LHCb computing activity and 2022/2023 requests

Stefano Perazzini
INFN Bologna

Riunione referaggio calcolo – 9 Settembre 2021

Outline

- Usage of resources in 2020/2021
- News and current activity
- Reassessment of requests for 2022
- Look to projections for 2023 and 2024

Summary of computing activity in 2020/2021

- As usual CPU consumption has been dominated by MC productions
- Full reprocessing + stripping of particular dataset during 2020
  - 2015 proton collisions, in January 2020
  - 2016 proton collisions, in February 2020
  - 2017 proton-Neon (beam-gas) collisions, in April 2020
  - 2018 Pb-Pb collisions, in May 2020
  - 2018 Pb-Ar collisions, in June 2020

Require data staging from tape

- New restripping campaigns of p-p Run2 collisions during 2021
- Nevertheless users' jobs represent a non-negligible fraction of total activity
  - Complemented by centralised productions running users' code

Summary of computing activity in 2020/2021

- As usual CP ctions
- Full reproce
  - 2015 proton col
  - 2016 proton col
  - 2017 proton-Ne
  - 2018 Pb-Pb coll
  - 2018 Pb-Ar coll
- New restrip
- Neverthele f total activity
  – Compleme

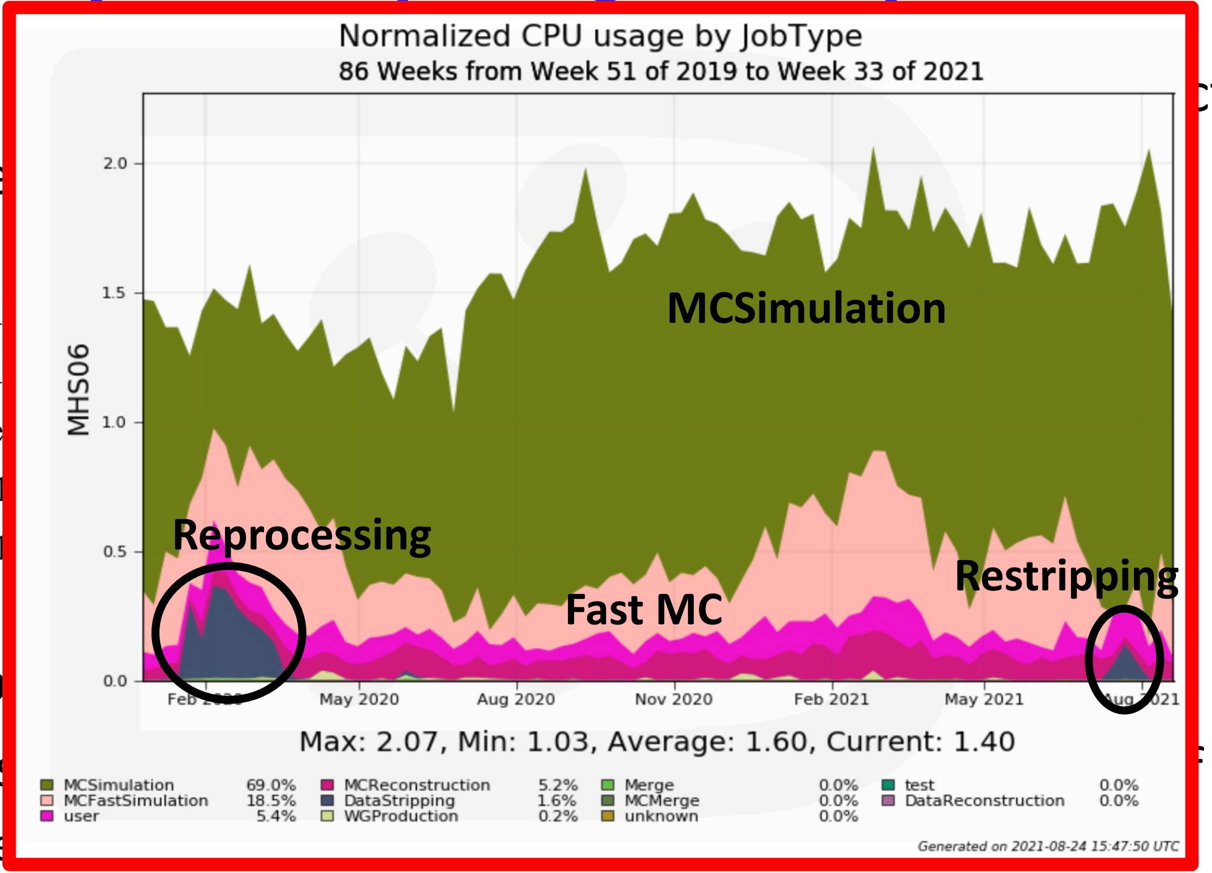

CPU usage

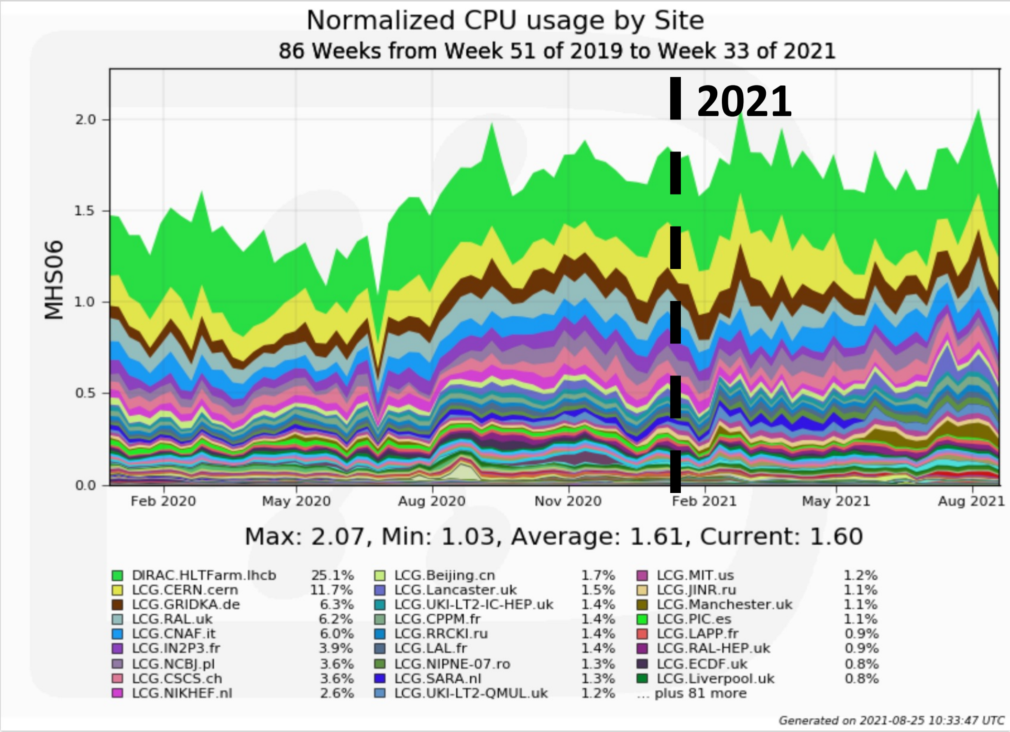

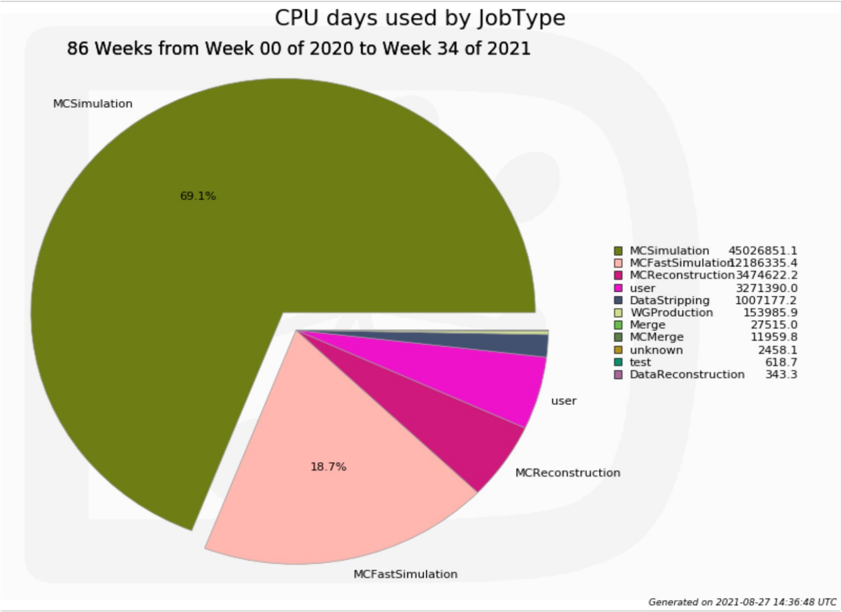

- About 90% of CPU power is used for MCSimulation
- About 25% of CPU power used comes from HLTFarm
  - **Not available during 2022**

