Computing Resources Scrutiny Group Report

Pekka K. Sinervo, C.M., FRSC

University of Toronto

For the Computing Resources Scrutiny Group

April 24-25, 2023

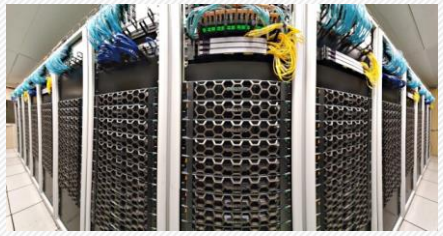

CERN-RRB-2023-013

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)	

- Have two new members nominated to the C-RSG who “shadowed” colleagues in this scrutiny:
  - Eric Fede (IN2P3) will replace Nadine Neyroud (France)
  - Markus Schulz (CERN) will replace Jan van Eldik (CERN)
- The C-RSG thanks Nadine and Jan for their significant contributions to the C-RSG.
- C-RSG thanks the experiment representatives and CERN management for their support.

Spring 2023 Scrutiny Process

- C-RSG met with LHCC WLCG referees and LHCC chair in advance
  - Identified a number of issues of common concern
- LHC experiments reported on
  - 2022 computing resource usage,
  - Computing activities and plans for 2023
  - Resource requests for 2024
- Computing activities in 2022 continued according to plans
- Revised 2023 and 2024 LHC have had impact on computing plans
  - Both ALICE and LHCb have had to make adjustments
  - ATLAS and CMS have only made modest changes

Resource Utilization in 2022

- 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 and 2024

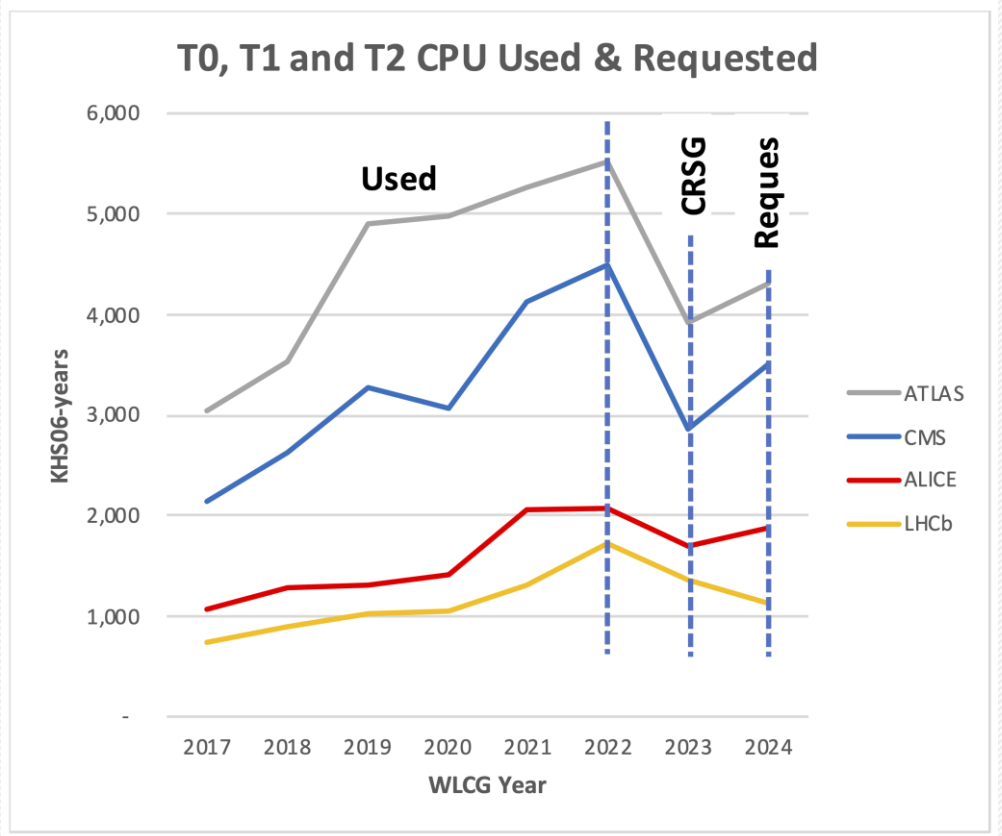

CPU Usage in 2022

- Total 2022 CPU utilization 12.6 MHS06-years
  - Roughly 1.2M cores used 24-7 through 2022
  - Half delivered by T2 sites, with T1 (30%) and T0 (20%)
- ALICE dominates CERN; ATLAS and CMS dominate T1 sites

