



CMS Update

D. Piparo (CERN), M. Klute (MIT) - LHCC, June 3rd, 2020 - Meeting with the LHCC referees



Usage of Resources During the Lockdown




Contribution to CERN Against COVID-19 Taskforce

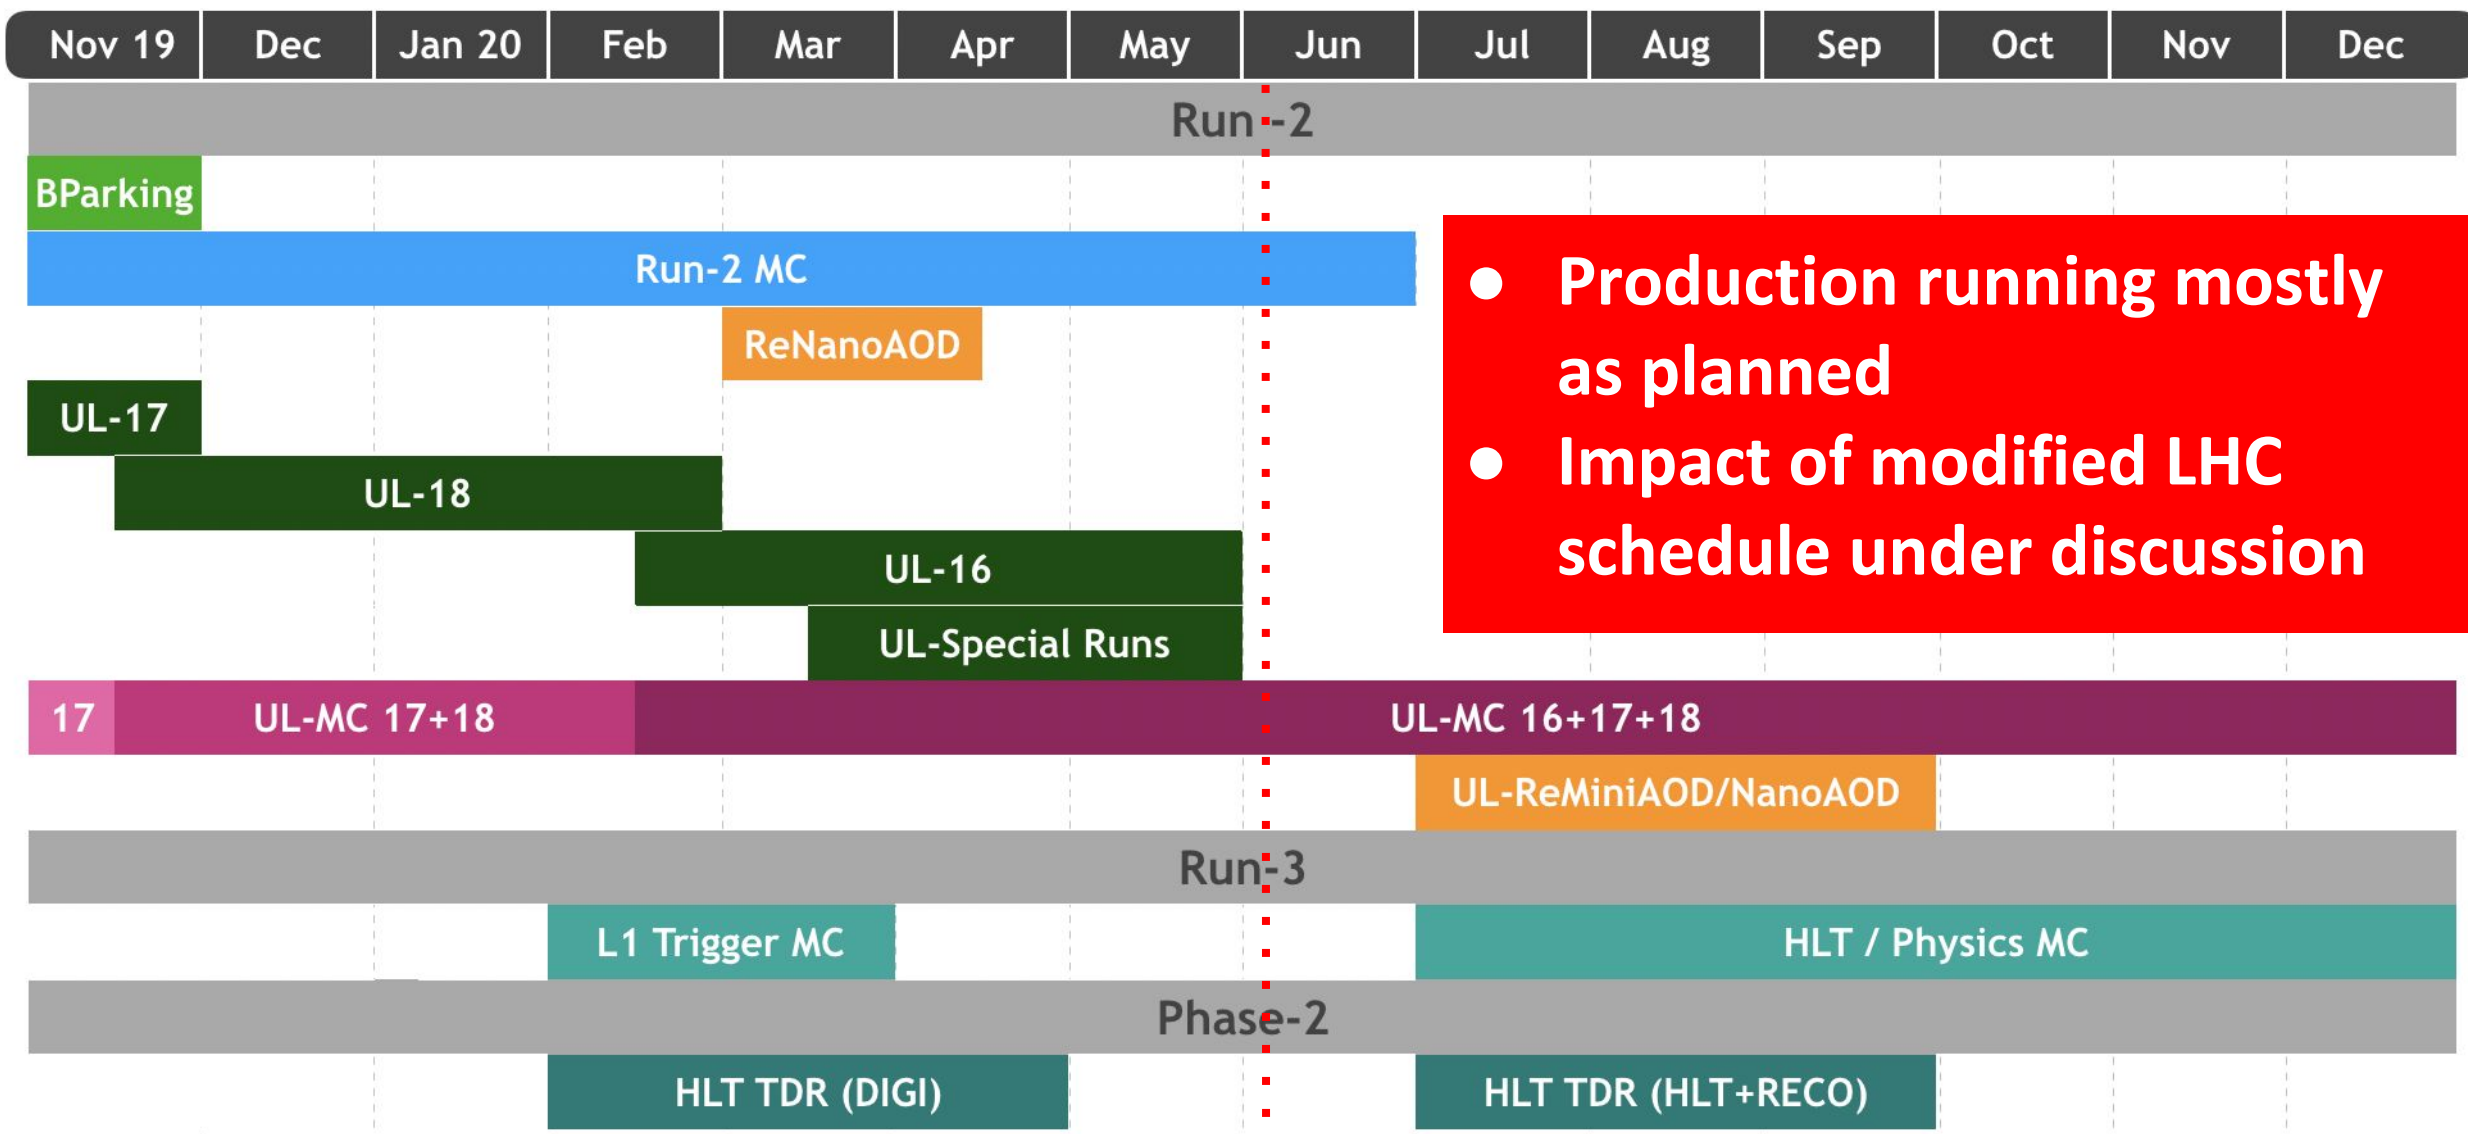
- Folding@Home: initiative chosen in the context of the CERN Against COVID-19 task force
- Running stably on about 64k job slots, divided in three main categories:
 - Opportunistic execution on machines dedicated to CMS internal services: ~700 cores
 - Grid distribution: EU sites, via opt-in agreement: ~2000 cores
 - HLT farm: running 30k physical cores (60k virtual)
- **Great flexibility demonstrated, lessons learned which are useful for future CMS processing!**
 - Direct injection via HTCondor, seamless handling of a non-CMS application
 - Accurate internal monitoring via standard CERN MONIT infrastructure
 - Potential application at backfilling Grid slots with CMS jobs (e.g. CMS@Home tasks) in the future
- **Plan to ramp down donation in the near future**