WLCG T0/T1 accounting for 2020

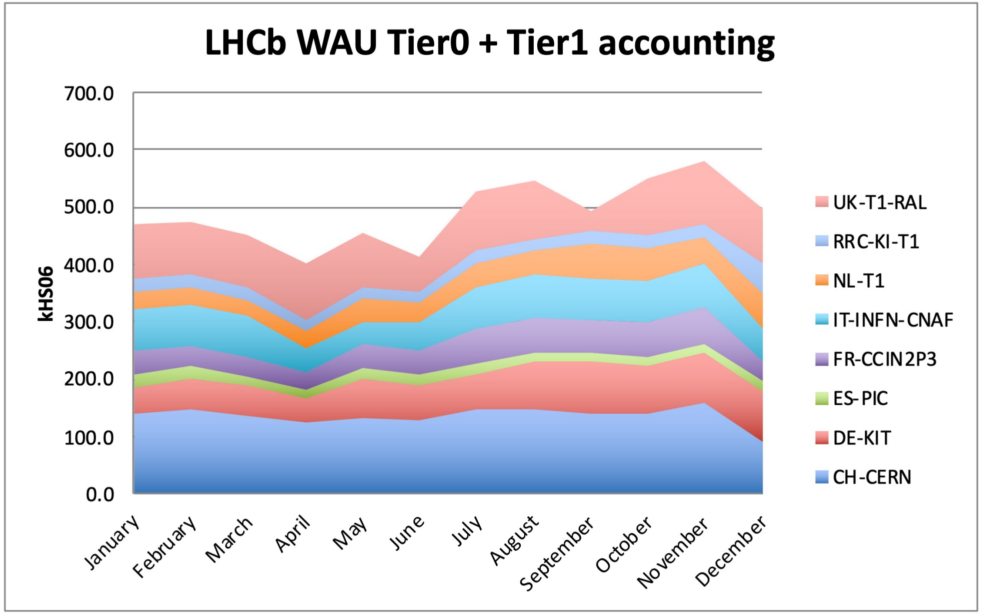

LHCb-PUB-2021-003

<Power>	Used (kHS06)	Pledge (kHS06)
CH-CERN	136.0	98
DE-KIT	68.1	54.8
ES-PIC	17.6	13.1
FR-CCIN2P3	46.5	47.2
IT-INFN-CNAF	64.8	55.8
NL-T1	42.3	26.2
RRC-KI-T1	23.5	16.4
UK-T1-RAL	89.0	81.3
Total	**487.8**	**392.8**

WLCG T2 accounting for 2020

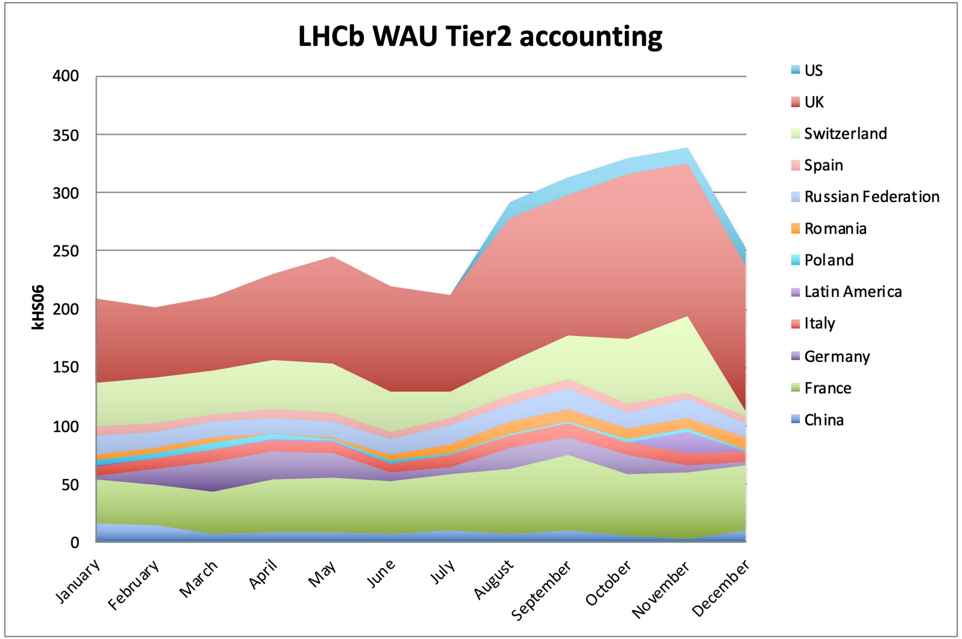

LHCb-PUB-2021-003

<Power>	Used (kHS06)	Pledge (kHS06)
China	9.1	18.0
Brazil	2.0	1.0
France	48.4	30.1
Germany	12.9	10.6
Israel	0.4	0
Italy	9.9	31.5
Poland	2.6	7.4
Romania	6.6	6.9
Russia	14.9	18.2
Spain	6.5	7.0
Switzerland	37.5	32.0
UK	98.1	31.6
US (*)	14.0	12.0
Total	**262.8**	**206.3**

- NOTE: most of CPU provided by Italian T2s is actually accounted in T1 (LHCb-T2 is at CNAF)

Italian CPU resources

- According to WLCG, sum of INFN-T1 + Italy-T2 is **slightly below pledge** (~10-15%)
  - Problems in finding proper settings after migration to HTCondor
  - Operational problems from LHCb
    - Misconfigurations in LHCbDirac settings
  - Delay in installation of new resources

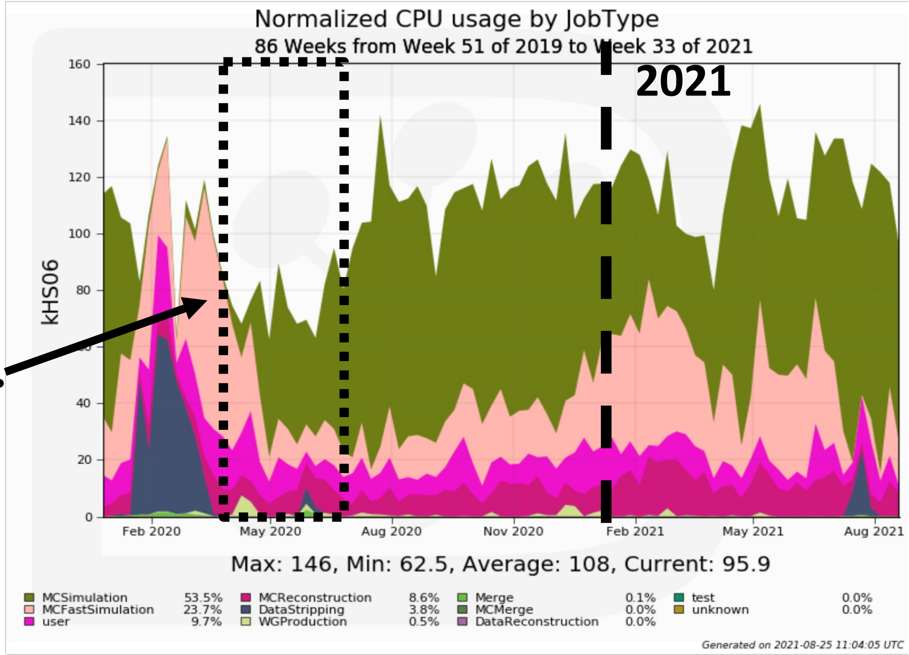

- About 90% of CPU provided by Italian sites comes from CNAF
  - Only site pledging resources for LHCb