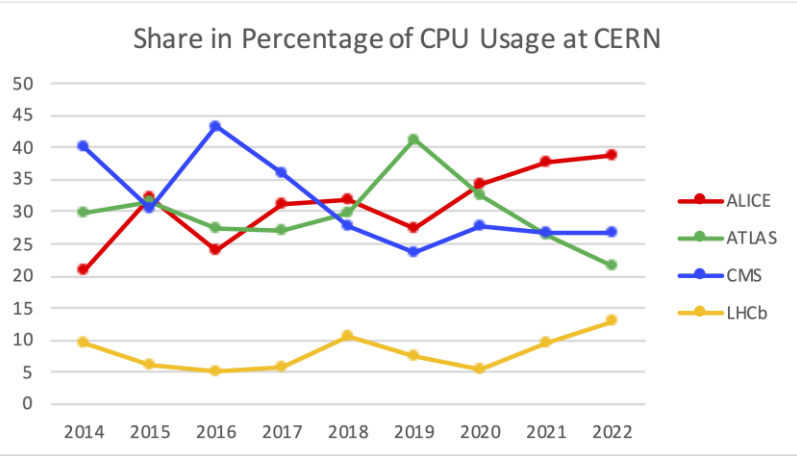

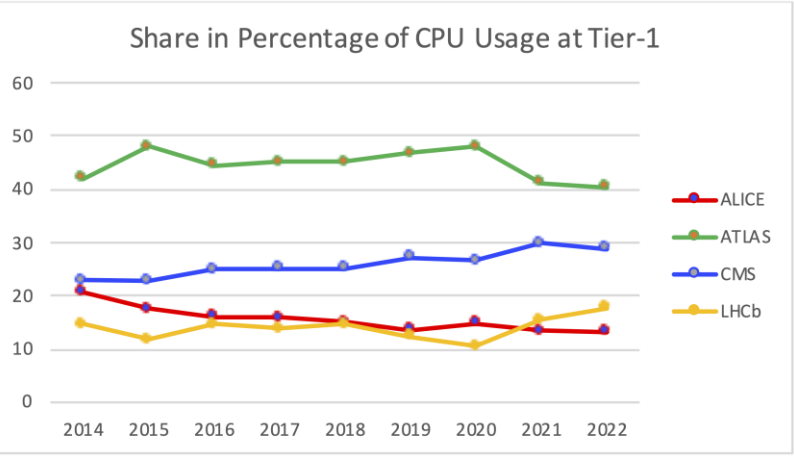

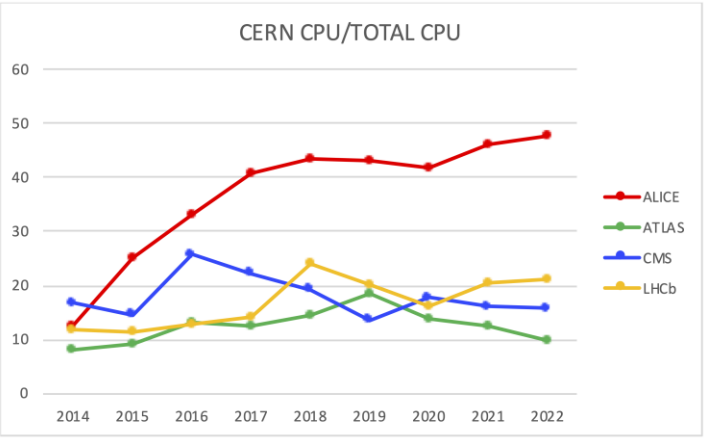

Disk Usage in 2022

- Disk usage tracked for CERN and T1 sites
- Trends from previous years have continued
- ALICE has increased share of CERN disk; ATLAS remains largest usage of T1 disk

