

Computing Resources Scrutiny Group Report

1

□

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

For the Computing Resources Scrutiny Group

October 24-25, 2022

CERN-RRB-2022-079

Pekka Sinervo, C.M.



October 25, 2022

C-RSG Membership

- C-RSG thanks the experiment representatives and CERN management for their support.
- C-RSG also thanks H. Meinhard for his 11 years of dedication to the C-RSG.


C Allton (UK)	J Hernandez (Spain)
N Neyroud (France)	J Kleist (Nordic countries)
J van Eldik (CERN)	A Valassi (CERN, scient. secr.)
P Christakoglou (Netherlands)	P Sinervo (Canada)
A Connolly (USA)	V Vagnoni (Italy)
T Mkrtchyan (Germany)	

Fall 2022 Scrutiny Process

- The four LHC experiments gave updates on computing and data processing plans,
 - Described computing activities for 2022 year
 - Updated computing plans for 2023 year
 - Initial resource estimates for 2024
- Run 3 is now dominating all computing activities
 - Using estimates of delivered integrated luminosity for 2023
 - There are significant uncertainties for 2023 and 2024
 - Effect of power reductions on running schedule
 - Longer-term effect of loss of Russian and Belarusian computing resources
- C-RSG believes that the 2024 estimates are needed to achieve the approved physics programs
 - Estimates reflect the Run 3 data-taking and physics programs
 - All collaborations working hard to make most effective use of the allocated resources

Resource Utilization in 2022

- 2022 is first year of Run 3
- Approved resources for 2023 are in-line with the needs for the Run 3 data processing and simulation
- Significant opportunistic CPU utilization by all experiments
 - Use of HPC resources continue to rise
 - Being offset in part by less utilization from internal resources (such as HLT and O2 systems)



ALICE Preliminary Request for 2024

ALICE		2022		2023			2024	
		C-RSG recomm.	Pledged	Request	2023 req. / 2022 C-RSG	C-RSG recomm.	Prelim Request	2024 req. / 2023 C-RSG
CPU	Tier-0	471	471	541	115%	541	622	115%
	Tier-1	498	448	572	115%	572	655	115%
	Tier-2	515	517	592	115%	592	683	115%
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	1484	1436	1705	115%	1705	1960	115%
	<i>Others</i>							
Disk	Tier-0	50.0	50.0	58.5	117%	58.5	67.5	115%
	Tier-1	55.0	49.7	63.5	115%	63.5	71.5	113%
	Tier-2	49.0	55.2	57.5	117%	57.5	66.5	116%
	Total	154.0	154.9	179.5	117%	179.5	205.5	114%
Tape	Tier-0	95.0	95.0	131.0	138%	131.0	167.0	127%
	Tier-1	63.0	71.8	82.0	130%	82.0	102.0	124%
	Total	158.0	166.8	213.0	135%	213.0	269.0	126%

- 2023 & 2024 priorities
 - Processing of p-p and Pb-Pb data collected in 2023
 - Reconstruction of p-Pb data taken in 2024
- 2024 reflects expected growth
 - Δ CPU = 15%
 - Δ Disk = +14%
 - Δ Tape = +26%
- Growth is driven by planned p-p and Pb-Pb running

ATLAS Preliminary Request for 2024

ATLAS		2022		2023			2024	
		C-RSG recomm.	Pledged	Request	2023 req. / 2022 C-RSG	C-RSG recomm.	Prelim Request	2024 req. / 2023 C-RSG
CPU	Tier-0	550	550	740	135%	740	850	115%
	Tier-1	1300	1300	1430	110%	1430	1501	105%
	Tier-2	1588	1588	1747	110%	1747	1834	105%
	HLT	n/a	n/a	n/a	n/a	n/a	0	n/a
	Total	3438	3438	3917	114%	3917	4185	107%
Others								
Disk	Tier-0	32.0	32.0	40.0	125%	40.0	46.0	115%
	Tier-1	116.0	116.0	136.0	117%	136.0	162.0	119%
	Tier-2	142.0	142.0	168.0	118%	168.0	198.0	118%
	Total	290.0	290.0	344.0	119%	344.0	406.0	118%
Tape	Tier-0	120.0	120.0	174.0	145%	174.0	205.0	118%
	Tier-1	272.0	272.0	353.0	130%	353.0	448.0	127%
	Total	392.0	392.0	527.0	134%	527.0	653.0	124%