Situation at CNAF in 2021 so far

1st April 2021

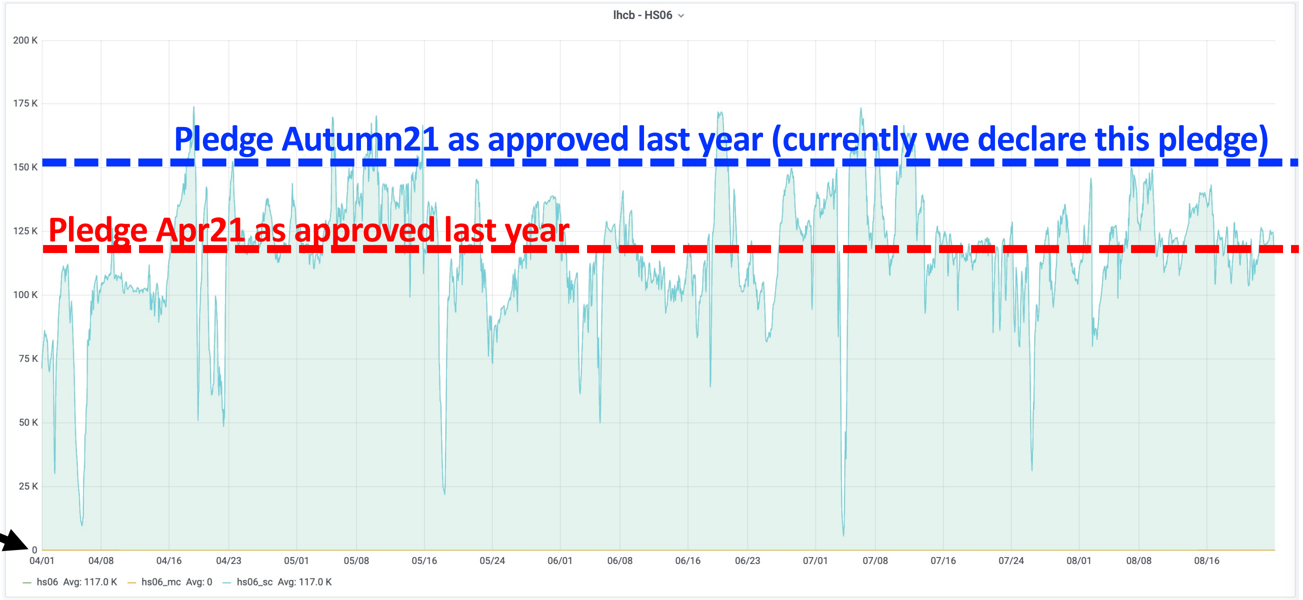

- Since April 2021 CNAF(T1+T2) provides an average of about 117 kHS06
  - Just slightly below pledged for April2021 agreed last year (~120 kHS06)
  - Nominal partition of **used resources is 77 kHS06 fot T1 and 43 kHS06 for T2**
- As approved last year, in autumn CNAF should reach about 152 kHS06 with the installation of SJ resources

Job efficiency

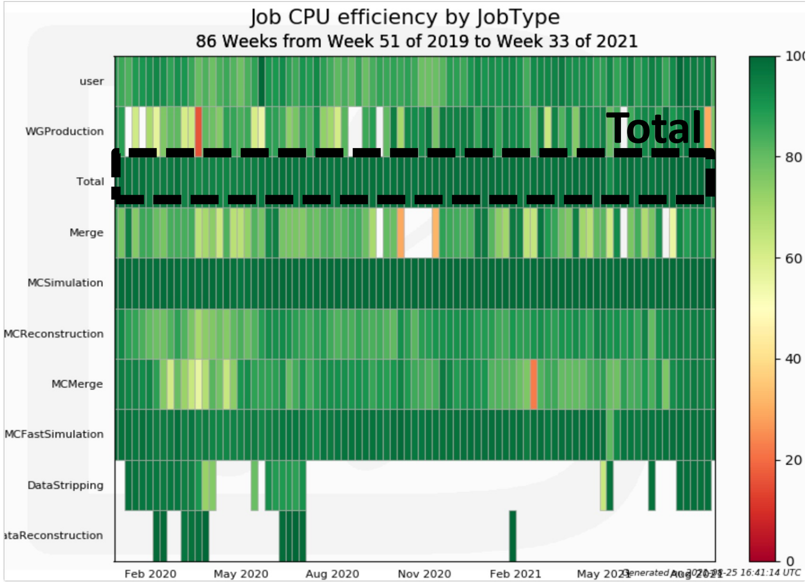

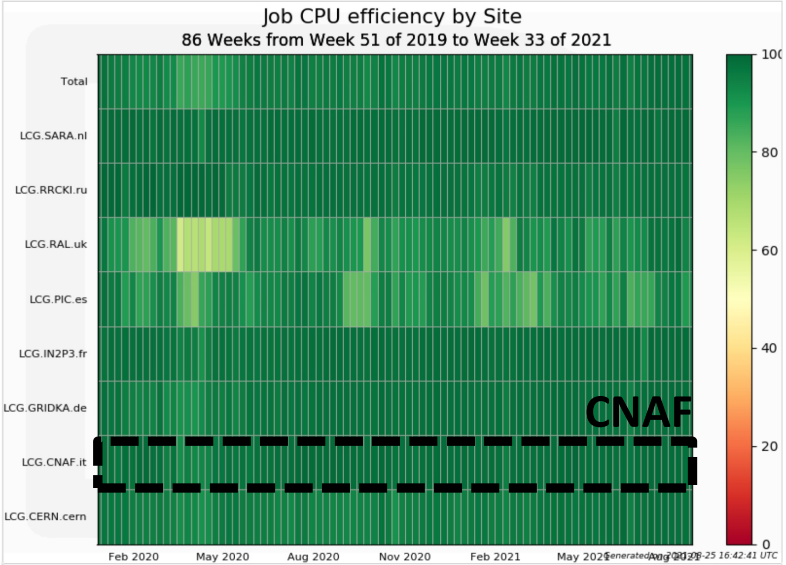

- Very good usage of CPU
  - CPU inefficiencies concentrated in jobs with large I/O → small fraction of total CPU
  - Excellent performance from CNAF

Job success rate

- Very high success rate
  - **Almost 97% of CPU time is used by successful jobs**
- Most of failed jobs are “stalled”
  - Killed by batch system because out of queue limits
  - Wrong determination of time needed for completion
- **About 30% of failed jobs are users’ jobs ➔ users’ errors**

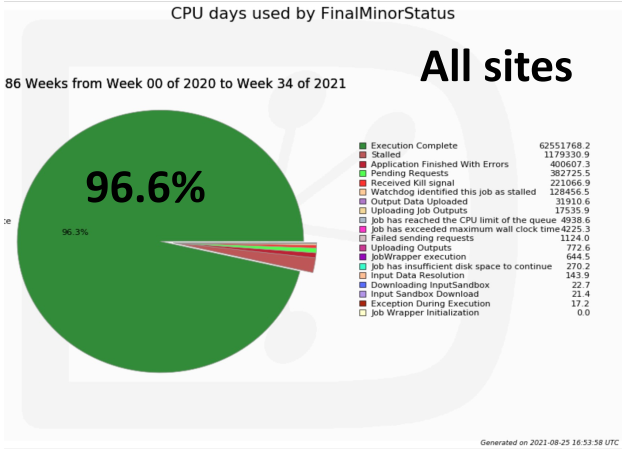

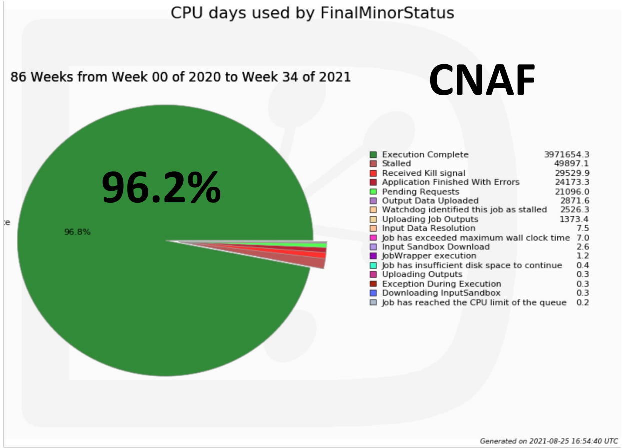

Tape usage 2020

- During 2020
  - **Archived 2.9 PB** of old MC productions and RDST from reprocessing campaigns
- Total 74.6 PB on tape (31st December 2020)
  - 38.4 PB (RAW) + 13.7 PB (RDST) + 22.5 PB (ARCHIVE)
  - Lower than pledge → in part recovered during first part of 2021
- Total of 10.7 PB of tape at CNAF → **19% of Tier1s** (for LHCb policy should be 24%)
  - 3.4 PB (RAW) + 2.6 PB (RDST) + 4.7 PB(ARCHIVE)

LHCb Pledges for 2020

2020	CPU (kHS06)	Disk (PB)	Tape (PB)
Tier 0	98	17.2	36.1
Tier 1	295	31.7	56.0
Tier 2	206	4.3	
Total WLCG	**599**	**53.2**	**92.1**
Non WLCG	**20**		