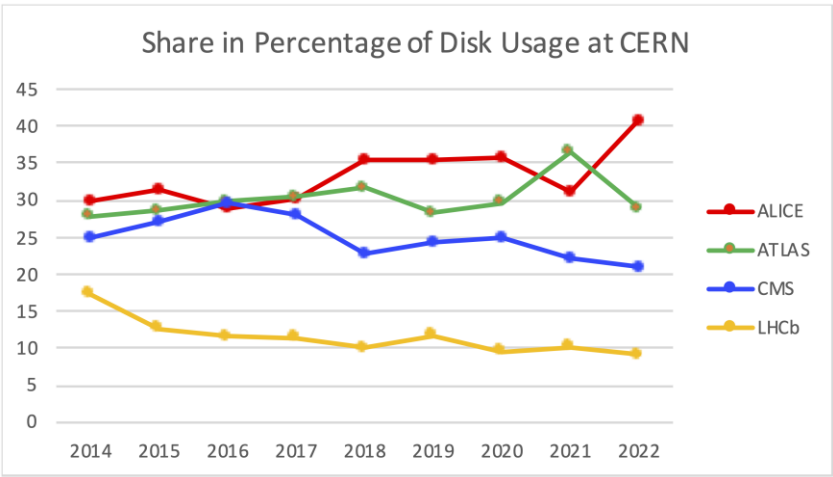

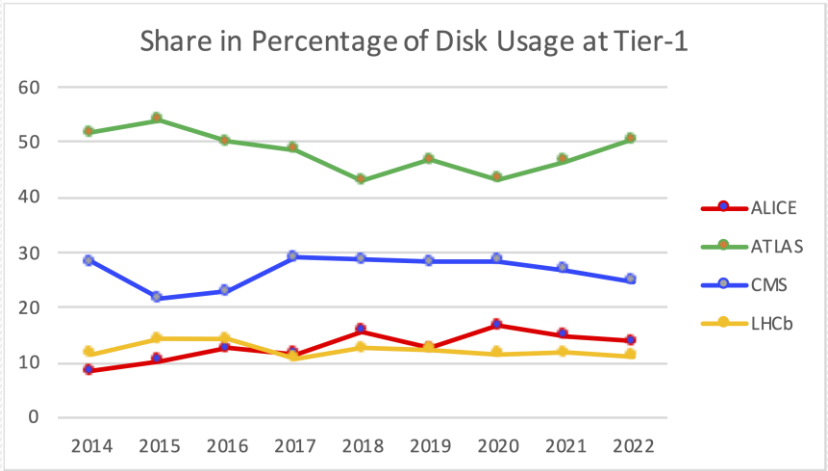

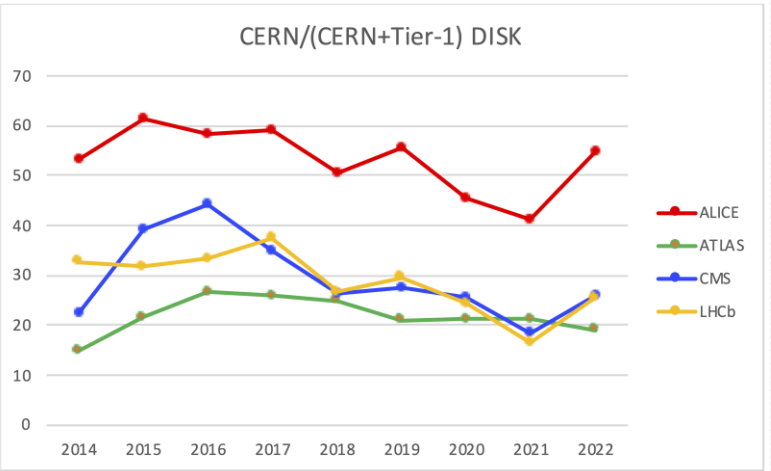

Tape Usage in 2022

- Total of 708 PB (down from 788 PB in 2021)
- Reflects archiving of raw and superseded physics datasets
- About 70-80% of allocated storage volume used

