

# Computing Resources Scrutiny Group Report

1

**Pekka K. Sinervo, C.M., FRSC**  
**University of Toronto**

**For the Computing Resources Scrutiny Group**

April 22-23, 2024

**CERN-RRB-2024-013**

Pekka Sinervo, C.M.



April 22-23, 2024

## C-RSG membership

C Allton (UK)	J Hernandez (Spain)
E Fede (France)	J Kleist (Nordic countries)
M Schulz (CERN)	A Valassi (CERN, scient. secr.)
P Christakoglou (Netherlands)	P Sinervo (Canada)
A Connolly (USA)	D Elia (Italy)
T Mkrtchyan (Germany)	

2

- The C-RSG membership has remained the same since the Fall 2023 meeting
- C-RSG thanks the C-RSG Scientific Secretary, A Valassi, the collaboration computing representatives and CERN management for their support.



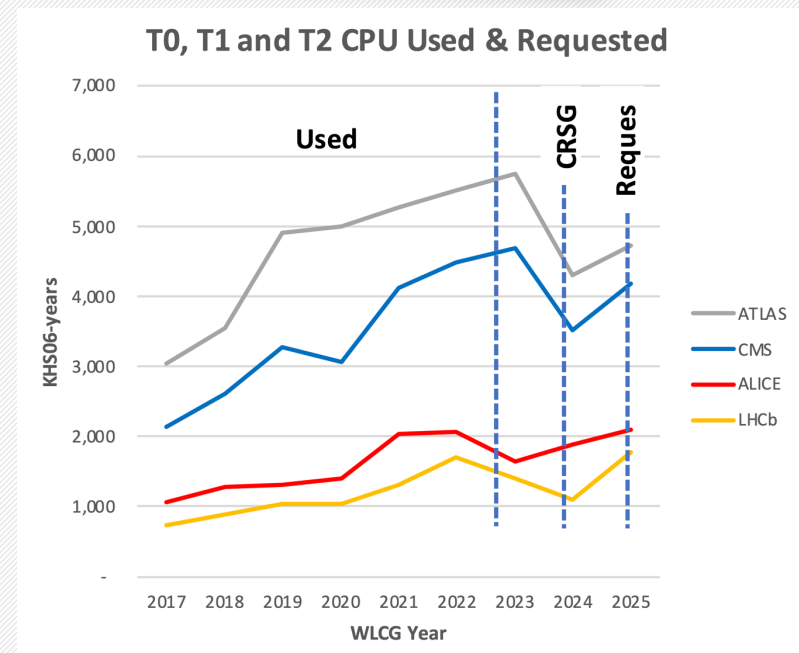
# Spring 2024 Scrutiny Process

3

- C-RSG met with LHCC WLCG referees and LHCC chair in advance
  - Identified a number of issues of common concern
- LHC Collaborations submitted
  - Report on 2023 computing resource usage,
  - Computing activities and plans for 2024
  - Resource requests for 2025
- C-RSG responded with written questions
  - Met with Collaboration computing representatives
  - Provided Collaborations with draft report for any corrections of facts
- Made recommendations regarding requests and plans



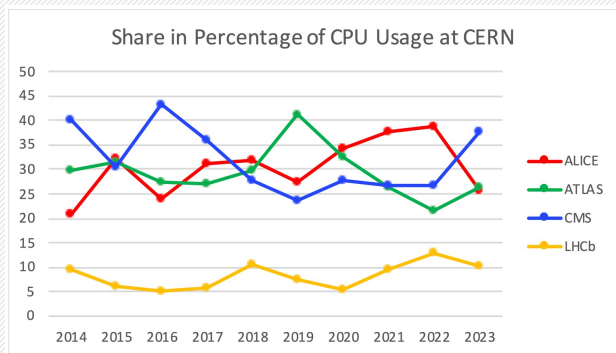
# Resource Utilization in 2023

- CPU utilization dominated by simulation & analysis
- All the Collaborations had taken advantage of opportunistic CPU well above pledged values
- Trends for ALICE and LHCb reflect changes in running plans for 2023, 2024 and 2025



