WLCG Status Report

S. Campana (CERN)

Acknowledgements

Many thanks to the WLCG experiments' computing coordinators, the ROOT and GEANT4 representatives and the WLCG software liaisons for the input and useful discussion

The WLCG Collaboration - March 2023

66 MoU's
165 sites
42 countries

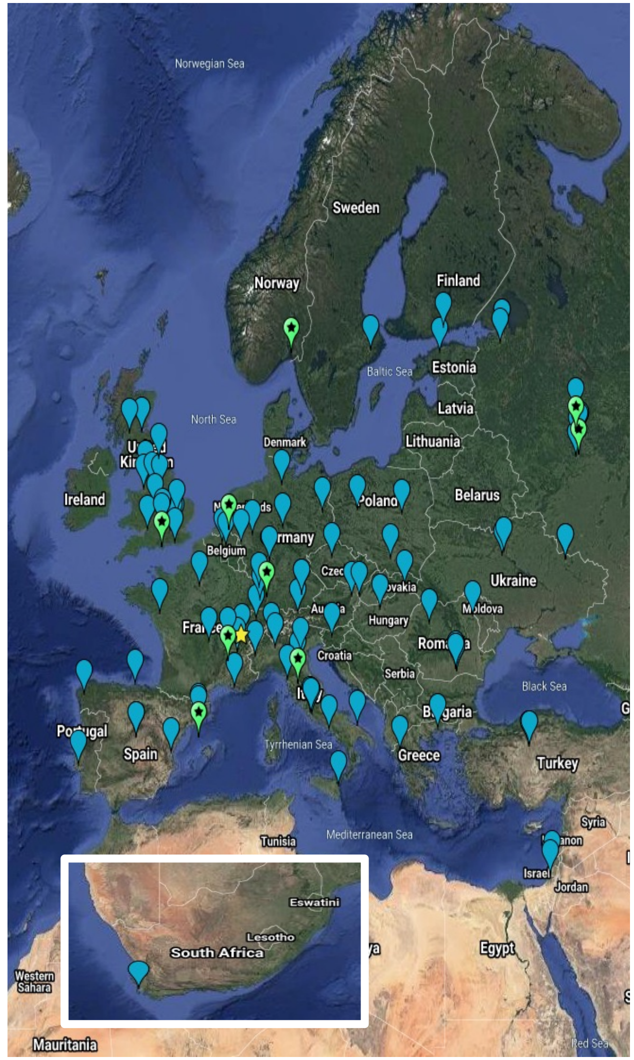

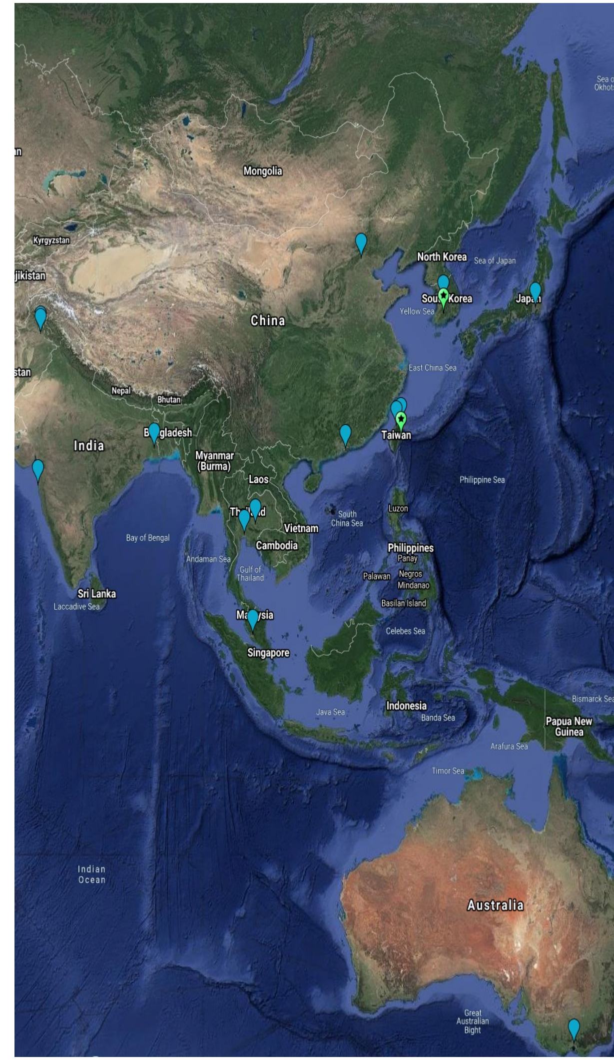

Run-3 data taking

Data written in the CERN tape storage per month (since 2008)

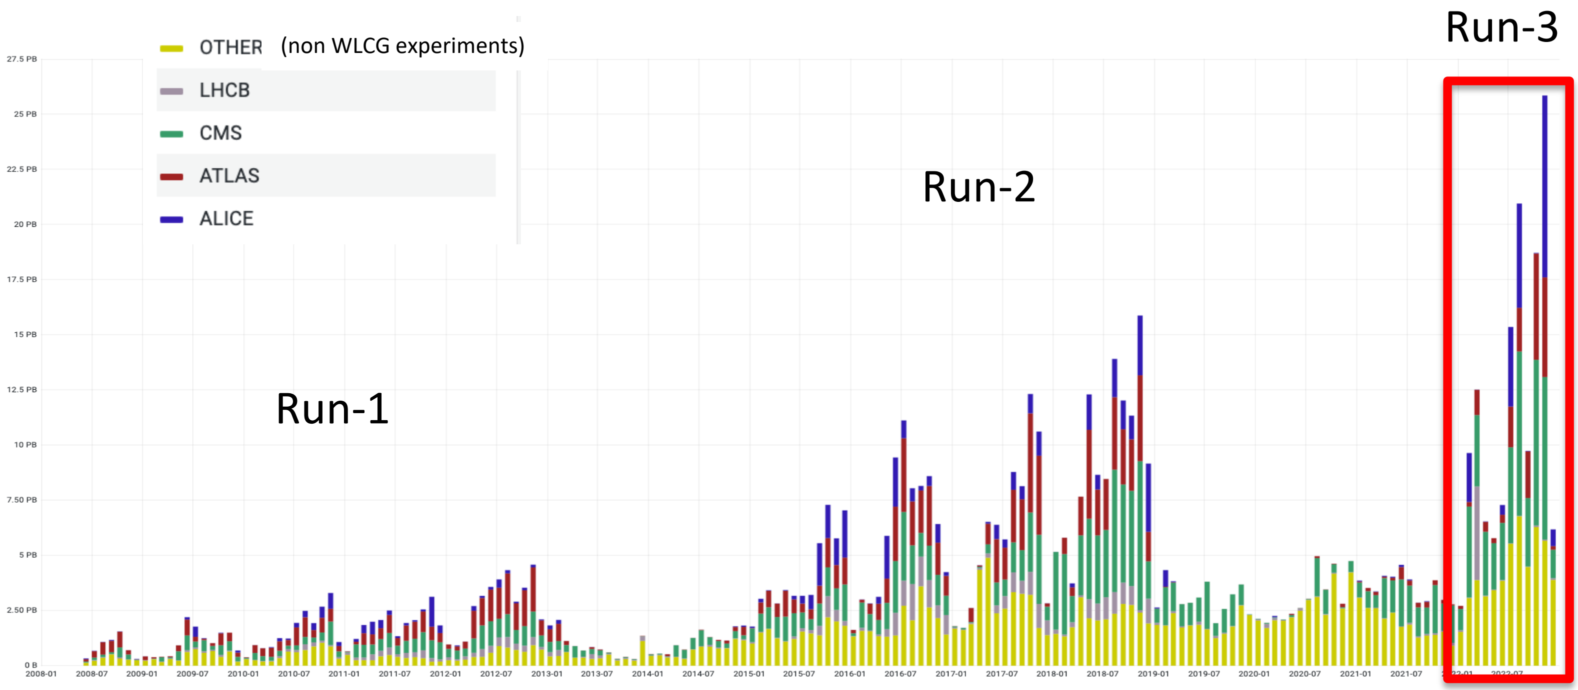

26 PB of data written on tape in November 2022

Data Export from CERN - last 2 years

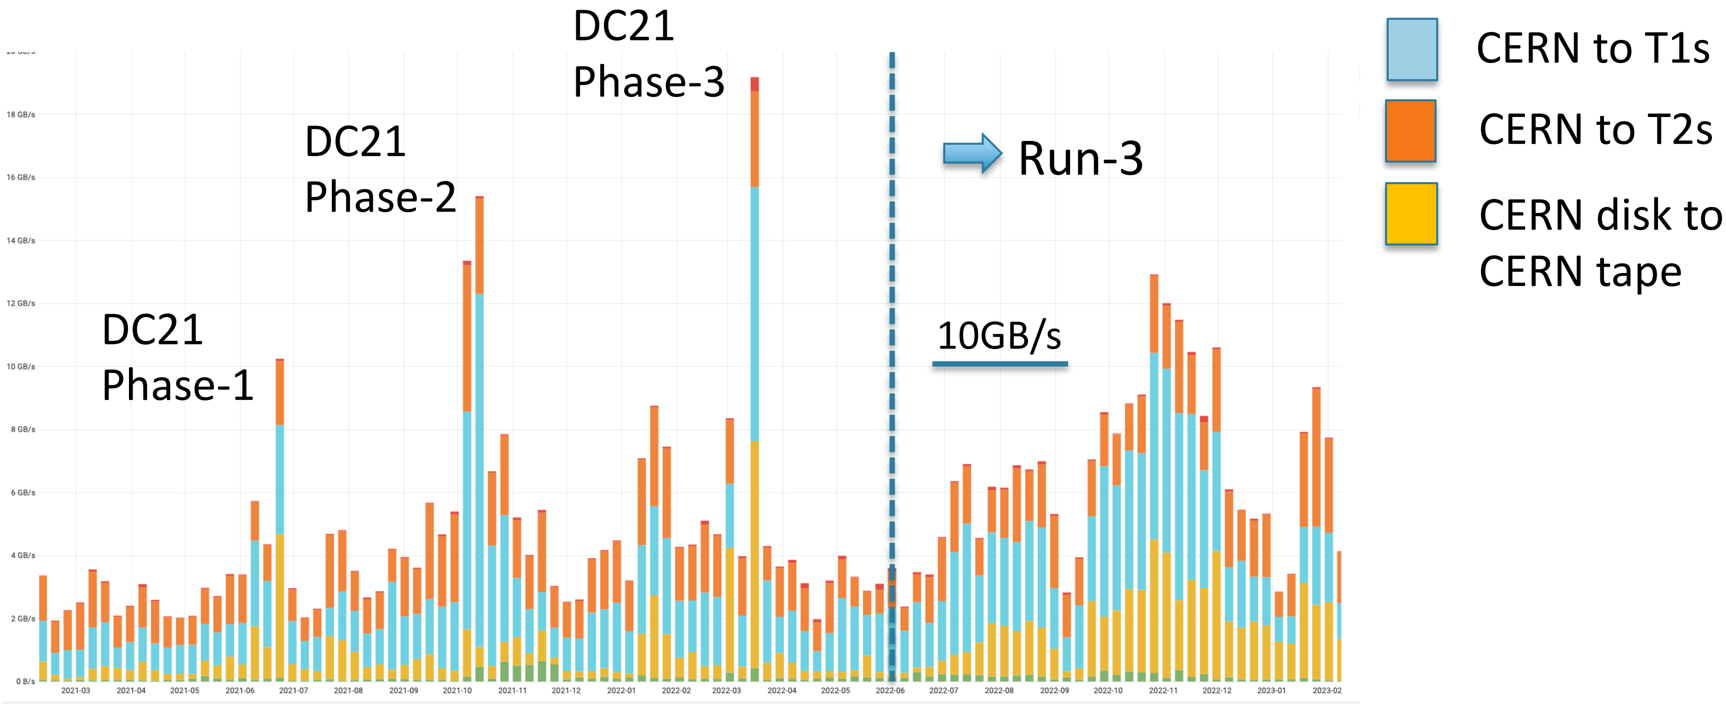

Data Challenge in 2021/2022 to prepare for Run-3:

Phase-1: focus on T0 workflows

Phase-2: worldwide network challenge

Phase-3: focus on archive storage

WLCG data transfers – last 2 years

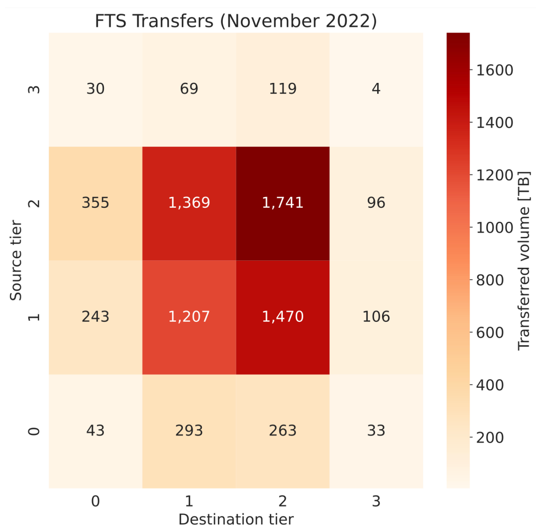

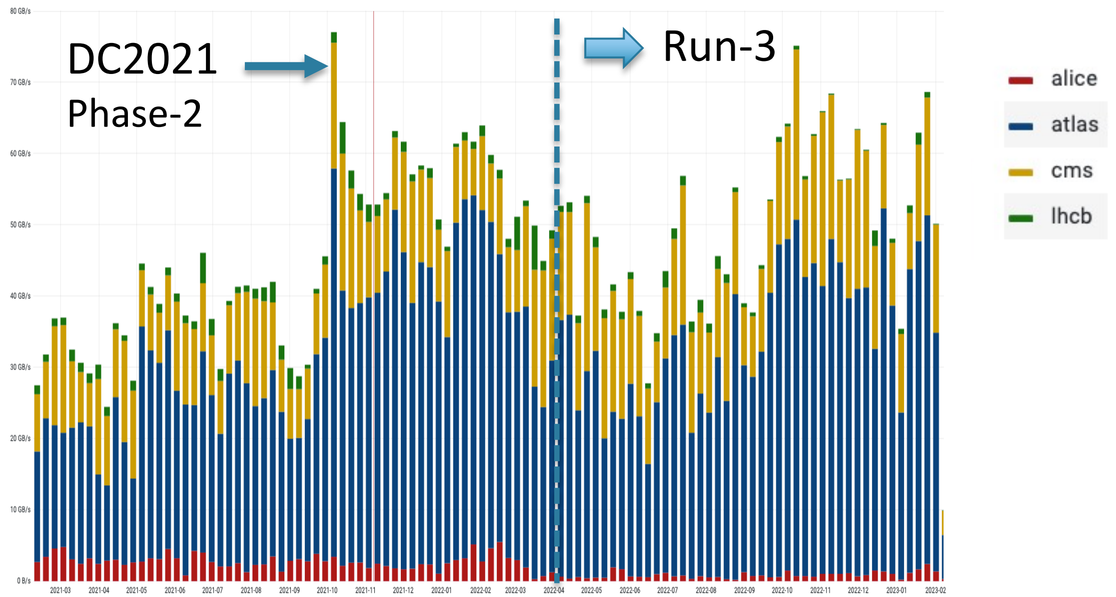

The WLCG transfer activities increased after the start of Run-3, recently reaching the Data Challenge 2021 level.

LHCb resource usage in 2022

CPU usage = 1.55x WLCG pledges in 2022

Disk (50%) and tape (35%) usage way below requirements / pledges / deployed capacities

For LHCb, 2022 has been a commissioning year instead of a full data-taking year

~80pb^{-1} collected (to be quality-checked)

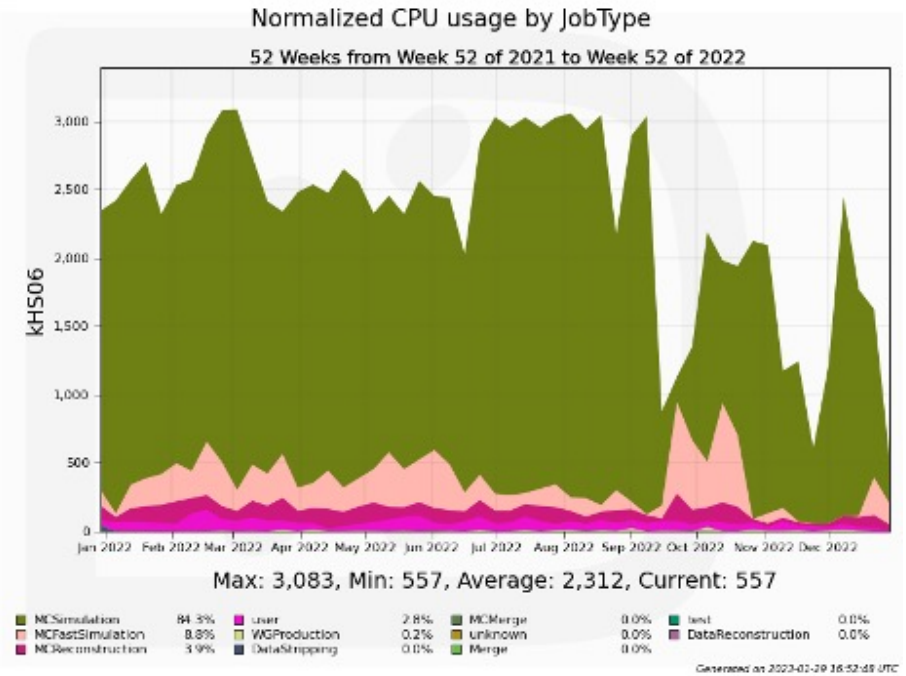

Offline computing infrastructure ready for data-taking: successful data challenges have been performed, exceeding target levels

ALICE: first taste of Pb-Pb in 2022

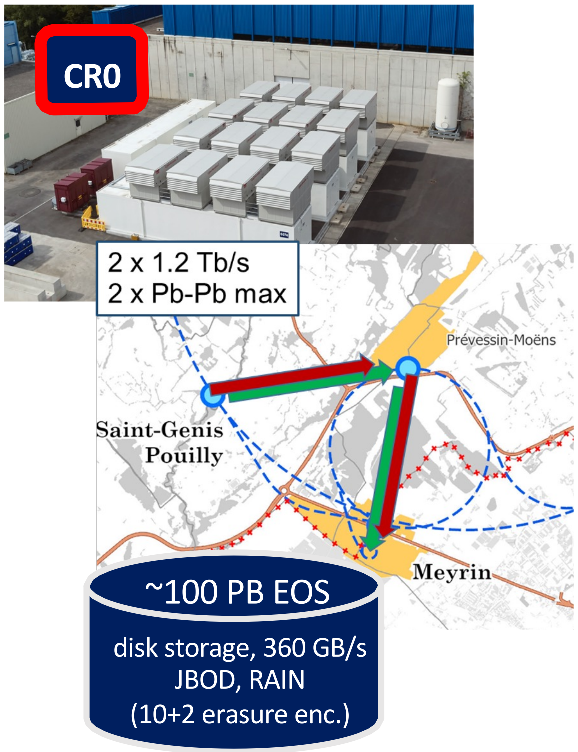

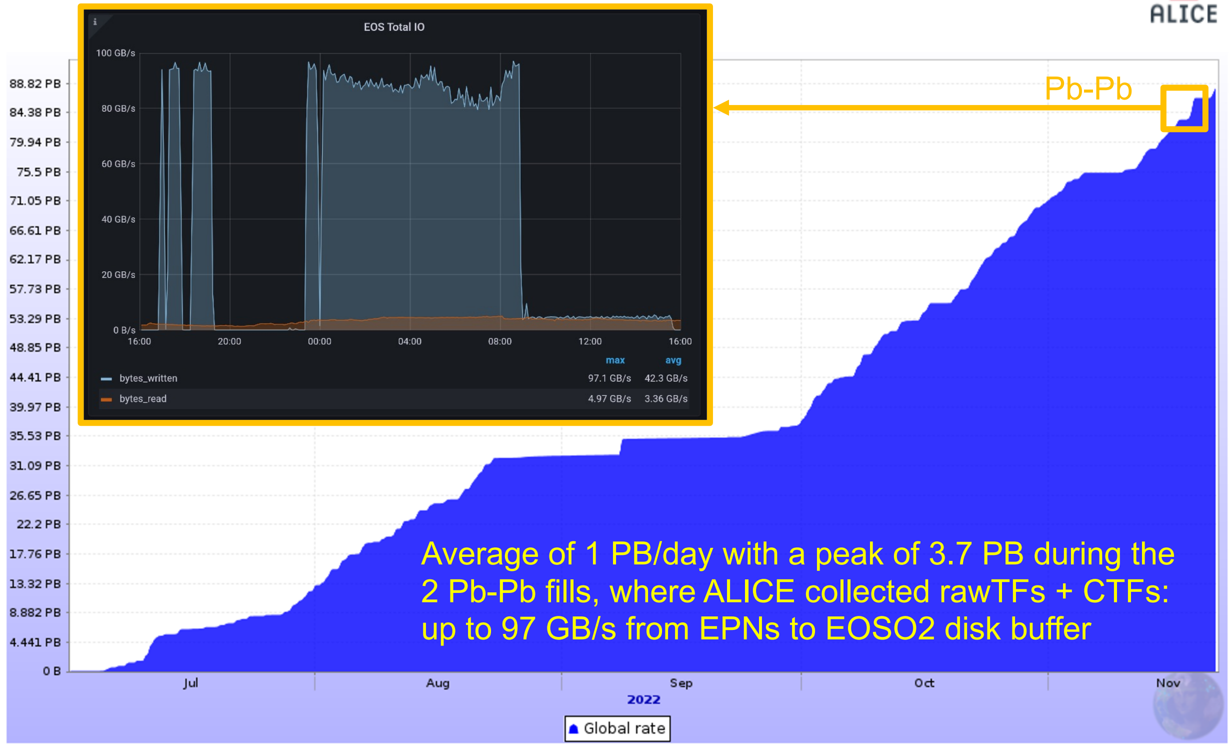

800Gbps in 2022 to be compared with 1200 Gbps max expected. A bit more than a taste …

Resources from the Russian sites

Russian sites provide 10% of WLCG T1 resources (not evenly spread across experiments)

Russia continues being part of WLCG and providing CPUs to the LHC experiments. These are efficiently used

CPU capacity from Russian sites (last 12 months)

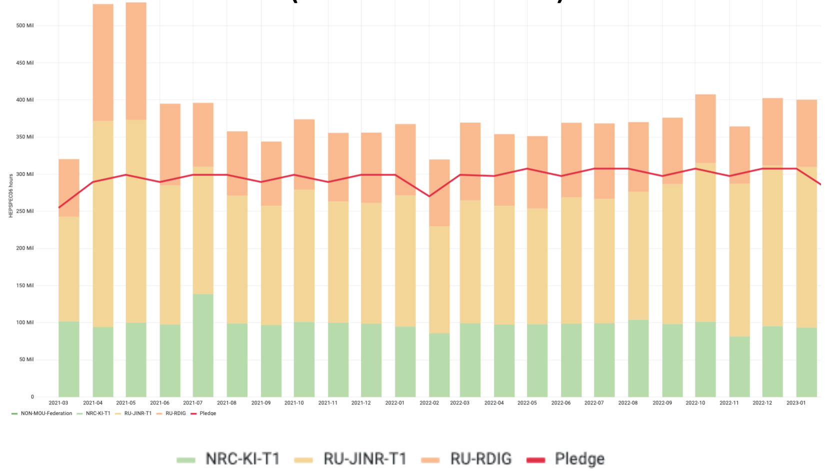

Data transfer INTO Russia (last 12 months)

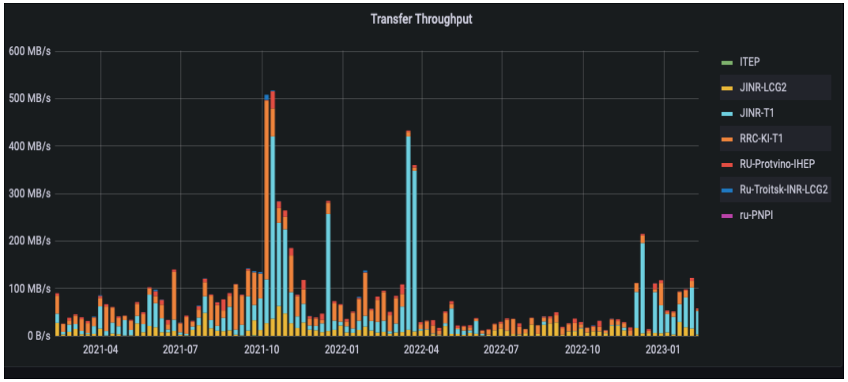

Data transfers decreased by ~ a factor 2 in the last months: tactical use of storage by the experiments.