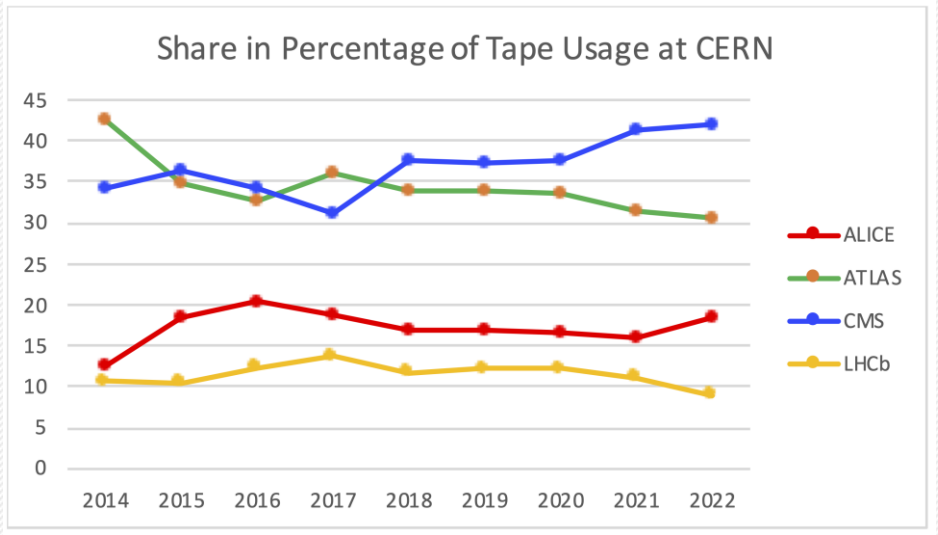

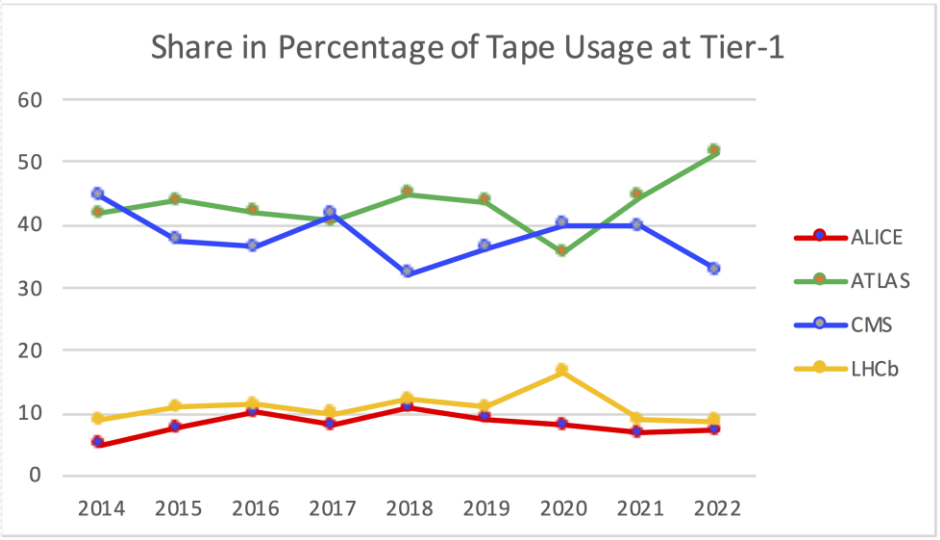

2023 Resource Utilization and 2024 Requests

- ”Big picture” view for 2023
  - Continuing data collection and processing
  - Increased focus on Run 3 physics analyses
  - ALICE expected to acquire large dataset
  - LHCb making best use of 2023 detector configuration
- “Big picture” view for 2024
  - Increases ~ 10% for CPU, ~18% for disk, ~25% for tape
  - ALICE, ATLAS and CMS large increases in disk and tape
    - Run 3 data-taking and physics analysis
  - LHCb expects no change from 2023

