LHCb computing activities and requests for 2023/24

Lucio Anderlini
INFN Firenze

2022-09-06
Università “La Sapienza”, Roma

Istituto Nazionale di Fisica Nucleare
SEZIONE DI FIRENZE

Outline

- Usage of resources in 2021/22
- The upgrade computing model and other developments
- Plans for 2023 and beyond

Usage of resources in 2021/22

Requests and pledges for 2021

The LHCb requests for 2021, scrutinized by the cRSG, were reported in LHCb-PUB-2020-001

Resource usage for 2021 is detailed in the public document LHCb-PUB-2022-011

Pledges from different sites were taken from WLCG-CRIC

2021	CPU (kHS06)			Disk (PB)			Tape (PB)		
	Req.	Pledged	Pled./Req.	Req.	Pledged	Pled./Req.	Req.	Pledged	Pled./Req.
Tier 0	175	175	100.0%	18.8	18.8	100.0%	44.0	44.0	100.0%
Tier 1	574	470	81.9%	37.6	33.9	90.1%	76.0	64.7	85.1%
Tier 2	321	289	90.0%	7.2	6.0	83.0%	n/a	n/a	-
Total WLCG	1070	934	87.3%	63.7	58.7	92.2%	120.0	108.7	90.6%
HLT EFF		50							
Total	**1070**	**984**	**92.0%**	**63.7**	**58.7**	**92.2%**	**120.0**	**108.7**	**90.6%**

Usage of CPU resources

In 2021 we had limited reprocessing needs, most CPU (> 93%) devoted to Simulation jobs. Significant contribution (21.6%) from HLT EFF (*not available for simulation next year*).

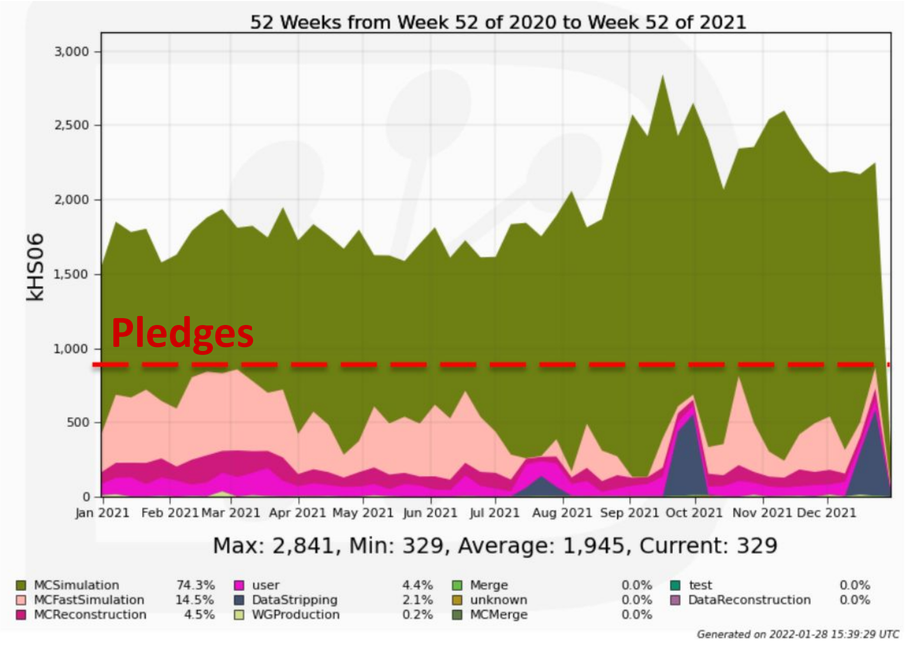

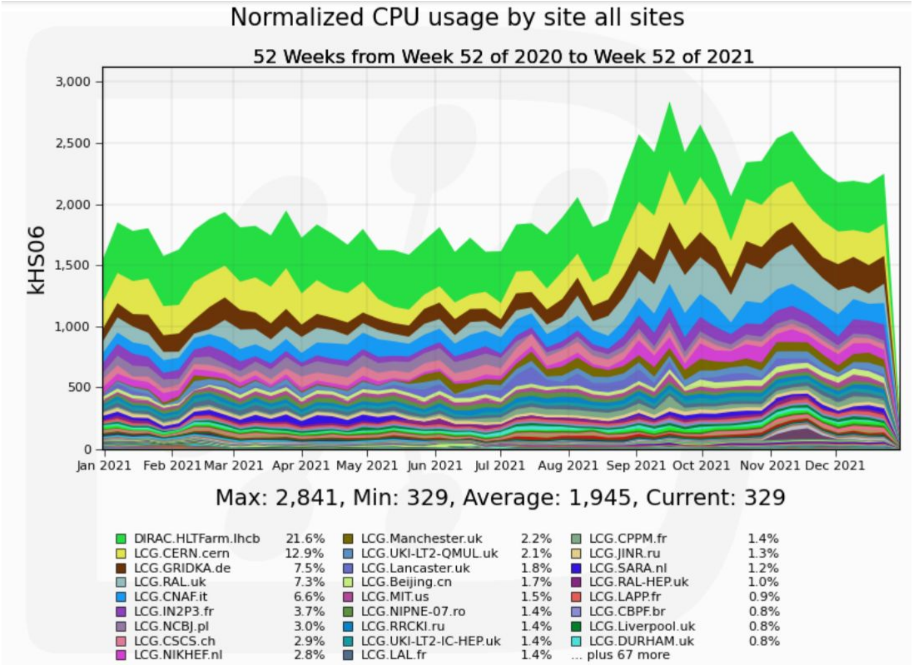

WLCG accounting (WAU)

Average power used:

at Tier-0 and Tier-1 sites, **122%** of the pledges;

at Tier-2, **160%** of the pledges.

In terms of operations, Simulation jobs can run independently on **Tier-0**, **Tier-1** and **Tier-2** sites.

In addition, in Italy, Tier-2 CPU is mostly allocated at CNAF, and is identical to Tier-1 resources.

In total, LHCb used **1145 kHS06** from WLCG resources plus **380 kHS06** from non-WLCG sites (mainly HLT EFF).

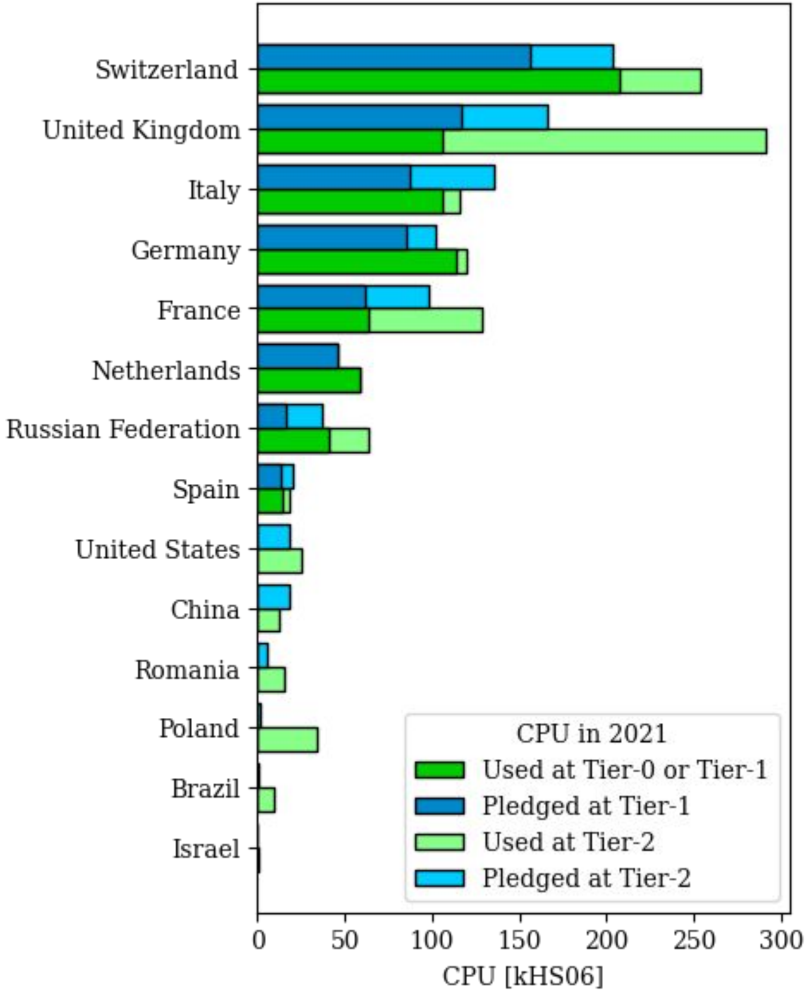

CPU usage at CNAF

LHCb running slightly under pledge in 2021 and first half of 2022, before getting CINECA CPU up and running, with which we are recovering quickly.

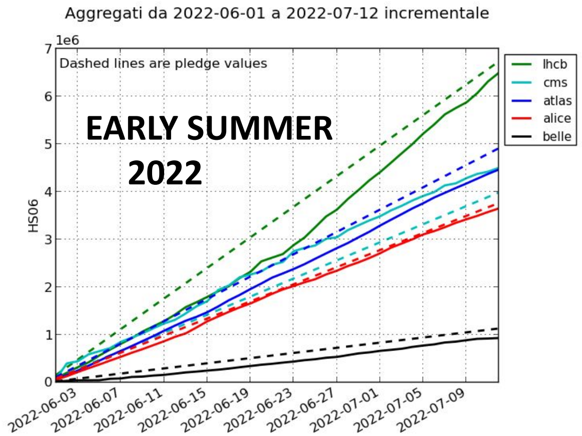

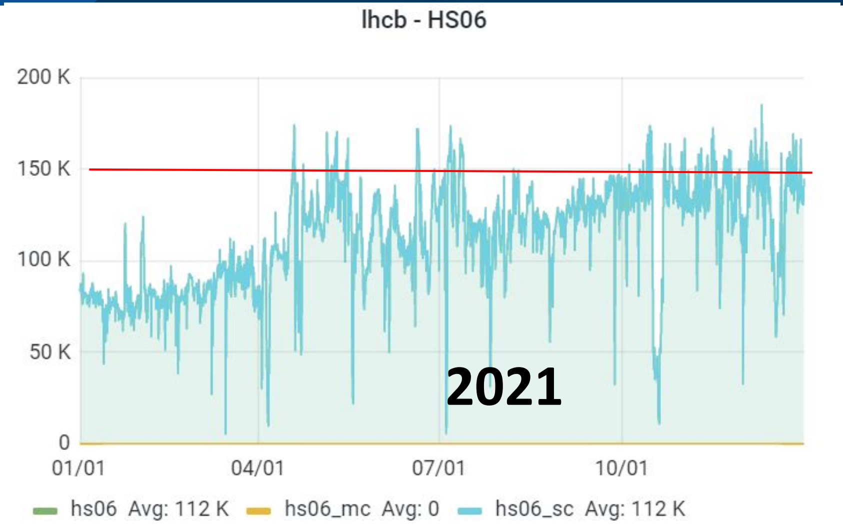

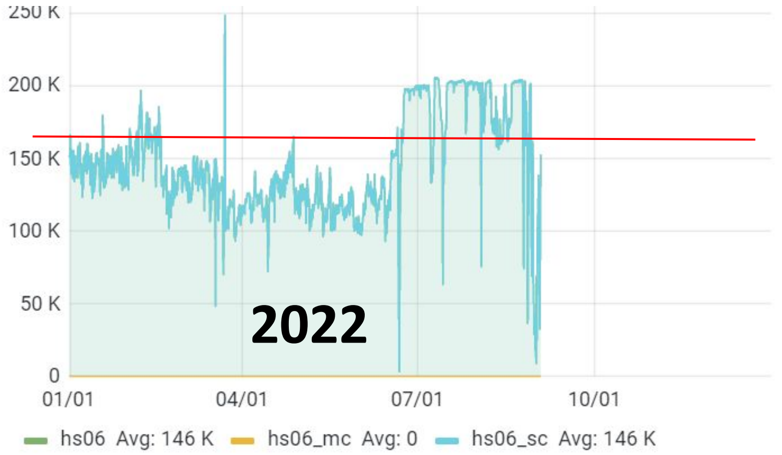

Tape usage as of December 2021

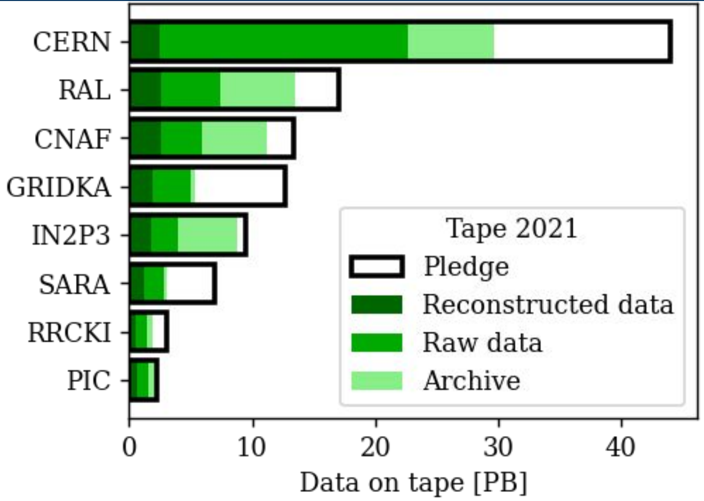

Minor increase in tape storage during 2021, mostly due to steady archiving of less-popular datasets.