LHCb-PUB-2021-003

Tape usage 2021 so far

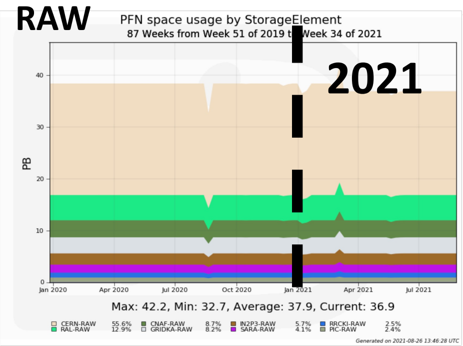

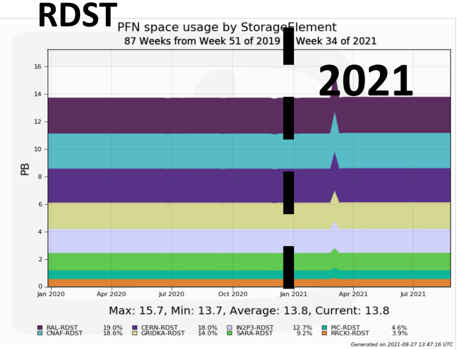

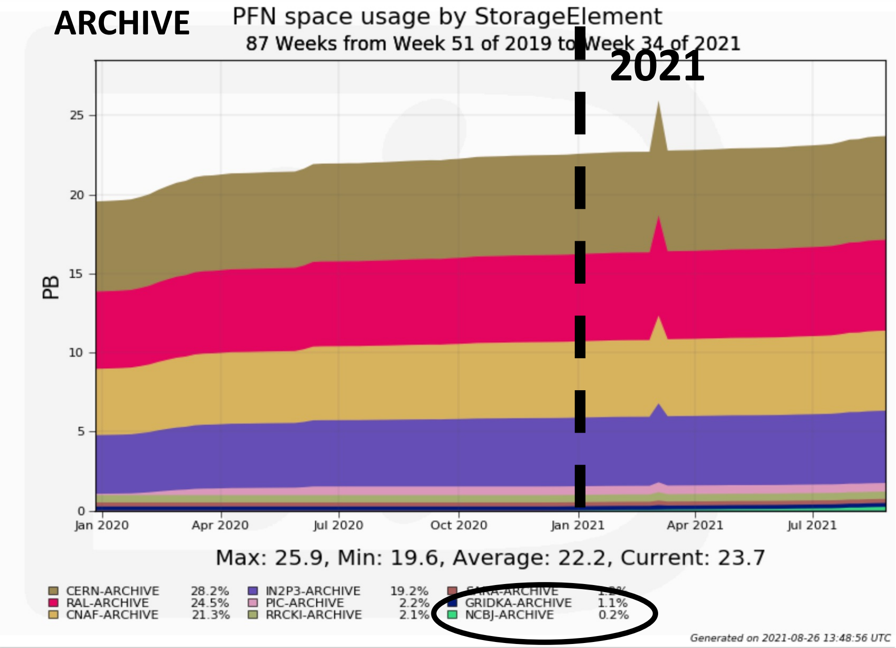

- RAW and RDST are stable
  - No data taking and no reprocessing
- Steady increase for ARCHIVE
  - **Additional 1.2 PB during first half of 2021 and about a factor 3 by the end of the year** to clean space for Run1+Run2 incremental strippings
- **Start exploiting some tape resources in Poland (NCBJ) for archival**
  - Just a "pilot test" (about 50 TB for now) but discussion is ongoing to increase this quota and have a substantial pledge of tape resources in the coming years **(not yet rubber stamped)**

Disk usage T0+T1 in 2020

Disk (PB)	CERN	Tier1s	CNAF	GRIDKA	IN2P3 (*)	PIC	RAL	RRCKI	SARA
LHCb accounting	**7.91**	**22.62**	**4.65**	**3.87**	**3.35**	**1.25**	**5.61**	**1.88**	**2.01**
WSSA disk used	8.03	23.56	4.69	3.88	3.36	1.25	6.46	1.89	2.03
WSSA disk free	0.37	7.53	1.81	1.35	1.44	0.09	1.91	0.46	0.48
WSSA tape buffer (used+free)[3]	(**)	2.65	2.30	0.15	0.03	0.05	(**)	0.09	0.03
WSSA total[4]	**8.40**	**33.75**	**8.80**	**5.38**	**4.83**	**1.39**	**8.37**	**2.44**	**2.54**
Pledge '20	*17.2*	*31.7*	*6.87*	*5.55*	*4.65*	*1.33*	*8.37*	*2.30*	*2.645*

**: CERN and RAL do not publish buffer information

LHCb-PUB-2021-003

- About **8 PB of disk free at the end of 2020**
  - **Partially filled** with data reprocessing at beginning of 2021
  - Still to come new data from incremental stripping campaigns of Run1+Run2
  - **Underusage of disk resources mainly driven by delay in LHC schedule**

Disk usage T0+T1 in 2021

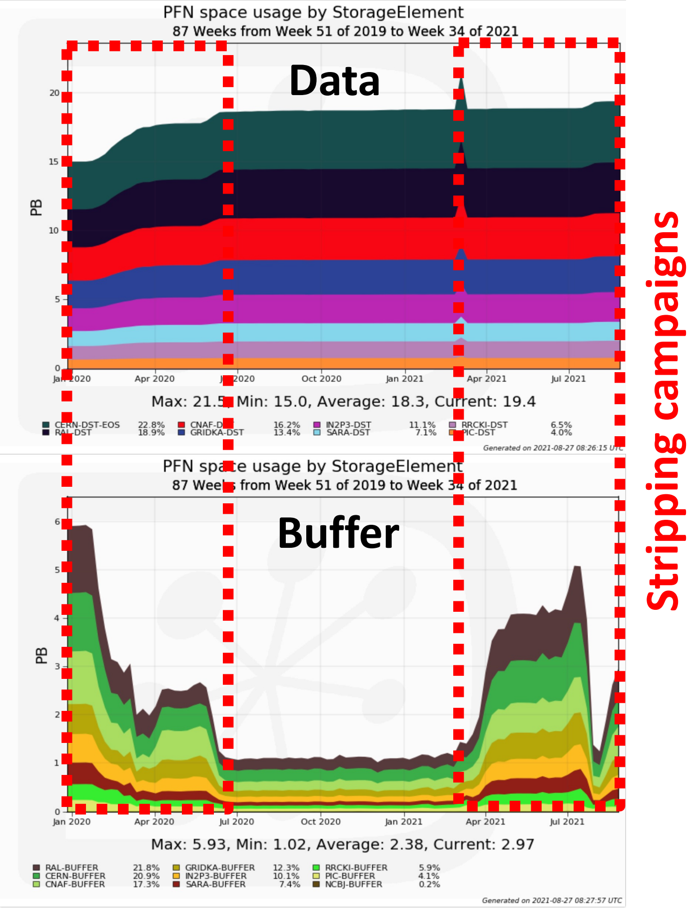

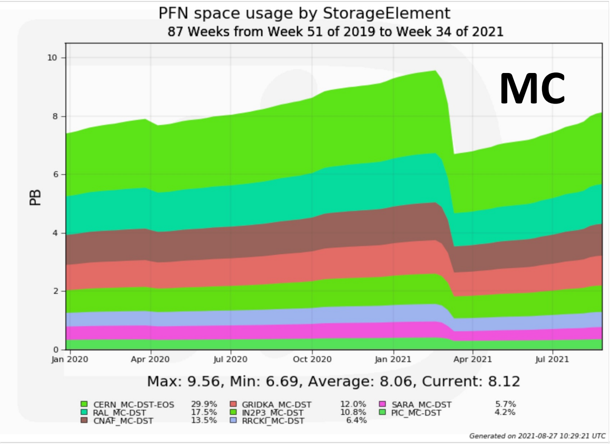

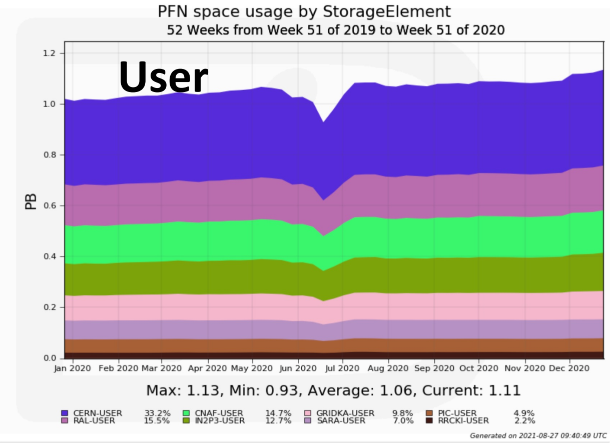

- Steady increase of MC samples
  - Despite archival of old MC productions
- Increase of data size due to stripping campaigns
  - PbPb and fixed target data processed beginning of 2020
  - Incremental stripping of Run1+Run2 just started
  - Note that up to 4 PB of buffer is needed to restrip one year of data taking
- **CNAF provides its fair share of about 15% of Tier1 (should be 19% according to LHCb policy)**

