Computing Resources Scrutiny Group Report

Pekka K. Sinervo, C.M., FRSC

University of Toronto

For the Computing Resources Scrutiny Group

October 25, 2021

C-RSG membership

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

- C-RSG thanks the experiment representatives and to CERN management for their support.

Fall 2021 Scrutiny Process

- The four LHC experiments gave updates on computing and data processing plans,
  - Described computing activities for 2021 year (April 2021 – March 2022)
  - Updated estimates for 2022 year (April 2022 – March 2023)
  - Initial resource estimates for 2023
- Effect of COVID-19 pandemic on computing activities has been largely mitigated
  - Firm LHC Run 3 schedule assists the planning for 2022
  - Computing efforts have proceeded in line with spring 2021 plans
    - Continued Run 2 data processing and scientific analysis
    - Preparing for Run 3 with new algorithms, data formats and higher data rates
  - Computing needs for 2022 and 2023 have been adjusted due to LHC schedule
- C-RSG believes that the 2023 estimates are well-supported by requirements of the approved physics programs
  - The group recognizes its role as recommending the minimum necessary resource allocations for the LHC experiments to meet their approved physics programs
  - We work in a context where computer resources are constrained by funding capacity and the need to continue to minimize our computing carbon footprint

Resource Requirements for 2022 and Estimates for 2023

- 2022 will be first full year of Run 3 data-taking
  - Pledged resources appear adequate to support data-processing
- Computing model is changing for LHCb and ALICE
  - Evolution in resource requirements for 2022 onwards
  - Changes for 2023 are approximately in-line with expectations
- Strategy for delaying some resource deployment has not had a deleterious effect
  - Some concerns over changes in resources mid-year that are inconsistent with pledges

