Computing: challenges and future directions

S. Campana (CERN)

Introduction

For HEP software and computing the time horizon of future challenges is the next 15 years

The main contributor to those challenges is HL-LHC, both in terms of volume and complexity. The largest needs come from ATLAS and CMS

The LHC computing resources are provided by the WLCG infrastructure. Other HEP experiments will share a large part of such an infrastructure. Other sciences will use many of the same facilities

The HL-LHC challenge statement at ECFA 2016

In a nutshell:

- Projections assume constant funding every year for LHC computing
- Technology improvements will bring ~20% more resources every year

And this was the initial conclusion

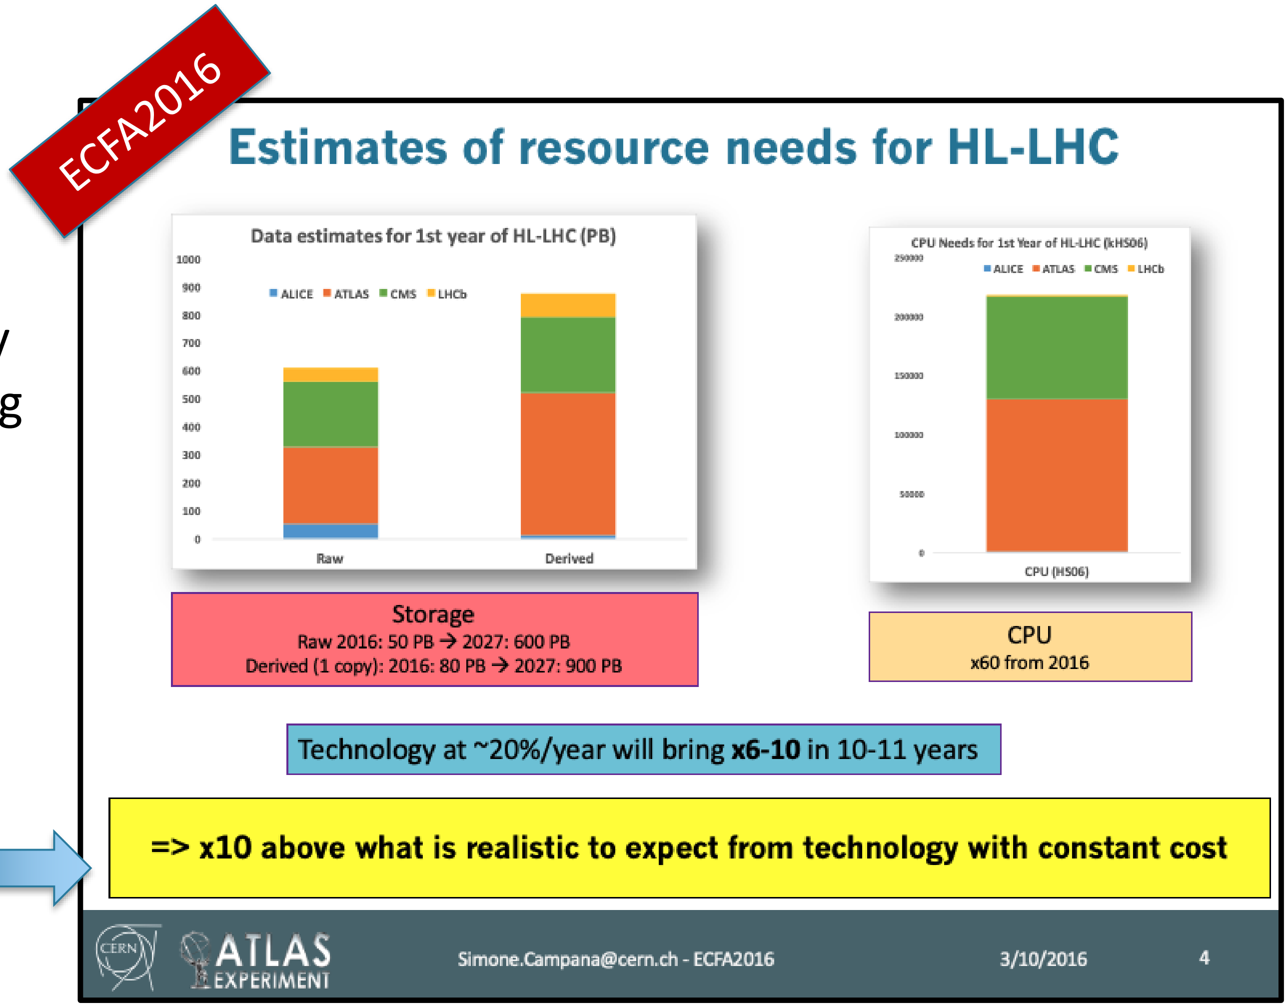

The HL-LHC computing roadmap process

HEP Software Foundation Community Whitepaper: a bottom-up exercise. Identify the areas of work to address the HEP challenges of the 2020s