Strong contribution substantially driven by unpledged resources, lots learned for future CMS operations!



Production update since Nov LHCC




- Production running mostly as planned
- Impact of modified LHC schedule under discussion




Progress on Legacy Production

- The Run-2 Legacy program was started two years ago
- Since August '19, **40B events have been produced**
 - including 32B events of data, covering standard pp 13 TeV collisions, special runs (not including ~12B B-Parking events)
 - 8B events of MC, so far, for 2017-2018, used by Physics Object Groups (POG) to produce final energy corrections and efficiency scale factors. More MC to come.



NanoAOD(Sim) Production

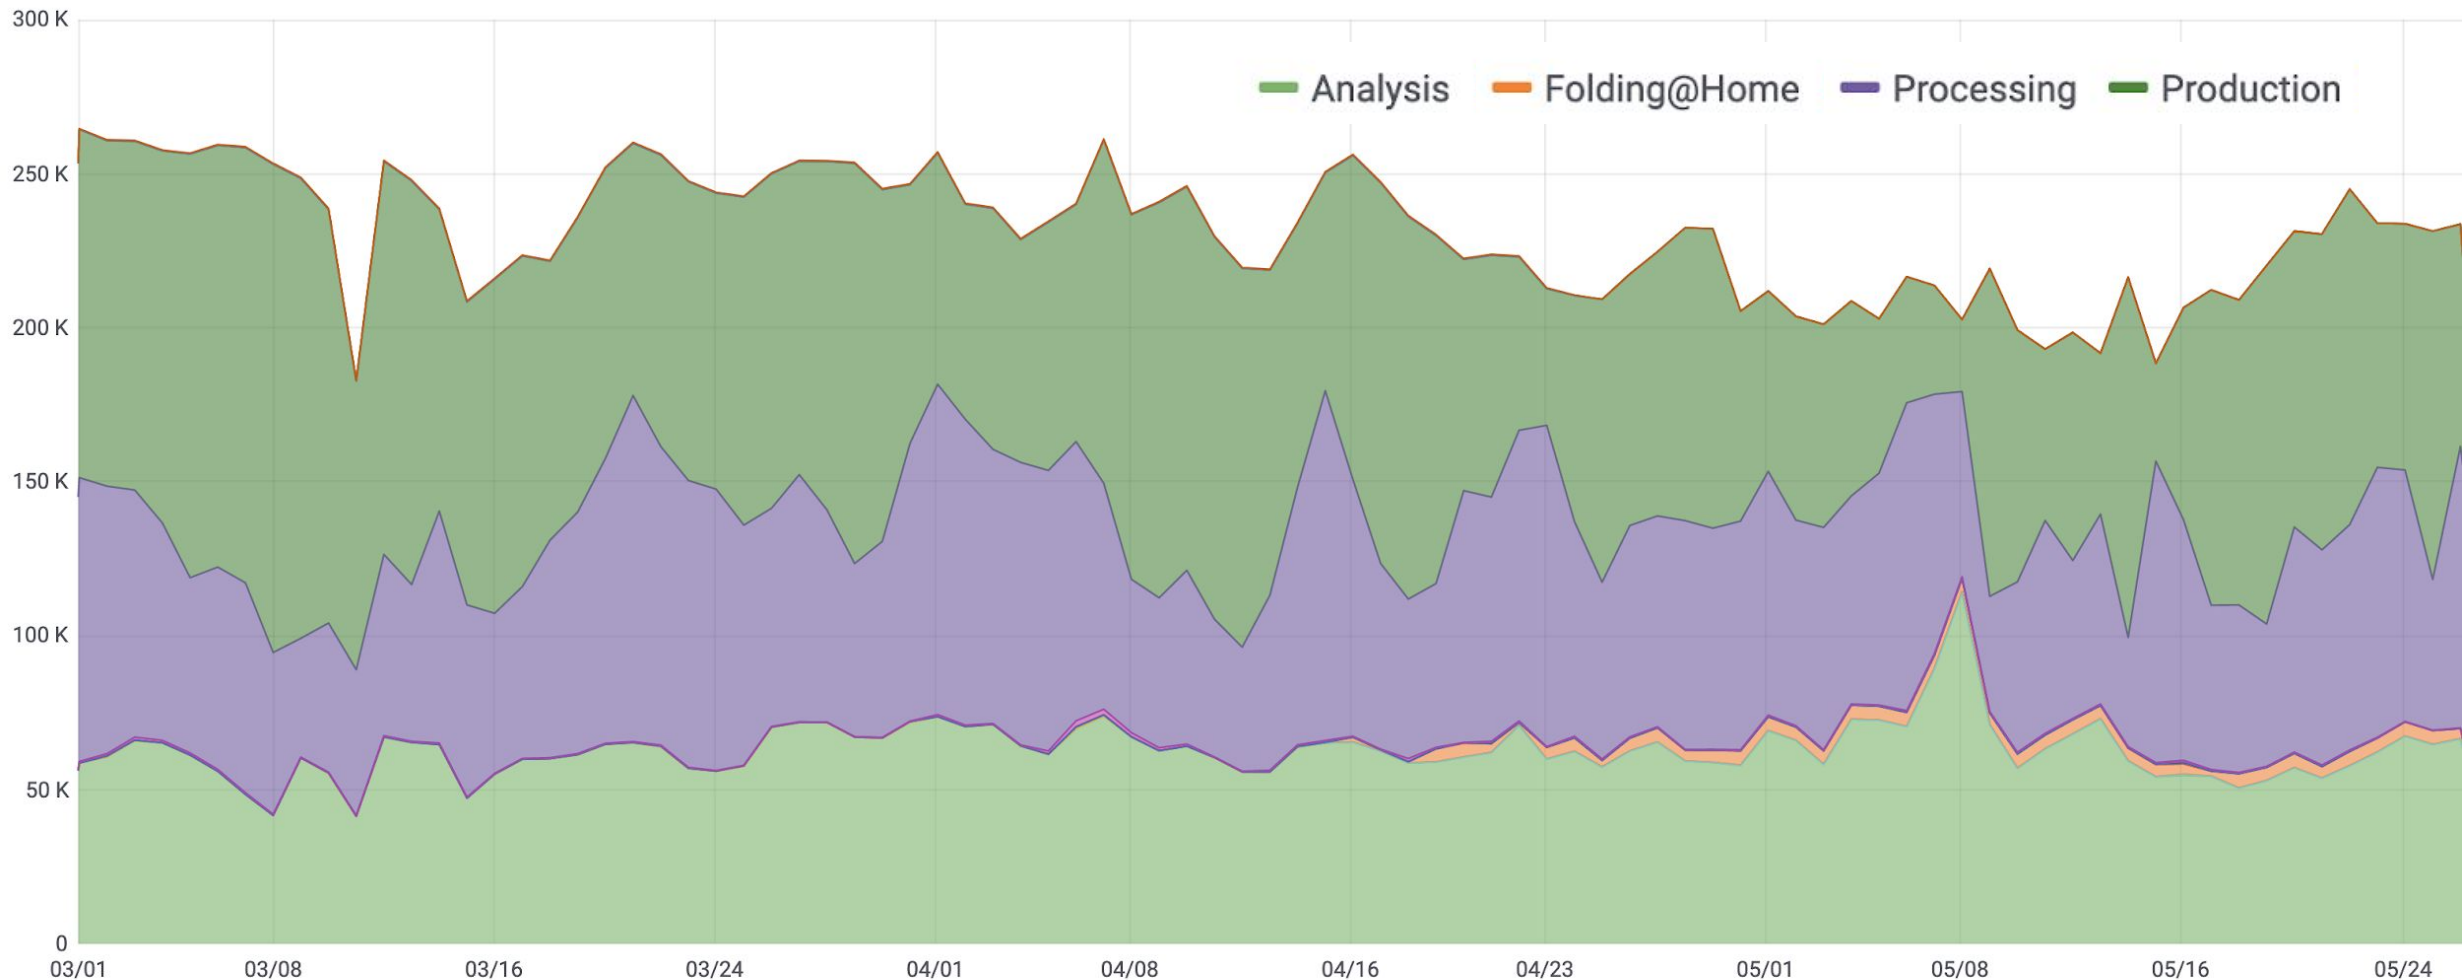
- Previously shown 4B/day possible. Peak of 25 B/day achieved
- Lasted campaign of 60B non-legacy Run 2 events unveiled limits in production system that are being addressed
- Four-threaded jobs for producing NanoAOD