No use of storage for data retention

Increasing cost of energy

Different initiatives to face the increasing cost of energy

- Reduction of capacity in critical months or for the full year
- (R&D) Real time power load shedding to compensate peaks of energy demand

Germany: load shedding based on energy cost

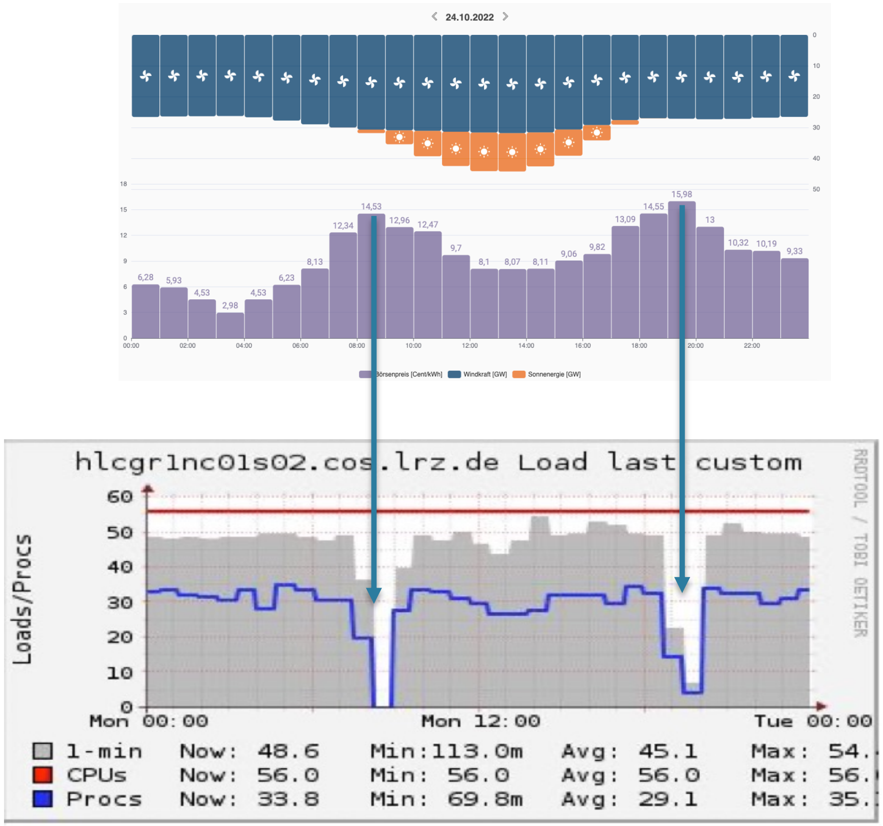

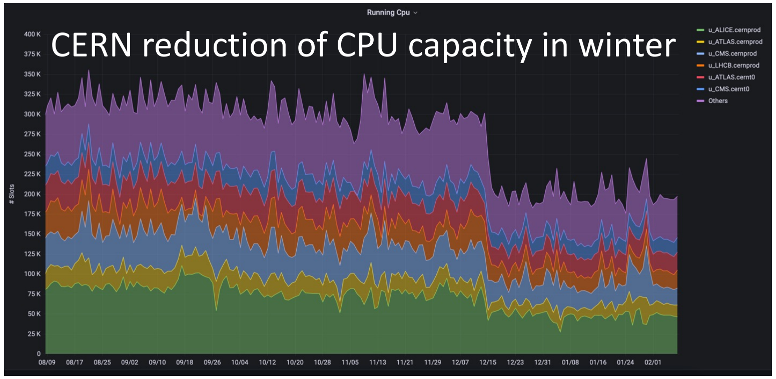

Hardware cost and time of delivery

The hardware cost is roughly at the pre-COVID (early 2020) level

In the ”flat budget model” of WLCG we expect a 15%/year decrease of the hw cost. This did not materialise in the last two years. On the positive side, some models predicted even higher prices for longer times

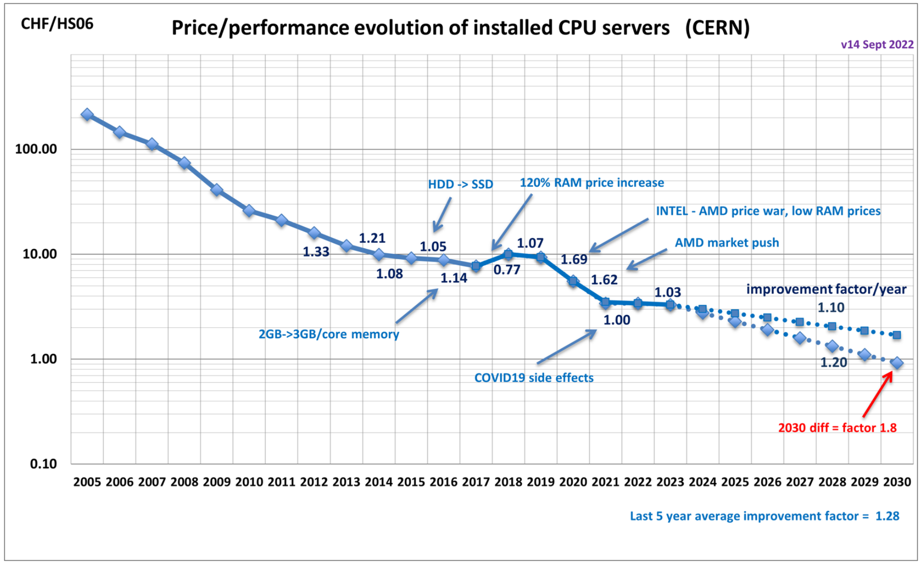

We still have strong problems in the hardware supply chain after COVID

The hardware time of delivery increased by x2 with respect to the pre-COVID times. More problematic for some hardware type (e.g. network components)

WLCG processing

Reduction of processing activities at WLCG sites in the last few months

- partially due to the external factors mentioned before
- partially to a lower experiments' activity

All experiments still benefit from more resources than pledged at WLCG sites

- plus, extra from HPCs, HLTs

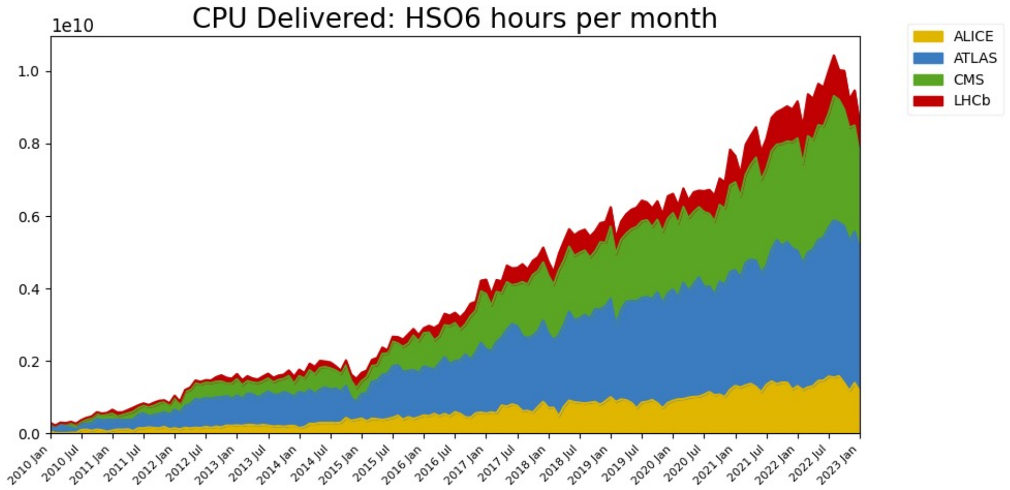

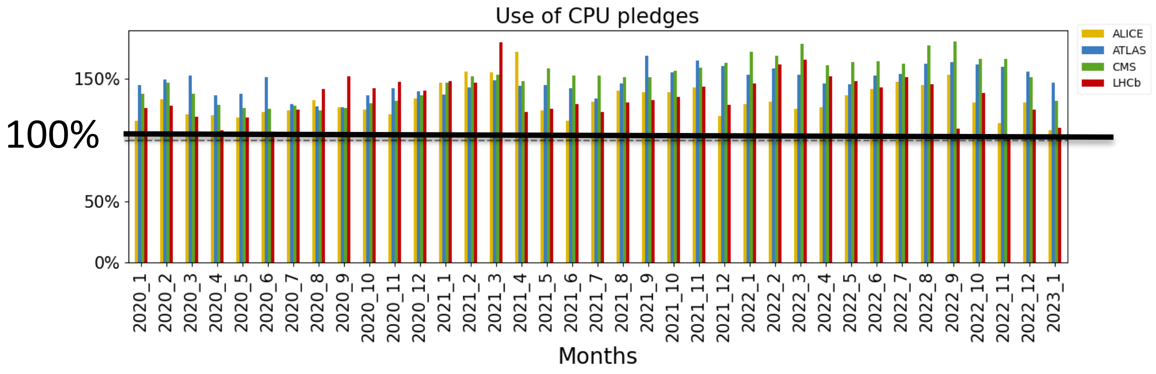

Opportunistic resources - HPCs

The use of HPC facilities increased considerably in the last years. Challenges and constraints:

- diversity in access, usage policies, the services available and the system architectures.
- HPCs live for a relatively short time (3-5 years) compared to WLCG sites, and then the work might need to be done again.

WLCG computing is done most economically on the sort of hardware used on the Grid, but can be usefully harmonized with other national computing priorities. Access and usage are improving as both sides gain experience; HPC will be a significant part of the future WLCG in some countries.

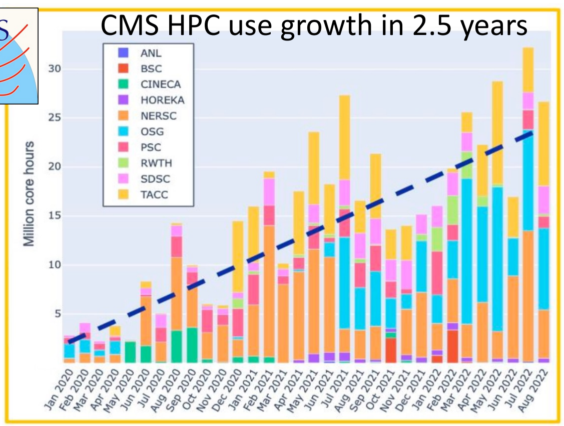

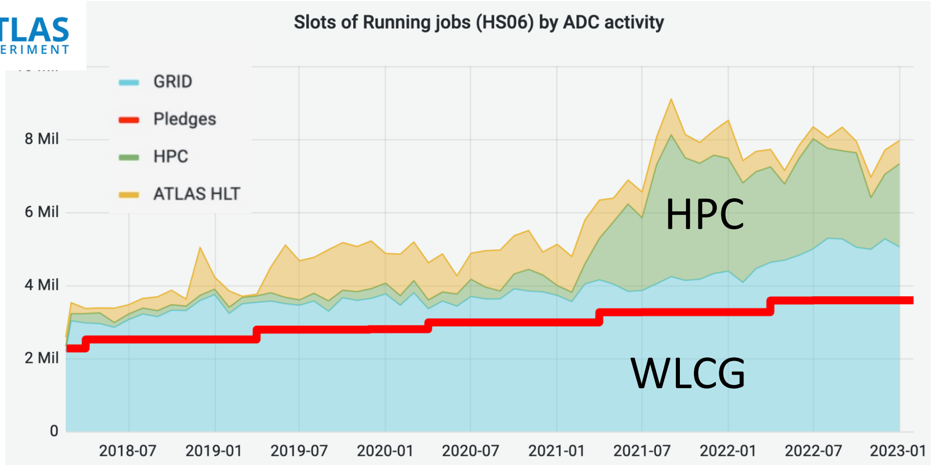

Run-3 forecast

	2022 (what happened)	2023 (planning for computing)	2024 (planning for computing)
ATLAS/CMS luminosity	<40/fb	<90/fb	<110/fb
ATLAS/CMS average pile-up	<45	<62 (peak 65)	<62 (peak 65)
LHCb luminosity	<0.1/fb	10/fb	15/fb
ALICE luminosity (pp)	<20/pb	50/pb	100/pb
Running time pp	<4 M seconds	4.2 M seconds	5.2 M seconds
Running time ions	No HI run	1.7 M seconds (Pb-Pb)	1.7 M seconds (Pb-Pb)

Considerable increase between 2022 and 2023:

- ATLAS/CMS: x2 more luminosity, 40% higher pileup; ALICE: x2.5 more pp luminosity
- LHCb: mostly commissioning in 2022
- 5 weeks of HI run in 2023, none in 2022

WLCG resources in 2023 and 2024

Run-3 will be challenging in terms of resources: unprecedented LHC running conditions (pile-up, extended HI run)

Risks:

- Increasing cost of energy in various regions impacting resource availability
- Market trends affected by post-COVID effects and trading difficulties
- Political situation

