutiny rounds, experiments submit **preliminary requests** for resources to be provided 18 months later
  - Opportunity for early assessment of requests and iterations with experiments
- 2027 is an LHC Long-Shutdown-3 year
  - No new experimental data will be collected in 2027
    - ➡ **Computing resource requirements are expected to remain flat or grow only modestly**
  - The HLT farms and Tier0 computing resources will be available for offline use
    - ➡ CPU demands on WLCG sites are not expected to increase
  - Storage requirements are foreseen to grow slowly
    - ➡ Driven primarily by data reprocessing campaigns and new MC simulations.
- WLCG has revised **downwards** its estimates of “flat-budget” purchasing power
  - Annual resource increases of approximately 10% in CPU and tape capacity, and 5% in disk capacity, are projected to be achievable within the same budget envelope
  - Disk storage is the most expensive infrastructure at WLCG sites
    - ➡ Optimizing disk usage remains essential

2027 preliminary computing resource requests

Experiments have established comprehensive plans for computing resource utilization in 2027, along with the estimated computing resources required for the planned processing, simulation and analysis of the collected data

- Extensive reprocessings of Run 3 data
- Enlarging simulation samples

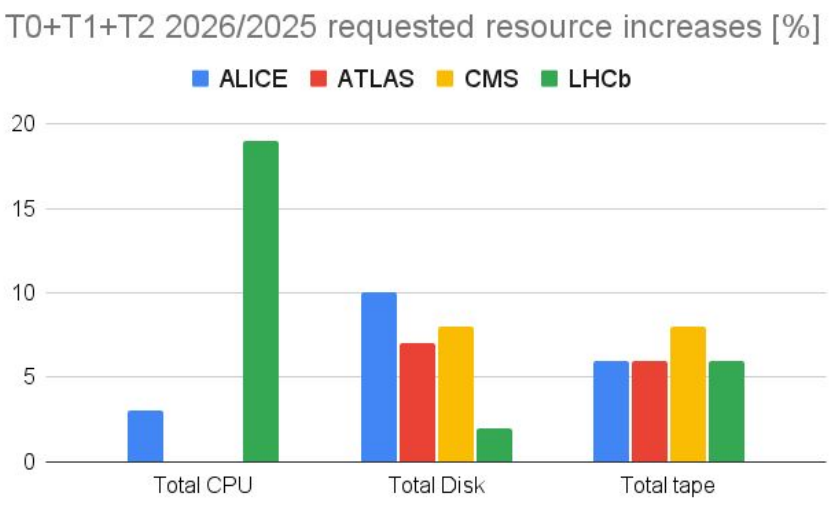

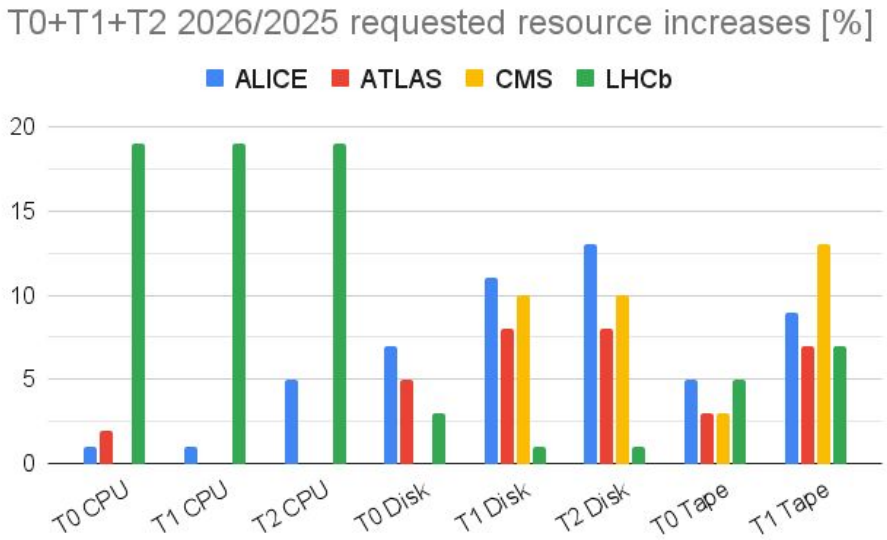

ALICE preliminary requests for 2027

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							
	Others							
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	Tier0	220.0	220.0	292.0	133%	292.0	306.5	105%
	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

ALICE recommendations

1. The C-RSG recommends to include into future reports **more detailed background information**, so that the reasoning behind decisions and requests can be realized [...]
2. The C-RSG recommends that, in addition to the reported resources needed to produce one MC event, the corresponding **information on the resources required for the other processing steps** also be provided, to ensure a complete understanding of the total resource requests.
3. The C-RSG recommends that future reports include a **more detailed description of activities** targeted at reducing resource requirements, such as the ongoing work to move from strategy A to strategy B for data compression.
4. The C-RSG recommends to share a **more detailed insight into the use of the different opportunistic resources**, including an assessment of the potential impact on the physics program if these resources were no longer available.
5. The C-RSG recommends that the experiment **reassess** whether the **requested CPU increases**, albeit moderate, are fully justified, taking into account the availability of the online EPN farm in 2027 for offline processing.