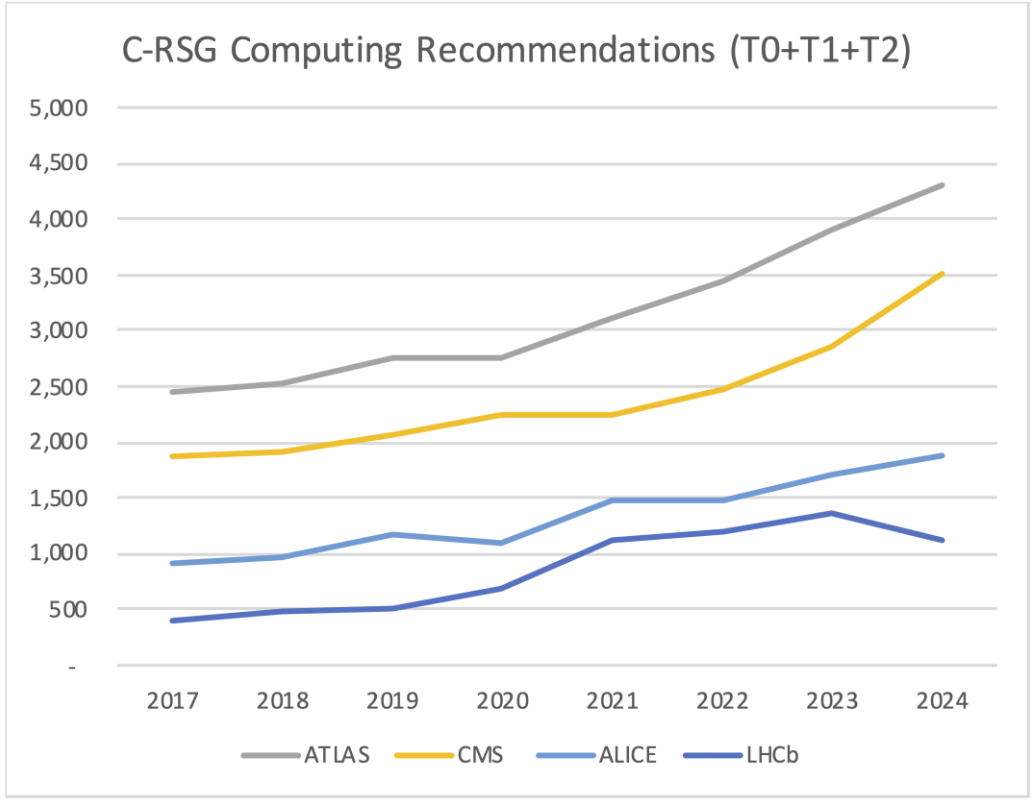

Alice Usage for 2022 and Request for 2024

ALICE		2022			2023		2024		
		C-RSG recomm.	Pledged	Used	C-RSG recomm.	Pledged	Request	2024 req. /2023 C-RSG	C-RSG recomm.
CPU	Tier-0	471	471	921	541	541	600	111%	600
	Tier-1	498	448	505	572	506	630	110%	630
	Tier-2	515	517	507	592	567	650	110%	650
	HLT	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Total	**1484**	**1436**	**1933**	**1705**	**1614**	**1880**	**110%**	**1880**
	Others			139					
Disk	Tier-0	50.0	50.0	46.6	58.5	58.5	67.5	115%	67.5
	Tier-1	55.0	49.7	38.3	63.5	57.6	71.5	113%	71.5
	Tier-2	49.0	55.2	40.3	57.5	60.4	66.5	116%	66.5
	Total	**154.0**	**154.9**	**125.2**	**179.5**	**176.5**	**205.5**	**114%**	**205.5**
Tape	Tier-0	95.0	95.0	61.4	131.0	131.0	181.0	138%	181.0
	Tier-1	63.0	71.8	39.5	82.0	87.7	107.0	130%	107.0
	Total	**158.0**	**166.8**	**100.9**	**213.0**	**218.7**	**288.0**	**135%**	**288.0**

- 2022 and 2023 priorities
  - Analyze 2022 pp data
  - Take and analyze larger PbPb data sample in 2023

- 2024 reflects expected growth
  - Accommodate pp, PbPb and pPb data
- Will still be working to implement more aggressive compression
- Increase in MC/data by x4