Core Usage for the CMS Physics Program


- Excellent usage of resources during lockdown: we are thankful to the sites supporting CMS!
 - HLT farm was off and then used for f@h - not monitored here



Use of HPCs for Monte Carlo production

- Not only pledged resources were used in recent months
- Example: **Marconi Supercomputer at CINECA**
 - Excellent utilization of this KNL machine
 - Site extension mechanism: HPC not visible to CMS operations, CNAF T1 capacity transparently increased
- Other HPCs also in production

**Site Extension: exploitation of HPCs
Transparent to computing operations**



Preparations for Run 3



Transitioning from PhEDEx to Rucio

- **PhEDEx: data management tool used by CMS for more than a decade**
 - It served CMS very well, especially in conjunction with the [Dynamo DDM](#), but ...
 - ... an in-house product, expensive to sustain for Run 3 and beyond
- Rucio has been identified in 2018 as the technology to replace PhEDEx in CMS
 - A community project adopted, among others, by ATLAS and DUNE
- A transition is ongoing since 2019
- **Rucio took ownership of NanoAOD(Sim) datasets:** start of PhEDEx phaseout
 - Transparent for CMS physicists, e.g. data discovery and analysis jobs!
- **Next steps**
 - Rucio take over data location services for all data tiers by mid-summer
 - Transitioning additional data tiers to be fully managed by Rucio starting in the fall
 - Turn off PhEDEx after that



Major Milestone: CMS uses Rucio to manage NanoAOD(Sim) datasets

Tape Usage and Incremental Processing

- **Excellent tape throughput in the last 12 months**


- Average over 1 year > 350 TB/day
- Peaks above 450 TB/day

- **Incremental processing possible since 2017**

- No need to stage entire datasets on disk
- Block-wise processing
- Deletion of blocks after processing

- Crucial feature for the simultaneous (re)-reconstruction of B-parked and Heavy Ions dataset in 2019!