Disk usage at T2-D

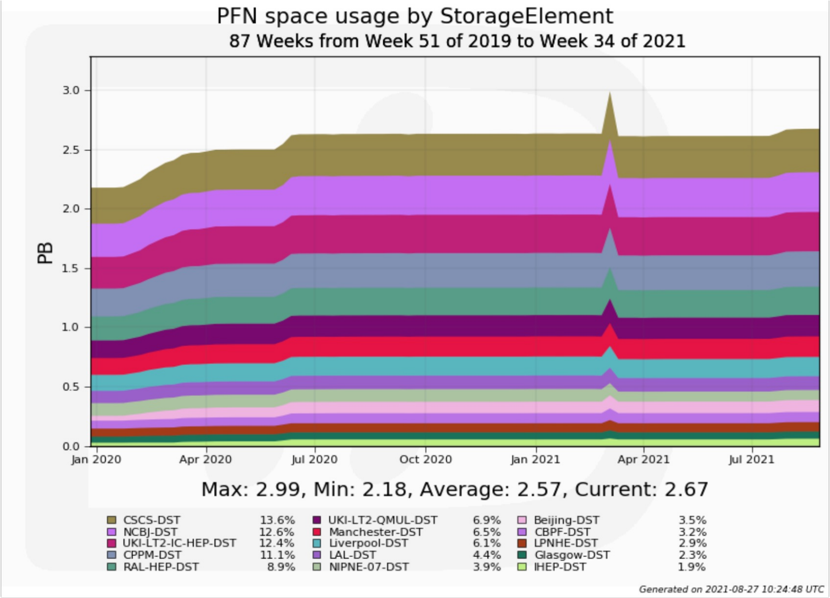

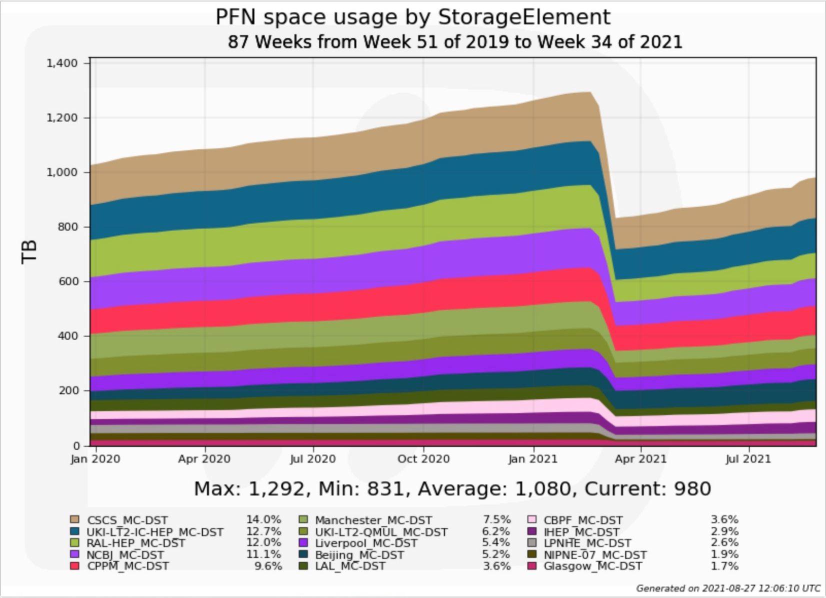

- Since few years several T2s also pledge disk resources
  - Very useful additional resources
  - Similar behaviour of disk usage as for T1s

Data popularity

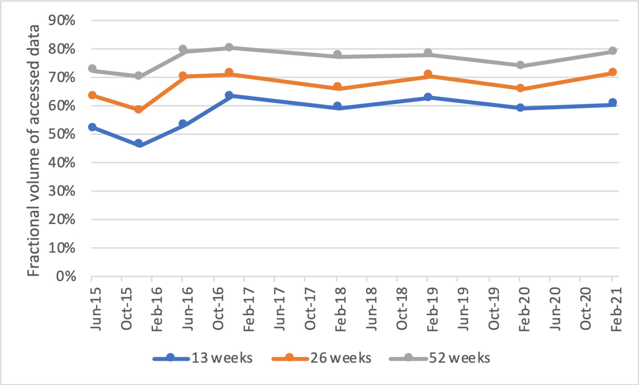

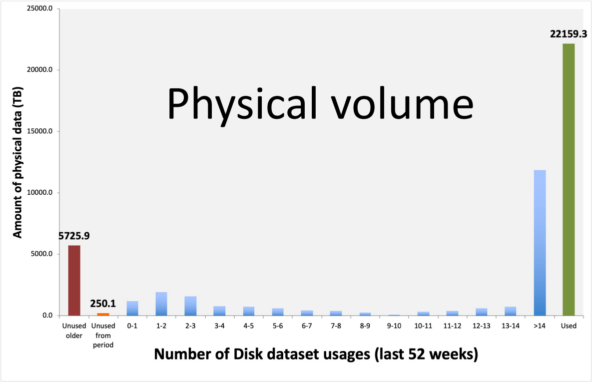

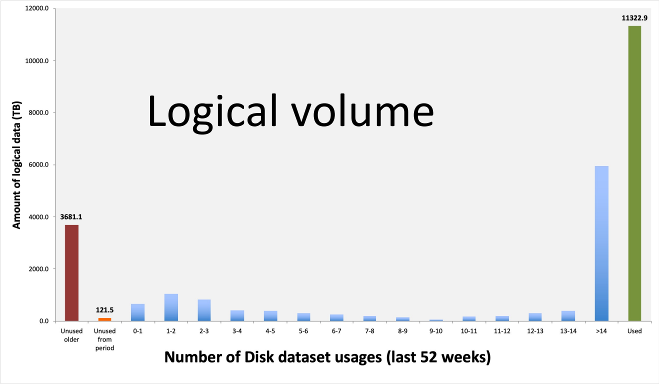

- Optimisation of datasets present on disk
  - Analysis allows to “optimise” number of replicas for unpopular datasets
  - General increase in old dataset usage during the last year due to legacy analyses going on
  - Archiving additional copies of all unused dataset in the last year we can expect to gain up to 1.1 PB
- No major news since last year:
  - Automatic preparation of weekly reports to be consulted by the Data Managers
  - Automatic removal of “cold” data not yet implemented but in the pipeline

Fast MC productions

- Extensive usage of different flavour of FastMC continued in 2020/2021
  - Not viable solution for all analyses
  - **More than ~85% of MC events produced in the last year are from some FastMC**
- Simulated events are filtered to save disk
  - Start of production for Run3 trigger optimisation requires unfiltered samples

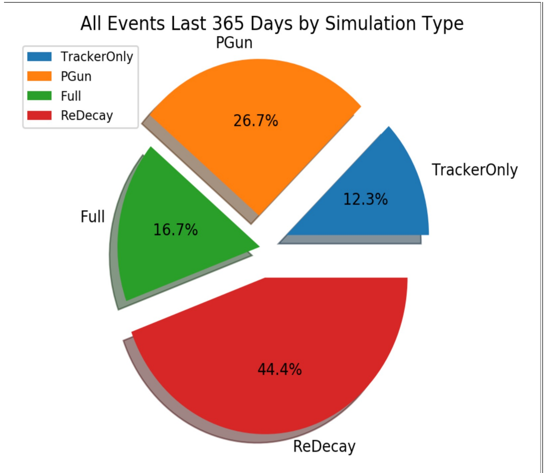

Distribution of events produced or requested (not yet produced) in the last 365 days.

Year	Simulated events (10^9)	Stored events (10^9)	Ratio	CPU work kHS06.y	CPU per event kHS06.s	LFS TB
2017	10.3	4.2	40.3%	817	2.50	640
2018	12.0	3.0	25.3%	1009	2.65	550
2019	45.0	6.9	15.2%	1290	0.90	1110
2020	53.0	16.8	31.7%	1357	0.81	2010

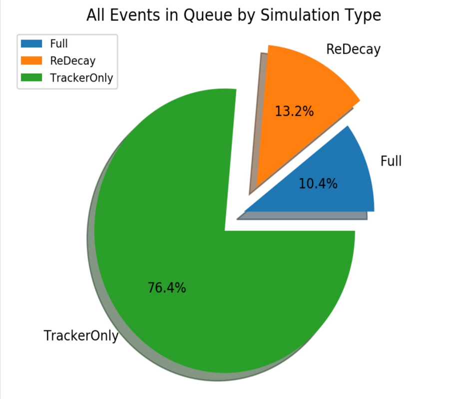

Distribution of Simulation Type over all events currently requested.

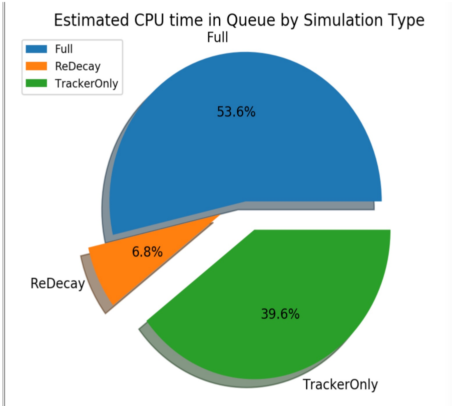

Estimate on distribution of CPU time to produce the events currently requested.