Delays in the LHC schedule resulted into a tape occupancy 37% lower than the original 120 PB requested.

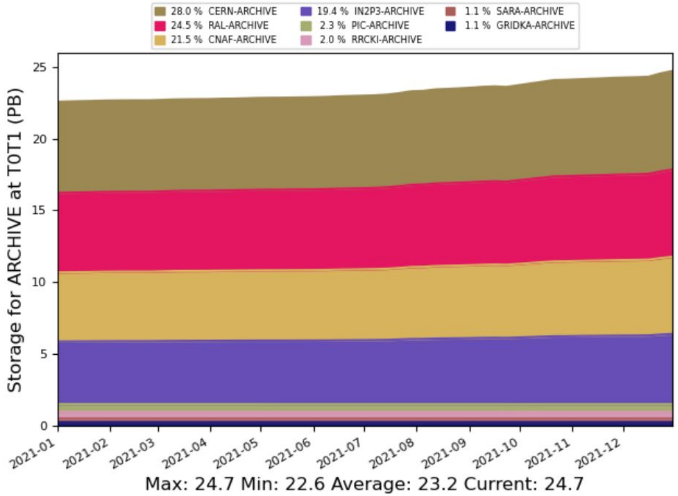

Tape usage

Increase in pledge should be followed by an increase of stored data, that will hopefully start flushing soon.

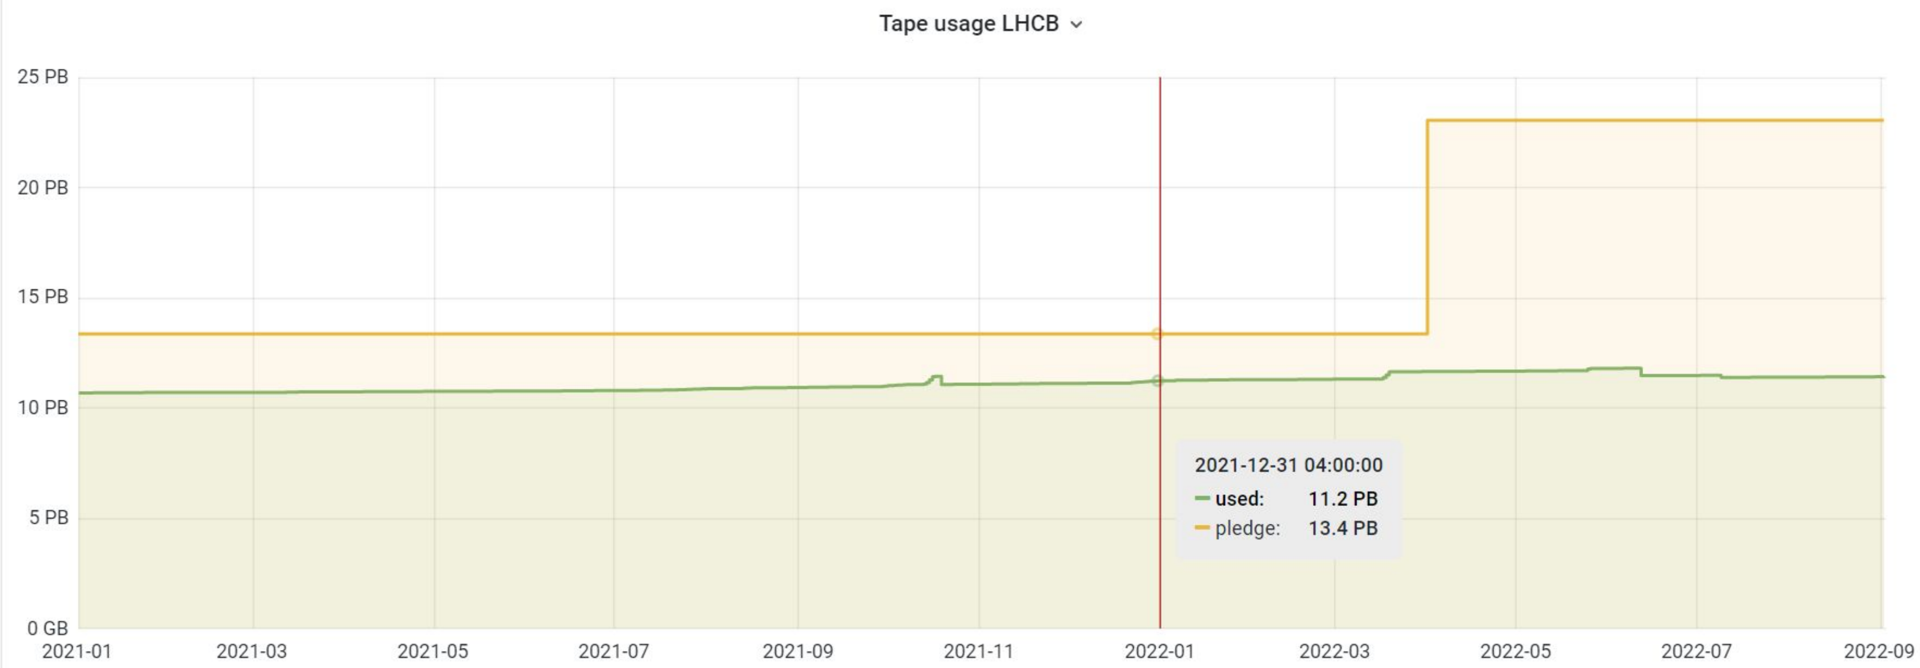

Disk usage

- Analysis-level datasets [Data]
- Simulated datasets [Simulation]
- Analysis-level nTuples [User]
- Buffer space for temporary files [Buffer]

EOS is also used to store large datasets for interactive analysis of relevance for a whole Physics Analysis Working Group.

In the end of 2021 calendar year, LHCb was using **39.2 PB** of disk space, 28% lower than the available disk space of **54.6 PB**.

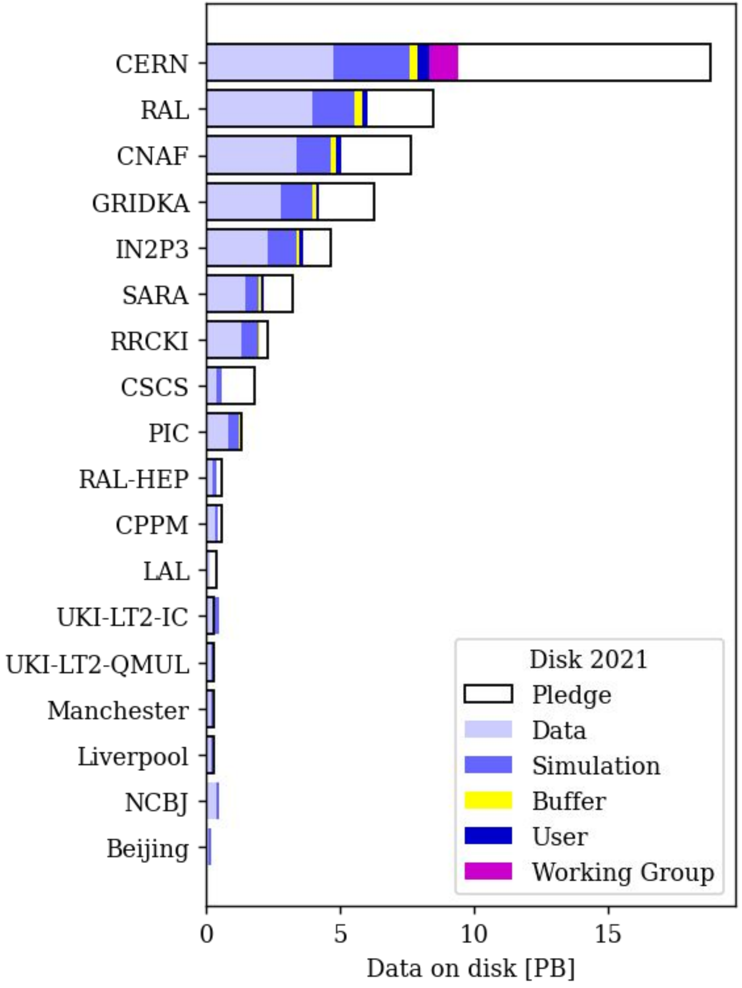

Disk usage: *Real data*

Mild increase of the disk usage due to incremental restripping campaigns.