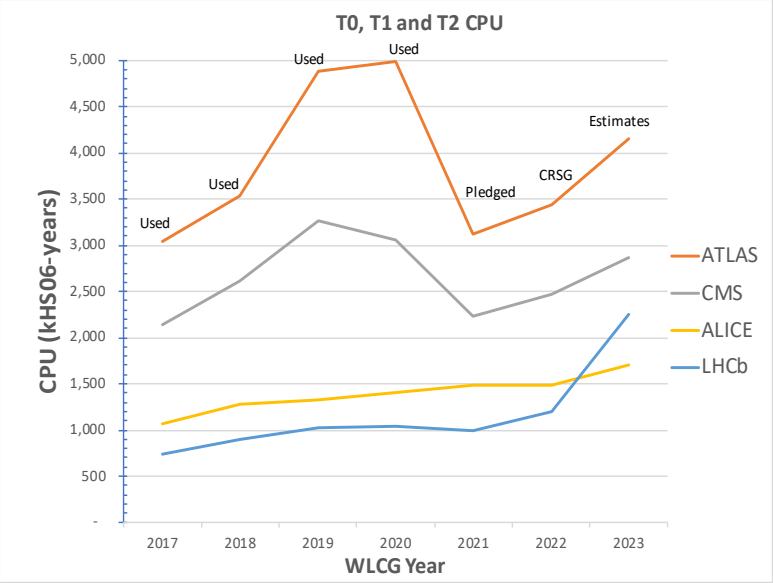

ALICE Computing Activities

Run 2 data processing dominates current activities

- MC production completed August 2021
- Now in Run 3 pp/PbPb MC production

AOD preparation through to Apr 2022

- Produce AODS in new format

Significant utilization of HPC resources

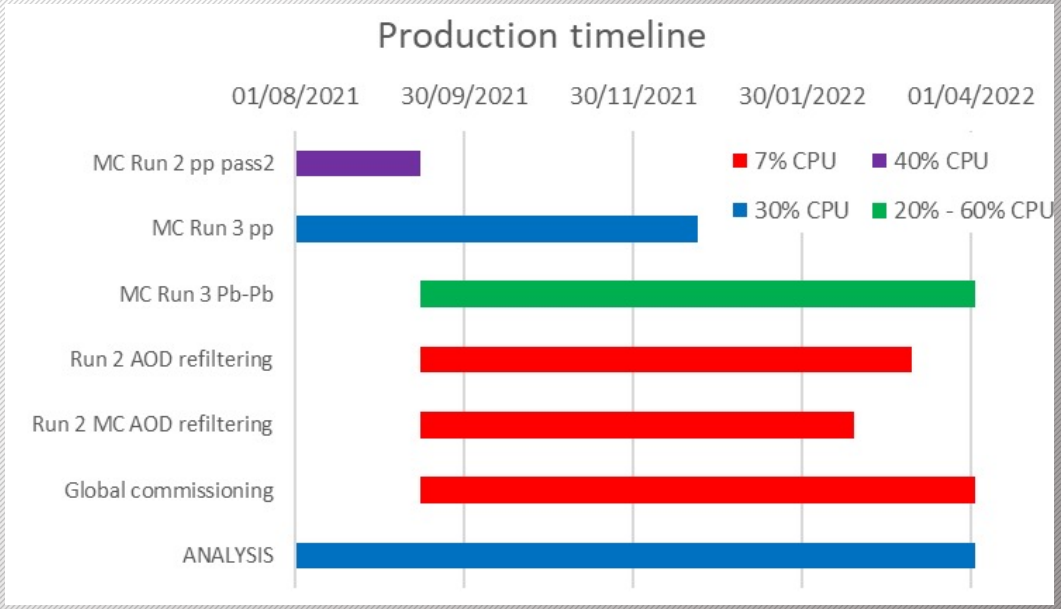

Alice Resources for 2022 and Estimates for 2023

ALICE		2021		2022			2023	
		C-RSG recomm.	Pledged	Request	2022 req. /2021 C-RSG	C-RSG recomm.	Preliminary Request	2023 req. /2022 C-RSG
CPU	Tier-0	471	471	471	100%	471	541	115%
	Tier-1	498	412	498	100%	498	572	115%
	Tier-2	515	481	515	100%	515	592	115%
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	**1484**	**1364**	**1484**	**100%**	**1484**	**1705**	**115%**
	Others							
Disk	Tier-0	45.5	45.5	50	110%	50.0	58.0	116%
	Tier-1	53.3	45.4	55	103%	55.0	63.0	115%
	Tier-2	44.8	50.4	49	109%	49.0	56.0	114%
	Total	**143.6**	**141.3**	**154**	**107%**	**154.0**	**177.0**	**115%**
Tape	Tier-0	86.0	86.0	95.0	110%	95.0	126.0	133%
	Tier-1	57.0	59.8	63.0	111%	63.0	79.0	125%
	Total	**143.0**	**145.8**	**158.0**	**110%**	**158.0**	**205.0**	**130%**

- No increase estimated for 2022 relative to 2021 C-RSG recommendations
  - All Pb-Pb and p-p running in 2022
  - Pb-Pb running is primary resource driver
- Focus for 2023 will be 2022 data-processing
  - 2 full reconstruction passes and simulations
  - Results in 44 PB of data on disk & tape
  - Underpledges may result in resource contention during pp running
  - Expect “critical point” to be Apr/May 2023
- N.B. O2 will be an opportunistic resource
  - Need to understand GPU accounting

ATLAS Computing Activities

- 2021 computing focused on Run 2 analysis
  - Continued Run 2 analyses
  - Production of intermediate datasets
  - Run 3 simulation and software framework
- Continued to
  - 4126 kHS06-y used vs 3265 kHS06-y pledged
  - Disk utilization 249 PB out of 272 PB pledged
  - Campaign to free up tape space shows up as a short-term under-utilization

Compute	Pledged	Tier-0 CPU used (pledged)	467 (525) kHS06
		Tier-1 CPU used (pledged)	1280 (1243) kHS06
		Tier-2 CPU used (pledged)	2379 (1497) kHS06
		CPU efficiency at Tier-1 and Tier-2	82%
	Opportunistic	HLT CPU used (to end of June)	829 kHS06
		HPC	704 kHS06
		HPC special	144 kHS06
		BOINC	385 kHS06
Storage		Tier-0 disk used (pledged)	26 (29) PB
		Tier-1 + Tier-2 disk used (pledged)	223 (243) PB
		Tier-0 tape used (pledged)	84 (95) PB
		Tier-1 tape used (pledged)	179 (241) PB

ATLAS Resources for 2022 and Estimate for 2023

ATLAS		2021		2022			2023	
		CRSG recomm.	Pledged	Request	2022 req. /2021 C-RSG	C-RSG recomm.	Preliminary Request	2023 req. /2022 C-RSG
CPU	Tier-0	525	525	550	105%	550	740	135%
	Tier-1	1170	1243	1356	116%	1300	1536	118%
	Tier-2	1430	1497	1656	116%	1588	1877	118%
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	**3125**	**3265**	**3562**	**114%**	**3438**	**4153**	**121%**
	Others							
Disk	Tier-0	29.0	29.0	32.0	110%	32.0	40.0	125%
	Tier-1	105.0	116.3	121.0	115%	116.0	136.0	117%
	Tier-2	130.0	127.2	148.0	114%	142.0	167.0	118%
	Total	**264.0**	**272.5**	**301.0**	**114%**	**290.0**	**343.0**	**118%**
Tape	Tier-0	95.0	95.0	120.0	126%	120.0	174.0	145%
	Tier-1	235.0	241.2	272.0	116%	272.0	353.0	130%
	Total	**330.0**	**336.2**	**392.0**	**119%**	**392.0**	**527.0**	**134%**