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

ATLAS Usage for 2022 and Request for 2024

ATLAS		2022			2023		2024		
		C-RSG recomm.	Pledged	Used	C-RSG recomm.	Pledged	Request	2024 req. /2023 C-RSG	C-RSG recomm.
CPU	Tier-0	740	544	512	740	740	936	126%	936
	Tier-1	1300	1349	1545	1430	1520	1516	106%	1516
	Tier-2	1588	1616	3151	1747	1841	1852	106%	1852
	HLT	n/a	n/a	311	n/a	n/a	n/a	n/a	n/a
	Total	**3628**	**3509**	**5519**	**3917**	**4101**	**4304**	**110%**	**4304**
	Others			2161					
Disk	Tier-0	32.0	32.0	33.0	40.0	40.0	49.0	123%	49.0
	Tier-1	116.0	130.0	139.0	136.0	150.5	163.0	120%	163.0
	Tier-2	142.0	142.0	141.0	168.0	160.0	200.0	119%	200.0
	Total	**290.0**	**304.0**	**313.0**	**344.0**	**350.5**	**412.0**	**120%**	**412.0**
Tape	Tier-0	120.0	120.0	102.0	174.0	174.0	207.0	119%	207.0
	Tier-1	272.0	280.0	281.0	353.0	360.3	452.0	128%	452.0
	Total	**392.0**	**400.0**	**383.0**	**527.0**	**534.3**	**659.0**	**125%**	**659.0**

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

- 2022 and 2023 has Run 3 focus
  - Higher pileup -> 15% CPU increase
  - Moving to smaller data formats

- 2024 requests reflect full year of running to collect 110 fb^{-1} of data
  - CPU increase reflects improved simulation performance
  - Doesn't account for HI run -- OK
- Porting to other architectures
  - Both ARM and GPUs
  - Continued use of opportunistic resources
- HL-LHC work remains priority

CMS Usage for 2022 and Request for 2024

CMS		2022			2023		2024		
		C-RSG recomm.	Pledged	Used	C-RSG recomm.	Pledged	Request	2024 req. /2023 C-RSG	C-RSG recomm.
CPU	Tier-0	540	540	638	720	720	980	136%	980
	Tier-1	730	730	1107	800	916	930	116%	930
	Tier-2	1200	1200	2302	1350	1313	1600	119%	1600
	HLT	n/a	n/a	347	n/a	n/a	n/a	n/a	n/a
	Total	**2470**	**2470**	4394	**2870**	**2949**	**3510**	**122%**	**3510**
	Others			649					
Disk	Tier-0	35.0	35.0	24.0	45.0	45.0	54.0	120%	54.0
	Tier-1	83.0	83.0	68.0	98.0	96.8	122.0	124%	122.0
	Tier-2	98.0	98.0	87.0	117.0	109.7	149.0	127%	149.0
	Total	**216.0**	**216.0**	**179.0**	**260.0**	**251.5**	**325.0**	**125%**	**325.0**
Tape	Tier-0	155.0	155.0	140.0	228.0	228.0	320.0	140%	320.0
	Tier-1	260.0	260.0	179.0	316.0	303.7	380.0	120%	380.0
	Total	**415.0**	**415.0**	**319.0**	**544.0**	**531.7**	**700.0**	**129%**	**700.0**

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

- 2022 and 2023 focused on Run 3
  - Increased T0 resources due to higher pileup
  - Continued migration to nano-DST

- 2024 requests reflect higher pileup and larger dataset
- T0 increase result of larger data set and higher
- Increased effort on HL-LHC simulation

LHCb Usage for 2022 and Request for 2024

LHCb		2022			2023		2024		
		C-RSG recomm.	Pledged	Used	C-RSG recomm.	Pledged	Request	2024 req. /2023 C-RSG	C-RSG recomm.
CPU	Tier-0	189	189	305	215	215	174	81%	174
	Tier-1	622	515	676	707	598	572	81%	572
	Tier-2	345	333	470	391	434	319	82%	319
	HLT	50	50	271	50	50	50	n/a	50
	Total	**1206**	**1087**	**1722**	**1363**	**1297**	**1115**	**82%**	**1115**
	Others			53					
Disk	Tier-0	26.5	26.5	10.5	30.3	30.3	30.6	101%	30.6
	Tier-1	52.9	47.8	30.6	60.5	54.7	61.2	101%	61.2
	Tier-2	10.2	6.9	4.0	11.6	7.9	11.8	102%	11.8
	Total	**89.6**	**81.2**	**45.1**	**102.4**	**92.9**	**103.6**	**101%**	**103.6**
Tape	Tier-0	81.0	81.0	29.8	91.0	91.0	117.1	129%	117.1
	Tier-1	139.0	116.0	47.1	157.0	133.7	133.3	85%	133.3
	Total	**220.0**	**197.0**	**76.9**	**248.0**	**224.7**	**250.4**	**101%**	**250.4**

