Computing Resources Scrutiny Group Report

Pekka K. Sinervo, C.M., FRSC

University of Toronto

For the Computing Resources Scrutiny Group

October 24, 2023

CERN-RRB-2023-115

C-RSG Membership

- C-RSG thanks the experiment representatives and CERN management for their support.
- C-RSG also thanks A Valassi and H Meinhard for their support of the scrutiny group.

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)	

Fall 2023 Scrutiny Process

- The collaborations gave updates on computing and data processing plans
  - Described computing activities for 2023 year
  - Updated computing plans for 2024 year
  - Initial resource estimates for 2025, based on LHC schedule
- Run 3 analysis and simulation production comprise bulk of computing activities
  - Using estimates of delivered integrated luminosity for 2024 and 2025
  - The computing requirements for 2023 met the needs of the experiments
    - Modified by shorter pp run and somewhat longer heavy-ion run
    - Uncertainties in running plans and WLCG computing resources being mitigated
- C-RSG believes that the 2025 estimates are a reasonable basis for further discussion and refinement

Resource Utilization in 2023 and 2024

- Available resources for 2023 meet the needs
  - Especially given the change in running schedule
- Approved resources for 2024 are sufficient to support approved physics program
- Continued availability of opportunistic CPU utilization by all experiments
  - Use of HPC resources continue to rise
  - That may be changing as some of the HPC resources are incorporated into pledges
- Note: new CPU benchmark
  (HEPSPEC-06 -> HEPSCORE-23)

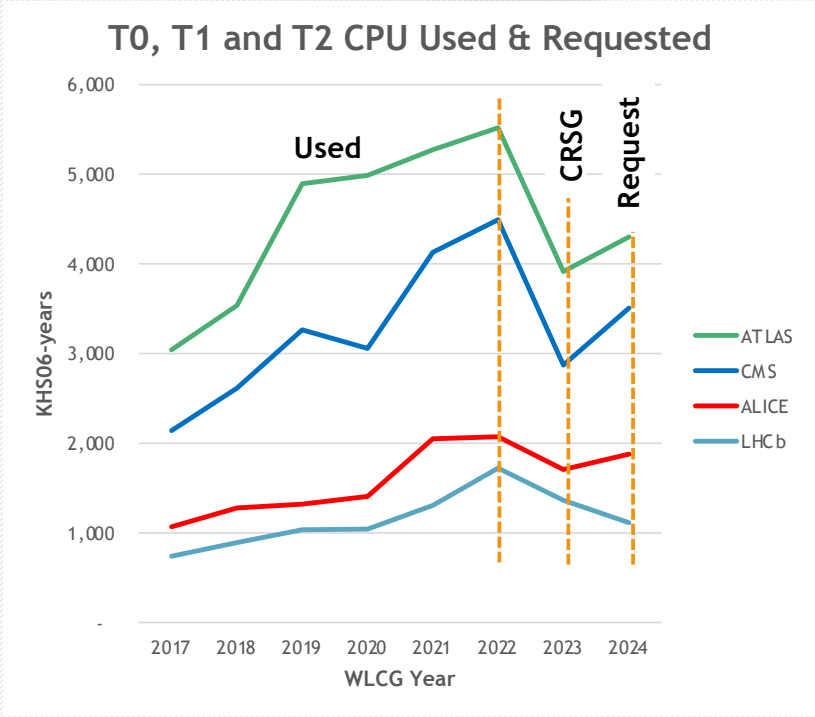

ALICE Preliminary Estimates for 2025

ALICE		2023		2024			2025	
		C-RSG recomm.	Pledged	Request	2024 req. / 2023 C-RSG	C-RSG recomm.	Prelim Request	2025 req. / 2024 C-RSG
CPU	Tier-0	541	541	600	111%	600	690	115%
	Tier-1	572	506	630	110%	630	725	115%
	Tier-2	592	567	650	110%	650	750	115%
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	**1705**	**1614**	**1880**	**110%**	**1880**	**2165**	**115%**
	Others							
Disk	Tier-0	58.5	58.5	67.5	115%	67.5	78.5	116%
	Tier-1	63.5	57.6	82.5	130%	71.5	82.5	115%
	Tier-2	57.5	60.4	77.5	135%	66.5	77.5	117%
	Total	**179.5**	**176.5**	**238.5**	**133%**	**205.5**	**238.5**	**116%**
	Others							
Tape	Tier-0	131.0	131.0	181.0	138%	181.0	226.0	125%
	Tier-1	82.0	87.7	107.0	130%	107.0	135.0	126%
	Total	**213.0**	**218.7**	**288.0**	**135%**	**288.0**	**361.0**	**125%**
	Others							