- **Plan after migration to Rucio is completed: optimise further tape throughput**



Steadily Improving in Preparation of Run 3 - I/2

■ Geant 4: CMS officially adopted version 10.6, released in January

- Quite some work to achieve excellent physics performance, now done
- **~10% speedup measured with early release confirmed**
- Further performance improvements in the pipeline




■ ROOT: CMS will incorporate version 6.22 as soon it's out

- **All tests are ok using most recent versions of their git repo**
- ROOT 7 enabled: all new functionality will be offered to Physicists (RNTuple, new graphics, new PyROOT)



■ Migration to DD4Hep for the detector description towards completion

- Community tool which will be used by LHCb in Run 3, among others
- Thread friendly, modern and more sustainable than previous in-house solution



DD4Hep Migration Progress Map

	CSC	DT	RPC	GEM	ME0	Pixel	TRK	ECal	HCal	HGCal	MTD	PPS	MF
XML description	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done
Algorithms	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	In progress	Done
Overlaps	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	In progress	Done
DB reader	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	In progress	Done
Sim app	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	In progress	Done
G4 geometry	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	In progress	Done
Reco geometry	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	In progress	Done
Reco / Sim numbering	Done	Done	Done	Done	Done	In progress	In progress	In progress	In progress	In progress	In progress	In progress	In progress
Production cuts	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	In progress	Done
SD classes in SimG4CMS	Done	Done	In progress	In progress	In progress	In progress	In progress	In progress	In progress	In progress	In progress	In progress	In progress
Validation	In progress	In progress	In progress	In progress	In progress	In progress	In progress	In progress	In progress	In progress	In progress	In progress	In progress
Alignment	Done	In progress	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done	Done