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%
	Total	**4733**	**5036**	**5269**	**111%**	**5269**	**5289**	**100%**
	HLT			438			1000	
	Others						800	
Disk	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%
	Total	**469.0**	**461.6**	**507.0**	**108%**	**507.0**	**544.0**	**107%**
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%
	Total	**819.0**	**825.6**	**994.0**	**121%**	**994.0**	**1053.0**	**106%**

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

ATLAS recommendations

1. The C-RSG recommends that ATLAS examine the factors preventing a **more rapid transition to fast simulation**. Extensive use of fast simulation remains essential to reduce the CPU needs.
2. Considering the availability of the HLT farm and Tier0 in 2027, which may be repurposed for offline data processing and simulation, the C-RSG recommends that the experiment carefully **reassess** whether the **requested CPU and disk increases at the Tier0 level**, albeit moderate, are genuinely necessary.

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	**4180**	**4176**	**4550**	**109%**	**4550**	**4550**	**100%**
	HLT			350			860	
	Others							
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	**387.0**	**363.4**	**443.0**	**114%**	**443.0**	**478.0**	**108%**
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	**887.0**	**853.5**	**1055.0**	**119%**	**1055.0**	**1140.0**	**108%**

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

- CPU requirements during the LHC shutdown are expected to remain stable
  - This expectation is fulfilled by all experiments except LHCb
- For storage resources, experiments request only modest increases, consistent with the revised flat funding scenario
  - Nevertheless, the C-RSG encourages further optimization of storage usage
- C-RSG encourages CMS and ALICE to include expected opportunistic CPU resources beyond the HLT farms in the CPU planning
  - Already done by ATLAS and LHCb
- ATLAS and CMS allocate a substantial share of their computing resources to processing and storing the so-called parked or delayed data streams
  - CMS: 60% of CPU and disk resources; 45% of tape
  - ATLAS: 20% of total raw data
  - ATLAS require more resources globally than CMS (~15% more CPU and disk)
  - Beyond the scope of this committee to judge the relevance of those datasets

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 expected during 2026

Overall findings and recommendations summary

- Several tasks with a direct impact on resource requirements are currently limited by a **shortage of available effort**
  - ATLAS: derivation and approval of performance recommendations needed to enable the use of fast simulation
  - CMS: monitoring and purging of cold data on disk.
  - ALICE: resolve the calibration issues of the TPC that currently hinder the transition to a more compressed raw data format
  - LHCb: work required to accelerate the full simulation
- **Strengthening support** for these activities would help optimize resource usage and reduce future computing demands

Overall recommendations

1. To better assess resource requirements and harmonize resource request information across experiments, the C-RSG requests that all future resource request reports include **summary tables** showing the percentage of requested resources along two dimensions, data-taking era and processing activity for CPU or data format type for storage.

2. Following feedback from the experiments, the agreed **data popularity metric** for assessing the volume of cold data on disk is defined as the total volume of disk-resident data that has not been accessed within 12 months of its placement on disk. The C-RSG requests that this metric be included in all future resource request reports.

Overall recommendations

3. In preparation for the final computing resource requests in the next scrutiny round, the C-RSG recommends evaluating potential **optimizations in storage usage** to help limit projected increases. These include more proactive deletion of cold data from disk, systematic replacement of analysis datasets with newly produced versions, and removal of obsolete data from tape.

4. In view of the forthcoming long shutdown 3, where computing requirements are not expected to increase significantly, and the start of the Run 4 that will presumably require a sharp increase in resources, the C-RSG recommends that the experiments include in their next report **long-term projections of computing resource needs**. Such projections will assist the funding agencies in defining the most effective investment strategy.