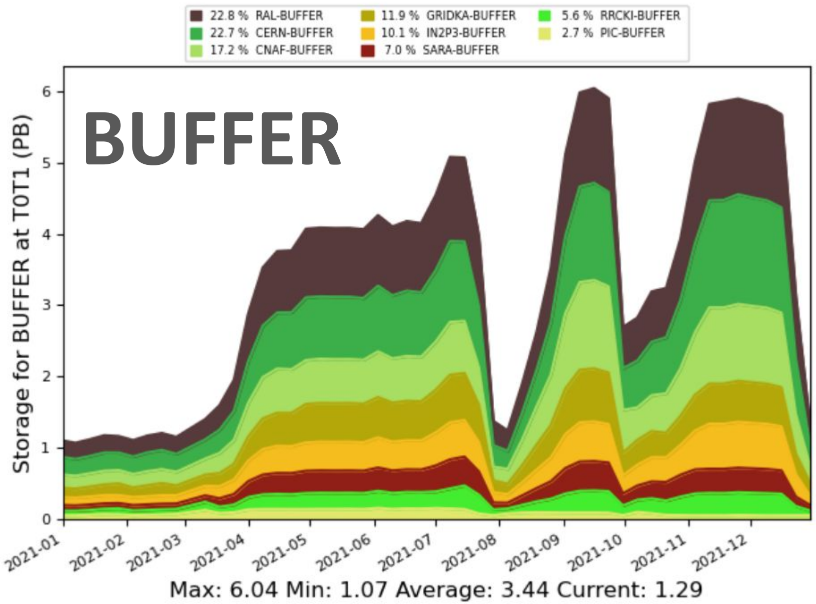

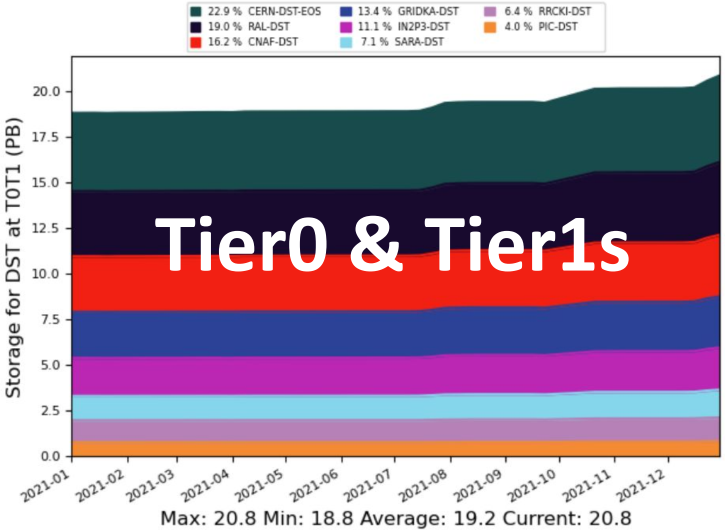

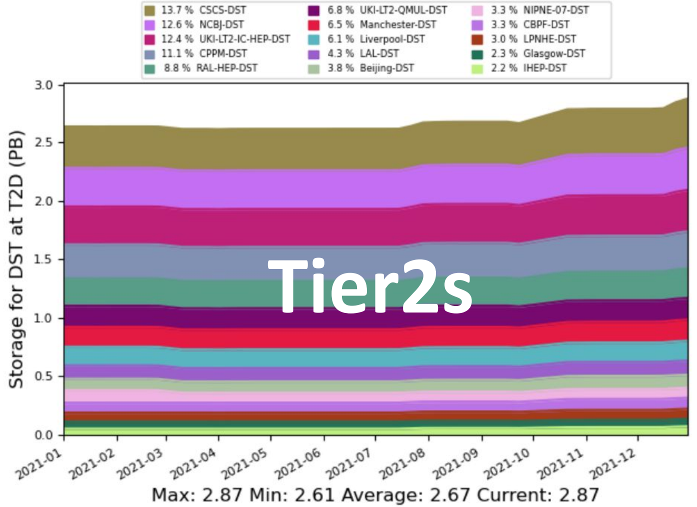

Simulated datasets

Volume of simulated datasets increases steadily as new events are simulated.

Major cleanup campaign performed in March 2021 perfectly balanced the newly generated events (3 PB).

A small fraction of the deleted datasets was retrieved from tape later in a semi-automatic way.

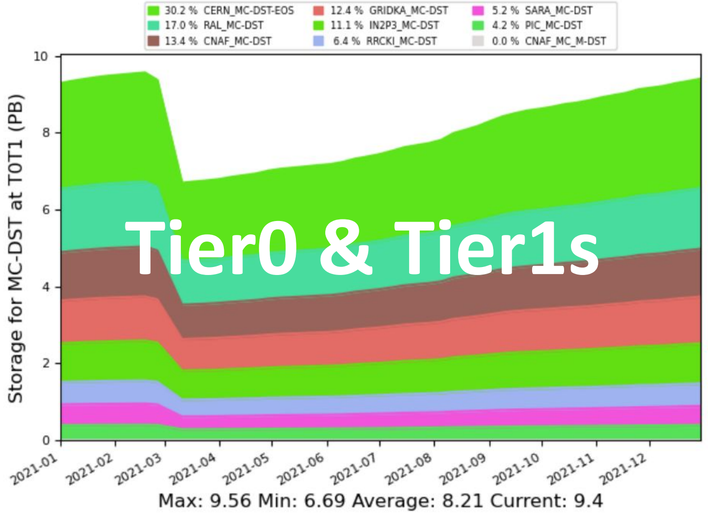

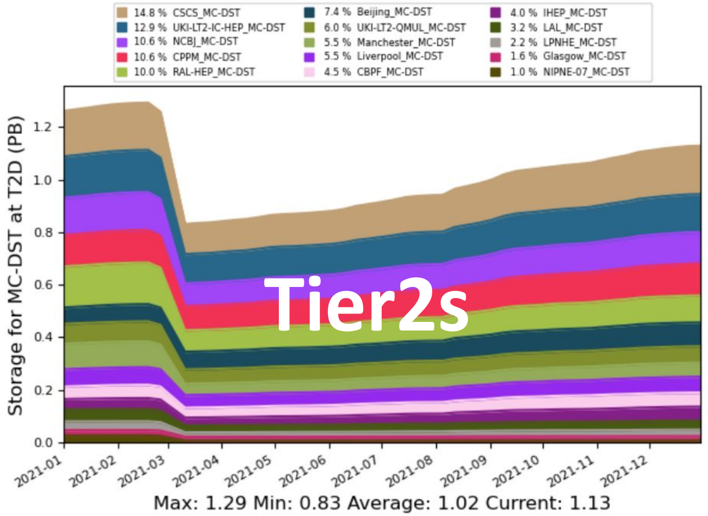

Disk usage at CNAF

As for tape, we expect a disk usage to rapidly increase with new data flushing in.

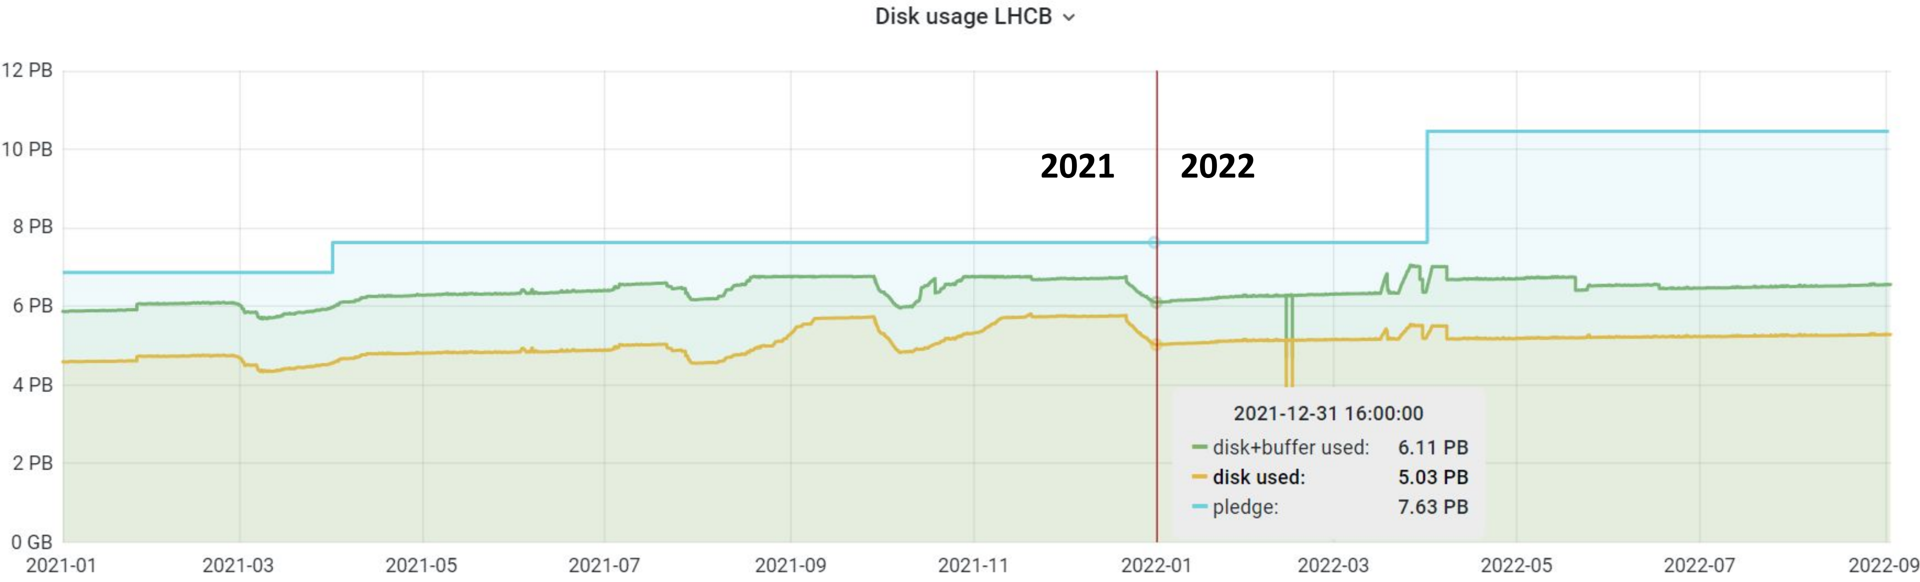

The upgrade computing model and other developments

Upgrade computing model

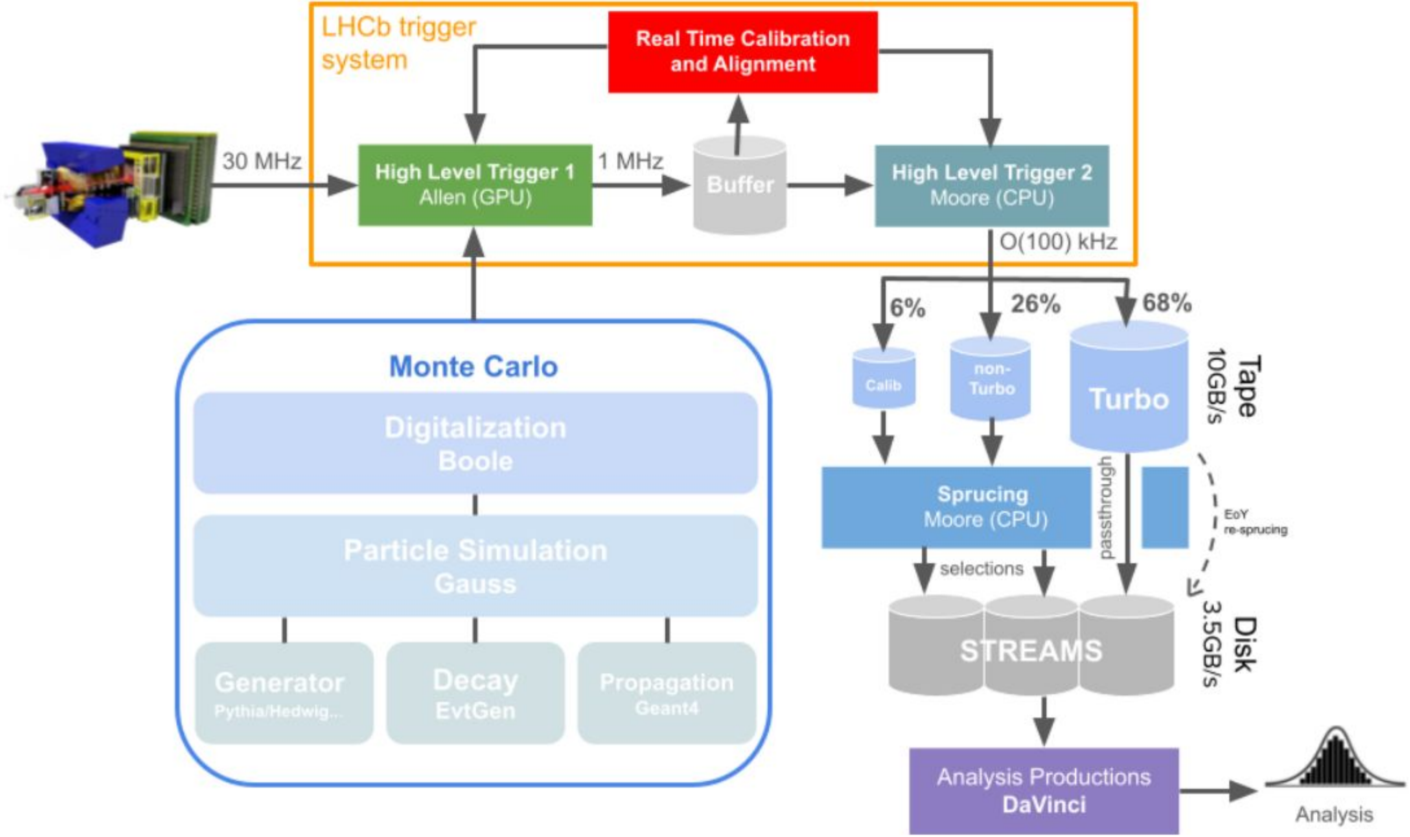

Trigger system

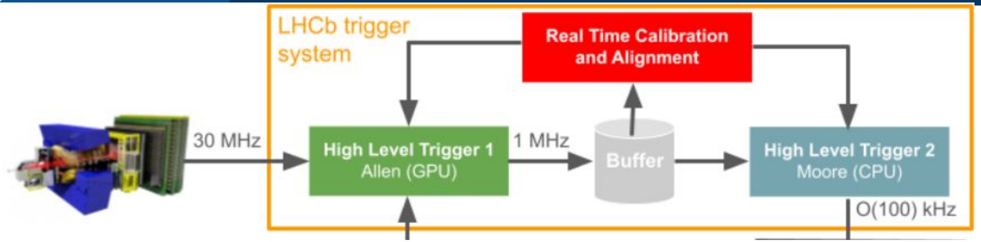

LHCb will run without a hardware trigger stage, using a farm of GPUs to perform an upfront "fast" reconstruction at 30 MHz. Software Project: *Allen*

Events passing a first loose preselection are stored on the local disk (**Buffer**) of the trigger farm and are used for online *calibration* and *alignment*.