Mitigations:

- The data parking strategy of the experiments to smoothen peaks of resource needs
- A more active use of tape media commissioned through the tape challenges
- Flexibility for the computing models commissioned through the network challenges
- Enable the use of more heterogeneous computing architectures (e.g. ARM, GPUs)

The shortage of storage should be avoided, particularly for tape. We encourage the sites and the Funding Agencies to ensure enough storage can be provided

ATLAS reconstruction for Run-3

The ATLAS offline reconstruction software is continuously improved to better deal with high pileup

Naively one assumes exponential scaling of CPU time/event as a function of pileup due to combinatorial effects in tracking

Preliminary studies with Run-3 software at Tier-0 show ~ linear scaling

Good news considering 2023 and 2024 running conditions ("only" a +60% in reconstruction time from 2022 to 2023 pileup conditions)

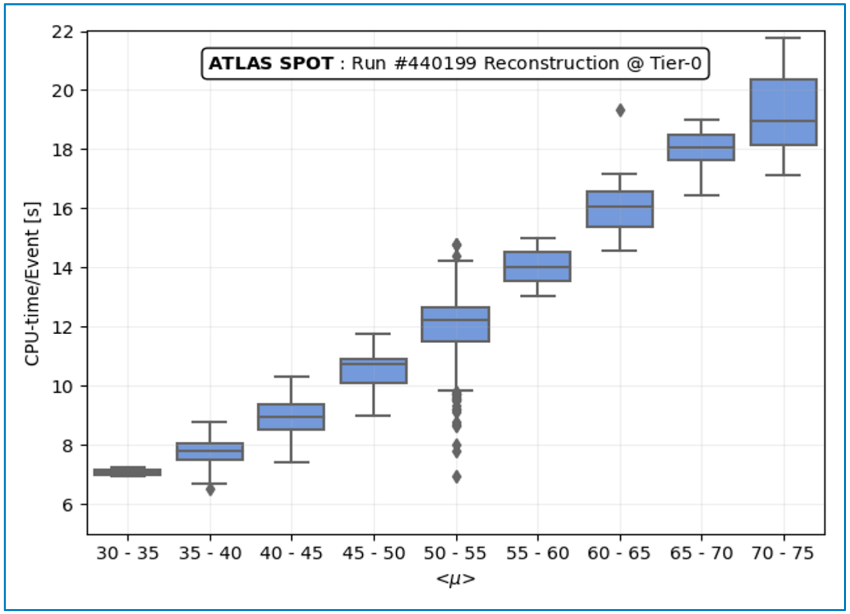

In the plot: CPU time per event during a 2022 data-taking run, as a function of <μ>. The boxes indicate the coverage of the middle two quartiles of data, with a line at the median. The whiskers indicate the location of the data point furthest from the median, within 1.5 times the difference between the 3rd and 1st quartiles (the IQR). The points indicate outliers.

CMS accelerated HL-LHC sw already in Run-3

Use from Run 3 to deploy innovations targeting Phase-2: e.g. pileup conditions approaching the ones at the beginning of Run 4 (levelling at PU 65 in 2023)

High Level Trigger: CPU + GPU (Nvidia T4, 2 per node) solution. CMSSW data processing framework already capable to orchestrate computations to accelerators

- 40% of evt. processing offloaded, +70% throughput, +50% per kWh

CMS is working to profit from accelerator-ready code for offline processing already in Run 3

- Solid plan ahead: aim to 10% offload this year
- Re-invention of algorithms: e.g. new vertexing
- One code-base for all platforms: performance portability through Alpaka
- Support for AMD GPUs backend added this year

Follow industry trends, save energy in some configurations, exploit more allocations at HPCs, faster code also for CPUs: ultimately, avoid becoming *legacy computing*.

ALICE asynchronous reconstruction on GPU and CPU

- Asynch. reconstruction of 2022 pp data with CPU on GRID (8 cores per job) and with CPU+GPU in O2
- The concurrent productions on GRID and on O2 allowed for the physics validation of the reconstruction with GPU and to speed up the processing
- Efficient reconstruction with GPU requires tailored job tuning, taking into account the GPU and CPU models and memory.
- Generalized Grid submission cannot currently use the GPU to the full extent.

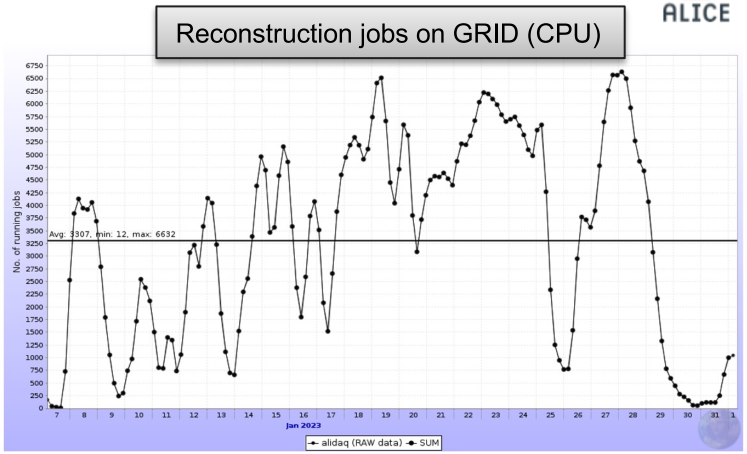

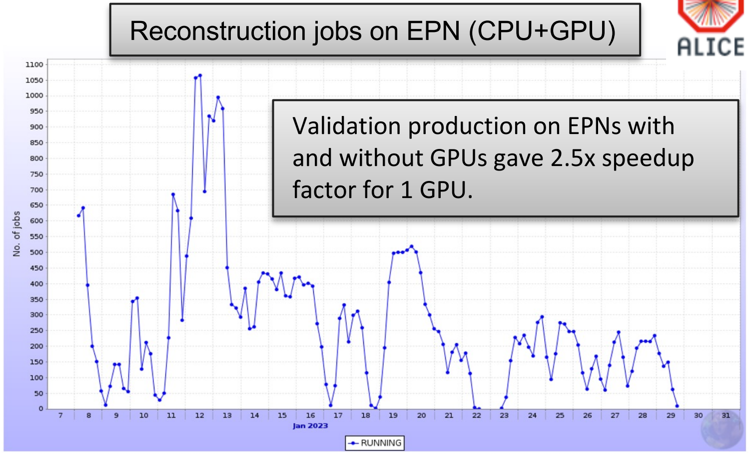

More non-x86 platforms

It is not just about GPUs but also non-X86 CPUS

- Power CPUs: CMS software validated in 2022 on the INFN M100 HPC at CINECA during 2022. Machine now used in production

- ARM64: most of the LHC workloads ported to ARM. ATLAS software successfully validated. [Promising results](#) in terms of energy per processed event

Give options to WLCG sites to diversify their CPU architectures, based on overall cost of ownership considerations

IBM Power CPU architecture for the first time at LHC

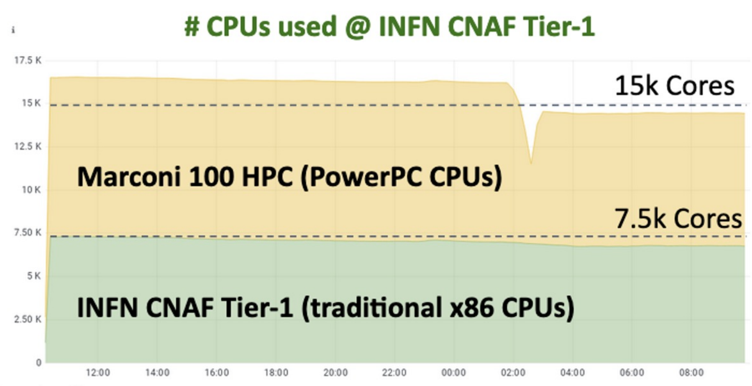

- Measured energy-used per unit WL-Score for 2 ATLAS and 3 CMS HEP-Score workloads.
- Three test machines:
  - x86_64: Single AMD EPYC 7003 series (SuperMicro)
  - 2*x86_64: Dual AMD EPYC 7513 series Processors (DELL)
  - arm64: Single socket Ampere Altra Processor (SuperMicro)
- Normalised to (first) x86. Lower is better!

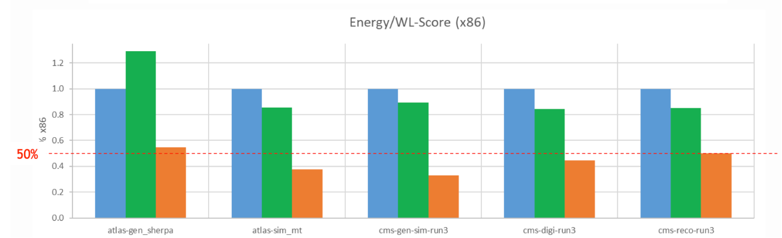

Geant4 simulation for HL-LHC

Geant4 is the common workhorse of detector simulation. Its evolution follows a three-prong strategy, as recommended during the 2021 LHCC common software review

- Improvements in the physics description. A continuous process, based on the dialog with the communities

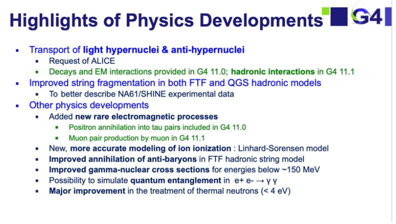

Highlights of Physics Developments

- Transport of **light hypernuclei & anti-hypernuclei**
  - Request of ALICE
  - Decays and EM interactions provided in G4 11.0; **hadronic interactions** in G4 11.1
- Improved string fragmentation in both FTF and QGS hadronic models
  - To better describe NA61/SHINE experimental data
- Other physics developments
  - Added **new rare electromagnetic processes**
    - Positron annihilation into tau pairs included in G4 11.0
    - Muon pair production by muon in G4 11.1
  - New, **more accurate modeling of ion ionization** : Linhard-Sorensen model
  - **Improved annihilation of anti-baryons** in FTF hadronic string model
  - **Improved gamma-nuclear cross sections** for energies below ~150 MeV
  - Possibility to simulate **quantum entanglement** in $e^+ e^- \to \gamma\gamma$
  - **Major improvement** in the treatment of thermal neutrons (< 4 eV)

- Speed up of Full Simulation: same accuracy in the physics description, but faster. Again a continuous process, bringing still very considerable benefits e.g. a recent 20% improvement for ATLAS simulation

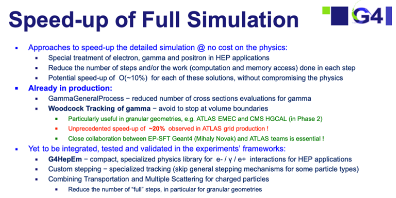

Speed-up of Full Simulation

- Approaches to speed-up the detailed simulation @ no cost on the physics:
  - Special treatment of electron, gamma and positron in HEP applications
  - Reduce the number of steps and/or the work (computation and memory access) done in each step
  - Potential speed-up of O(~10%) for each of these solutions, without compromising the physics
- **Already in production:**
  - GammaGeneralProcess − reduced number of cross sections evaluations for gamma
  - **Woodcock Tracking of gamma** − avoid to stop at volume boundaries
    - Particularly useful in granular geometries, e.g. ATLAS EMEC and CMS HGCAL (in Phase 2)
    - Unprecedented speed-up of **~20%** observed in ATLAS grid production !
    - Close collaboration between EP-SFT Geant4 (Mihaly Novak) and ATLAS teams is essential !
- Yet to be integrated, tested and validated in the experiments' frameworks:
  - **G4HepEm** − compact, specialized physics library for e- / γ / e+ interactions for HEP applications
  - Custom stepping − specialized tracking (skip general stepping mechanisms for some particle types)
  - Combining Transportation and Multiple Scattering for charged particles
    - Reduce the number of "full" steps, in particular for granular geometries

- R&D
  - Integration of Fast Simulation techniques in G4 including ML
  - GPU prototypes: AdePT and Celeritas