# CPU Usage in 2023




- Total 2023 CPU utilization 13.9 MHS06-years
  - Roughly 1.5M cores used 24-7 through 2023
  - Half delivered by T2 sites, with T1 (30%) and T0 (20%)
- CMS has been largest T0 user; ATLAS and CMS dominate T1 sites







# Disk Usage in 2023

- Disk usage tracked for CERN and T1 sites
- Trends been consistent, though total CERN disk has increased relative to others



# Tape Usage in 2023


- Total of 1150 PB up from 880 PB in 2022)
- Increase reflects volume of Run 3 data recorded
- Allocated resources are keeping up with data increase



# 2024 Utilization and 2025 Requests

- The 2024 plans are based on the Run 3 schedule
  - Continuing data collection and processing
  - Focus largely on Run 3 analyses
- The 2025 plans are based on the current Run 3 schedule

2025 Requests	ALICE	ATLAS	CMS	LHCb	Total
CPU	+12%	+10%	+19%	+60%	+18%
Disk	+14%	+14%	+19%	+57%	+20%
Tape	+19%	+24%	+27%	+46%	+27%





# Alice Usage for 2023 and Request for 2025

ALICE		2023			2024		2025		
		C-RSG recomm.	Pledged	Used	C-RSG recomm.	Pledged	Request	2025 req. / 2024 C-RSG	C-RSG recomm.
CPU	Tier-0	541	541	660	600	600	680	113%	680
	Tier-1	572	506	446	630	540	690	110%	690
	Tier-2	592	567	555	650	641	730	112%	730
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	<b>Total</b>	<b>1705</b>	<b>1614</b>	<b>1661</b>	<b>1880</b>	<b>1781</b>	<b>2100</b>	<b>112%</b>	<b>2100</b>
	<i>Others</i>			222					
Disk	Tier-0	58.5	58.5	51.6	67.5	67.5	78.0	116%	78.0
	Tier-1	63.5	57.6	50.2	71.5	61.9	79.0	110%	79.0
	Tier-2	57.5	60.4	47.3	66.5	69.8	77.0	116%	77.0
	<b>Total</b>	<b>179.5</b>	<b>176.5</b>	<b>149.1</b>	<b>205.5</b>	<b>199.2</b>	<b>234.0</b>	<b>114%</b>	<b>234.0</b>
Tape	Tier-0	131.0	131.0	115.0	181.0	181.0	220.0	122%	220.0
	Tier-1	82.0	87.7	42.5	107.0	102.4	123.0	115%	123.0
	<b>Total</b>	<b>213.0</b>	<b>218.7</b>	<b>157.5</b>	<b>288.0</b>	<b>283.4</b>	<b>343.0</b>	<b>119%</b>	<b>343.0</b>

- 2023 utilization
  - Pb-Pb run in 2023 more than doubled Alice dataset
  - Continued work to increase compression of CTF
  - Less 2022 data reflects T1 utilization

- 2025 plan is 15 days of Pb-Pb run
  - Reduction of about 20%
- Continuing work to calibrate TPC and implement higher compression
- Working on understanding decrease in efficiency

# ATLAS Usage for 2023 and Request for 2025

ATLAS		2023			2024		2025		
		C-RSG recomm.	Pledged	Used	C-RSG recomm.	Pledged	Request	2025 req. / 2024 C-RSG	C-RSG recomm.
CPU	Tier-0	740	693	675	936	936	1100	118%	1100
	Tier-1	1430	1462	1613	1516	1514	1635	108%	1635
	Tier-2	1747	1841	3457	1852	2074	1998	108%	1998
	HLT	n/a	n/a	467	n/a	n/a	n/a	n/a	n/a
	<b>Total</b>	<b>3917</b>	<b>3996</b>	<b>6212</b>	<b>4304</b>	<b>4524</b>	<b>4733</b>	<b>110%</b>	<b>4733</b>
	<i>Others</i>			1667					
Disk	Tier-0	40.0	40.0	36.4	49.0	49.0	56.0	114%	56.0
	Tier-1	136.0	140.9	149.0	163.0	163.1	186.0	114%	186.0
	Tier-2	168.0	160.6	155.0	200.0	194.0	227.0	114%	227.0
	<b>Total</b>	<b>344.0</b>	<b>341.5</b>	<b>340.4</b>	<b>412.0</b>	<b>406.1</b>	<b>469.0</b>	<b>114%</b>	<b>469.0</b>
Tape	Tier-0	174.0	174.0	127.0	207.0	207.0	258.0	125%	258.0
	Tier-1	353.0	360.3	346.0	452.0	460.0	561.0	124%	561.0
	<b>Total</b>	<b>527.0</b>	<b>534.3</b>	<b>473.0</b>	<b>659.0</b>	<b>667.0</b>	<b>819.0</b>	<b>124%</b>	<b>819.0</b>