- 2022 driven by pp and PbPb data-taking
  - Prompt processing of newly collected data
  - Partial reprocessing of 2022 data
  - 25×10^9 MC events
    - Drives resource requirements
    - Plans for more fast simulation/new pileup model/code improvements
- 2023 resource estimates driven by Run 3
  - Expects to record 10×10^9 events
  - Will need about 30×10^9 MC events
    - 60% fast simulation
  - Need 10 PB of disk for PbPb running
- 2023 estimates are significant increases
  - CPU +21%, disk +18% and tape +34%

CMS Computing Activities

- The main 2021 computing activities for CMS are:
  - Preparation for Run 3
  - Analysis of Run 2 data continued at the same pace
  - Implementation of Rucio as for data management
  - Transition to CTA for tape management
  - More effective use of HPC resources
- Planned software upgrades underway
  - Improvements to detector-specific software
  - Implementing more efficient codes
  - Updating CMS software stack to latest releases of underlying codes (e.g. Root, Geant4)

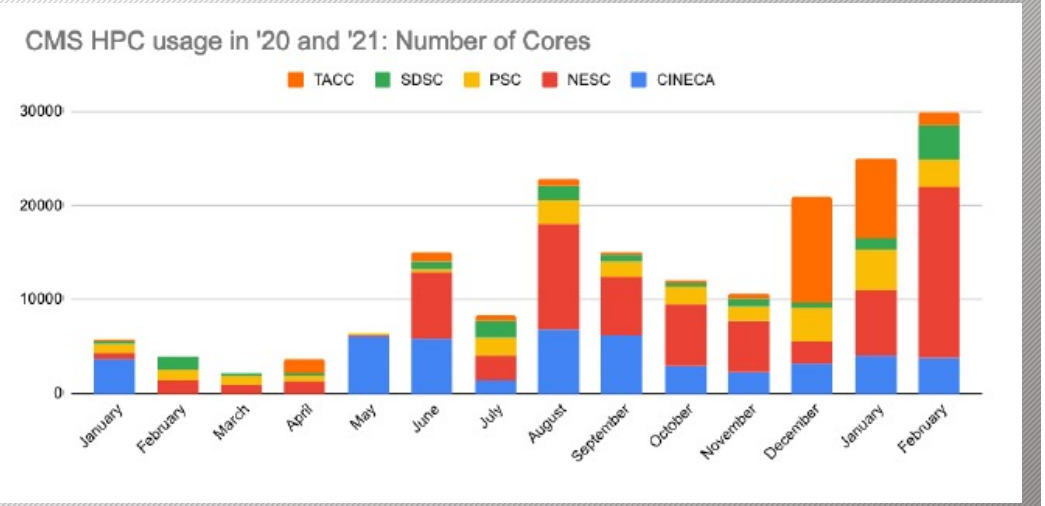

CMS Resources for 2022 and Estimates for 2023

CMS		2021		2022			2023	
		C-RSG recomm.	Pledged	Request	2022 req. /2021 C-RSG	C-RSG recomm.	Preliminary Request	2023 req. /2022 C-RSG
CPU	Tier-0	500	500	540	108%	540	720	133%
	Tier-1	670	764	730	109%	730	800	110%
	Tier-2	1070	1151	1200	112%	1200	1350	113%
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	**2240**	**2415**	**2470**	**110%**	**2470**	**2870**	**116%**
	Others							
Disk	Tier-0	30.0	30.0	35.0	117%	35.0	45.0	129%
	Tier-1	77.0	76.0	83.0	108%	83.0	98.0	118%
	Tier-2	92.0	96	98.0	107%	98.0	117.0	119%
	Total	**199.0**	**202**	**216.0**	**109%**	**216.0**	**260.0**	**120%**
Tape	Tier-0	120.0	120.0	155.0	129%	155.0	228.0	147%
	Tier-1	230.0	219.0	260.0	113%	260.0	316.0	122%
	Total	**350.0**	**339**	**415.0**	**119%**	**415.0**	**544.0**	**131%**

- 2022 requests driven by
  - Run 3 data-taking and simulations
  - HL-LHC work and Run 2 analysis
- 2023 increases are driven by Run 3 data-taking conditions and analysis
  - Higher pileup and integrated luminosity
  - Estimate 29×10^9 MC events for 100 fb^{-1}
  - ”Ambitious” HI run, requiring large min-bias sample
  - Implementing ”data parking” and “data scouting” strategies