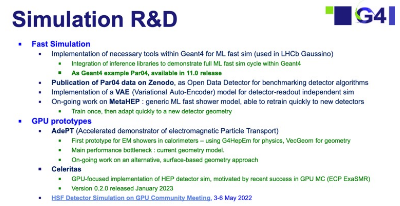

Simulation R&D

- **Fast Simulation**
  - Implementation of necessary tools within Geant4 for ML fast sim (used in LHCb Gaussino)
    - Integration of inference libraries to demonstrate full ML fast sim cycle within Geant4
    - **As Geant4 example Par04, available in 11.0 release**
  - **Publication of Par04 data on Zenodo**, as Open Data Detector for benchmarking detector algorithms
  - Implementation of a **VAE** (Variational Auto-Encoder) model for detector-readout independent sim
  - On-going work on **MetaHEP** : generic ML fast shower model, able to retrain quickly to new detectors
    - Train once, then adapt quickly to a new detector geometry
- **GPU prototypes**
  - **AdePT** (Accelerated demonstrator of electromagnetic Particle Transport)
    - First prototype for EM showers in calorimeters – using G4HepEm for physics, VecGeom for geometry
    - Main performance bottleneck : current geometry model.
    - On-going work on an alternative, surface-based geometry approach
  - **Celeritas**
    - GPU-focused implementation of HEP detector sim, motivated by recent success in GPU MC (ECP ExaSMR)
    - Version 0.2.0 released January 2023
  - HSF Detector Simulation on GPU Community Meeting, 3-6 May 2022

More details in the backup slides

ROOT foundation layer for HL-LHC

ROOT provides the common I/O layer for the LHC experiments.
The re-design of the internal ROOT data format (TTree to RNTuple) is a major upgrade in view of HL-LHC

- 10-20% smaller files, x3 to x5 better single core performance
- enables fast adaptation to modern technologies, like object stores

The 2021 and 2022 milestones from the HL-LHC computing review were achieved

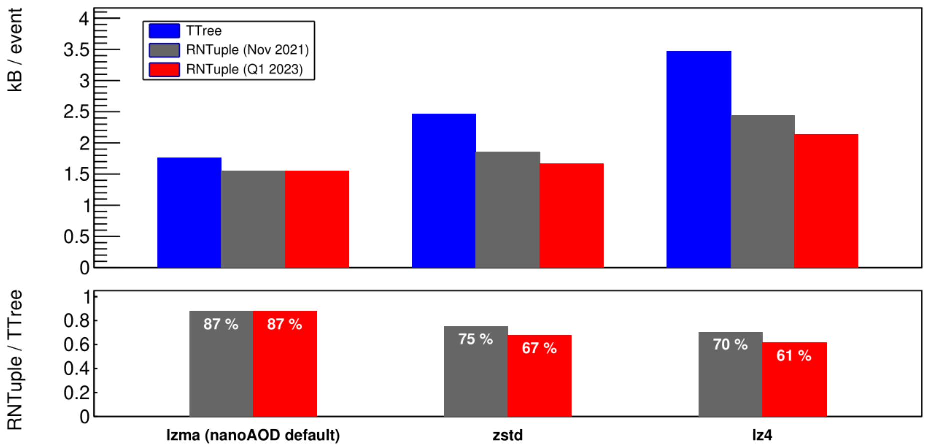

2024 goal: release RNTuple 1.0 supporting the Event Data Models of all experiments for analysis data and preserves backward compatibility. See backup slides for more details

ATLAS and CMS CPU power needs for HL-LHC

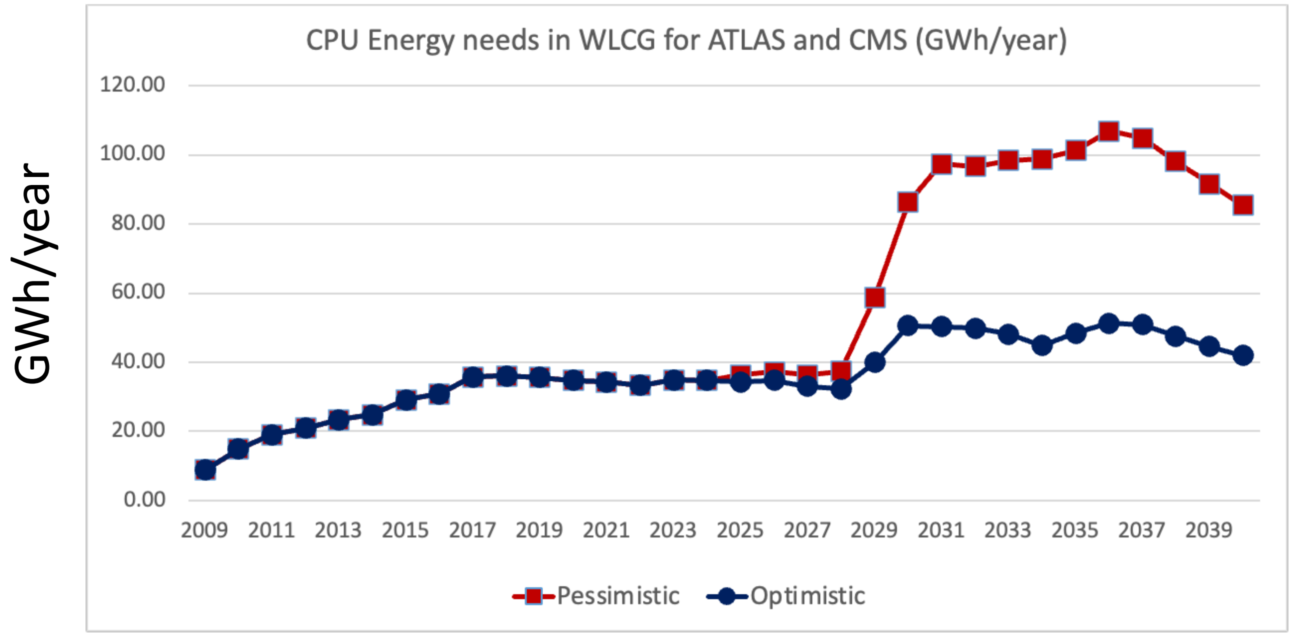

Discussed at the 2022 WLCG workshop

Based on the ATLAS and CMS resource projections made in 2022. More details here.

The peak of energy need happens in 2036 (start of Run-5): 300% higher than 2022 in the pessimistic scenario and 50% higher in the optimistic scenario.

GWh/fb^{-1} for WLCG

In WLCG GWh/fb^{-1} represent the energy needed to **analyse** the data

- look at each run rather than each year and assume that in the shutdown periods the processing of the previous run is completed

The scale on the left **(RED)** shows the energy and the scale on the right **(BLUE)** shows GWh/fb^{-1} (log!)

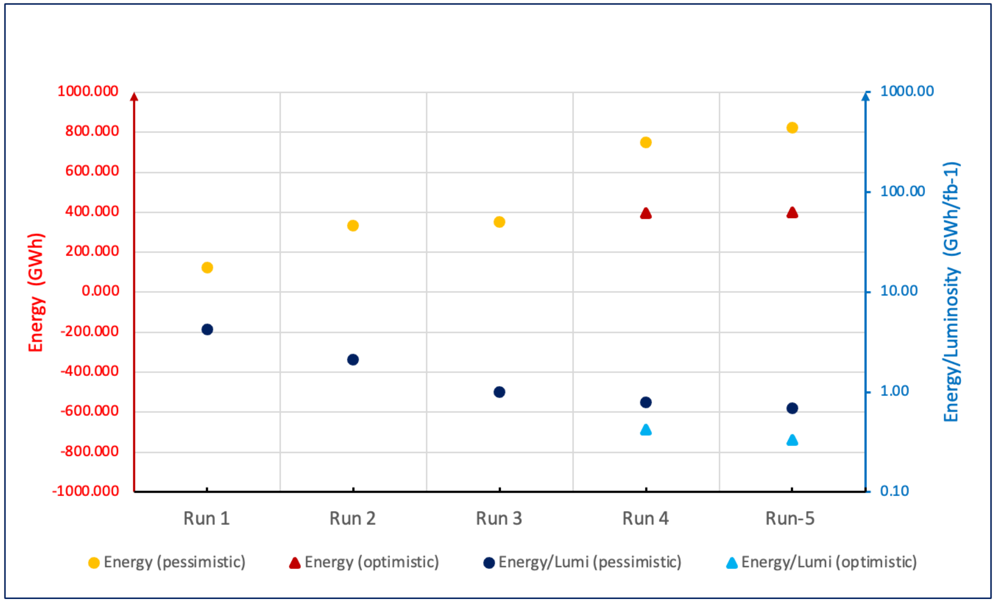

Energy needs in Run-4 and Run-5: +100% compared to Run-2 in the pessimistic scenario, only +10% in the **optimistic** scenario

GWh/fb^{-1} decreases a factor 10 between Run-1 and Run-5.
In Run-5, GWh/fb^{-1} in the **optimistic** scenario is half compared to the **pessimistic** scenario

Training in Software Developement

HEP C++ COURSE AND HANDS-ON TRAINING

550 STUDENTS TRAINED IN 2 ½ YEARS

- 3 day courses organised from within HEP for HEP
  - supported by the HSF training WG & SIDIS
- 6th course running this week
  - https://indico.cern.ch/e/CppSpring23
- More courses coming up this year
  - May @ JLAB / US
  - Sept @ Manchester / UK
  - Oct @ CERN

Organising team of the 6th course

WLCG partners and collaborators

Many communities collaborate with WLCG, share some of the same technologies, services and resources

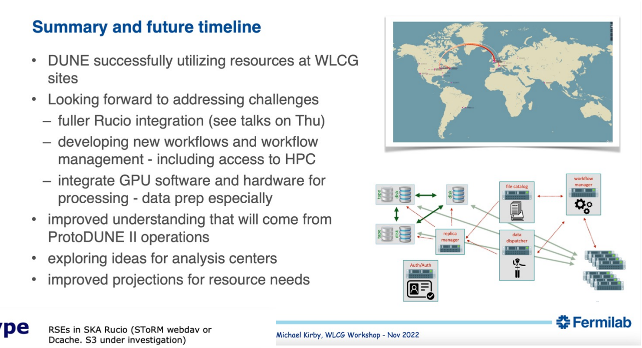

DUNE, Belle-2, JUNO and (recently) VIRGO are now WLCG "observers"

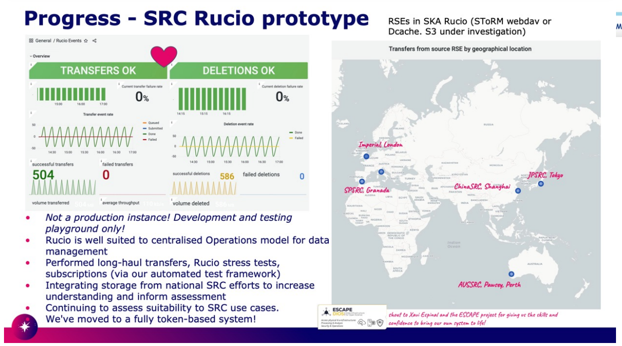

Collaboration with Astronomy in the context of the [ESCAPE Open Collaboration]

Snowmass 2021

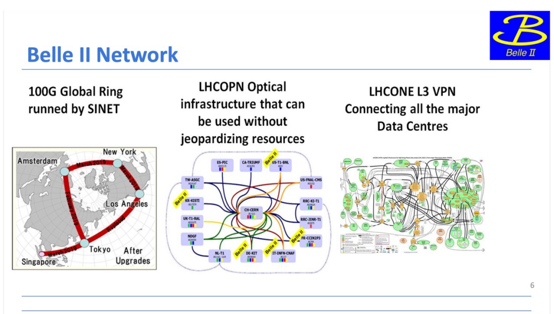

Fostering this collaboration is essential for the sustainability of the WLCG infrastructure. See the WLCG/DUNE/Belle-2 Snowmass [contribution]

WLCG data challenges

In 2021/2022, WLCG ran data challenges to commission the infrastructure for Run-3 and start preparing for HL-LHC

We are discussing a common 2024 data challenge involving WLCG, Belle-2 and DUNE

For WLCG this will be at 25% of the HL-LHC nominal rate. Other experiments will fix their targets

Also, an opportunity to commission new features at scale: in 2024 demonstrate the use of tokens for bulk data transfer and some of the most recent network R&D

Progress needs to be made in the next year

Statistics by Tag	Pass	Fail	Pass / Fail / Skip
critical	369	39	
not-critical	8	26	
se-cern-eos	20	6	
se-cnaf-amnesiac-storm	24	2	
se-florida-xrootd	23	3	
se-florida-xrootd-redir	23	3	
se-fnal-dcache	26	0	
se-infn-t1-xfer-storm	24	2	
se-nebraska-xrootd	20	6	
se-nebraska-xrootd-redir	18	8	
se-prague-dcache	14	12	
se-prague-xrootd	24	2	
se-prometheus-dcache	26	0	
se-ral-test-xrootd	22	4	
se-ubonn-xrootd	24	2	
se-ucsd-xrootd	23	3	
se-ucsd-xrootd-redir	22	4	
se-wisconsin-xrootd	22	4	
se-wisconsin-xrootd-redir	22	4	