Finally, the selected events are processed with the full reconstruction and selection algorithms on the CPU *Event Filter Farm*. Software Project: *Moore*

10 GB/s

The trigger output is a set of collection of events organized in streams:

- ***Turbo***: *decay candidates* completely processed at trigger level
- ***Non-Turbo***: (parts of) events that require *additional combinatorics* before analysis, *e.g.* inclusive lines
- ***Calib***: full events targeting *detector studies* and calibration of the simulation

CNAF, Ferrara, Firenze

Offline data processing

Non-Turbo and *Calib* streams are processed (**Sprucing**)

Software project: *Moore*

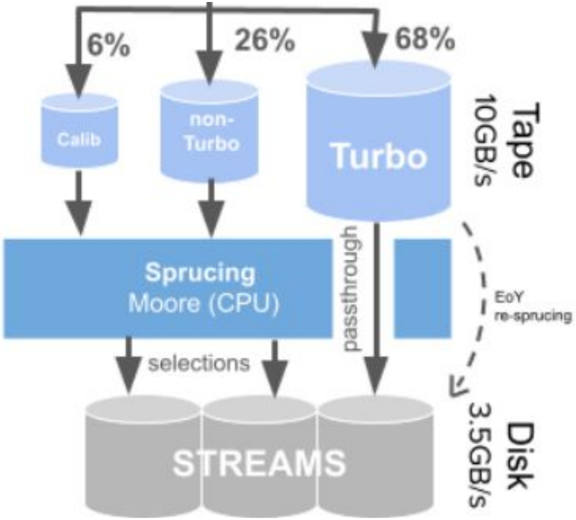

Sprucing runs at Tier-0 and at Tier-1s, during data taking.
While sprucing input is **archived on tape** in two copies:
Tier-0 + one Tier-1 to allow for legacy reprocessing.

Spruceing output (3.5 GB/s) is kept on disk where produced.

Bandwidths in GB/s

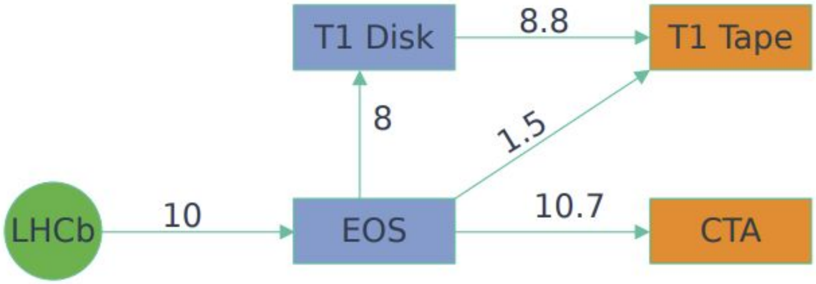

This results in **severe constraints** on both the amount of **available tape resources** and the bandwidth for **WRITE-ON-TAPE** operations.

Offline data processing at Tier-1s

During data-taking:

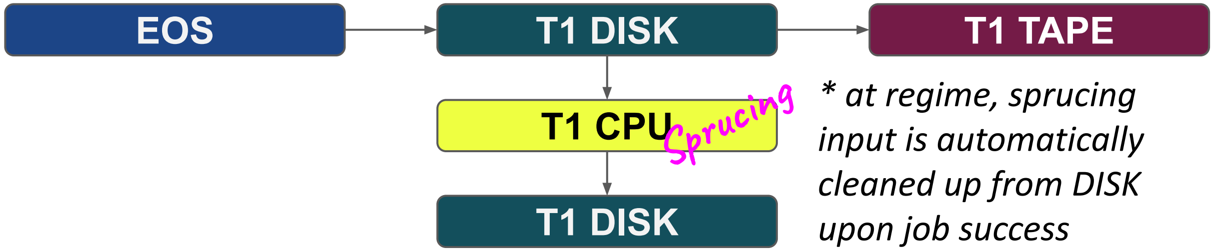

During re-sprucing:

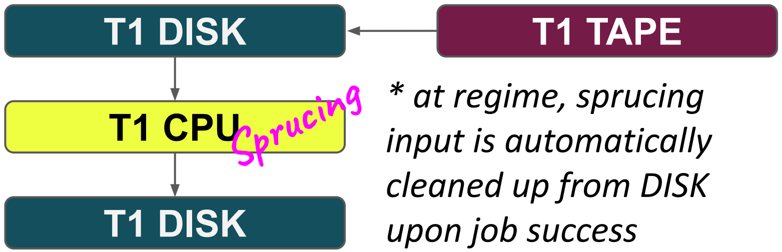

LHCb bandwidth expectations

	WRITE GB/s	READ GB/s
CERN	11	
CNAF	1.72	1.35
GRIDKA	2.23	1.36
IN2P3	1.25	0.98
NCBJ	1.32	0.91
PIC	0.2	0.17
RAL	2.96	1.93
RRCKI	0.25	0.21
SARA	1.07	0.74

Setup at CNAF: Data flow EOS to CNAF-DISK

FTS handles large numbers of simultaneous copies

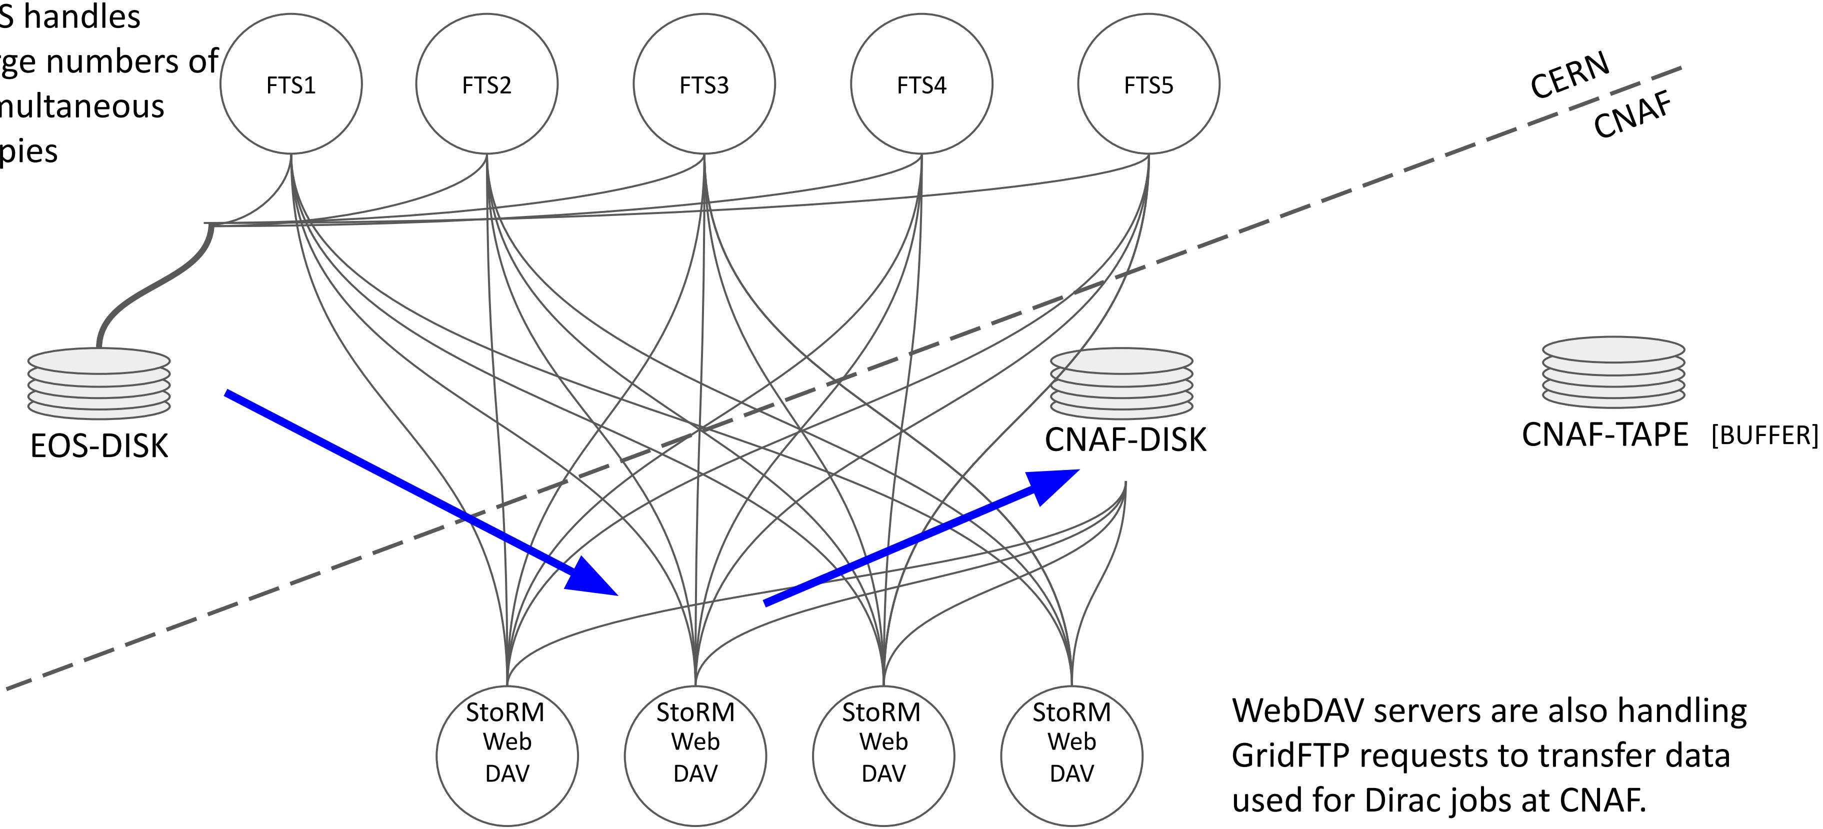

WebDAV servers are also handling GridFTP requests to transfer data used for Dirac jobs at CNAF.