Ultrafast MC productions

- Development of a fully-parametric simulation framework LAMARR continues

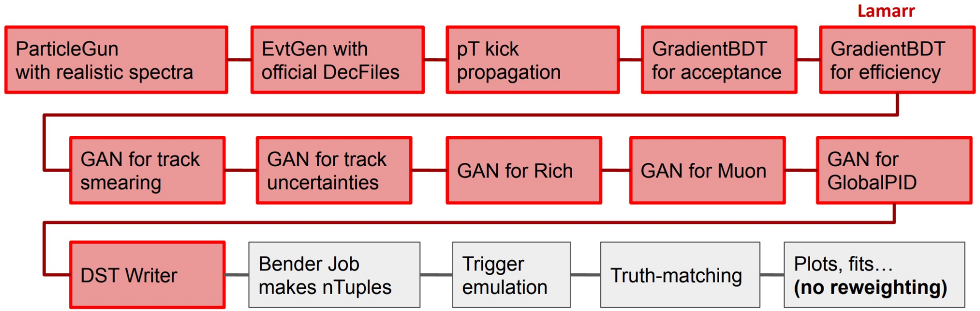

- With respect to last year several improvements and additions
  - **Much better particle propagator, tracking simulation and inclusion of particleID**
  - **All based on GANs ➔ would benefit from dedicated GPU resources meant for training**

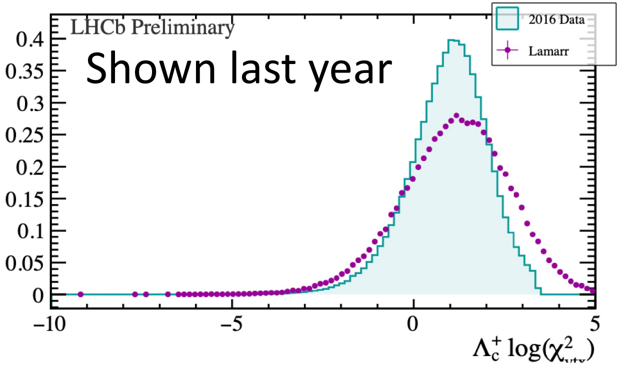

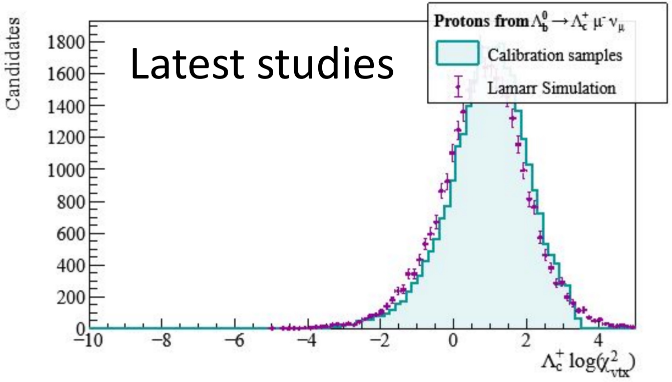

- Important contributions from Italian community (L. Anderlini, B. G. Siddi)

HPC resources

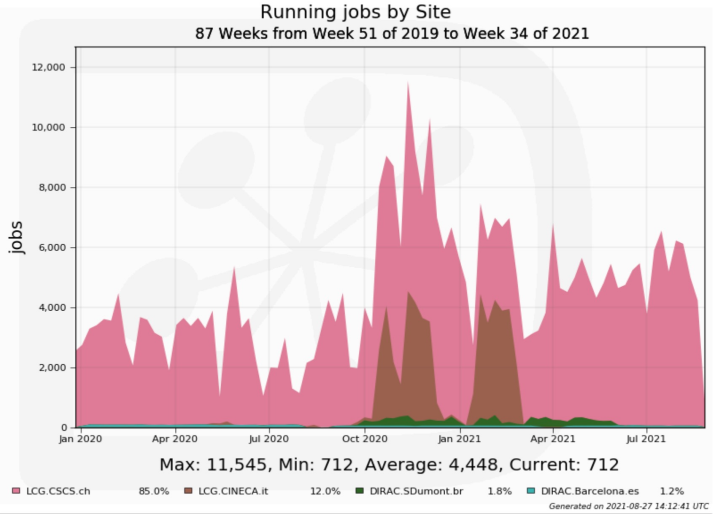

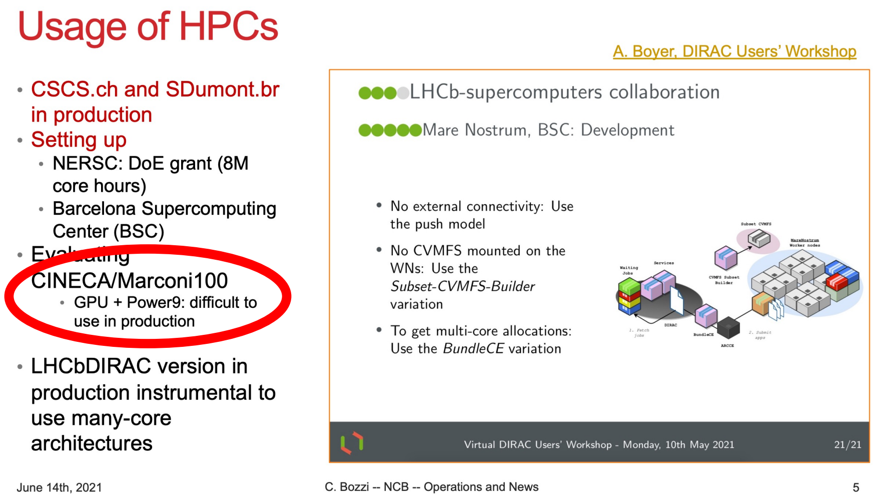

- Work continue to exploit other resources and in particular HPC centres
- Aftern nice experience with MarconiA2, started to investigate the usage of Marconi100 cluster CINECA
  - Not suitable for production jobs ➔ LHCb software relies on LCG dependencies not available for PowerPC
  - GPUs constitute the real power of the cluster ➔ Very interesting for NNet training and users' applications
  - Major LHCbDirac version released during the summer allowing to submit to the cluster
  - Several use cases prepared to start the testing phase (**Padova, Firenze, Cagliari, Milano**)

MarconiA2@CINECA

- As part of PRACE Project Access Grant together with the other LHC experiments
- KNL farm: 68-processor Xeon-Phi
  - Up to 4x hyper-threading, 96 GB DDR4 RAM
- Resurrection and development of MP/MT application for LHCb simulation (Gaussino)
  - Code is now validated and working
  - Vary important to validate MP/MT LHCb simulation to be put in production soon

Summary slides from Tommaso Boccali

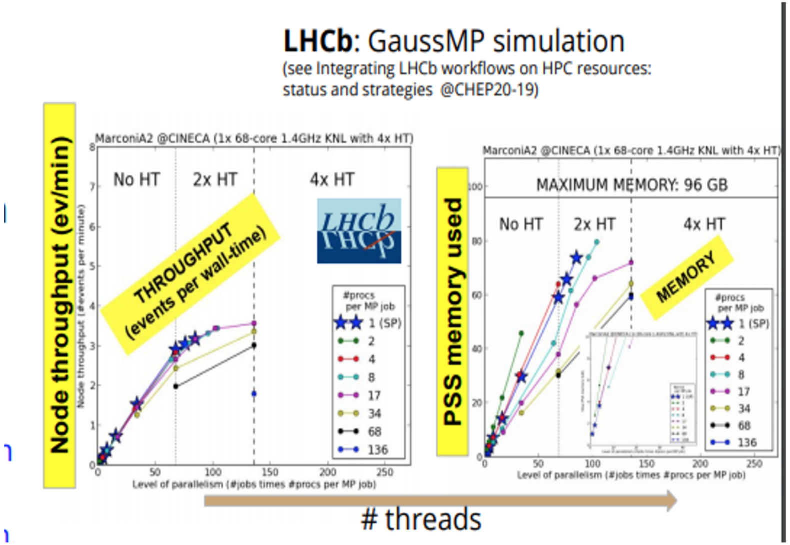

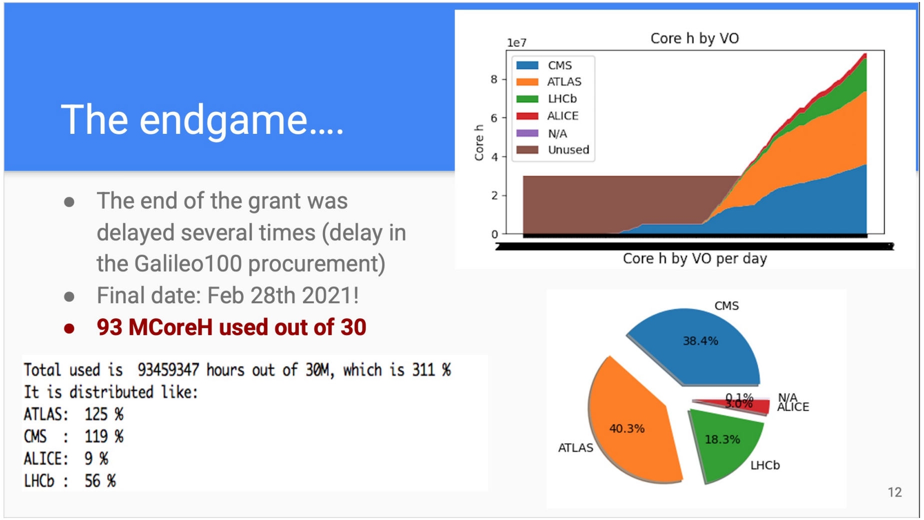

Requests for 2022

- Some history
  - Requests for 2021 were still based on the assumption of 2021 being a year with "physics"
  - Requests where split in two parts:
    - April2021 ➔ minimal resources considering no data taking in 2021
    - Autumn2021 ➔ deploy resources as originally planned (sort of anticipation for 2022)
- In light of the new schedule of LHC requests for 2022 have been reviewed (LHCb-PUB-2021-002)
  - Assuming **operations will start on 1st Feb. 2022**
  - Total of **6×10^6 s (pp collisions) + 1.2×10^6 s (ion collisions)**
  - Total integrated luminosity of **10 fb^{-1}** (integrated luminosity in **Run2 was 2 fb^{-1} per year**)