The first WLCG strategy toward HL-LHC document: a top-down high-level prioritization of the whitepaper, for the LHC needs

The LHCC review series of HL-LHC computing: a multistep process tracking the progress towards HL-LHC

- May 2020: review of ATLAS and CMS plans, Data Management (DOMA), offline software, the WLCG collaboration and infrastructure. Documents
- November 2021: update from ATLAS and CMS, common software activities (generators, simulation, foundation software, analysis, DOMA). Charge

CPU needs

At HL-LHC:

CMS CPU needs are dominated by reconstruction of data and MC (>80%)

ATLAS CPU needs: > 20% reconstruction, ~30% simulation (Geant4 + Fast Sim), 18% event generators

Event generators, GEANT4 and ROOT are common software libraries used by both ATLAS and CMS

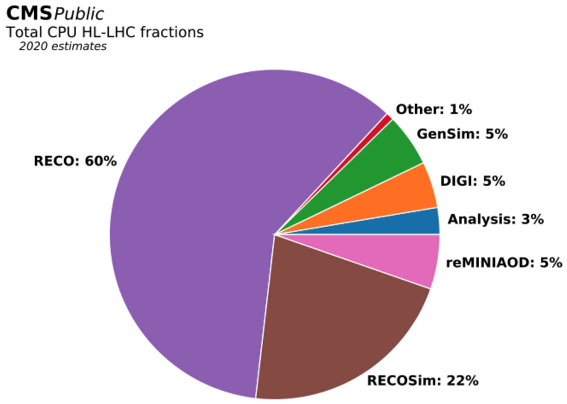

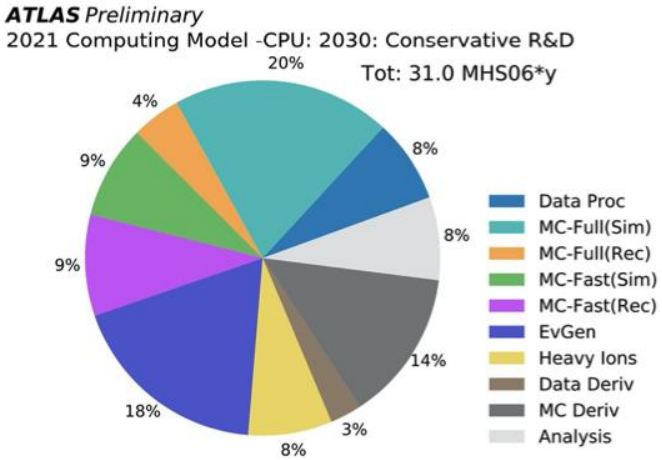

Event Generators

The HSF generators WG identified the main areas of work in this paper.
Some highlights:

- R&D activities for moving part of the code to GPUs and vectorization. E.g. Matrix Element calculations: small code base, suitable for parallelization

- Reducing the fraction of events with negative weights (example from Madgraph5_aMC@NLO)

Challenges and risks:

- No current real estimate of Event Generation CPU cost at higher precision (e.g. NNLO)

- Funding and career opportunities for performance work on generators remain problematic

	MC@NLO 111	MC@NLO-Δ Δ-441
$pp \to e^+e^-$	6.9% (1.3)	2.0% (1.1)
$pp \to e^+\nu_e$	7.2% (1.4)	2.3% (1.1)
$pp \to H$	10.4% (1.6)	0.5% (1.0)
$pp \to Hb\bar{b}$	40.3% (27)	31.3% (7.2)
$pp \to W^+j$	21.7% (3.1)	7.4% (1.4)
$pp \to W^+t\bar{t}$	16.2% (2.2)	11.5% (1.7)
$pp \to t\bar{t}$	23.0% (3.4)	7.7% (1.4)

Simulation

The workhorse of LHC simulation is Geant4, complemented with fast simulation, including Machine Learning techniques

The strategy towards HL-LHC:

- Improve the description of the physics processes, adequately to the HL-LHC precision needs
- Adiabatically improve the CPU performance
- Leverage fast simulation techniques and improve their accuracy for larger use
- R&D to investigate novel techniques and architectures, e.g. GPUs