- 2023 and 2024 activities
 - Increasing use of fast simulation and smaller data format
- 2024 requests reflect full year of running to collect 2×10^{10} events
 - Δ CPU = +7%
 - Δ Disk = +18%
 - Δ Tape = +24%
- Evolving data format/architectures
 - Implementing DAOD_PHYS
 - Continued development of a more compact format for Run 4

CMS Preliminary Request for 2024

CMS		2022		2023			2024	
		C-RSG recomm.	Pledged	Request	2023 req. / 2022 C-RSG	C-RSG recomm.	Prelim Request	2024 req. / 2023 C-RSG
CPU	Tier-0	540	540	720	133%	720	750	104%
	Tier-1	730	730	800	110%	800	860	108%
	Tier-2	1200	1200	1350	113%	1350	1500	111%
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	2470	2470	2870	116%	2870	3110	108%
Others								
Disk	Tier-0	35.0	35.0	45.0	129%	45.0	52.0	116%
	Tier-1	83.0	83.0	98.0	118%	98.0	108.0	110%
	Tier-2	98.0	98.0	117.0	119%	117.0	130.0	111%
	Total	216.0	216.0	260.0	120%	260.0	290.0	112%
Tape	Tier-0	155.0	155.0	228.0	147%	228.0	293.0	129%
	Tier-1	260.0	260.0	316.0	122%	316.0	370.0	117%
	Total	415.0	415.0	544.0	131%	544.0	663.0	122%

- 2023 and 2024 activities focus on Run 3
 - Increasing use of nanoDST
 - Using legacy HLT system for offline
- 2024 requests driven by expectation of 100 fb⁻¹ dataset
 - Δ CPU = 8%
 - Δ Disk = 12%
 - Δ Tape = 22%
- Use of nanoAOD format reduces increase in disk requirements

LHCb Preliminary Request for 2024

LHCb		2022		2023			2024	
		C-RSG recomm.	Pledged	Request	2023 req. / 2022 C-RSG	C-RSG recomm.	Prelim Request	2024 req. / 2023 C-RSG
CPU	Tier-0	189	189	215	114%	215	344	160%
	Tier-1	622	515	707	114%	707	1128	160%
	Tier-2	345	333	391	113%	391	630	161%
	HLT	50	50	50	n/a	50	50	n/a
	Total	1206	1087	1363	113%	1363	2152	158%
Others								
Disk	Tier-0	26.6	26.5	30.3	114%	30.3	46.8	154%
	Tier-1	52.9	47.8	60.5	114%	60.5	93.6	155%
	Tier-2	10.2	6.9	11.6	114%	11.6	18.1	156%
	Total	89.7	81.2	102.4	114%	102.4	158.5	155%
Tape	Tier-0	81.0	81.0	91.0	112%	91.0	188.5	207%
	Tier-1	139.0	116.0	157.0	113%	157.0	199.7	127%
	Total	220.0	197.0	248.0	113%	248.0	388.2	157%

- 2023 and 2024 priorities
 - Commission new detector and computing system
 - Increased computing support
- 2024 request reflects expected increases
 - Δ CPU = 58%
 - Δ Disk = 55%
 - Δ Tape = 57%
- Resource under-pledges
 - Shortfall in 2023 (10-15% at T1 level)
 - If trend continues, 2024 might become problematic