- Simulation dominates LHCb 2022 and 2023 CPU usage
  - Disk needed for data sprucing and analysis
  - Reduced data set in 2023

- 2024 request reduced due to status of detector
  - Data used for commissioning and some physics studies
- Working to improve simulation in high-pileup environment
- Developing new T1 sites

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

Summary of Resource Requests for 2024

- Expect 2024 to be a full-year of data-taking
  - Increases in computing requirements driven by data volume
  - Simulation scales comparably to data
- The requested resources reflect increased Run 3 data volume
  - ALICE requests driven by HI running plans
  - LHCb requests reflect expected detector status in 2023
- Effort continues to optimize computing resource utilization
  - Effort continues to reduce analysis CPU and event size
  - More efficient use of multi-core nodes
  - Porting to new CPU architectures (e.g., ARM, GPU)
- Following recommendations are synopsis of those in report

ALICE Recommendations

ALICE-1 The C-RSG recommends that the ALICE Collaboration's resource requests for 2024 be approved.

ALICE-2 ALICE Collaboration is encouraged to focus on CTF compression strategies.

ALICE-3 The number of MC events that will be generated by ALICE has increased by a factor of four increasing the MC/event. More detailed justification requested.

ALICE-4 The C-RSG requests that ALICE further encourages their users to migrate to the new AOD format to take advantage of their faster throughput.

ALICE-5 ALICE is encouraged to increase its opportunistic resources beyond T0… other Tiers and on HPCs.

ATLAS Recommendations

ATLAS-1 The C-RSG recommends that the ATLAS Collaboration's resource requests for 2024 be approved.

ATLAS-2 The C-RSG would like to understand whether increased use of the DAOD_PHYS and DAOD_PHYSLITE data formats would influence future Grid resource requests, specifically disk space.

CMS Recommendations

CMS-1 The C-RSG recommends that the CMS Collaboration's resource requests for 2024 be granted.

CMS-2 The C-RSG would like to understand if the increased popularity of the nano-AOD datasets will impact the Grid resource requests, in particular disk space.

CMS-3 The C-RSG recognises the effort to increase disk utilization on the T0 and T1 sites, and encourages CMS to raise this to the 90% limit considered optimal.

CMS-4 The C-RSG requests that the CMS Collaboration continue reporting on efforts to reduce the causes behind the relatively low 2022 CPU efficiency.

LHCb Recommendations

LHCb-1 The LHCb Collaboration requests no increase in resources for 2024. However, less than 90% of the T1 and T2 disk approved by the C-RRB for 2023 have been pledged. The C-RSG encourages the FAs to provide the LHCb collaboration with these resources.

LHCb-2 The C-RSG notes that the zero growth of LHCb computing resources in 2024 will be followed by almost 100% growth in 2025. The experiment considers it acceptable to stagger this growth over 2024.

LHCb-3 The C-RSG recommends that LHCb continue work to decrease CPU consumption in simulations and to reduce storage demands.

Overall Recommendations

ALL-1 The C-RSG requests that the collaborations report high-level summaries of the manner in which their disk space is utilised and optimized… and should identify the space used for:

i. persistent datasets (differentiating between primary copies and additional replicas),
ii. cached storage used to hold datasets for short periods of time, and
iii. buffering space used, for example, to transfer data from one storage media to another.

This information will assist the C-RSG in better understanding the pressures on disk utilization and help identify best practice.

ALL-2 The C-RSG recognises the significant efforts … to identify additional T1 sites. The addition of such sites is essential to mitigate the effects on the physics programmes of the loss of any existing T1 facility.

Some Final Comments

- C-RSG thanks all collaborations for effective and collegial engagement
- Potential loss of Russian and Belarusian resources challenging,
  but being mitigated in short term and possibly longer term
- The support of the funding agencies in meeting these resource needs is greatly appreciated by the collaborations and the C-RSG