LHCb Computing Activities

- Preparation for Run 3 and analysis of Run 2 dominate LHCb 2021 resource requirements
  - Simulation dominates CPU requirements
  - Delay in running has mitigated some of the shortfall in 2021 resources
  - Commissioning new data-taking and computing model

CPU Work in WLCG year (kHS06.years)	2021	TDR
First pass sprucing	70	100
End-of-year sprucing	70	100
Simulation	760	447
Core and distributed computing infrastructure	10	0
User Analysis and working group productions	260	216
Total Work (kHS06.years)	**1170**	**863**

Disk storage usage forecast (PB)		2021		TDR
Real data	Run1+Run2 pp data	15.2	37.8	36.3
	Run1+Run2 PbPb + SMOG	2.7		
	Run3: FULL	3.4		
	Run3: TURBO	7.6		
	Run3: TURCAL	0.9		
	Run3: Minimum bias	2.4		
	Run3: PbPb + SMOG2	5.6		
Simulated data	Run1+Run2 Simulated Data	8.7	10.0	10.0
	Run3 simulated data	1.3		
Other	User data	4.8	15.9	19.3
	Buffers	11.1		
Total			**63.7**	**65.7**

LHCb Resources for 2022 and Estimates for 2023

LHCb		2021		2022			2023	
		C-RSG recomm.	Pledged	Request	2022 req. /2021 C-RSG	C-RSG recomm.	Preliminary Request	2023 req. /2022 C-RSG
CPU	Tier-0	175	175	189	108%	189	361	191%
	Tier-1	574	470	622	108%	622	1185	191%
	Tier-2	321	292	345	107%	345	657	190%
	HLT	50	50	50	100%	50	50	100%
	Total	1120	987	1206	108%	1206	**2253**	**187%**
	Others		50	50			50	
Disk	Tier-0	18,8	18,8	26,5	141%	26,5	42,2	159%
	Tier-1	37,6	33,9	52,9	141%	52,9	84,4	160%
	Tier-2	7,3	6,1	10,2	140%	10,2	16,2	159%
	Total	63,7	58,8	89,6	141%	89,6	**142,8**	**159%**
Tape	Tier-0	43,8	43,8	81	185%	81,0	131,6	162%
	Tier-1	75,9	64,7	139	183%	139,0	227,7	164%
	Total	119,7	108,5	220	184%	220,0	**359,3**	**163%**

- 2022 usage driven by Run 3 data-taking
  - "Sprucing" of Run 3 data
  - Simulation for Run 2 and Run 3 physics
- 2023 resources needed for Run 3 data processing and simulation needs
  - Simulation takes 80% of CPU resources
  - CPU & disk doubling between 2021 and 2023, as predicted by TDR
  - 120 PB of tape for 2 copies of raw data
  - Computing support is area of concern

C-RSG Summary

- Overall picture for 2021 and 2022 is consistent with plans
  - Plan for flexible implementation of pledges in 2021 has been successful
  - Experiments have been able to continue Run 2 analysis while preparing for Run 3
- Resource requirements for 2023 coming into focus
  - Alice will be taking both PbPb and pp data with new computing model
    - GPUs playing increasing role
  - ATLAS and CMS anticipate significant increases in CPU/data driven by integrated luminosity
    - Continue to use significant opportunistic resources
  - LHCb has fast ramp – ambitious but credible
    - Personnel limitation remain a concern

Comments and Recommendations

ALL-1 The C-RSG thanks all four experiments for the responses to the Spring 2021 recommendations. It appreciates the constructive discussions … with the computing coordinators …

ALL-2 The C-RSG notes that computing resource allocations are usually made at the beginning of the year, providing a well-defined and committed resource for the subsequent 12 months. The group is concerned that reductions in the WLCG MOU commitment mid-year compromise the experiments' ability to complete their physics programs.

ALL-3 The use of GPUs as a source of T1 and T2 computing is increasing, but the benchmarks necessary to account for the GPU computing resources have not been established. The C-RSG recommends that a consistent interim approach to GPU accounting be established…

ALL-4 Given that the C-RSG sees a future where processor architectures may evolve, the group suggests the experiments continue to assess the long-term outlook for architectures alternative to the current X86 platform… and anticipate porting implications.

ALL-5 The C-RSG notes the increasing importance of networking in view of the changing computing models for the LHC experiments.. The group requests that experiments report on any risks to their computing plans from possible network constraints.