Summary of Preliminary Requests for 2024

- Anticipate 2024 being a full-year of data-taking
- The preliminary requests reflect
 - Full year of Run 3 data collection, processing and analysis
 - Full implementation of new computing models for Alice and LHCb
- Continued work to develop new tools and optimize resource utilization
 - More compact data formats and faster simulation
 - Continuing to Innovate techniques for maximizing physics
 - Continued increased use of HPC opportunistic resources
- Contingency planning in the event of significant reduction in computing resources appears robust

ALICE Recommendations

- ALICE-1** The C-RSG suggests that for future Autumn scrutinies ALICE reduce the length of the report by focusing on the requests for future computing resources and a justification for these resources. The C-RSG recommends reducing the detail provided in Section 2 and excluding Section 3.
- ALICE-2** To evaluate the robustness of the predictions for 2024 resource requests, the C-RSG asks that for the next scrutiny ALICE provide a comparison of the predicted performance of their simulations (in terms of required CPU and disk) that were derived from Run-2 benchmarks with the actual performance that is being achieved using the O2 framework today.

ATLAS & CMS Recommendations

11

- ATLAS-1** The C-RSG recommends that ATLAS continue increasing the adoption of the more compact DAOD_PHYS and DAOD_PHYSLITE data formats for its physics data analyses.
- CMS-1** The C-RSG applauds CMS for their continuous efforts in making their software and computing environment more efficient in order to minimise their resource needs.
- CMS-2** The C-RSG recommends to the CMS collaboration that it investigate the economic trade-off between raw data compression CPU overhead versus the reduction of tape storage requirements at T0. This information would be helpful for understanding the overall benefits of such a data

LHCb Recommendations (I)

- LHCb-1** In view of the completion and commissioning of the upgraded LHCb detector, the C-RSG expects the experiment to reassess its 2024 resource estimates in the next scrutiny round.
- LHCb-2** The experiment reported that the full simulation of Run-3 events is currently significantly slower than the expected performance. The C-RSG asks LHCb to report ... on the actual performance of the Run-3 simulation ...
- LHCb-3** Simulation of pile-up collisions in the full simulation is currently done by generating every underlying minimum-bias collision together with every signal collision. Considering that expected CPU requirements are dominated by full MC simulation for Run-3 data, ... the C-RSG recommends LHCb explore alternative pile-up modelling strategies to reduce CPU requirements.

LHCb Recommendations (II)

13

- LHCb-4** In view of the projected increase in data storage, the C-RSG encourages LHCb to explore more reduced data formats for analysis as an avenue to significantly reduce the disk storage footprint.
- LHCb-5** Despite the progress the LHCb Collaboration has made in integrating and exploiting HPC opportunistic resources at various centres, the CPU time contributed by HPCs in 2024 is not planned to increase The C-RSG encourages LHCb to continue looking for increased allocations in HPC centres.
- LHCb-6** The C-RSG congratulates the LHCb Collaboration for the efforts made to address the shortage of personnel available for computing activities. The C-RSG reiterates the importance of having additional LHCb institutions contributing increased efforts to computing.

Overall Recommendations

14

- ALL-1** The C-RSG notes that there may be reductions in pledged resources due to power availability. It recommends that the WLCG collaboration identify the risks arising from this development and, where appropriate, develop contingency plans.
- ALL-2** Several experiments will be placing increased demands on T1 tape resources. The C-RSG recommends that an assessment be made of the effect on tape storage capability of another loss of a T1 facility.
- ALL-3** Given the shorter 2023 data-taking period, the C-RSG recommends all experiments review the effect this has on the estimated 2024 resources and report on any changes that arise from the LHC schedule adjustments in the Spring 2023 scrutiny.

Final Comments

15

- C-RSG thanks all collaborations for effective and collegial engagement
- Collaborations effectively using the 2022 computing resources
- There are significant uncertainties for 2023, but pledges remain appropriate
- The 2024 preliminary requests are robust, and will be refined over next 6 months
- Face-to-face scrutinies are much appreciated!