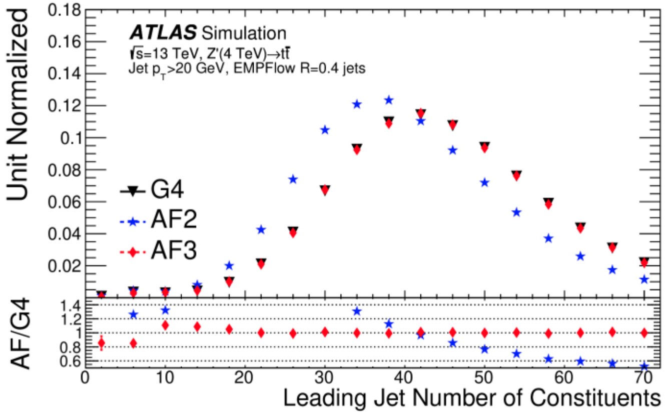

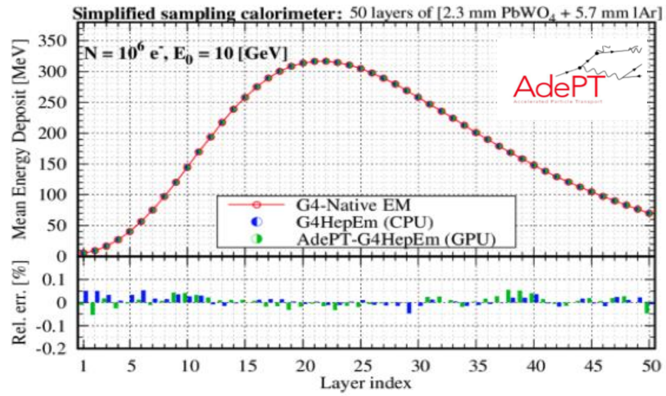

Reconstruction

Reconstruction (data and Monte Carlo) has a very large impact on CPU needs. It is very experiment dependent

- Continuous improvement of the Run-4 reconstruction code through tuning and partial re-engineering
- Adoption of Fast Reconstruction techniques where the discrepancy with the standard reconstruction is tolerable
- R&D to offload part of the reconstruction code to accelerators

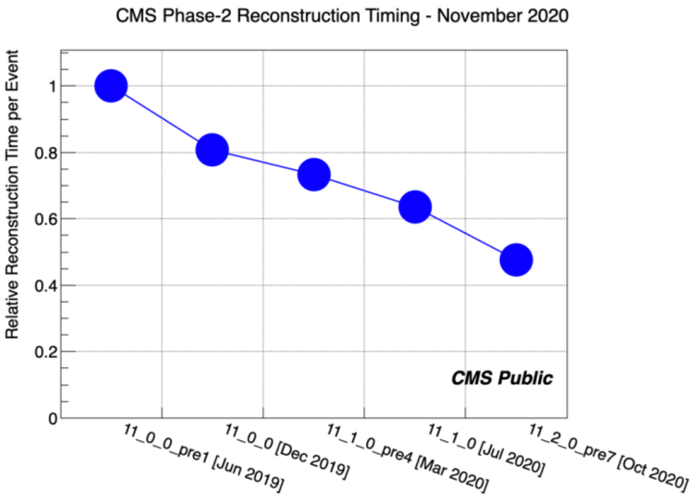

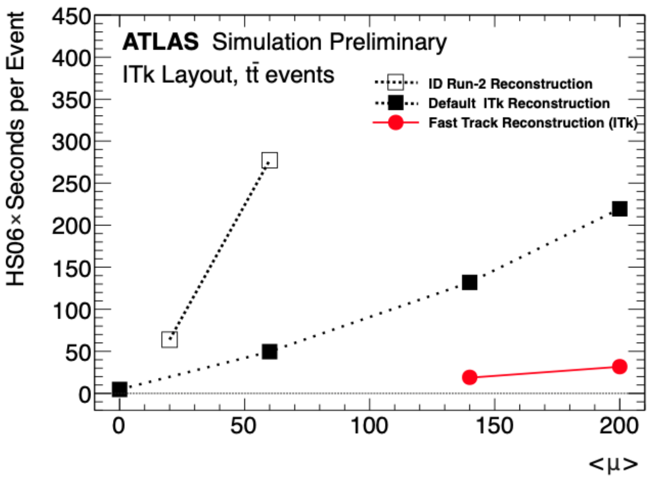

Storage needs

At HL-LHC:

Archive storage needs (tape) are dominated by the amount of RAW data

- Possible gains with compression/suppression, but moderate
- Be prepared to invest in the tape volume we need and optimize the rest

Disk storage needs are dominated by the amount of reconstructed data, in different formats and versions. The strategy focuses on:

- Reduced analysis formats
- Data carousels

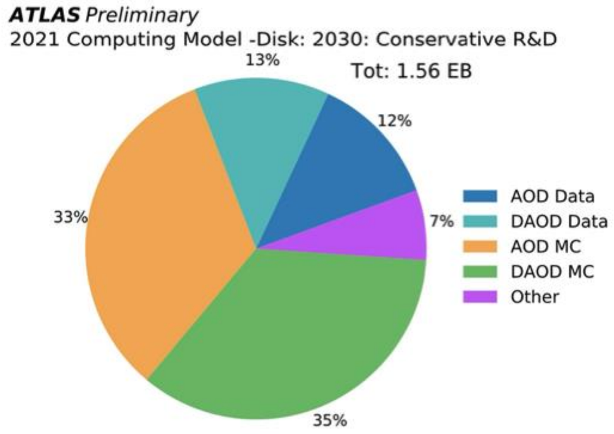

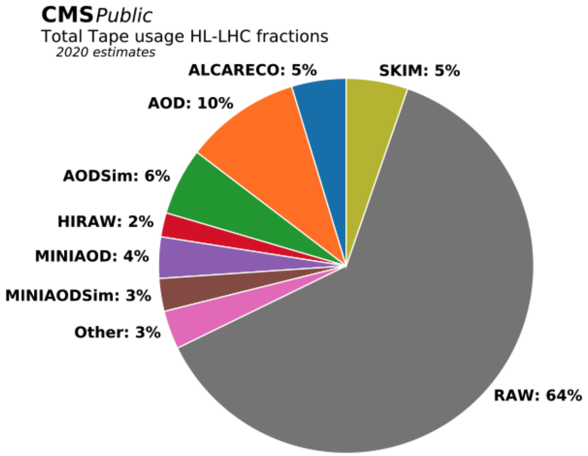

Reduced analysis formats

Both ATLAS and CMS developed derived formats from AODs

- A reduced amount of information, but still sufficient for most analyses

MiniAODs were the main analysis format in CMS in LHC Run-2

- Some analyses also used NanoAODs

The Run-4 strategy of ATLAS and CMS foresees the adoption of the slimmer format for a sizeable (~50%) number of analyses

This will reduce considerably the disk storage needs, assuming the larger formats will not be generally needed on disk

NanoAOD:
- 1kB in Run-3
- <5kB in HL-LHC

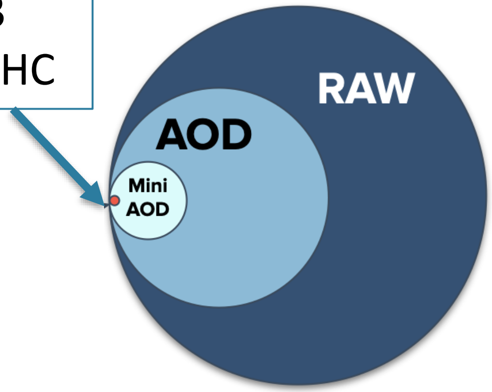

Data Carousels

Tape in our infrastructure is not just an archive media. Data is recalled regularly for further processing

- Organized and marshalled activities, nothing to do with the old use of hierarchical mass storages

Active use of tape will play an even more important role in HL-LHC

- Leverage the lower cost of tape **MEDIA** w.r.t. disk. Rely on the capability of tape systems to recall data fast enough