Setup at CNAF: Data flow EOS to CNAF-DISK

FTS handles large numbers of simultaneous copies

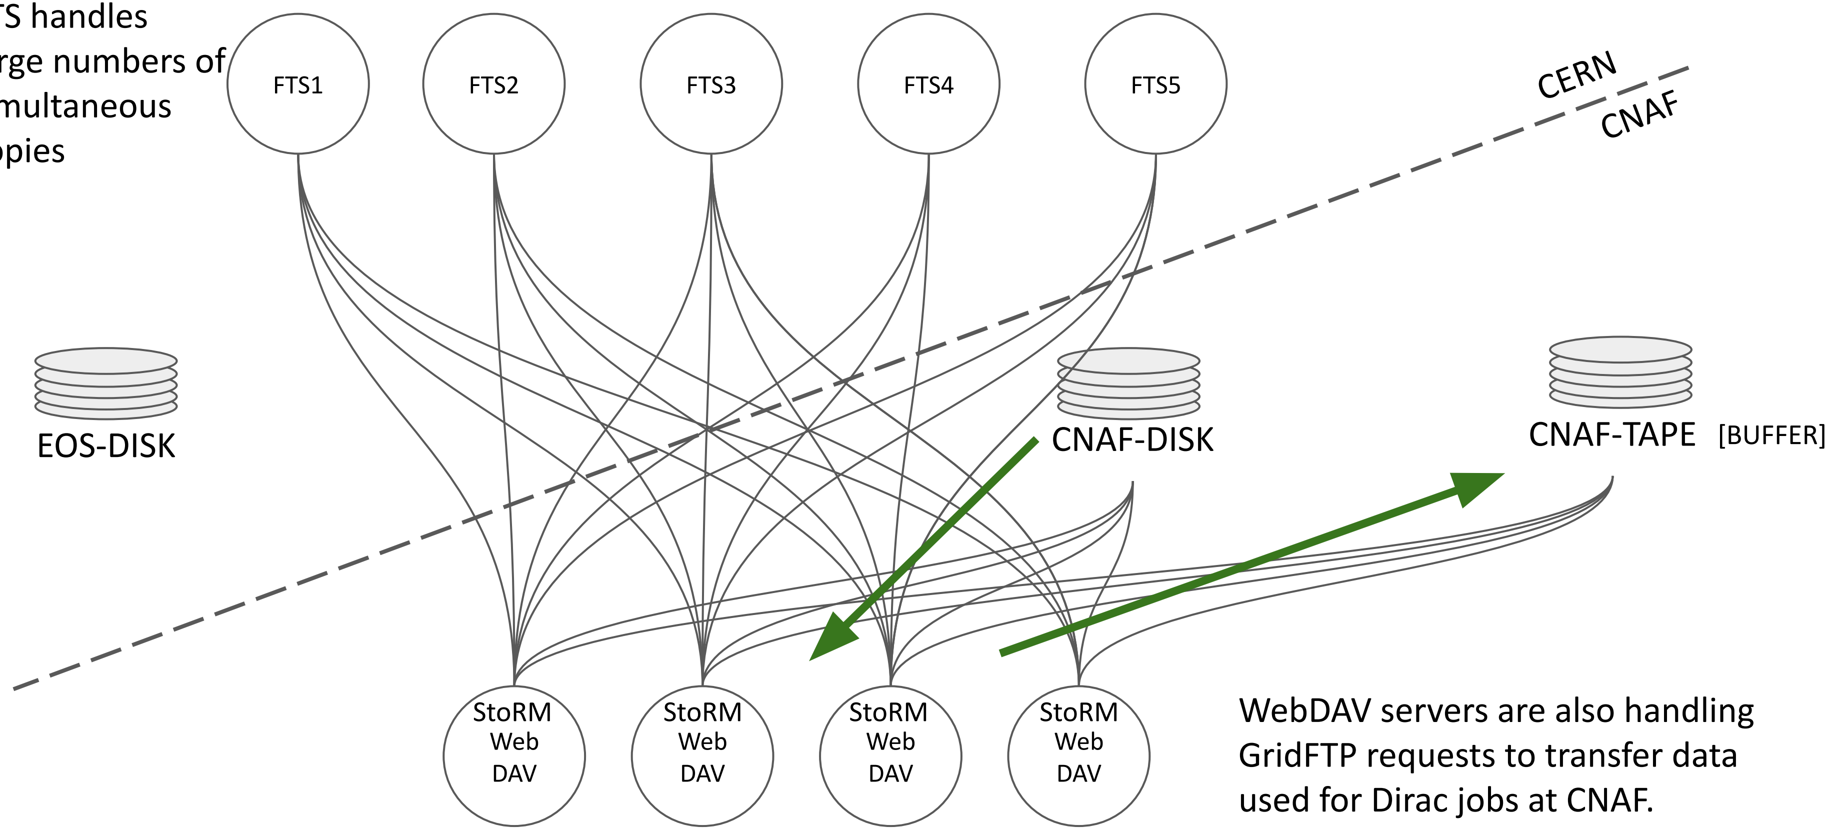

WebDAV servers are also handling GridFTP requests to transfer data used for Dirac jobs at CNAF.

Validation runs (*tape challenges*)

LHCb bandwidth expectations		
	WRITE GB/s	READ GB/s
CNAF	1.72	1.35

A **first validation** run was performed in **Dec. '21** copying manually (without WebDAV) data to disk and tape.
CNAF well above requirements (up to 5.16 GB/s).

Second tape challenge, involving the full setup, in **March '22**.
Read bandwidth immediately reached without issues.
Write bandwidth was basically acceptable, but with a **huge error-rate** (75% copies failed)

Third dedicated run, involving the full setup with max. verbosity, in **June'22**.
Confirmed the bottleneck is in the WebDAV endpoints that gets contacted multiple times per copy to handle the authentication and authorization process.

CNAF, Firenze

Readiness of CNAF

WebDAV endpoints are our bottleneck at CNAF.

Ongoing effort to move data access from the Condor jobs to POSIX to lighten WebDAV.

Dedicating nodes to WebDAV, only, and assigning all other protocols to a **single node** results into a significant improvement.

Connections with **FTS** and **WebDAV** developers seeking for **improvements software-side**. Several packages were updated during the summer.

Improvements of the hardware are also being considered.

A further test would be beneficial, we are trying to fit it in the LHCb & CNAF schedules.

Readiness of other sites

Both read and write tests adequate for 2022 data taking.

At regime, limitations at RAL, NCBJ and CNAF may become relevant.

The LHCb Computing Project is following up with these sites to monitor the evolution of the infrastructure.

Site	expected Speed (GB/s)	Target achieved
CERN	11	Yes
CNAF	1.72	~Yes
GRIDKA	2.23	Yes
IN2P3	1.25	Yes
NCBJ	1.32	No
PIC	0.2	Yes
RAL	2.96	~No
RRCKI	0.25	Yes
SARA	1.07	Yes

Site	Expected Speed (GB/s)	Success
CNAF	1.35	Yes
GRIDKA	1.36	Yes
IN2P3	0.98	Yes
NCBJ	0.91	No
PIC	0.17	Yes
RAL	1.93	No
RRCKI	0.21	Yes
SARA	0.74	Yes

Modernization of the analysis software

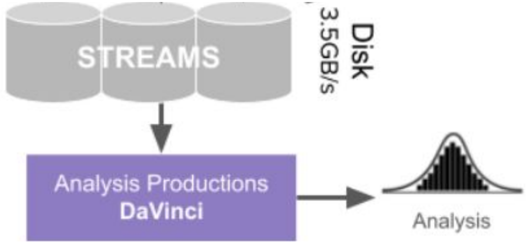

Analysis software package (DaVinci) rewritten from scratch.

Focus on centralized nTuple production, improving:

- CI/CD-based job validation
- reusability of the output
- data & analysis preservation
- job monitoring

LHCb-FIGURE-2020-016

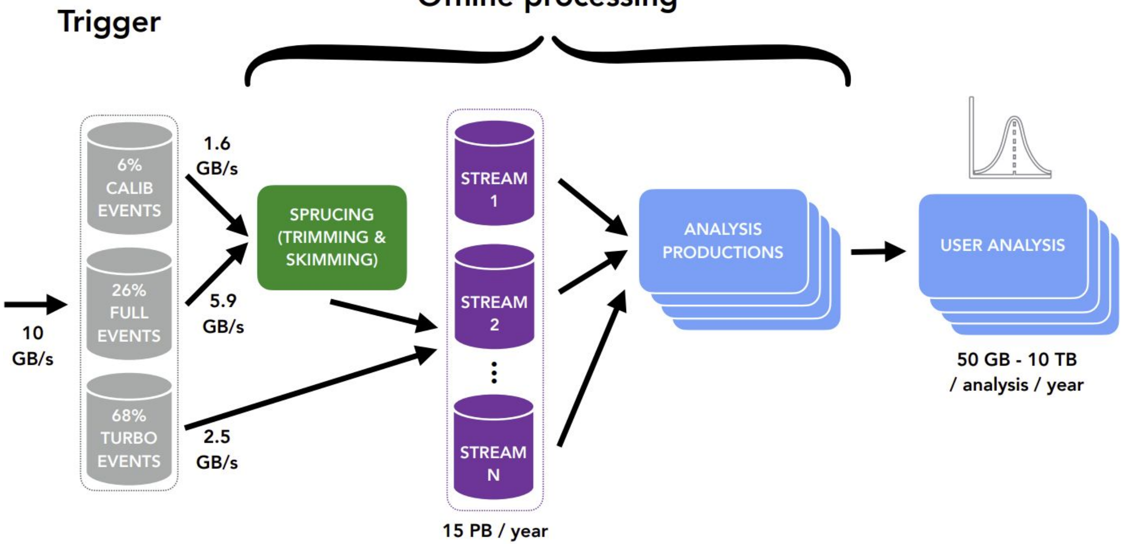

Standardizing the **output data format to nTuples** and extending the application of **functors**, the new DaVinci will be *easier to maintain*, more *open to contributions* and *manageable* as other Computing Operation tasks.

Getting ready for new technologies

GPUs

Padova, Firenze, Ferrara

Ongoing effort to access Marconi100 GPUs with DIRAC. Two ways explored:

- Custom DIRAC Computing Element
- Dedicated CNAF Computing Element submitting the jobs to CNAF (thanks to CNAF, CMS and ATLAS colleagues for setting this up)

Workflow management systems based on **kubernetes** and **snakemake** on INFN Cloud are also being explored to employ GPU nodes for a subset of tasks, only.

Urbino

Quantum technologies

Quantum-computing algorithms can play a role even before having access to a real quantum computer.

For example, a *b*-jet charge identification classifier based on quantum machine learning was submitted in February.

https://arxiv.org/abs/2202.13943

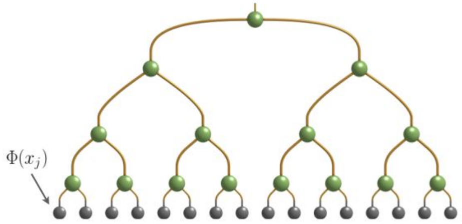

Padova

Firenze, Pisa, Ferrara

Faster Simulations

⚠️ CPU pledges for producing simulated samples won't be sufficient to cope with the MC requests expected for Run3: **need for speed-up the simulation production**.

Developments split in two pillars (more details here):

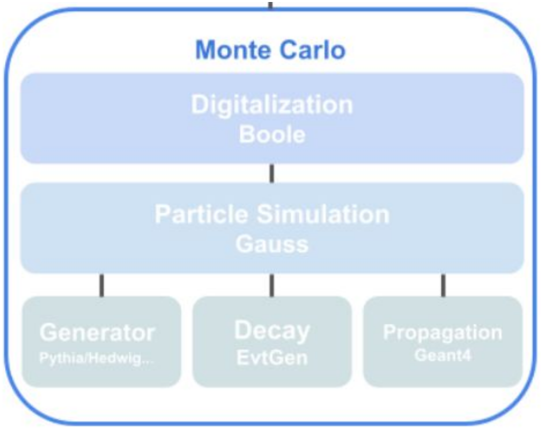

[FastSim] Accelerate Geant4-based simulation

- Upgrade of the simulation framework
- **Drop sub-detectors** from simulation (e.g. track-only simulation)
- Avoid simulating parts of the event (**ReDecay**)
- **Parameterize** some energy deposits instead of relying on Geant4 (e.g. hit library, CaloGAN)

[UltraFastSim] Avoid using Geant4 in simulation

Parameterize the **reconstructed quantities** given the MCTrue values. Provide both:

- a pipeline of parameterizations describing the whole event (the new **Lamarr** framework)
- tools for completing or cross-checking Fast Simulation options

Requests for 2023

LHCb-PUB-2022-010

LHC running scenario and LHCb plans

Further **delay of the LHC due to the pandemic situation** resulted into a reassessment of the needs for computing resources in 2022.

In particular, ***storage resources for 2022 were overestimated*** as the expected integrated luminosity at the end of 2022 is **5.7 times lower than previous expectations**.
CPU (and to some extent, disk) resources will also reduce as the need for Run3-simulated samples, driving the requests, is postponed.