Conclusions

- Despite the complications due to external factors, the first year of Run-3 operations for WLCG was very successful. Ingredients of this success are:
    - WLCG activities run 24/7/365 also during shutdowns, so the system is always maintained in operation
    - Data-taking specific workflows were extensively tested in several data challenges in 2021 and early 2022
    - Continuous (and possibly growing) commitment from our funding agencies
    - A very good collaboration between sites, experiments and software providers
- Computing for the remaining of Run-3 will not be “just more of the same” and will increase in scale and complexity. Risk management is in place
- There is continuous progress with the HL-LHC computing roadmap. Many solutions can be early adopted and play a role already during Run-3

Backup

Network use Px⇔CERN Data Center

Network Px => CERN Data Center (Gb/s)

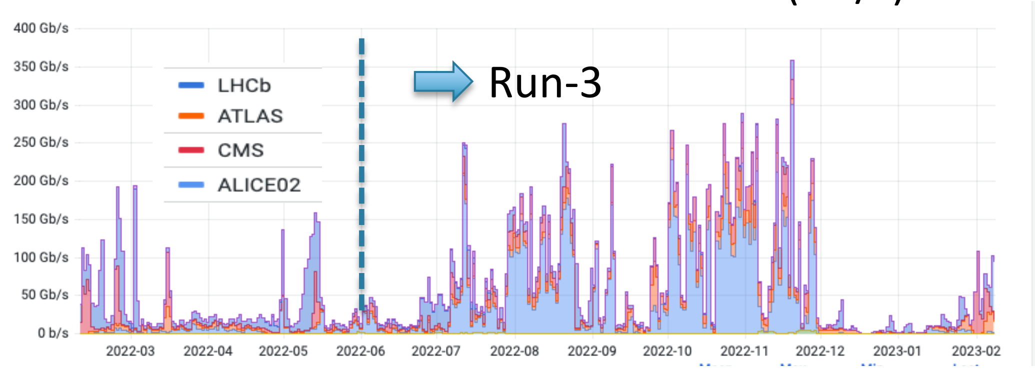

Network CERN Data Center => Px (Gb/s)

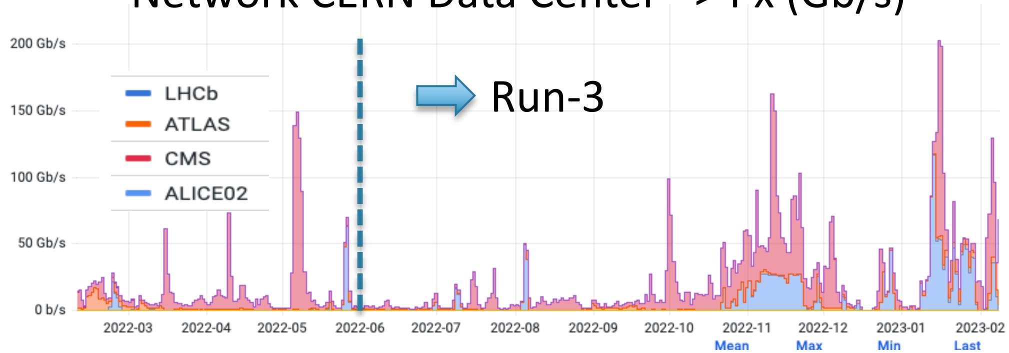

Commissioning for Run-3: networks

We executed data challenges testing the WLCG network and archive storages in preparation for Run-3 (Oct 2021)

Targets were met: the nominal rate was sustained, and the peak transfer rate were reached

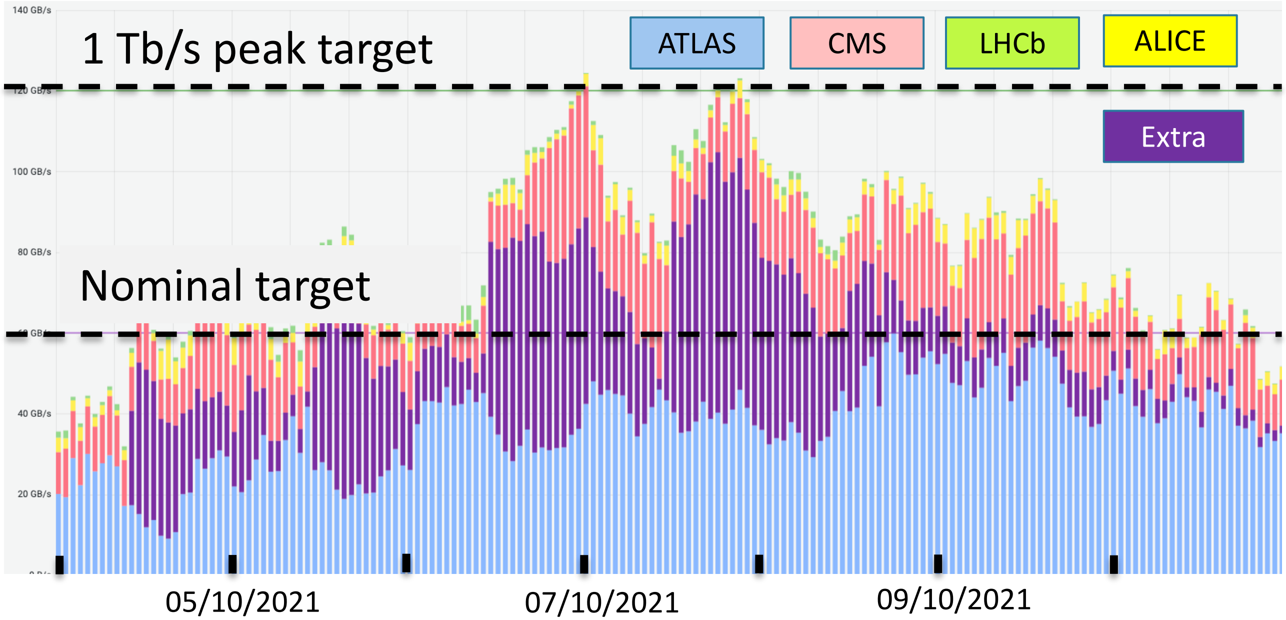

The data challenge are part of a longer-term process to prepare for the HL-LHC needs

Impact of the sanctions on Russia

The implementation of the 6th EU sanctions package (31st July 2022) affected the use of Russian resources in WLCG.

The European Council bans providing dual-use technologies to a list of Russian institutes, including some WLCG sites. Details here and links therein.

WLCG is not providing services/tools to Russian sites. However EGI is providing services to WLCG (e.g. accounting, GGUS) and that includes the Russian sites. EGI stopped providing those services to two Russian T2s at the end of Jan 2023.

The impact on WLCG is very moderate. The future availability of Russian resources remains very uncertain and unpredictable.

2023 Running Conditions for Computing estimates including contingency

• ATLAS/CMS luminosity:	<90/fb
• ATLAS/CMS average pile-up:	62 (peak PU=65)
• LHCb luminosity:	<10/fb
• ALICE luminosity (pp):	<50/pb
• Running time pp:	4.2×10^6 seconds
• Running time ions (PbPb):	1.7×10^6 seconds

Assumes first LHC beam on March 27 and end of run on October 30

2024 Running Conditions for Computing estimates including contingency

• ATLAS/CMS luminosity:	<110/fb
• ATLAS/CMS average pile-up:	62 (peak PU=65)
• LHCb luminosity:	<15/fb
• ALICE luminosity (pp):	<100/pb
• Running time pp:	5.2×10^6 seconds
• Running time ions (PbPb):	$<1.7 \times 10^6$ seconds

Assumes that the ion run in 2024 and/or 2025 could be extended to 5 weeks. Currently, 4 weeks of PbPb is foreseen for both years, but a short p-Pb run in one year and longer PbPb run in the other year is also a possibility.

Overall Run 3 Plan

Tentative updated Run 3 running plan (not to scale):

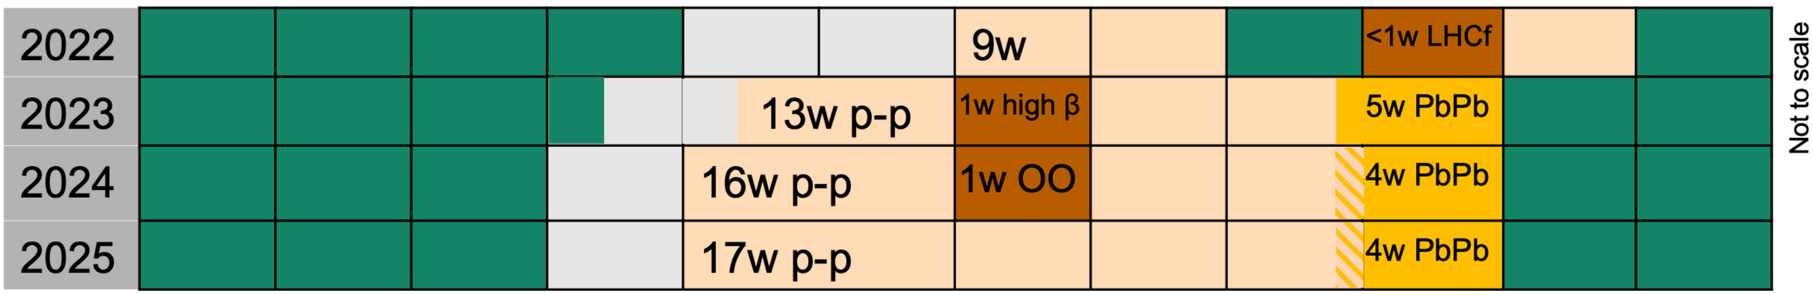

Each ion run includes a pp reference run. Currently, 4 weeks of PbPb is foreseen in 2024/2025, but a short p-Pb run in one year and longer PbPb run in the other year is also a possibility.

WLCG pledges: trend

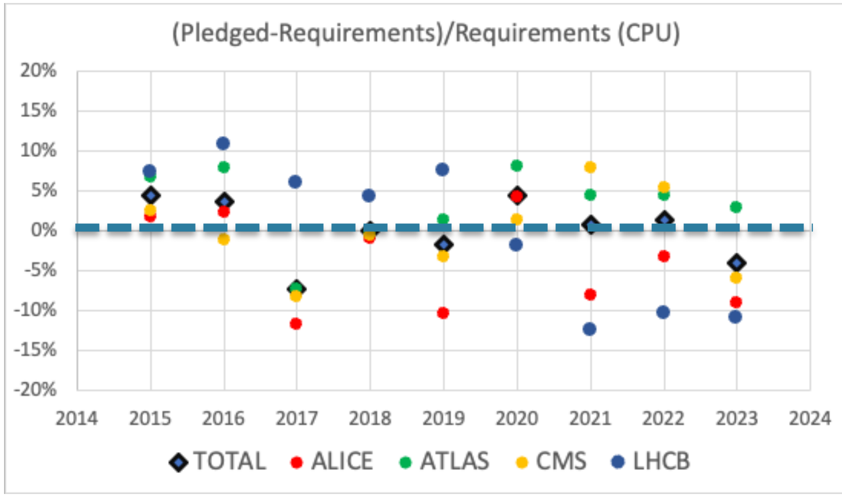

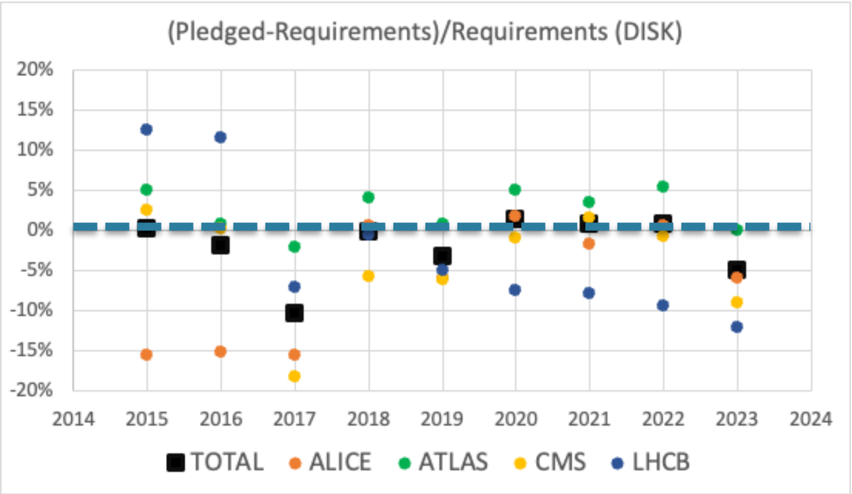

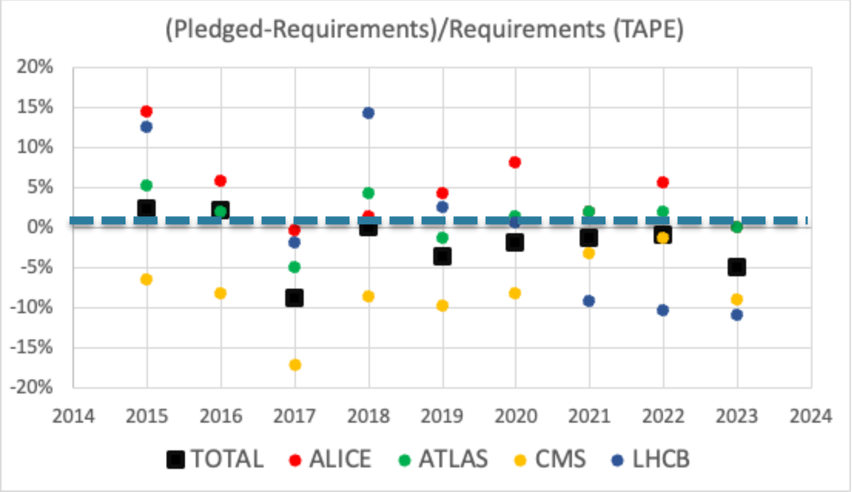

In the plots the contribution of Russian resources in 2023 is **NOT** considered as part of the pledge