Tape bandwidth has also a cost that needs to be monitored and optimized

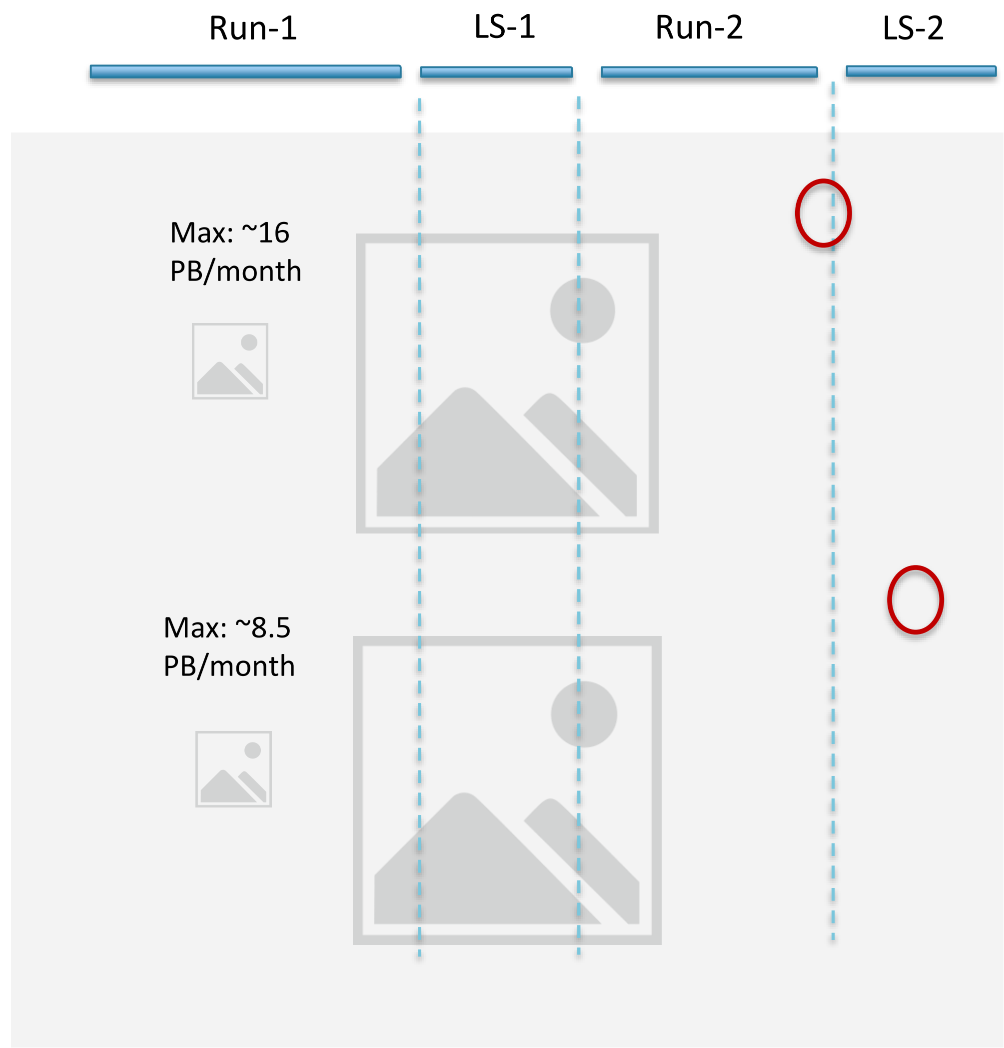

Networks

We established an incremental process to prepare for the HL-LHC network needs, through a regular dialog between the network providers, the experiments and the facilities

Regular Data Challenges of increasing complexity are part of this process.