For **2023**, given the LHC schedule, we expect:

- **12 fb^{-1} of proton-proton collision data**
- requests for ion-ion collisions have been shifted from 2022 to 2023, but their amount is unchanged

CPU requests

Sprucing CPU demand assuming consistent with Run2-stripping

Parametric simulation expected to contribute negligibly in 2023

CPU Work in WLCG year (kHS06.years)	2022 LHCB-PUB-2021-002	2023 LHCB-PUB-2021-008	2023 THIS DOCUMENT
First pass sprucing	80	80	80
End-of-year sprucing	80	80	80
Simulation	870	1800	1010
Core and distributed computing infrastructure	10	10	10
User Analysis and working group productions	220	330	240
Total Work (kHS06.years)	**1260**	**2300**	**1410**
LHCb-TDR-018	**1580**	**2750**	

Core and distributed computing infrastructure used mainly for for CI/CD.

User analysis and WG productions expected to scale as $\frac{\mathrm{CPU}(Run3)}{\mathrm{CPU}(Run2)} = \frac{1}{2}\frac{\mathcal{L}(Run3)}{\mathcal{L}(Run2)}$

thanks to the important effort in code optimization and centralized WG productions.

Overall, a 24% increment over the 2022 pledges is requested.

CPU requests

We assume HLT farm will be partly available during the winter shutdowns and similar contributions from opportunistic resources as in the past.

The corresponding computing power is subtracted from the requests to WLCG.

Distribution through different Tiers remains the same as described in LHCb-PUB-2013-002.

CPU Power (kHS06)	2022	2023
Tier 0	189	215
Tier 1	622	705
Tier 2	345	395
Total WLCG	**1156**	**1315**
HLT farm	50	50
Opportunistic	50	50
Total non-WLCG	**100**	**100**
Grand total	**1256**	**1415**

Disk requests

Run1-2 data will not evolve.
Same for simulation.

Most resources to keep analysis-level datasets (FULL + TURBO) readily available for nTupling (WG productions).

Disk storage usage forecast (PB)		2022 LHCb-PUB-2021-002		2023 LHCb-2021-008		2023 This document	
Real data	Run1+Run2 pp data	10.2	65.9	10.2	114.6	10.2	78.8
	Run1+Run2 PbPb + SMOG						
	Run3: FULL	13.7		27.4		16.2	
	Run3: TURBO	30.3		60.6		35.7	
	Run3: TURCAL	3.7		7.4		4.4	
	Run3: Minimum bias	2.4		0.0		0.0	
	Run3: PbPb + SMOG2	5.6		8.9		11.2	
Simulated data	Run1+Run2 Sim. Data	8.7	11.0	8.7	14.9	8.7	11.5
	Run3 simulated data	2.2		6.2		2.8	
Other	User data	1.8	12.8	2.4	13.4	2.4	13.4
	Buffers	11.0		11.0		11.0	
Total			**89.6**		**142.9**		**102.5**
	LHCb-TDR-018		111.0		159.0		

11 PB of disk space reserved for staging 44 PB of data from tape during “re-sprucing”. Corresponding to a two-month procedure with 2-week contingency.

Today: 6 PB over a period of 5 months.

Disk (PB)	2022	2023
Tier0	26.5	30.3
Tier1	52.9	60.5
Tier2	10.2	11.7
Total	**89.6**	**102.5**

Tape requests

Tape request dominated by the need of preserving a two copies of the raw data, at CERN and at Tier-1s.

Tape storage usage forecast (PB)		2022 LHCb-PUB-2021-002		2023 LHCb-2021-008		2023 This document	
Run1 + Run2	RAW data (pp+HI+fixed target)	38.4	82.1	38.4	83.1	36.9	78.4
	RDST data (pp+HI+fixed target)	13.7		13.7		13.8	
	ARCHIVE	30.0		31.0		27.7	
Run3	pp data (FULL+TURBO+TURCAL)	120.1	137.8	240.1	276.2	141.4	169.9
	minimum bias / no-bias	0.6		0.6		0.6	
	Heavy Ion Data + fixed target	5.6		10.6		11.2	
	ARCHIVE (data+MC)	11.5		25.4		16.7	
Total			**219.9**		**359.3**		**248.3**
	LHCb-TDR-018		243.0		345.0		

To cope with limited buffer space at EOS, data must be transferred to Tier-1s will happen during the data-taking, representing a significant challenge for the infrastructure.

Tape (PB)	2022	2023
Tier0	81	91
Tier1	139	157
Total	**220**	**248**

CERN-RRB-2022-044

Share through countries (Tier1s+Tier2s)

CERN researchers cannot contribute to computing resources beyond what is allocated in Tier-0.

Countries not hosting a Tier-1 cannot contribute to Tape.

LHCb requests distribute the missing resources through the remaining sites, according to fraction of PhD equivalents.

$$\mathrm{CPU}(\bar{c}) = \frac{\mathrm{PhD}(\bar{c})}{\sum_{c\neq\mathrm{CERN}} \mathrm{PhD}(c)} \times \mathrm{CPU}_{\mathrm{cRSG}} \qquad \mathrm{DISK}(\bar{c}) = \frac{\mathrm{PhD}(\bar{c})}{\sum_{c\neq\mathrm{CERN}} \mathrm{PhD}(c)} \times \mathrm{DISK}_{\mathrm{cRSG}}$$

$$\mathrm{TAPE}(\bar{c}) = \frac{\mathrm{PhD}(\bar{c})}{\sum_{c\in\{\mathrm{w/\ Tier1}\}} \mathrm{PhD}(c)} \times \mathrm{TAPE}_{\mathrm{cRSG}}$$

Such sharing scheme is considered unfair by some funding agency, including INFN, which contributes by its fair-share (**16%**)

2023	PhD eq. total/ funding auth.	%
AUSTRALIA	3	0.5
BRAZIL	19	3.3
COLOMBIA	2	0.3
FRANCE	41	7.1
BMBF GERMANY	29	5.0
MPG, GERMANY	2	0.3
IRELAND	1	0.2
INFN ITALY	93	16.0
HUNGARY	1	0.2
NETHERLANDS	22	3.8
P. R. CHINA	45	7.8
POLAND	22	3.8
HHNIPNE ROMANIA	5	0.9
RUSSIA INST.	40	6.9
RUSSIA UNI.	15	2.6
SPAIN	24	4.1
SWEDEN	1	0.2
SWITZERLAND	26	4.5
UKRAINE	1	0.2
UK	88	15.2
USA	32	5.5
CERN	68	11.7
TOTAL	580	100.0

Requests and pledges for 2022

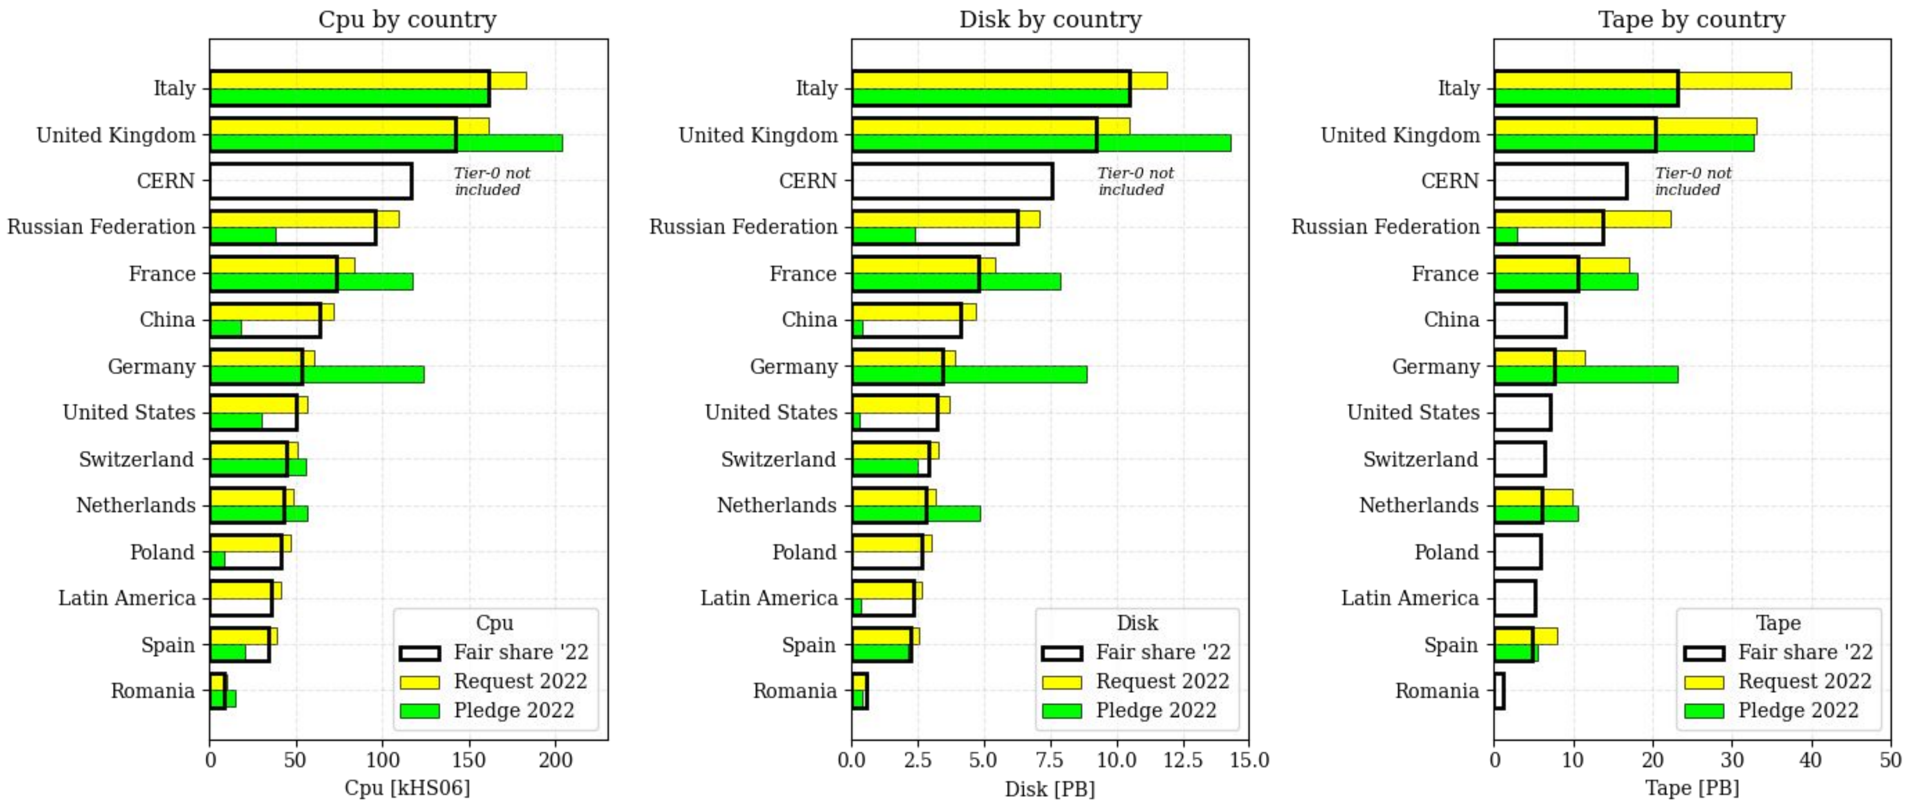

Reassessed requests for 2022 (due to LHC delays)