WLCG pledges for 2023: detailed view

	ALICE			ATLAS			CMS			LHCb		
	2022	2023	2023 w/o Ru	2022	2023	2023 w/o Ru	2022	2023	2023 w/o Ru	2022	2023	2023 w/o Ru
Tier-0												
CPU	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Disk	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Tape	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Tier-1												
CPU	-10%	-12%	-17%	6%	6%	4%	17%	15%	-15%	-17%	-15%	-18%
Disk	-10%	-9%	-16%	12%	11%	7%	0%	-1%	-15%	-10%	-10%	-13%
Tape	14%	7%	0%	3%	2%	0%	-2%	-4%	-15%	-16%	-15%	-17%
Tier-2												
CPU	0%	-4%	-10%	4%	5%	3%	1%	-3%	-3%	-4%	-5%	11%
Disk	13%	5%	0%	0%	-5%	-6%	-2%	-7%	-7%	-32%	-37%	-35%
Total												
CPU	-3%	-5%	-9%	4%	5%	3%	5%	3%	-6%	-10%	-10%	-11%
Disk	1%	-2%	-6%	5%	2%	0%	-1%	-3%	-9%	-9%	-10%	-12%
Tape	6%	3%	0%	2%	1%	0%	-1%	-2%	-9%	-10%	-9%	-11%

Good match between requests and commitments for ATLAS

WLCG pledges for 2023: detailed view

	ALICE			ATLAS			CMS			LHCb		
	2022	2023	2023 w/o Ru	2022	2023	2023 w/o Ru	2022	2023	2023 w/o Ru	2022	2023	2023 w/o Ru
Tier-0												
CPU	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Disk	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Tape	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Tier-1												
CPU	-10%	-12%	-17%	6%	6%	4%	17%	15%	-15%	-17%	-15%	-18%
Disk	-10%	-9%	-16%	12%	11%	7%	0%	-1%	-15%	-10%	-10%	-13%
Tape	14%	7%	0%	3%	2%	0%	-2%	-4%	-15%	-16%	-15%	-17%
Tier-2												
CPU	0%	-4%	-10%	4%	5%	3%	1%	-3%	-3%	-4%	-5%	11%
Disk	13%	5%	0%	0%	-5%	-6%	-2%	-7%	-7%	-32%	-37%	-35%
Total												
CPU	-3%	-5%	-9%	4%	5%	3%	5%	3%	-6%	-10%	-10%	-11%
Disk	1%	-2%	-6%	5%	2%	0%	-1%	-3%	-9%	-9%	-10%	-12%
Tape	6%	3%	0%	2%	1%	0%	-1%	-2%	-9%	-10%	-9%	-11%

Shortage of resources for ALICE and CMS, particularly at T1s. The possible loss of Russian resources was not fully mitigated by the other funding agencies.

WLCG pledges for 2023: detailed view

	ALICE			ATLAS			CMS			LHCb		
	2022	2023	2023 w/o Ru	2022	2023	2023 w/o Ru	2022	2023	2023 w/o Ru	2022	2023	2023 w/o Ru
Tier-0												
CPU	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Disk	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Tape	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
Tier-1												
CPU	-10%	-12%	-17%	6%	6%	4%	17%	15%	-15%	-17%	-15%	-18%
Disk	-10%	-9%	-16%	12%	11%	7%	0%	-1%	-15%	-10%	-10%	-13%
Tape	14%	7%	0%	3%	2%	0%	-2%	-4%	-15%	-16%	-15%	-17%
Tier-2												
CPU	0%	-4%	-10%	4%	5%	3%	1%	-3%	-3%	-4%	-5%	11%
Disk	13%	5%	0%	0%	-5%	-6%	-2%	-7%	-7%	-32%	-37%	-35%
Total												
CPU	-3%	-5%	-9%	4%	5%	3%	5%	3%	-6%	-10%	-10%	-11%
Disk	1%	-2%	-6%	5%	2%	0%	-1%	-3%	-9%	-9%	-10%	-12%
Tape	6%	3%	0%	2%	1%	0%	-1%	-2%	-9%	-10%	-9%	-11%

LHCb continues seeing a ~10% shortage as in 2022. Little impact of the situation with Russia

GPUs in Event Generators

A great progress porting event generators code to use accelerators (GPUs). Details in here

The example in this slide is about MadGraph and particularly the Matrix Element calculation in it

For the process below the total time to generate one event reduces by ~60 times

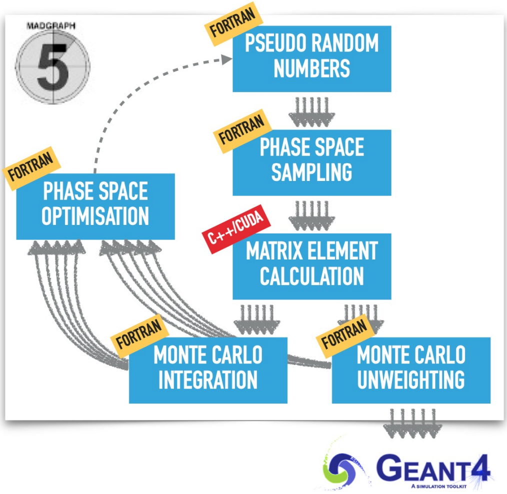

Next steps: integrate in the production environment and demonstrate the impact for the experiments

		madevent		
CUDA grid size		8192		
$gg \to t\bar{t}ggg$	MEs precision	$t_{TOT} = t_{Mad} + t_{MEs}$ [sec]	N_{events}/t_{TOT} [events/sec]	N_{events}/t_{MEs} [MEs/sec]
Fortran	double	1228.2 = 5.0 + 1223.2	7.34E1 (=1.0)	7.37E1 (=1.0)
CUDA	double	19.6 = 7.4 + 12.1	4.61E3 (x63)	7.44E3 (x100)

NVidia V100, Cuda 11.7, gcc 11.2

Geant4 contribution to LHCC software review

Marc Verderi & Alberto Ribon, for the Geant4 Collaboration

Geant4 & Experiments, Geant4 Evolution

- **Close collaboration with experiments remains a line of action:**
  - Support to experiments, responses to expressed needs (functionalities, physics, performances, etc.), tight collaboration on specific subjects (eg: ATLAS speed-up), etc.
  - **Joint events:** specific ones –like LPCC simulation workshop (Nov. 2021) dedicated to fast simulation and AI-based one– or general ones –technical forum, HSF simulation working group–

- **Geant4 evolution:**
  - **Software Evolution**:
    - Major release 11.0 (dec. 2021)
      - **New parallelism scheme :** multi-threading (core-bound) replaced by **tasking** (logical layer for tasks, with choice of backend to execute the task, investigated now for seamless bridge for hybrid computing)
      - C++17 as minimum standard
      - *To be noted* : new example demonstrating ML-based Fast Simulation
    - Minor release 11.1 (dec. 2022)
      - New or improved physics capabilities (see after)
  - **Collaboration Evolution**:
    - Actions to address the **"generation transfer knowledge"** challenge
      - Strong warning to funding agencies on vital need for **new, <u>stable</u>, young generation of developers**
      - Creation of an **"Early Career Researcher Committee"**, with representative at Steering Board
      - And also : internal seminars, detailed minutes, WG meetings open to all, etc.
    - Creation of **"contributor"** status:
      - **O(10) contributors** this past and first year; status seems to lower the barrier to join the Collaboration
    - **New website**, with design initiated by "new generation"

Highlights of Physics Developments

- Transport of **light hypernuclei & anti-hypernuclei**
  - Request of ALICE
  - Decays and EM interactions provided in G4 11.0; **hadronic interactions** in G4 11.1
- Improved string fragmentation in both FTF and QGS hadronic models
  - To better describe NA61/SHINE experimental data
- Other physics developments
  - Added **new rare electromagnetic processes**
    - Positron annihilation into tau pairs included in G4 11.0
    - Muon pair production by muon in G4 11.1
  - New, **more accurate modeling of ion ionization** : Linhard-Sorensen model
  - **Improved annihilation of anti-baryons** in FTF hadronic string model
  - **Improved gamma-nuclear cross sections** for energies below ~150 MeV
  - Possibility to simulate **quantum entanglement** in $e^+ e^- \rightarrow \gamma\gamma$
  - **Major improvement** in the treatment of thermal neutrons (< 4 eV)

Speed-up of Full Simulation

- Approaches to speed-up the detailed simulation @ no cost on the physics:
  - Special treatment of electron, gamma and positron in HEP applications
  - Reduce the number of steps and/or the work (computation and memory access) done in each step
  - Potential speed-up of O(~10%) for each of these solutions, without compromising the physics
- **Already in production:**
  - GammaGeneralProcess − reduced number of cross sections evaluations for gamma
  - **Woodcock Tracking of gamma** − avoid to stop at volume boundaries
    - Particularly useful in granular geometries, e.g. ATLAS EMEC and CMS HGCAL (in Phase 2)
    - Unprecedented speed-up of **~20%** observed in ATLAS grid production !
    - Close collaboration between EP-SFT Geant4 (Mihaly Novak) and ATLAS teams is essential !
- Yet to be integrated, tested and validated in the experiments' frameworks:
  - **G4HepEm** − compact, specialized physics library for e- / γ / e+ interactions for HEP applications
  - Custom stepping − specialized tracking (skip general stepping mechanisms for some particle types)
  - Combining Transportation and Multiple Scattering for charged particles
    - Reduce the number of "full" steps, in particular for granular geometries

Simulation R&D

- **Fast Simulation**
  - Implementation of necessary tools within Geant4 for ML fast sim (used in LHCb Gaussino)
    - Integration of inference libraries to demonstrate full ML fast sim cycle within Geant4
    - **As Geant4 example Par04, available in 11.0 release**
  - **Publication of Par04 data on Zenodo**, as Open Data Detector for benchmarking detector algorithms
  - Implementation of a **VAE** (Variational Auto-Encoder) model for detector-readout independent sim
  - On-going work on **MetaHEP** : generic ML fast shower model, able to retrain quickly to new detectors
    - Train once, then adapt quickly to a new detector geometry
- **GPU prototypes**
  - **AdePT** (Accelerated demonstrator of electromagnetic Particle Transport)
    - First prototype for EM showers in calorimeters – using G4HepEm for physics, VecGeom for geometry
    - Main performance bottleneck : current geometry model.
    - On-going work on an alternative, surface-based geometry approach
  - **Celeritas**
    - GPU-focused implementation of HEP detector sim, motivated by recent success in GPU MC (ECP ExaSMR)
    - Version 0.2.0 released January 2023
  - HSF Detector Simulation on GPU Community Meeting, 3-6 May 2022

Responses to recommandations (1/2)

- **General recommendations:**
  - The experiments have strong relationships with Geant4 that should continue. More coordination and collaboration on fast simulation would be beneficial.
    1) LPCC simulation workshop (Nov. 2021) on fast simulation, including AI-based ones, is along this line
    2) But experience shows that even AI-based fast-simulation remains quite experiment specific
       - Collaboration is essentially through building & sharing a common background on these techniques
  - We urge the experiments to consider whether their existing project management processes are sufficient to address the following:
    1) The (evolving) degree of certainty that the overall project will deliver the required performance in time for HL-LHC.
    2) Recognizing the level of risk to the project from internal and external dependencies, assumptions, and unexpected changes.
    3) Establishing the decision point when choices need to be made, priorities established or changed.
       - The three above points are considered and are part of "routine" management
    4) Person-power requirements and projections.
       - This remains by far the highest risk for Geant4 –and for the users' community–, with a high risk of expertise disruption. Funding agencies have been strongly warned about this risk, for Geant4 in general, and for each of the Geant4 working groups.
  - [About additional collaboration mechanisms with experiments] We encourage Geant4 and the experiments to reinforce such collaborative efforts that not only streamline communication but also facilitate sharing effort towards common goals such as, for example, AI-based fast Monte Carlo projects.
    1) Any opportunity to strengthen collaborative efforts are taken, we believe.
    2) For example, but not limited to:
       - LPCC simulation workshop, dedicated to fast-simulation, including AI-based ones
       - Large ATLAS performance increase, obtained thanks to tight collaboration