- 2024 and 2025 priorities
  - Extended PbPb and pPb runs
  - Two possible scenarios for pPb and PbPb running – comparable CPU requirements
  - Improving calibration procedures and data compression

- 2025 reflects expected growth
  - Δ CPU = 15%
  - Δ Disk = +16%
  - Δ Tape = +25%
- Growth driven by pp and PbPb running

ATLAS Preliminary Estimates for 2025

ATLAS		2023		2024			2025	
		C-RSG recomm.	Pledged	Request	2024 req. / 2023 C-RSG	C-RSG recomm.	Prelim Request	2025 req. / 2024 C-RSG
CPU	Tier-0	740	740	936	126%	936	1100	118%
	Tier-1	1430	1520	1516	106%	1516	1661	110%
	Tier-2	1747	1841	1852	106%	1852	2030	110%
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	**3917**	**4101**	**4304**	**110%**	**4304**	**4791**	**111%**
	Others							
Disk	Tier-0	40.0	40.0	49.0	123%	49.0	56.0	114%
	Tier-1	136.0	150.5	163.0	120%	163.0	186.0	114%
	Tier-2	168.0	160.0	200.0	119%	200.0	227.0	114%
	Total	**344.0**	**350.5**	**412.0**	**120%**	**412.0**	**469.0**	**114%**
Tape	Tier-0	174.0	174.0	207.0	119%	207.0	264.0	128%
	Tier-1	353.0	360.3	452.0	128%	452.0	567.0	125%
	Total	**527.0**	**534.3**	**659.0**	**125%**	**659.0**	**831.0**	**126%**

- 2024 and 2025 activities
  - Increase in use of smaller data formats
  - Raw event size increased by 20%
  - Anticipating greater use of HPC
    (NB: not all workflows run efficiently – or at all – on HPCs)

- 2025 estimates reflect full year of running
  - Δ CPU = +11%
  - Δ Disk = +14%
  - Δ Tape = +26%
- Smaller increases than previous years reflect increasing use of fast simulation

CMS Preliminary Estimates for 2025

CMS		2023		2024			2025	
		C-RSG recomm.	Pledged	Request	2024 req. / 2023 C-RSG	C-RSG recomm.	Prelim Request	2025 req. / 2024 C-RSG
CPU	Tier-0	720	720	980	136%	980	1180	120%
	Tier-1	800	916	930	116%	930	1100	118%
	Tier-2	1350	1313	1600	119%	1600	1900	119%
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	**2870**	**2949**	**3510**	**122%**	**3510**	**4180**	**119%**
	Others							
Disk	Tier-0	45.0	45.0	54.0	120%	54.0	64.0	119%
	Tier-1	98.0	96.8	122.0	124%	122.0	142.0	116%
	Tier-2	117.0	109.7	149.0	127%	149.0	175.0	117%
	Total	**260.0**	**251.5**	**325.0**	**125%**	**325.0**	**381.0**	**117%**
Tape	Tier-0	228.0	228.0	320.0	140%	320.0	420.0	131%
	Tier-1	316.0	303.7	380.0	120%	380.0	452.0	119%
	Total	**544.0**	**531.7**	**700.0**	**129%**	**700.0**	**872.0**	**125%**

- 2024 and 2025 activities focus on Run 3 and HL-LHC preparations
  - Run 3 data reprocessing
  - MC for Run 3 and HL-LHC

- 2025 estimates based on full year of running
  - Δ CPU = +19%
  - Δ Disk = +17%
  - Δ Tape = +25%