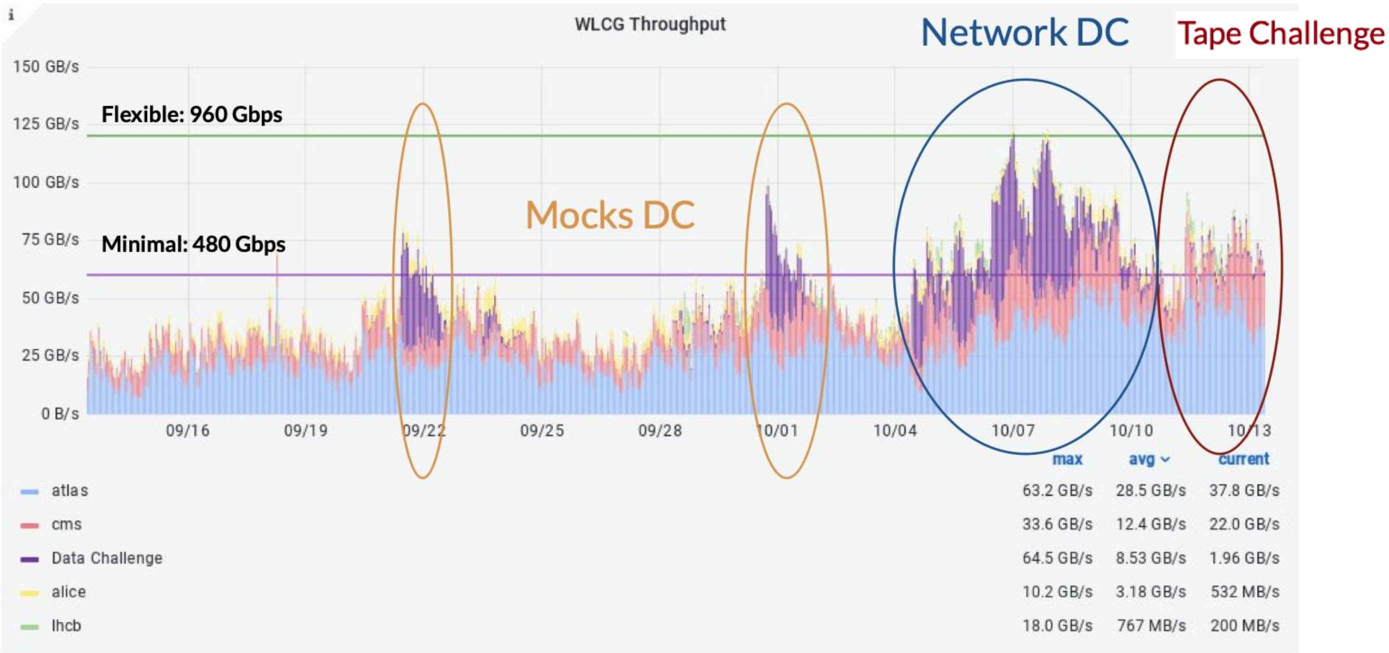

Metrics exist for a baseline (minimal) scenario and a flexible scenario

The first data challenge took place in October 2021, in preparation for LHC Run-3

Analysis

End user analysis will benefit from the re-design of the internal ROOT data format: from TTree to RNtuple

- Intent to adopt it for AODs and their derivatives by ATLAS and CMS

Tests indicate that

- RNTuple is 10-20% smaller than TTree, resulting in storage saving
- Read throughput improves by x3-x5 with RNtuple

Additionally, RNTuple enables fast adaptation to modern technologies, like object stores

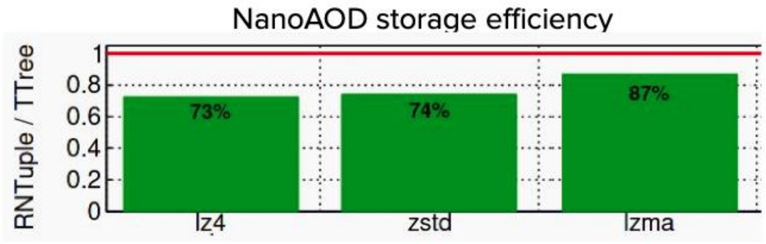

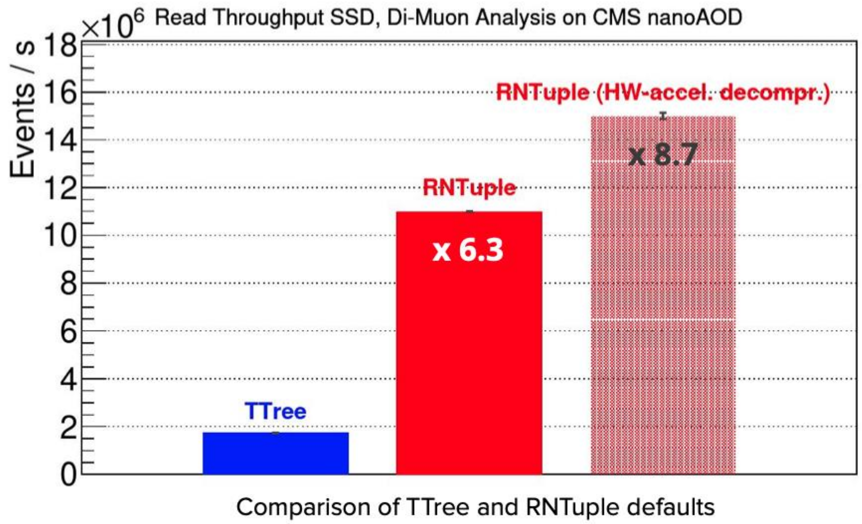

HL-LHC parameters for computing planning

The current scenarios for Run-4 and Run-5 were agreed with the LHC physics committee and are regularly reviewed

	Run-4 (2028-2030)	Run-5 (2032-2034)
ATLAS/CMS luminosity	<270/fb	<350/fb
ATLAS/CMS average pile-up	<140	<200
LHCb luminosity	15/fb	50/fb
ALICE luminosity (pp)	100/pb	
Running time pp	6 M seconds	8 M seconds
Running time ions	1.2 M seconds	