Responses to recommandations (2/2)

- **Specific recommendations to Geant4:**
  - We note that improvements in the speed of Geant4 do not automatically translate into equal gains within the experiments' workflows and we encourage Geant4 experts to reinforce their involvement in the tuning of Geant4 in the experiments' software stacks to maximize the benefits.
    1) Example of collaboration with ATLAS
  - [About fast simulation] We encourage Geant4 and the experiments to increase coordination on these developments.
    1) LPCC workshop mentioned before
  - In addition, a new role of Geant4 Contributor has been created to facilitate contributions from people who are not full members of the collaboration. We are keen to see if this new initiative adds flexibility and enables new engagement.
    1) About 10 contributors joined this past year
    2) Role looks to indeed lower the barrier for future membership (but need several years feedback for conclusion)
  - It is also crucial that the physics description of the detailed simulation continues to improve, and care should be taken in maintaining a sufficient core developer team with appropriate expertise.
    1) We fully support this conclusion.
    2) This is tightly linked to high risk of expertise disruption mentioned already
    3) And is tightly linked with the appropriate balance between adiabatic evolution and risky R&D.
  - The inclusion of FLUKA models in Geant4 may improve the overall precision but is a time-consuming task. We suggest that the experiments are involved in the prioritization of this activity in the overall project schedule.
    1) This is under discussion, but is also a way to respond to above point of improving the physics description for what hadronic physics is concerned.

ROOT Foundation Upgrade for HL-LHC

- **Major I/O upgrade** of the event data file format and access API: TTree → RNTuple
  - Less disk and CPU usage for same data content
    - **10-20% smaller** files, **at least x3-5 better single-core performance**
  - Give access to **novel and future storage technologies**
    - Native support for HPC and cloud object stores
    - Direct disk-to-GPU data transfers
  - Systematic use of checksumming and exceptions to prevent silent I/O errors

- **Generation hand-over** of I/O experts to ensure availability of I/O expertise compatible with the HL-LHC lifetime

- Est. 50 MCHF/year on storage in WLCG
  → strong incentive for **common, highly efficient I/O layer**

Supported by

EP R&D

iris hep

Coordination with Experiments

Approach to adding support for new EDMs

1. EDM representation in RNTuple (import / conversion support)
2. Integration in the experiment framework (writer / output module)
3. EDM R&D: optimizing EDMs for best performance with RNTuple

- The first few EDMs require substantial R&D and are onboarded one-by-one
- Support since 2021: CMS nanoAOD, LHCb final-stage ntuples (import only)

Milestone reached in 2022: Support for ATLAS xAOD (PHYS & PHYSLITE)

- Support for import as well as generation from ATLAS software framework
- Required support for *read rules* and *custom collections* in RNTuple
- Fruitful and very efficient collaboration of ROOT and ATLAS Core I/O teams to support RNTuple

Encouraging first results:

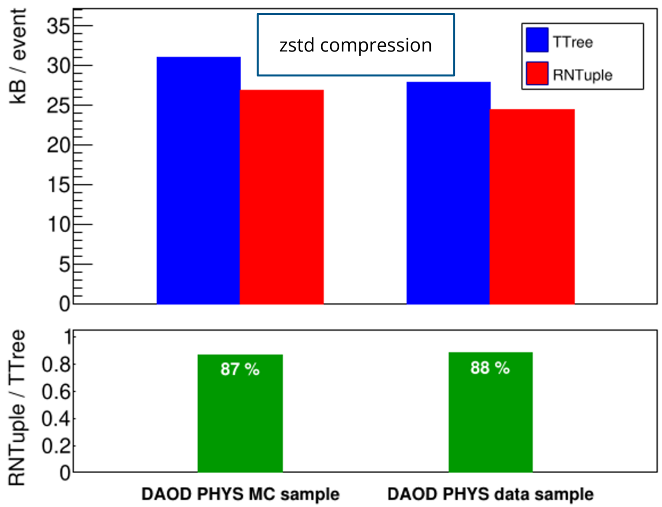

Goals for 2023:

- RNTuple support for all ATLAS & CMS AOD variants with initial framework integration
- Planning for onboarding of EDMs of LHCb and ALICE in 2024
- Validation of format forward-compatibility

Goal for 2024: Release of RNTuple 1.0, i.e. promise to keep backwards-compatibility
→ Requires firm understanding of requirements for RECO and AOD formats of all LHC experiments
(RAW formats have little impact because their format is experiment-specific)

ROOT RNTuple Technical Progress

Milestones reached

- TTree to RNTuple importer
- XRootD support
- Lossless compression R&D: encoding filters
- Full exploitation of HPC object store performance
- First validation of distributed analysis throughput

Progress on track

- Late schema extension
  - Essential for production-grade AOD writer modules in ATLAS & CMS software frameworks
- (Distributed) RDataFrame integration
- Lossy compression (`double` in memory, `float` on disk)

Moved to 2023/2024

- Schema evolution (initial design ready)
- Meta-data API & data set combinations (pending EP R&D funding)
- RNTupleLite: low-level RNTuple C API
- Validation on large storage pools (pending collaboration with infrastructure providers, e.g. ATLAS cloud R&D, CERN IT)

NanoAOD storage efficiency

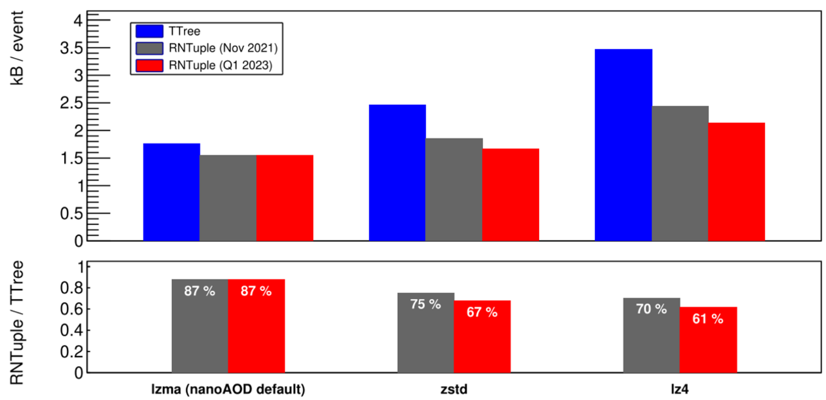

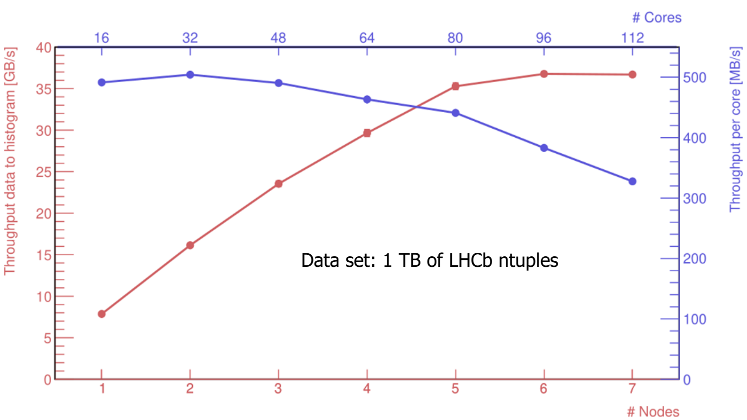

First scaling tests:

Distributed RDataFrame with RNTuple, data stored in HPC object storage (Intel DAOS)

Collaboration with other HEP experiments

WLCG presented a joint paper with DUNE and Belle-2 to the Snowmass 2021 process, which detailed the strategic directions needed to address the computing challenges of the experiments over the next decade. It complements the WLCG contribution to the European Strategy for Particle Physics in 2019.

Three strategic areas:

- Strengthen the backbone of core services and policies.
- Evolution of the infrastructure to integrate modern technologies and facilities.
- **Broadening the scope of the WLCG collaboration to create partnership with other HEP experiments**

Snowmass 2021

Physics > Computational Physics

[Submitted on 14 Mar 2022]

HEP computing collaborations for the challenges of the next decade

Simone Campana, Alessandro Di Girolamo, Paul Laycock, Zach Marshall, Heidi Schellman, Graeme A Stewart

Large High Energy Physics (HEP) experiments adopted a distributed computing model more than a decade ago. WLCG, the global computing infrastructure for LHC, in partnership with the US Open Science Grid, has achieved data management at the many-hundred-Petabyte scale, and provides access to the entire community in a manner that is largely transparent to the end users. The main computing challenge of the next decade for the LHC experiments is presented by the HL-LHC program. Other large HEP experiments, such as DUNE and Belle II, have large-scale computing needs and afford opportunities for collaboration on the same timescale. Many of the computing facilities supporting HEP experiments are shared and face common challenges, and the same is true for software libraries and services. The LHC experiments and their WLCG- partners, DUNE and Belle II, are now collaborating to evolve the computing infrastructure and services for their future needs, facilitated by the WLCG organization, OSG, the HEP Software Foundation and development projects such as HEP-CCE, IRIS-HEP and SWIFT-HEP. In this paper we outline the strategy by which the international HEP computing infrastructure, software and services should evolve through the collaboration of large and smaller scale HEP experiments, while respecting the specific needs of each community. We also highlight how the same infrastructure would be a benefit for other sciences, sharing similar needs with HEP. This proposal is in line with the OSG/WLCG strategy for addressing computing for HL-LHC and is aligned with European and other international strategies in computing for large scale science. The European Strategy for Particle Physics in 2020 agreed to the principles laid out above, in its final report.

Comments: contribution to Snowmass 2021
Subjects: **Computational Physics (physics.comp-ph)**
Cite as: arXiv:2203.07237 **[physics.comp-ph]**
(or arXiv:2203.07237v1 **[physics.comp-ph]** for this version)
https://doi.org/10.48550/arXiv.2203.07237

Power Usage Efficiency

The modernisation of the facilities generally reduces the PUE (overhead of cooling)

The new CERN Prevessin Computing Center (PCC) PUE will be ~1.1

The effect of the PCC on the energy needs for ATLAS and CMS CPUs at CERN is on the right plot

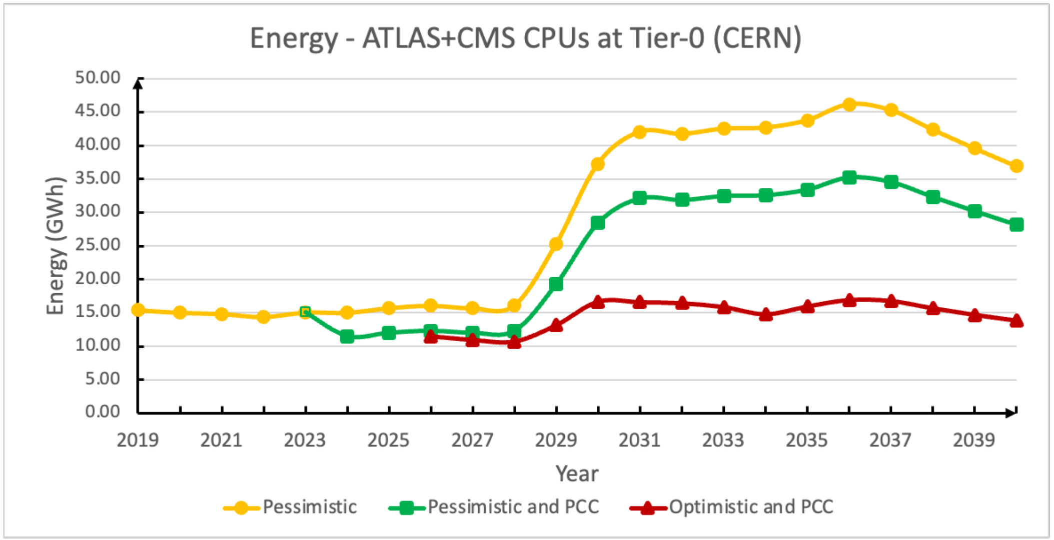

The introduction of the PCC reduces by 30% the CPU energy needs

The successful completion of the R&D program reduces the needs by another 50%