N.B.: Other CPU represents non-WLCG resources

- 2023 utilization focus was Run 3
  - Increasing use of smaller formats
  - Continued work on alternate architectures
- 2025 expected to be full year of running to collect  $120 \text{ fb}^{-1}$  of data
  - CPU & storage increases reflect improved simulation, smaller raw event size and lighter physics formats
- Porting to other architectures & platforms will continue
  - Opportunistic CPU expected to continue to play key role



# CMS Usage for 2023 and Request for 2025

CMS		2023			2024		2025		
		C-RSG recomm.	Pledged	Used	C-RSG recomm.	Pledged	Request	2025 req. / 2024 C-RSG	C-RSG recomm.
CPU	Tier-0	720	720	969	980	980	1180	120%	1180
	Tier-1	800	916	1173	930	1020	1100	118%	1100
	Tier-2	1350	1313	2544	1600	1484	1900	119%	1900
	HLT	n/a	n/a	0	n/a	n/a	n/a	n/a	n/a
	<b>Total</b>	<b>2870</b>	<b>2949</b>	<b>4686</b>	<b>3510</b>	<b>3484</b>	<b>4180</b>	<b>119%</b>	<b>4180</b>
	<i>Others</i>			0					
Disk	Tier-0	45.0	45.0	39.2	54.0	54.0	70.0	130%	70.0
	Tier-1	98.0	96.8	76.7	122.0	115.7	142.0	116%	142.0
	Tier-2	117.0	109.7	96.1	149.0	134.1	175.0	117%	175.0
	<b>Total</b>	<b>260.0</b>	<b>251.5</b>	<b>212.0</b>	<b>325.0</b>	<b>303.8</b>	<b>387.0</b>	<b>119%</b>	<b>387.0</b>
Tape	Tier-0	228.0	228.0	202.8	320.0	320.0	442.0	138%	442.0
	Tier-1	316.0	303.7	226.4	380.0	353.9	445.0	117%	445.0
	<b>Total</b>	<b>544.0</b>	<b>531.7</b>	<b>429.2</b>	<b>700.0</b>	<b>673.9</b>	<b>887.0</b>	<b>127%</b>	<b>887.0</b>

- 2023 focus was Run 3 data analysis
  - Successfully mitigated reductions in T1 availability
  - Impressive increase in opportunistic resources

- 2025 requests is needed for Run 3 data analysis
  - T0 changes arise from raw data repacking & buffers
- T1 and T2 resources used for simulation and periodic reprocessing of entire Run 3 dataset

N.B.: Other CPU represents non-WLCG resources



# LHCb Usage for 2023 and Request for 2025

LHCb		2023			2024		2025		
		C-RSG recomm.	Pledged	Used	C-RSG recomm.	Pledged	Request	2025 req. / 2024 C-RSG	C-RSG recomm.
CPU	Tier-0	215	215	258	174	174	283	163%	283
	Tier-1	707	598	652	572	542	928	162%	928
	Tier-2	391	434	492	319	394	518	162%	518
	HLT	50	50	0	50	0	50	n/a	50
	<b>Total</b>	<b>1363</b>	<b>1297</b>	<b>1402</b>	<b>1115</b>	<b>1110</b>	<b>1779</b>	<b>160%</b>	<b>1779</b>
	<i>Others</i>			26					
Disk	Tier-0	30.3	30.3	23.4	30.6	30.6	54.9	179%	54.9
	Tier-1	60.5	54.7	35.2	61.2	53.0	89.9	147%	89.9
	Tier-2	11.6	7.9	3.6	11.8	9.4	17.4	147%	17.4
	<b>Total</b>	<b>102.4</b>	<b>92.9</b>	<b>62.2</b>	<b>103.6</b>	<b>93.0</b>	<b>162.2</b>	<b>157%</b>	<b>162.2</b>
Tape	Tier-0	91.0	91.0	37.8	117.1	117.0	170.4	146%	170.4
	Tier-1	157.0	133.7	55.8	133.3	125.0	194.8	146%	194.8
	<b>Total</b>	<b>248.0</b>	<b>224.7</b>	<b>93.6</b>	<b>250.4</b>	<b>242.0</b>	<b>365.2</b>	<b>146%</b>	<b>365.2</b>