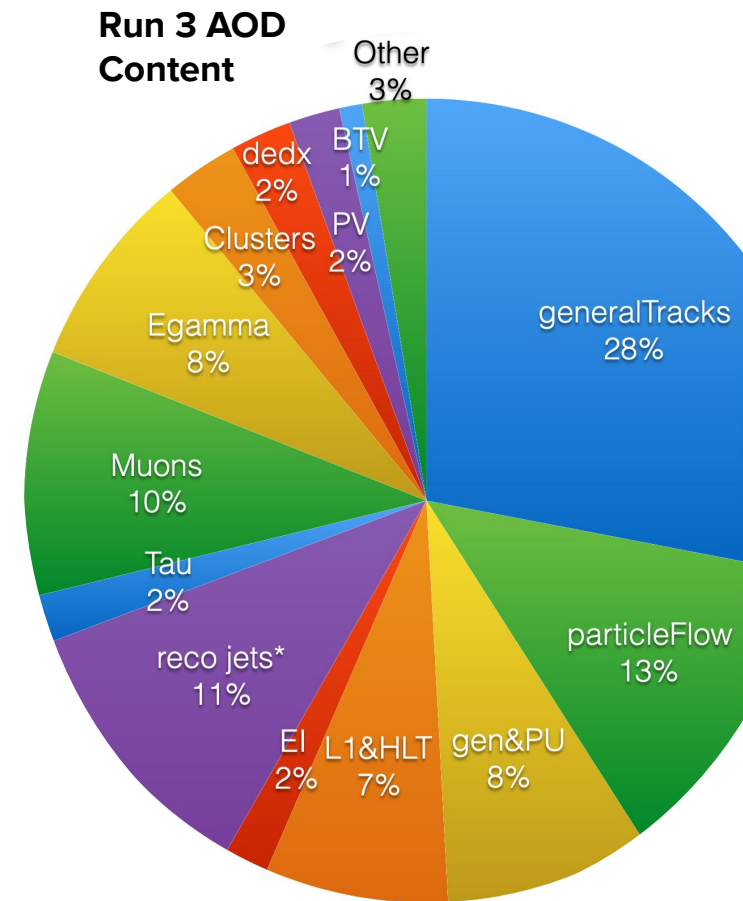
Done
In progress
Not needed



Steadily Improving in Preparation of Run 3 - 2/2

- **Possible to build packages of CMSSW for multiple instruction sets**
 - Before: sse3+vectorisation to fit all hardware on the Grid
 - Now: avx, avx2, fma, ...: select the right library once on the node, exploit the available hardware at best
 - Actively looking for pieces of code which will benefit (e.g. vertexing)
- **CMSSW Framework can now encapsulate thread unsafe code in processes**
 - Orchestrate forked processes coherently with the threads in the pool
 - E.g. allow to run multiple instances of thread unsafe Monte Carlo generators (e.g. EvtGen) inside a multistep MT job
- **Data tiers: striving to improve storage utilisation in Run 3**
 - 6% reduction of AOD, work ongoing to reduce this further through careful selection of information to be kept
 - NanoAOD: 20% of analyses adopted the smallest data tier

Continuous innovation of software and computing tools, targeting Run 3



Elements of HL-LHC Computing





- CMS prepared documentation highlighting areas of future investments and including estimation of resources needs for production year with agreed on parameters [[public plots](#)]
- **Thankful to the committee for their efforts and inputs**


Resource estimates follow bottom-up approach with well defined set of assumptions

- **10% rate year-by-year speedup for 3 years (2024-2027) of reconstruction software (from past experience)**
- **½ of CMS analyses based on NanoAOD**
- Practices
 - Studies before data taking, each year: 4B Monte Carlo events (commissioning/trigger menu)
 - Physics analysis, each year: 5B Monte Carlo events + 0.2 B per fb⁻¹ of luminosity collected
- **2024-2027**
 - Replicas on disk of the current processing version of each event:
 - **1.5 for MiniAOD, 3 for NanoAOD, <1 for AOD**
 - Replicas on tape:
 - **2 for RAW : one at Tier-1s and one at the Tier-0**
 - 1 for derived data samples (AOD, MiniAOD, NanoAOD)