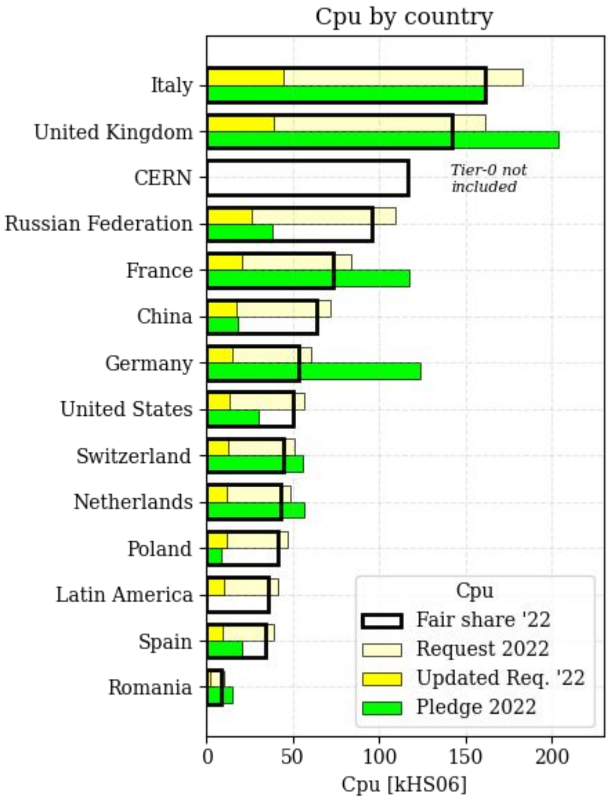

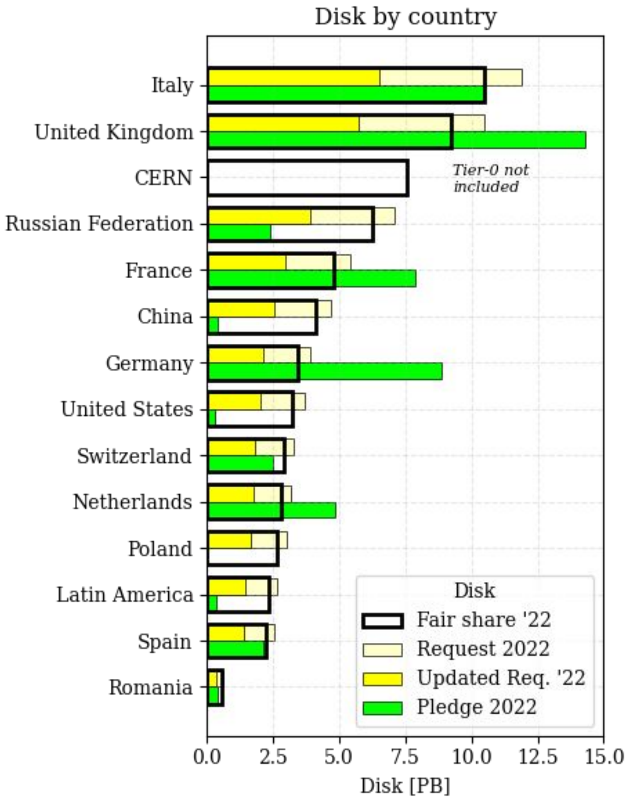

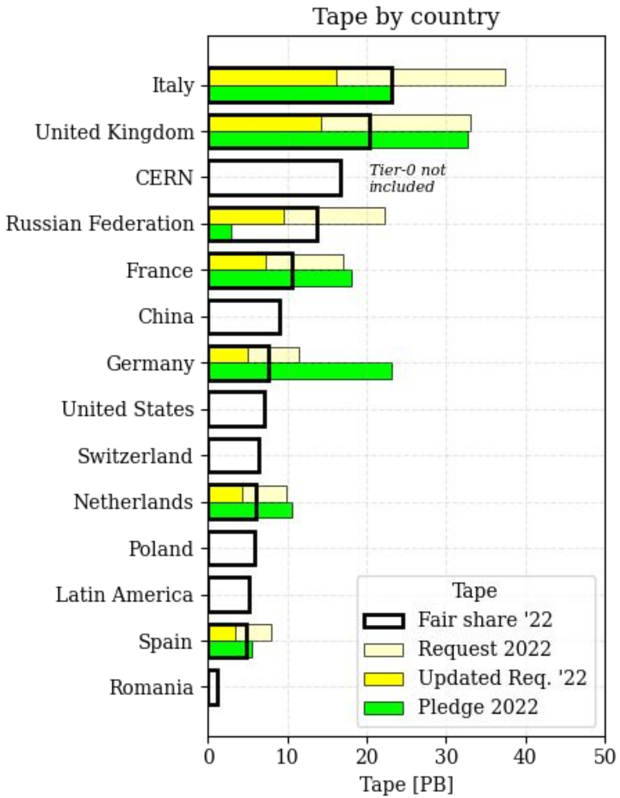

Requests for 2023

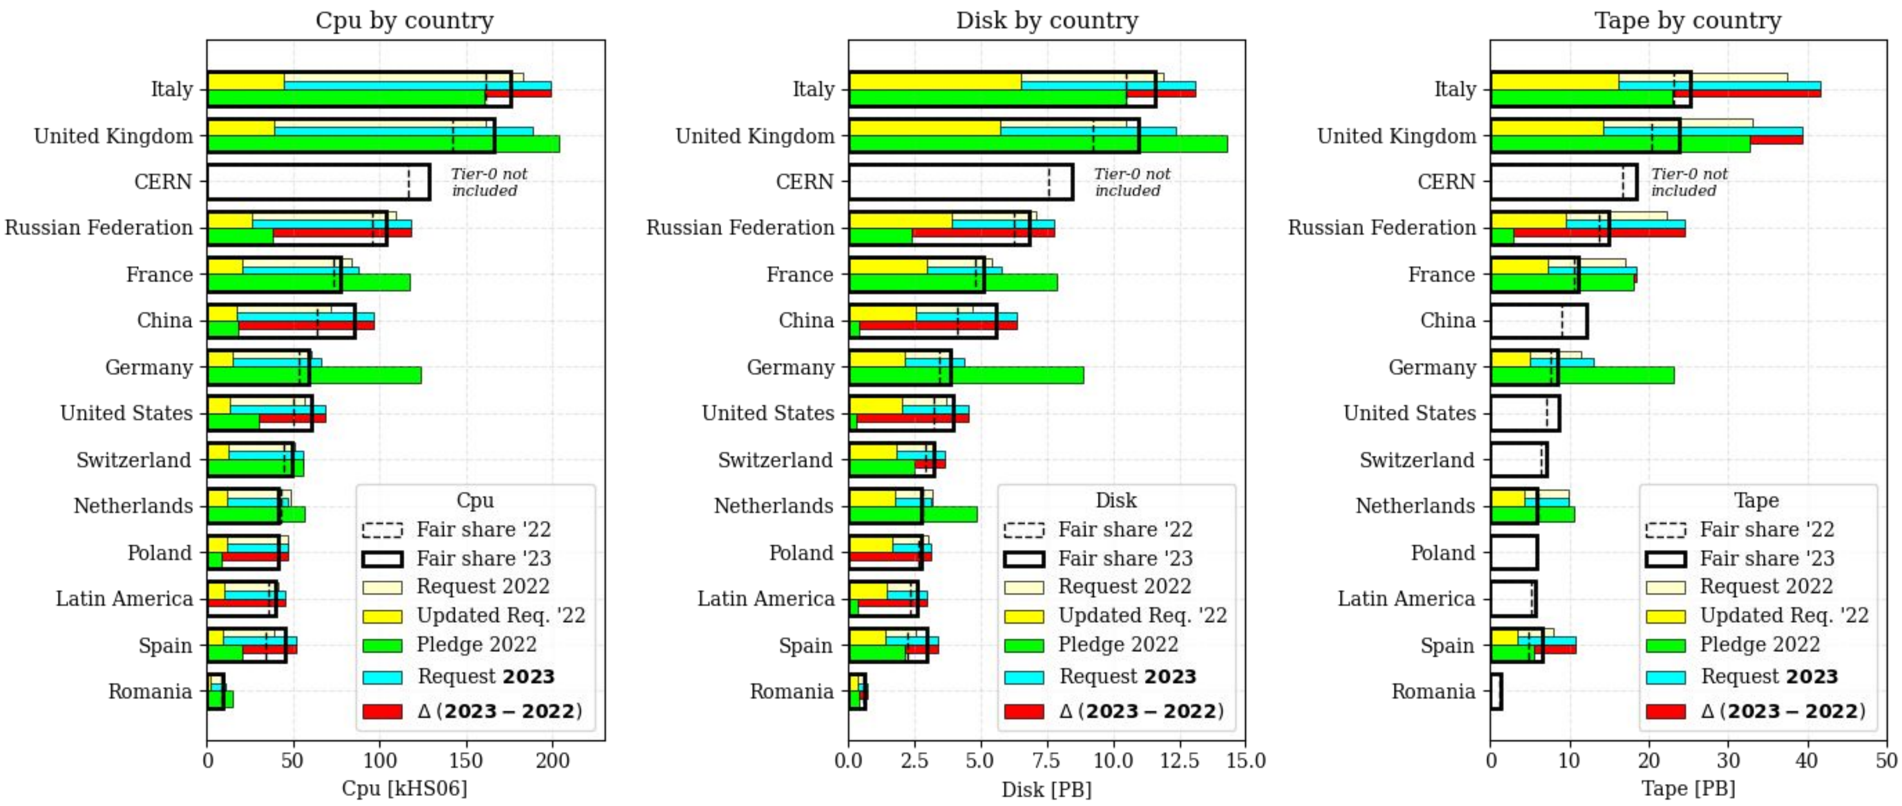

2021 Risk assessment

LHCb-PUB-2021-002

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	Respruncing 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	
Underestimation of analysis work	2	2	4	LHCb	more computing power needed	
Underestimation of simulation work	3	3	9	LHCb/G4/HSF	Cannot afford sufficient simulation; physics analysis suffers	

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

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

2022 tape challenge WRITE

Site	expected Speed (GB/s)	Target achieved
CERN	11	Yes
CNAF	1.72	~Yes
GRIDKA	2.23	Yes
IN2P3	1.25	Yes
NCBJ	1.32	No
PIC	0.2	Yes
RAL	2.96	~No
RRCKI	0.25	Yes
SARA	1.07	Yes

Final consideration on computing requests

Shortage of CPU and DISK resources may significant delay the physics output of the LHCb experiment, but we do not expect descoping of the program.

Lack of tape resources is worrisome, especially considering that increasing the available tape resources can hardly happen in an emergency.

Sticking to fair-share for CNAF tape resources represents a severe risk for LHCb data-taking.

Big effort ongoing in the Collaboration to add tape-enabled computing centers from other countries: NCBJ (Poland), IHEP (China), MGHPCC (USA)

LHCb - 2023 Requests to INFN-T1

Requests as approved by cRSG: LHCb-PUB-2022-010

LHCb		2023 Request	2023 req./ 2022 CRSG	2023 req. / 2022 pledge
WLCG CPU	Tier-0	215	114%	114%
	Tier-1	707	114%	137%
	Tier-2	391	113%	118%
	HLT	50	100%	100%
	Sum	1364	113%	126%
Others		50	100%	n/a
Total		**1,414**	**113%**	**124%**
Disk	Tier-0	30.3	114%	114%
	Tier-1	60.5	114%	127%
	Tier-2	11.6	114%	168%
	Total	**102.5**	**114%**	**126%**
Tape	Tier-0	91	113%	112%
	Tier-1	157	113%	135%
	Total	**248.3**	**113%**	**126%**

Pledge 2022 as from CRIC: wlcg-cric.cern.ch/core/pledge/list/

Federation	VO	Type	Pledge
IT-INFN-CNAF	LHCb	Tape	23044 TBytes
IT-INFN-CNAF	LHCb	CPU	103118 HEP-SPEC06
IT-INFN-CNAF	LHCb	Disk	10461 TBytes
IT-INFN-T2	LHCb	CPU	57197 HEP-SPEC06

Fair share INFN 2023

$$\frac{\mathrm{PhD}(\bar{c})}{\sum_c \mathrm{PhD}(c)} = 16.0\%$$

LHCb-share for INFN tape 2023

$$\frac{\mathrm{PhD}(\bar{c})}{\sum_{c\in\{\mathrm{w/\ Tier1}\}} \mathrm{PhD}(c)} = 26.4\%$$

Request	Pledge '23	Increment
CPU Tier-1 [HS06]	113428	+10310
DISK Tier-1 [TB]	11561	+1100
Tape Tier-1 [TB]	41624*	+18580
CPU Tier-2 [HS06]	62595	+5398

*Tape request computed according to the LHCb formula.

$$\mathrm{TAPE}(\bar{c}) = \frac{\mathrm{PhD}(\bar{c})}{\sum_{c\in\{\mathrm{w/\ Tier1}\}} \mathrm{PhD}(c)} \times \mathrm{TAPE}_{\mathrm{cRSG}}$$

Fair share would correspond to **25261 TB** (delta: **+2217 TB**)
Difference (extra request beyond FS): **+16363 TB**

Conclusion

An important effort is ongoing within LHCb to **modernize the software and the development cycle**. The Italian community is a very relevant player in this game.