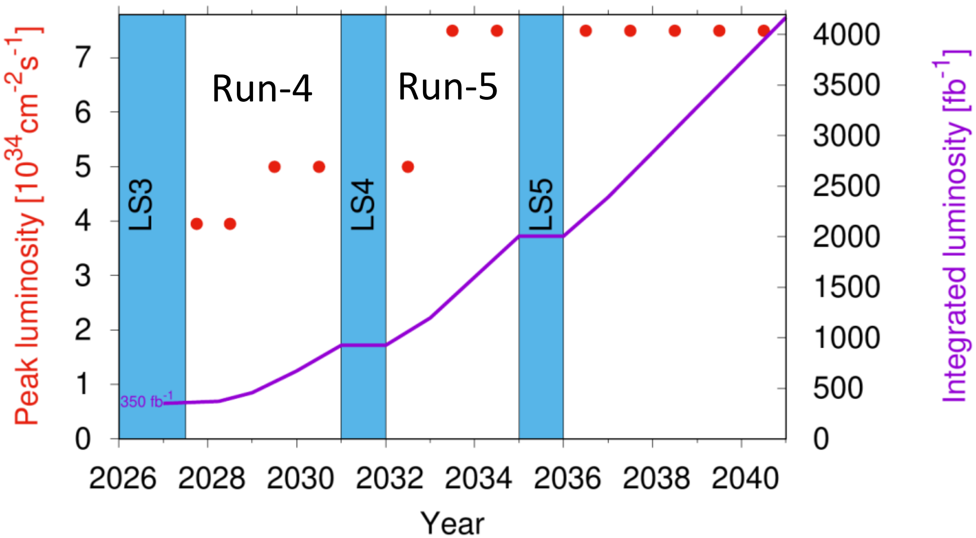

2027: commissioning year

Run-4 production years: 2028 to 2030

LS4: 2031

Run-5 production years: 2032 to 2034

ATLAS and CMS needs for HL-LHC

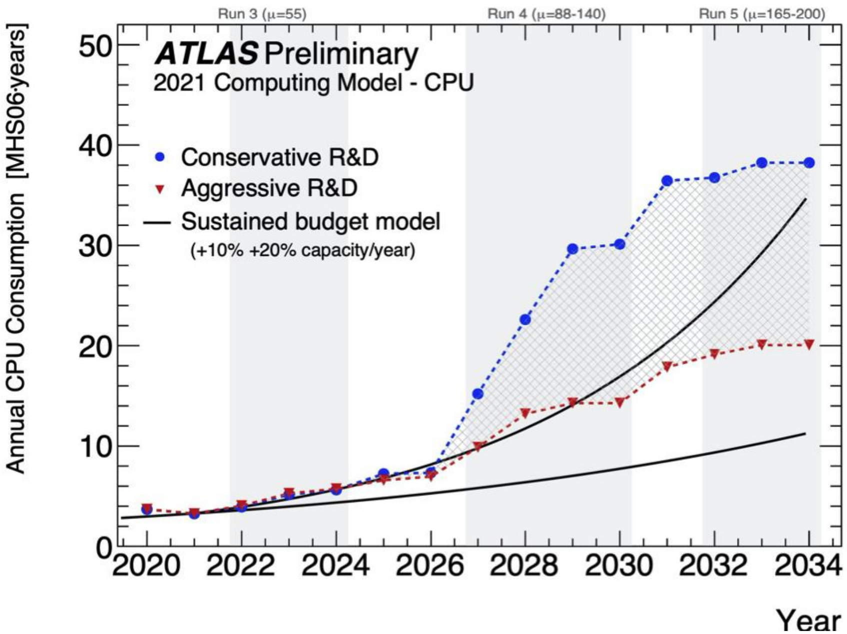

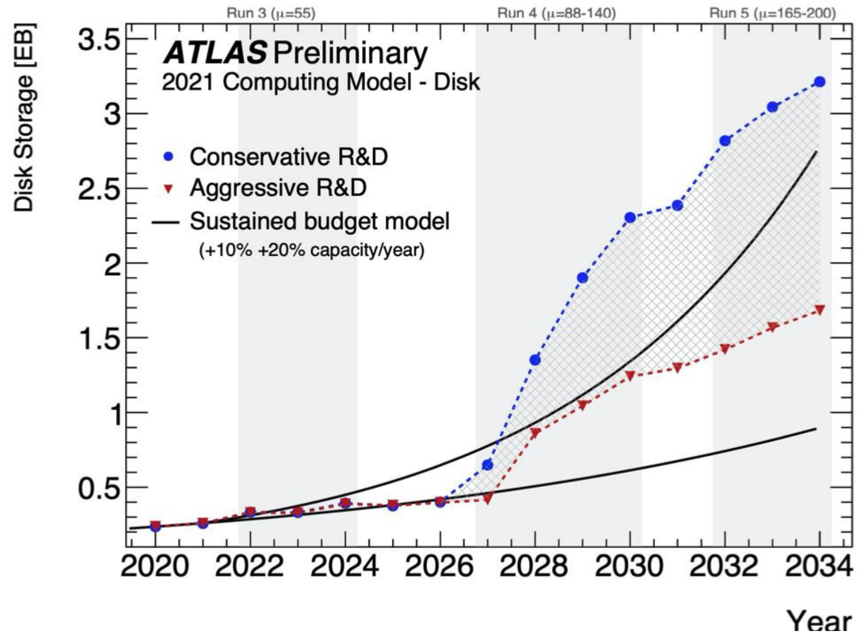

Those are the current estimates from ATLAS. Similar plots exist for CMS (will be public soon)

- The gap between available and needed resources is filling up, assuming the main R&D activities are successful
- Investing in further (identified) R&D activities would fill this gap further. Need more effort
- There are still large uncertainties