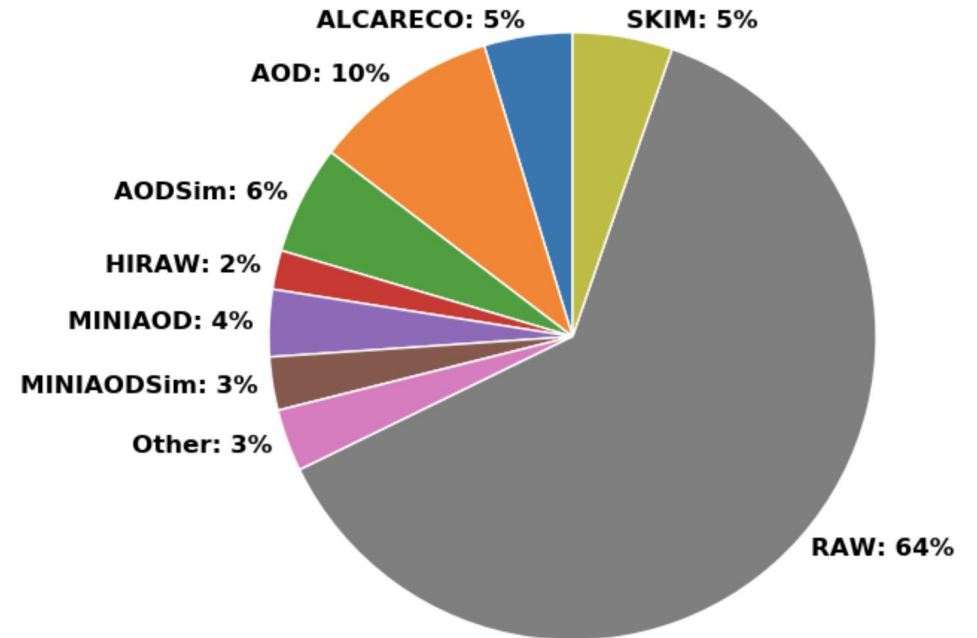
HL-LHC Resources: Top Consumers in 2029

CMS Public
Total CPU HL-LHC fractions
2020 estimates




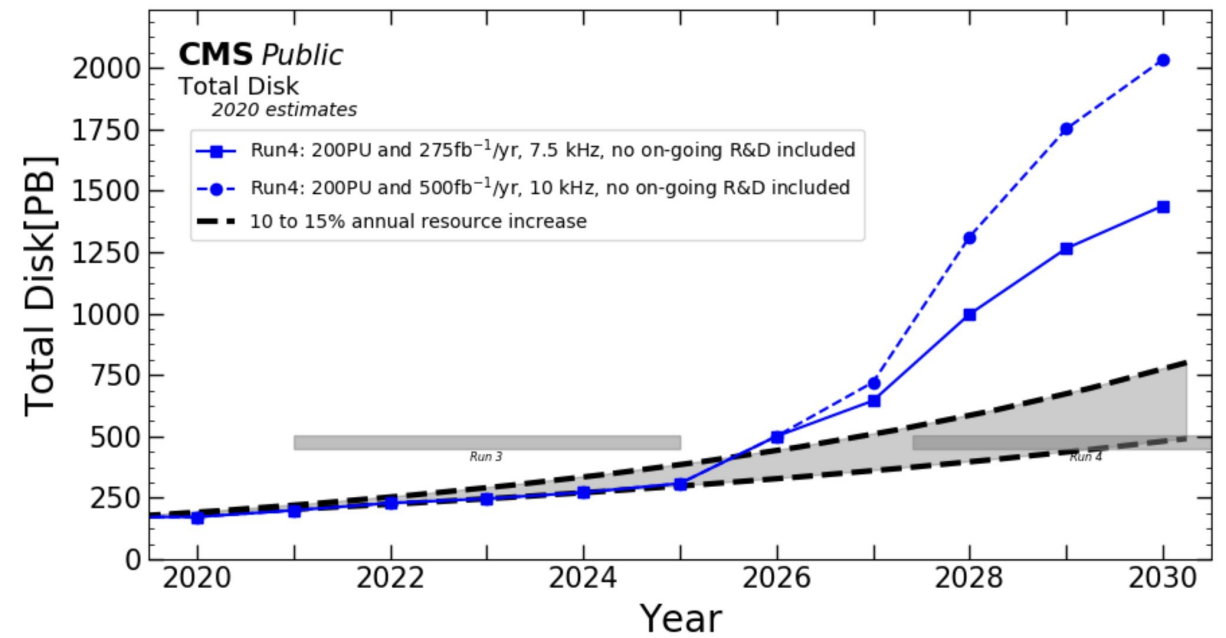

Compute: Reconstruction

CMS Public
Total Tape usage HL-LHC fractions
2020 estimates




Tape: RAW data storage

HL-LHC Resources by Year



Conclusions



Handing over the reins

From September 1st 2020 to August 31st 2022, **James Letts (UCSD)** will be **CMS Offline Software & Computing co-coordinator** replacing Markus Klute (MIT) and joining Danilo Piparo (CERN).



Conclusions

- **Excellent utilisation of computing resources during this difficult period**
 - Even with the entire team working from home
 - We thank the host lab and all of our sites for the excellent performance
- **Successful exploitation of an HPC machine via the *Site Extension* mechanism**
 - Regular workflows ran, executed transparently for the operations team
- **Intense preparation ongoing for Run 3**
 - **Adoption of Rucio in production for NanoAOD(Sim), transparent for users: phaseout of PhEDEx will follow**
 - Transition of the detector description technology to the DD4HEP community tool towards completion
 - Latest versions of ROOT and G4 included in the stack, validation completed, performance perks already there!
 - Can now handle in a multithreaded job thread unsafe code, such as some generators
- **First step of the HL-LHC Computing Model Review**
 - Projected resource needs published
 - Development and Research lines discussed, lots of useful input received!