Requests for 2022 (CPU)

LHCb-PUB-2021-002

CPU Work in WLCG year (kHS06.years)	2021	2022 LHCB-PUB-2020-005	2022 THIS DOCUMENT
First pass sprucing	70	160	80
End-of-year sprucing	70	160	80
Simulation	760	870	870
Core and distributed computing infrastructure	10	10	10
User Analysis and working group productions	260	335	220
Total Work (kHS06.years)	**1170**	**1535**	**1260**
LHCb-TDR-018 (2021 pledge)	**860 (934)**	**1580**	**1580**

- An **improvement of a factor 2 in the performances of sprucing** is assumed
  - Porting improvements from HLT algorithms to offline software
  - Same gain expected for user analysis algorithms
- **No resources for reconstruction of ion collisions is requested**
  - In case HLT farm will not be able to do that online, resources will be asked for 2023

Requests for 2022 (Disk)

LHCb-PUB-2021-002

Disk storage usage forecast (PB)		2021	2022 LHCb-PUB-2020-005		2022 This document	
Real data	Run1+Run2 pp data	37.8	17.9	73.7	10.2	65.9
	Run1+Run2 PbPb + SMOG					
	Run3: FULL		13.7		13.7	
	Run3: TURBO		30.3		30.3	
	Run3: TURCAL		3.7		3.7	
	Run3: Minimum bias		2.4		2.4	
	Run3: PbPb + SMOG2		5.6		5.6	
Simulated data	Run1+Run2 Simulated Data	10.0	8.7	10.9	8.7	10.0
	Run3 simulated data		2.2		2.2	
Other	User data	15.9	8.5	28.2	1.8	12.8
	Buffers		19.7		11.0	
Total		**63.7**	**112.7**		**89.6**	
	LHCb-TDR-018 (2021 pledge)	**66.0 (58.7)**	**111.0**		**111.0**	

- **Optimisation of Run1+Run2 number of copies**
- Reassessment of user data given **centralisation of user jobs management**
- Reassessment of buffer needed for end-of-the-year reprocessing requiring a 2-weeks contingency

Requests for 2022 (Tape)

LHCb-PUB-2021-002

Tape storage usage forecast (PB)		2021	2022	
Run1 + Run2	RAW data (pp+HI+fixed target)	81.1	38.4	82.1
	RDST data (pp+HI+fixed target)		13.7	
	ARCHIVE		30.0	
Run3	pp data (FULL+TURBO+TURCAL)	38.6	120.1	137.8
	minimum bias / no-bias		0.6	
	Heavy Ion Data + fixed target		5.6	
	ARCHIVE (data+MC)		11.5	
Total		**119.7**	**219.9**	
	LHCb-TDR-018 (2021 pledge)	**142.0 (108.7)**	**243.0**	

- No change with respect to previous estimation
  - **Tape is the most critical resource in the next years**

Requests for 2022 (BW to/from tape)

LHCb-PUB-2021-002

Country	Site	Tape Read BW (GB/s)	Tape Write BW (GB/s)
CERN		4.24	5.50
Tier1 sites			
France	CC-IN2P3	0.49	0.63
Germany	GridKA	0.86	1.12
Italy	CNAF	0.86	1.12
Netherlands	SARA/NIKHEF	0.34	0.44
Russia	RRCKI	0.34	0.44
Spain	PIC	0.23	0.29
UK	RAL	1.13	1.46
TOTAL Tier1 sites		4.24	5.50

Table 3-4: minimum required read and write bandwidths to the tape systems at CERN and Tier1 sites. A disk buffer corresponding to at least two weeks must be provided as well for data staging.

- Tape writing will have to sustain the **10 GB/s of data coming from HLT farm**
- Tape reading is evaluated assuming end-of-year reprocessing to finish in ~2 months
- Buffer disk must provide a contingency of at least two weeks

Summary of requests for 2022

- Reassessment of 2022 requests has been **scrutinised and approved by C-RSG** (CERN-RRB-2021-023)

LHCb		2022 Request	2022 req. /2021 C-RSG	C-RSG recomm.
CPU	Tier-0	189	108%	189
	Tier-1	622	108%	$\boxed{622}$
	Tier-2	345	107%	$\boxed{345}$
	HLT	50	100%	50
	Total	1206	108%	1206
	Others	50		
Disk	Tier-0	26.5	141%	26.5
	Tier-1	52.9	141%	$\boxed{52.9}$
	Tier-2	10.2	140%	$\boxed{10.2}$
	Total	89.6	141%	89.6
Tape	Tier-0	81	185%	81.0
	Tier-1	139	183%	$\boxed{139.0}$
	Total	220	184%	220.0

- In the red boxes the sum of resources that all T1s and T2s of the experiment will have to provide in 2022 for the proper working of LHCb
- To obtain INFN share total numbers must be multiplied by INFN share

Computation of INFN share

- By LHCb policy share should be computed in the following way
  - CPU and Disk: fraction of PhD Eq. of the country without CERN at the denominator ➔ 18.8% per INFN
  - Tape: fraction of PhD Eq. with only country with T1s at the denominator ➔ 26.7% per INFN
- In the next slides all numbers are computed assuming fair share ➔ **16.6 %** including CERN at the denominator
  - CPU: separated requests between Tier1 and Tier2
  - Disk: sum of T1+T2 share assigned to T1
  - Tape: only to T1

CERN-RRB-2021-047

2022	PhD eq. total/ funding auth.	%
AUSTRALIA	2	0.4
BRAZIL	19	3.4
COLOMBIA	2	0.4
FRANCE	43	7.6
BMBF GERMANY	29	5.1
MPG, GERMANY	2	0.4
HUNGARY		
IRELAND	1	0.2
INFN ITALY	94	16.6
NETHERLANDS	25	4.4
P. R. CHINA	37	6.5
POLAND	24	4.2
HHNIPNE ROMANIA	5	0.9
RUSSIA	56	9.9
SPAIN	20	3.5
SWITZERLAND	26	4.6
UKRAINE	2	0.4
UK	83	14.6
USA	29	5.1
CERN	68	12.0
TOTAL	567	100.0

INFN resources for 2022

- LHCb-T1:
  - CPU: Total T1s multiplied by INFN share
    ➔ **622 x 0.166 = 103 kHS06**
  - Disk (*): Total T1s+T2s multiplied by INFN share
    ➔ **(52.9+10.2) x 0.166 = 10.5 PB**
  - Tape: Total T1s multiplied by INFN share
    ➔ **139 x 0.166 = 23 PB**
- LHCb-T2:
  - CPU: Total T2s multiplied by INFN share
    ➔ **345 x 0.166 = 57.3 kHS06**

LHCb		2022		
		Request	2022 req. /2021 C-RSG	C-RSG recomm.
CPU	Tier-0	189	108%	189
	Tier-1	622	108%	$\boxed{622}$
	Tier-2	345	107%	$\boxed{345}$
	HLT	50	100%	50
	Total	1206	108%	1206
	Others	50		
Disk	Tier-0	26.5	141%	26.5
	Tier-1	52.9	141%	$\boxed{52.9}$
	Tier-2	10.2	140%	$\boxed{10.2}$
	Total	89.6	141%	89.6
Tape	Tier-0	81	185%	81.0
	Tier-1	139	183%	$\boxed{139.0}$
	Total	220	184%	220.0

(*): No difference between T1 and T2 disk in the LHCb computing model. Since LHCb-T2 is a part of CNAF disk is historically assigned to T1.