Heterogeneous Architectures in WLCG
2025/12/04

https://indico.cern.ch/event/1526077/timetable/

Summary

	LHCb software readiness	resources readiness (availability to LHCb users)	distributed computing (Dirac+X)
ARM CPUs	Ready	via WLCG	Ready
GPUs for Sim accelerators	Ongoing. Needs double-precision GPUs	LHCb-owned, or via HPCs (opportunistically)	Ongoing
GPUs for ML	Ongoing	LHCb-owned, or via clouds	A very different paradigm
TPUs/NPUs/FPGA	Specific applications, online	LHCb-owned	Not planned

Preliminary HEPScore23-GPU Benchmark Results

*Input from M. Mascheroni

- Additionally, able to run jobs on GPUs using this benchmark as payload in the Global Pool

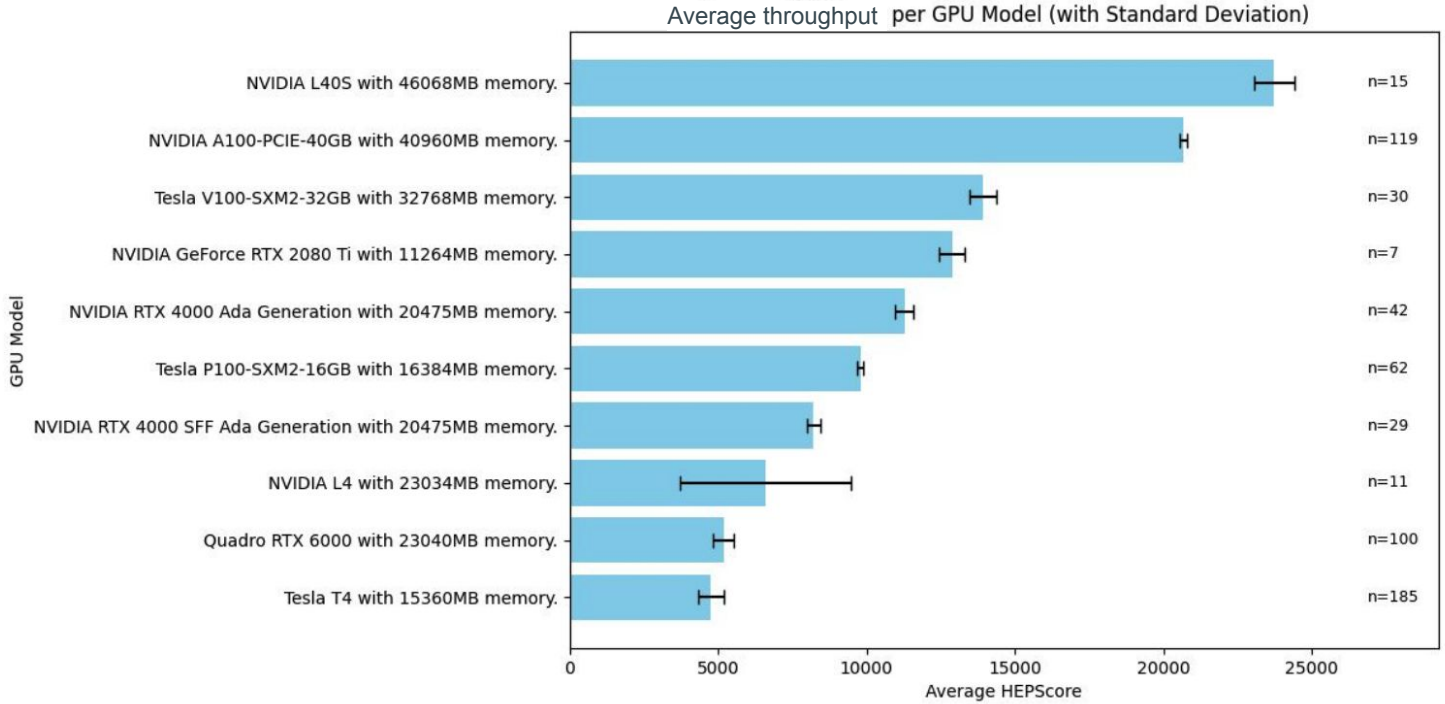

GPU Speedup

- In offline, TPC is 50% of the Pb–Pb processing time, 60% for pp.
  - **CPU-bound** (if GPUs are fast enough) → all **offloaded compute time** is **eliminated**.
  - **GPU speedup** on the **EPN** farm is **2x** for **Pb–Pb** and **2.5x** for **pp** (**exact match** of expectation when offloading **50%** / **60%**).

- Next step: Offload more processing steps to GPU:
  - **TPC Track Model Decoding in 2024 (see LHCP talk)**
  - **ITS tracking is the next big step**
  - **Ultimate plan is to offload the full barrel tracking to GPU: 80% of the workload, which should lead to 5x speedup.**

SONIC – **S**ervices for **O**ptimized **N**etwork **I**nference on **C**oprocessors

“As-a-service” approach to integration of hardware accelerators (GPU, FPGA, etc.):

- off-load heavy ML and non-ML computations to coprocessors
- “client–server” architecture with local or remote servers
- dynamically optimize CPU / coprocessor ratio for each workflow
- minimal disruption to existing workflows

SONIC developments are ongoing in CMS, ATLAS, IceCube, DUNE, LIGO