LHCb Preliminary Estimates for 2025

LHCb		2023		2024			2025	
		C-RSG recomm.	Pledged	Request	2024 req. / 2023 C-RSG	C-RSG recomm.	Prelim Request	2025 req. / 2024 C-RSG
CPU	Tier-0	215	215	174	81%	174	283	163%
	Tier-1	707	598	572	81%	572	928	162%
	Tier-2	391	434	319	82%	319	518	162%
	HLT	50	50	50	100%	50	50	n/a
	Total	**1363**	**1297**	**1115**	82%	**1115**	**1779**	**160%**
	Others			50			50	
Disk	Tier-0	30.3	30.3	30.6	101%	30.6	49.2	161%
	Tier-1	60.5	54.7	61.2	101%	61.2	98.5	161%
	Tier-2	11.6	7.9	11.8	102%	11.8	19.0	161%
	Total	**102.4**	**92.9**	**103.6**	101%	**103.6**	**166.7**	**161%**
Tape	Tier-0	91.0	91.0	117.1	129%	117.1	189.3	162%
	Tier-1	157.0	133.7	133.3	85%	133.3	208.1	156%
	Total	**248.0**	**224.7**	**250.4**	101%	**250.4**	**397.4**	**159%**

- 2024 and 2025 priorities
  - Full year of data taking with complete detector
  - First full year with new computing model
  - But available resources below 2023 pledges

- 2025 request reflects expected increase (shifted by 1 year)
  - Δ CPU = +60%
  - Δ Disk = +61%
  - Δ Tape =+59%
- Important that LHCb is provided these resources

Summary of Preliminary Estimates for 2025

- Based on a full year of data-taking in 2025
- The preliminary estimates of data processing needs appear reasonable
  - Last year of Run 3 data collection, with continued MC production and analysis
  - Continuing work to optimize use of resources
  - Compact data formats having increasingly beneficial effect
  - Continued increased use of HPC opportunistic resources
- Contingency planning in the event of significant reduction in computing resources appears robust
  - Work to mitigate uncertainties in T1 resources appears satisfactory

ALICE Observations and Recommendations

ALICE-1 The C-RSG requests that when different run scenarios are planned, ALICE Collaboration provide a high-level summary of the total amount of disk, tape and CPU for each run scenario. This will allow the C-RSG to better understand the potential resource requirements…

ALICE-2 When new processing, calibration or MC generation campaigns are undertaken by ALICE, the C-RSG requests that the ALICE Collaboration provide an assessment of the associated CPU, disk and tape requirements…

ALICE-3 The C-RSG requests that for the next scrutiny report the ALICE Collaboration provide an estimate of the contingency it holds….

ATLAS & CMS Observations and Recommendations

ATLAS-1 The C-RSG recommends that ATLAS continue increasing the adoption and usage monitoring of the more compact DAOD_PHYS and DAOD_PHYSLITE data formats for its physics data analyses.

ATLAS-2 The C-RSG is concerned by the 20% increase in RAW data event size. It encourages the ATLAS Collaboration to explore strategies to mitigate…

CMS-1 The C-RSG requests that the CMS Collaboration continue reporting on efforts to pinpoint & reduce the causes for the CPU efficiency.

CMS-2 The C-RSG recommends that CMS continue increasing the adoption and usage monitoring of the more compact NanoAOD data format for its physics data analyses.

LHCb Observations and Recommendations

LHCb-1 Considering that the CPU requirements are dominated by MC simulation for Run 3 data, the C-RSG encourages the LHCb Collaboration to explore strategies to reduce the simulation CPU footprint.

LHCb-2 A significant contributor to the CPU budget is user analysis. The C-RSG encourages the LHCb Collaboration to increase the utilization of the centrally managed analysis production system that has been shown to improve CPU efficiency.

LHCb-3 The physics data taken in 2023 during the pp and HI runs increase the storage requirements for 2025. These storage requests will need to be re-evaluated in the next scrutiny round in light of the actual physics data collected during 2023.

Collaborations have been very responsive

- The collaborations have been very responsive to previous recommendations and requests for additional information
  - Enabled in-depth and informed discussions of computing challenges
  - Some of the responses – such as CMS analysis of CPU efficiency – are very detailed and helpful
  - Deeply appreciated by scrutiny group
- C-RSG notes that all collaborations have been using resources efficiently
  - Observations and recommendations in report are made to support that effort

Summary

- C-RSG thanks the collaborations for their collegial engagement
- Collaborations efficiently using the 2023 computing resources
- The pledged 2024 resources are essential to pursue approved physics program
- The 2025 preliminary estimates appear robust, and will be refined over next 6 months