- 2023 processing reflected VELO incident
  - Underutilized disk and tape
  - CPU used for simulation of earlier data

- 2025 request reflects full year of data-taking in new model
  - Will need 125 PB of tape for data storage alone
  - 75% of CPU used for simulation
- Bringing new T1 sites online

# Summary of Resource Requests for 2025

- The requested resources are necessary to address the approved physics programs of the LHC Collaborations.
  - Increases in computing requirements driven by Run 3 data volume
  - Simulation scales comparably to data
- Several collaborations have had to adjust running plans
  - ALICE accommodating to changes in HI running schedule
  - LHCb will reach its planned data-taking rate
- Development effort continues to optimize computing utilization
  - Smaller data formats for physics analyses
  - Making use of alternative CPU architectures provided by HPC
- Following recommendations are synopsis of those in report



# ALICE Recommendations

- ALICE-1** The C-RSG recommends that, given substantial potential savings in tape utilization under Strategy B compression, the ALICE Collaboration prioritizes evaluation of the physics impact of adopting this improved compression.
- ALICE-2** The ALICE Collaboration tabulate their resource requests assuming both Strategy B and Strategy A.
- ALICE-3** Current efficiencies for have recently dropped. The C-RSG recommends allocating the personnel resources necessary to improve efficiencies.
- ALICE-4** The C-RSG recommends that the ALICE Collaboration evaluate the adoption of smaller derived AODs and MC AOD formats as a means to improve IO efficiencies, and to reduce storage needs.



# ATLAS Recommendations

15

- ATLAS-1** The C-RSG recommends that the ATLAS Collaboration continues its effort to reduce the raw event size.
- ATLAS-2** The C-RSG recommends that the ATLAS Collaboration continues its effort to adopt the DAOD\_PHYS and DAOD\_PHYSLITE data formats.



# CMS Recommendations

- CMS-1** The C-RSG recommends that the CMS Collaboration continues its efforts in adoption of the NanoAOD data format.
- CMS-2** The C-RSG requests that the CMS Collaboration continue reporting on efforts to pinpoint and reduce the causes behind the relatively low 2023 CPU efficiency.

The logo for GridPP, featuring the text "GridPP" in a bold, white, sans-serif font. The background of the logo is a blue-tinted image of a particle detector or server rack.

**GridPP**

Distributed Computing for Data-Intensive Research



# LHCb Recommendations

17

- LHCb-1** Estimates for the required CPU for data processing, simulation and analysis are based on extrapolations rather than on measurements on real data or simulations using the Run 3 nominal running conditions. The C-RSG recommends reassessing those parameters using the data collected in 2024 and the corresponding simulations.
- LHCb-2** The coordination of the software and computing areas has been reorganized within LHCb to optimize resources, increase efficiency, and improve communication. The C-RSG requests that the Collaboration provides a report in the next scrutiny round on the effectiveness of this new organization in addressing the shortage of personnel for computing and software support.



# Overall Recommendations

- ALL-1** The use of compact data formats has reduced pressure on disk resources while increasing throughput of physics analyses. The C-RSG encourages all the Collaborations to continue to focus development efforts in this area
- ALL-2** Increasingly the LHC Collaborations are relying on the use of HPC systems...[that] provide a large fraction of their capabilities in form of accelerators, such as GPUs, rather than CPUs. Their effective and efficient utilisation requires significant changes to the software architecture and adaptation of workloads and workflow management systems. Additional sustained investment in expert developers is essential and adequate mechanisms to fund these activities should be identified.

# Final Summary and Comments

19

- Clear that next few years will see changes in new architectures and facilities
- Critical need for increased system and user software development.
- Important for funding agencies to recognize and work to meet these resource needs
- C-RSG thanks all Collaborations for effective and collegial engagement