Currently installed resources

- Officially installed resources are declared on WLCG-CRIC
- Currently declared resources are the sum of 2021 scrutiny plus SJ unlocked in May
- **N.B.: Currently declared resources are not those actually intalled**

Site Topology Pledges Accounting Downtime Admin Logs Help Login

Export Filter Reload Columns 8/21 Pledge View

Federations list

Show 100 entries

Federation	Accounting name	Tier	Country	Type	DISK Q2 2021 LHCb	CPU Q2 2021 LHCb	TAPE Q2 2021 LHCb
IT-INFN-CNAF	INFN-T1	1	Italy	Official	7633	97580	13362
IT-INFN-T2	INFN T2 Federation	2	Italy	Official	0	54569	0
Federation	Accounting name	Tier	Country	Type	DISK Q2 2021 LHCb	CPU Q2 2021 LHCb	TAPE Q2 2021 LHCb

Showing 1 to 2 of 2 entries

Previous 1 Next

Summary of INFN resources for 2022

Variations computed according to declared resources

LHCb T1	2021 Pledge (Scrutiny 2021 + SJ)	2022 (Δ 2021)
CPU (HS06)	97580	103118 (+5538)
Disk (TB)	7633	10461 (+2828)
Tape (TB)	13362	23044 (+9682)

LHCb T2	2021 Pledge (Scrutiny 2021 + SJ)	2022 (Δ 2021)
CPU (HS06)	54569	57196 (+2627)

- Considerations on fair share
  - Using fair share instead of LHCb policy there is a shortage of resources of about: CPU ~ -22 kHS06, Disk ~ -1.4 PB, **Tape ~ -14 PB**
- **14 PB of tape corresponds to ~25 kHS06 of CPU (250 k€)**

Costs 2022

Simple multiplication by unitary costs

CPU = 10 €/HS06, Disk = 120 €/TB, Tape = 18 €/TB

LHCb T1	2021 Pledge (Scrutiny 2021 + SJ)	2022 (Δ 2021)	Costi
CPU (HS06)	97580	103118 (+5538)	55.4 kEuro
Disk (TB)	7633	10461 (+2828)	339 kEuro
Tape (TB)	13362	23044 (+9682)	174 kEuro
LHCb T2	**2021 Pledge (Scrutiny 2021 + SJ)**	**2022 (Δ 2021)**	
CPU (HS06)	54569	57196 (+2627)	26.3 kEuro

Total: 595 kEuro

Note: replacements for dismissed resources are not taken into account in these numbers

Risk assessment

Impact:
1: we can deal with it, no problem
2: a bit of a hassle but not too bad
3: can be managed, but with significant effort
4: crisis

Likelihood:
1: never expected to happen
2: could happen but very unlikely
3: could well happen
4: will probably happen

LHCb Risk Register

Risk	Risk Likelihood	Risk Impact	Risk Severity	Owner	Effect	Mitigation
Funding						
Tape shortage at CERN	2	4	8	CERN	Cannot store data coming from HLT. Descoping of physics program	Decrease HLT throughput to offline by migrating more lines from FULL to TURBO and/or tightening
Tape shortage at Tier1 sites	3	4	12	Tier1 sites		
Disk shortage	3	3	9	WLCG sites	Cannot analyse (part of) data taken. Impact on operations load and complexity, system resiliency. Delay in physics program	Park data on cold storage and activate if when feasible. Decide what to store on disk based on popularity.
Shortage of computing power	3	2	6	All sites	More time to get adequate simulation samples. Delay of physics analysis. Substantial changes in analysis model.	increase fraction of fast and parametric simulations. Increase pool of opportunistic resources
Operations / Technology						
Availability of tape write bandwidth	2	4	8	CERN + Tier1 sites	Cannot store data coming from HLT. Descoping of physics program	Decrease HLT trhoughput. Prioritize physics program
Availability of tape read bandwidth	2	4	8	CERN + Tier1 sites	Cannot recall data for end-of-year re-sprucing. Delay of physics analysis	Delay re-sprucing to EYETS or LS3
Underestimation of disk buffer	3	3	9	WLCG	Resprucing slowed-down and extending into data taking period	Delay re-sprucing or get space by temporarily removing other datasets
Software						
Underestimation of sprucing work	2	2	4	LHCb	More computing power needed. YETS re-sprucing is delayed	recover computing power by delaying simulation production
Availability of sprucing application	2	3	6	LHCb	delay in (re-)sprucing campaigns. Buffer space is not sufficient	park unspruced data on tape. Delay YETS re-sprucing
Underestimation of analysis work	2	2	4	LHCb	more computing power needed	recover computing power by delaying simulation production
Underestimation of simulation work	3	3	9	LHCb/G4/HSF	Cannot afford sufficient simulation; physics analysis suffers	find optimization point by studying interplay between full/fast/parametric simulation and technological developments in each ofthem

Outlook until LS3

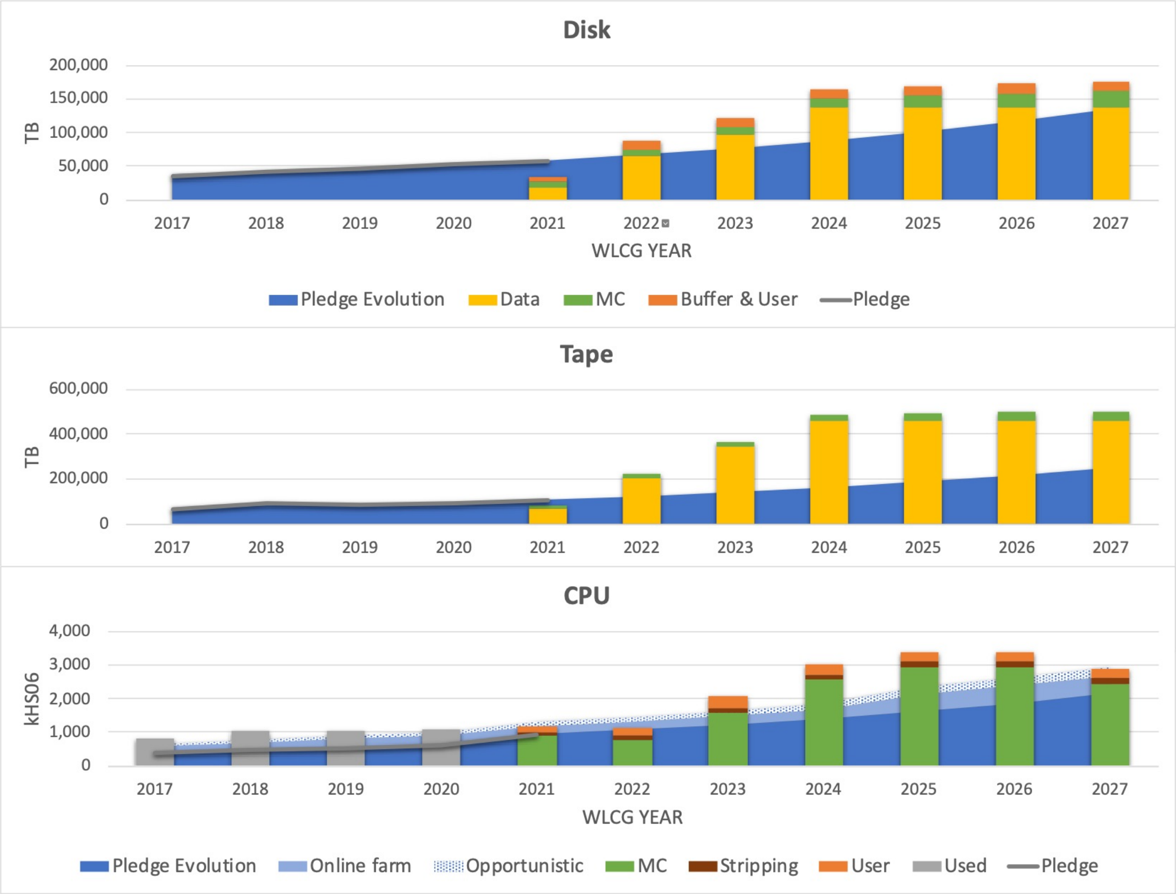

		DISK		
		PB	Increase	Average
2021	LS2	59		
2022		88	1.5	
2023	Run3	122	1.4	
2024		165	1.4	1.4
2025		169	1.0	
2026	LS3	173	1.0	
2027		177	1.0	1.2

		TAPE		
		PB	Increase	Average
2021	LS2	108.7	0	
2022		222	2.0	
2023	Run3	367	1.7	
2024		489	1.3	1.7
2025		495	1.0	
2026	LS3	498	1.0	
2027		501	1.0	1.3

		CPU		
		kHS06	Increase	Average
2021	LS2	933	0	
2022		1246	1.3	
2023	Run3	2097	1.7	
2024		3133	1.5	1.4
2025		3392	1.1	
2026	LS3	3392	1.0	
2027		2901	0.9	1.2

Conclusions

- **Good and efficient usage of computing resources** by LHCb in 2020 and 2021 so far
- Italian resources are provided in a very efficient way and are well used
  - **General underpledge of T1 impact for about 10-15% of CPU power provided by INFN to the collaboration** ➔ shortage will increase if further delay in CPU installation
- Work is ongoing to address most problematic points of LHCb computing model
- Delay in LHC schedule due to Covid-19 made necessary a reassessment of resource needs for 2022
  - Reduction of requests is based on some assumption on the capability to improve performances and efficiency of the computing model
- The likelihood and impact of potential resource shortages have been assessed
  - **shortage of storage resources, in particular tape, is at high risk and will impact physics programme ➔ need some extra resources**
- Hunger for resources will not stop in 2022