Covid-19 caused significant **delays to the start of Run-3** which set LHCb in a condition of *under-utilization of the computing resources*, in particular for the storage.
This is however expected to change very quickly with the start of the data-taking, given the **large bandwidth of the experiment** and *"debugging"* needs related to the **commissioning of a new experiment**.

CNAF is one of the four major Tier1s providing resources to LHCb and is technologically ready for the upcoming challenge (though with some bottleneck, well identified).

Pledging the resources requested by the experiment (especially tape) will consolidate the leading role of the Italian community in the experiment, contributing to its success.

Backup

Firenze (important synergy with *ML_INFN* and *INFN Cloud*)

Ultra-fast simulation: Lamarr

📯 Lamarr consists of a **pipeline of modular (ML-based) parameterizations** designed to replace the whole Simulation and Reconstruction steps.

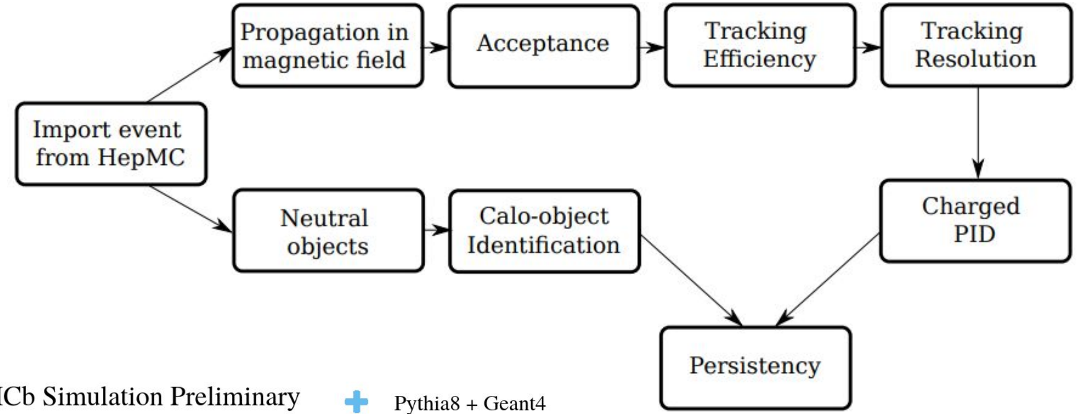

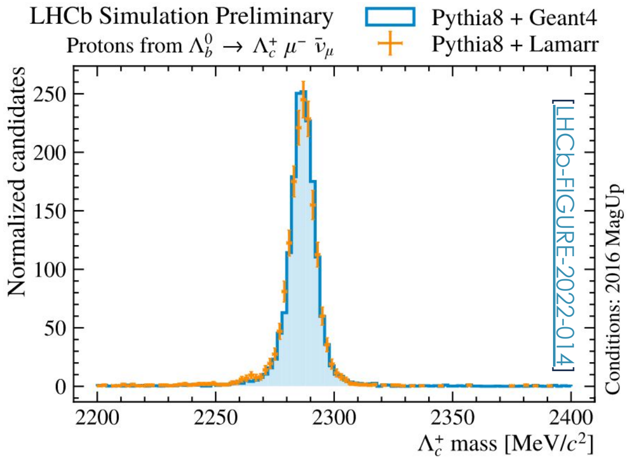

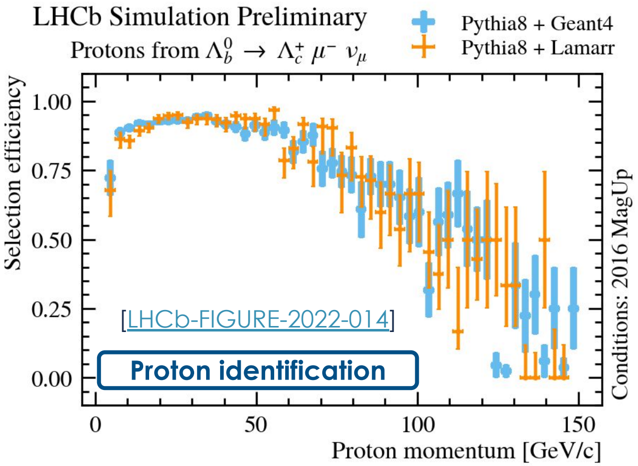

Development is focusing on Run-2 which will dominate simulation requests for the next few years.

Third tape challenge: gory details

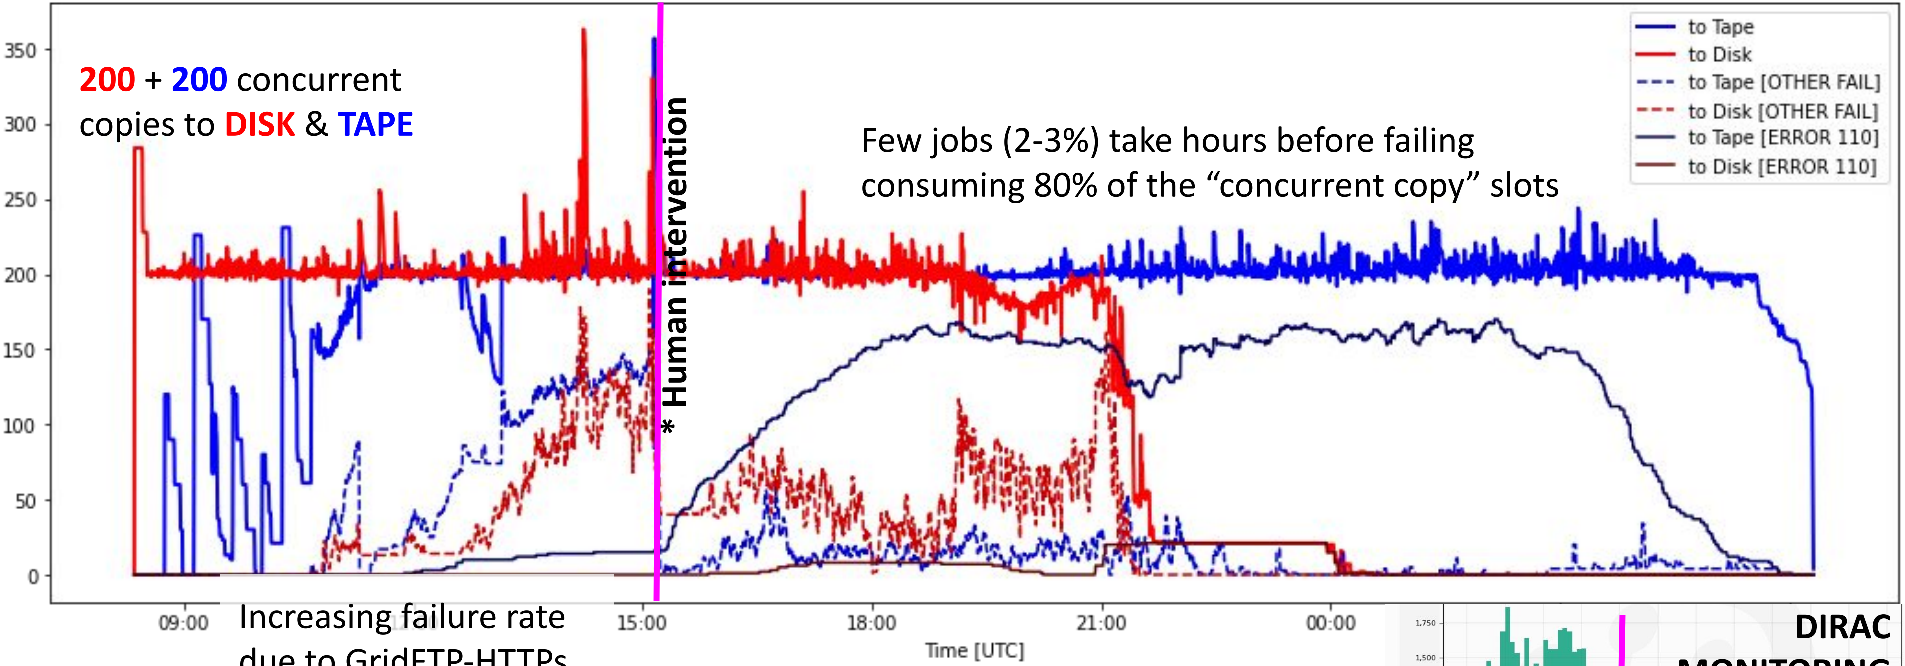

Increasing failure rate due to GridFTP-HTTPs interference

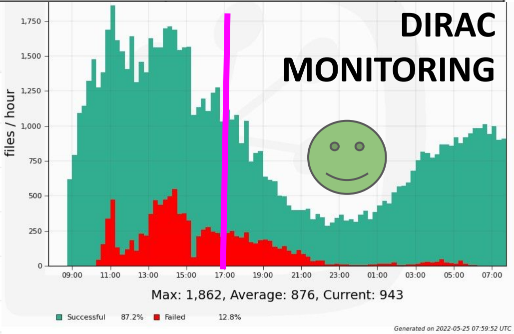

* Human intervention: split endpoints supporting WebDAV (3/4) and GFTP (1/4)

Effect of Fast & Filtered Simulations

As expected, during LS2 we had an explosion of requests for simulated samples, combining Run1 and Run2 data-taking conditions.

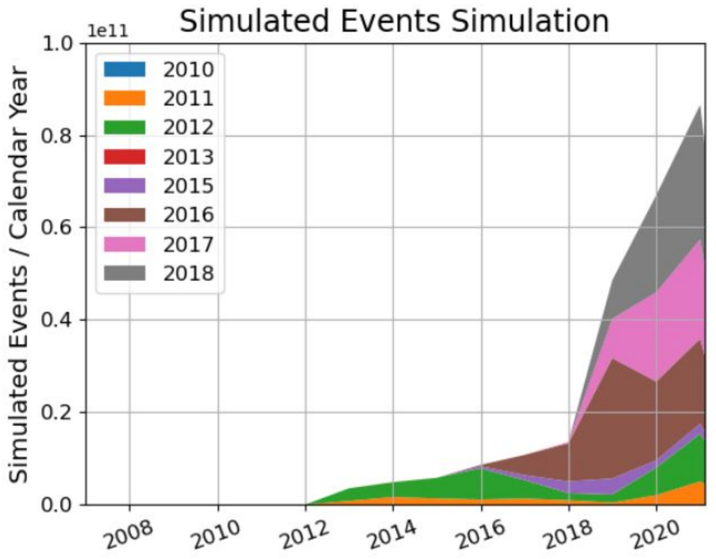

Year	Simulated events (10^9)	Stored events (10^9)	Ratio	CPU work kHS06.y	CPU per event kHS06.s	LFS TB
2017	10.3	4.2	40.3%	817	2.50	640
2018	12.0	3.0	25.3%	1009	2.65	550
2019	45.0	6.9	15.2%	1290	0.90	1110
2020	67.0	16.8	31.7%	1357	0.81	2010
2021	80.0	11.1	13.9%	1815	0.72	2030

*Logical File Size

Fast Simulation options (ReDecay) enabled a drastic reduction of CPU/event.

Filtered productions reduced the impact of simulation on storage resources.