HL-LHC for Alice and LHCb

I did not discuss ALICE and LHCb: Run-4 will be at a similar scale than Run-3. Plans for Run-5 and beyond are being developed

Early thinking in e.g. LHCb shows a challenging scenario (as ATLAS and CMS some years ago)

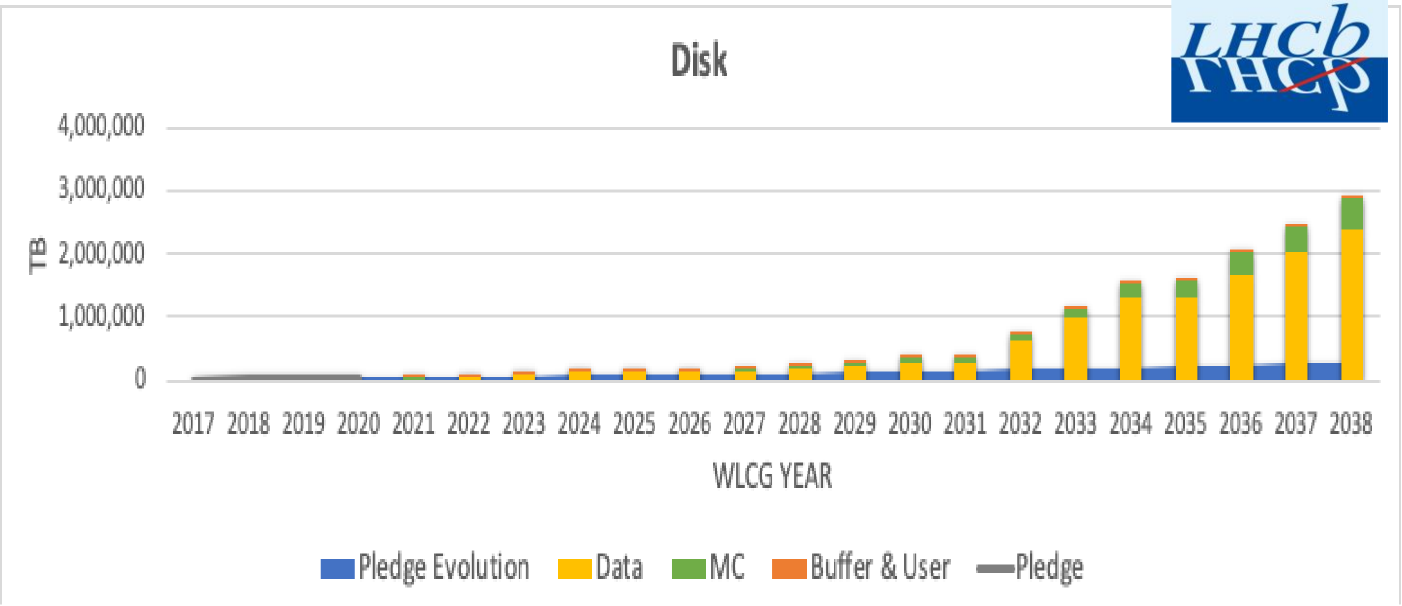

Hardware cost and market trends

Hardware cost is more and more dominated by market trends rather than technology

The assumption of +20%/year in storage and CPU for the same budget is still holding, but with large fluctuations (see plots below and notice the log scale)

It is not wise to expect better than this …

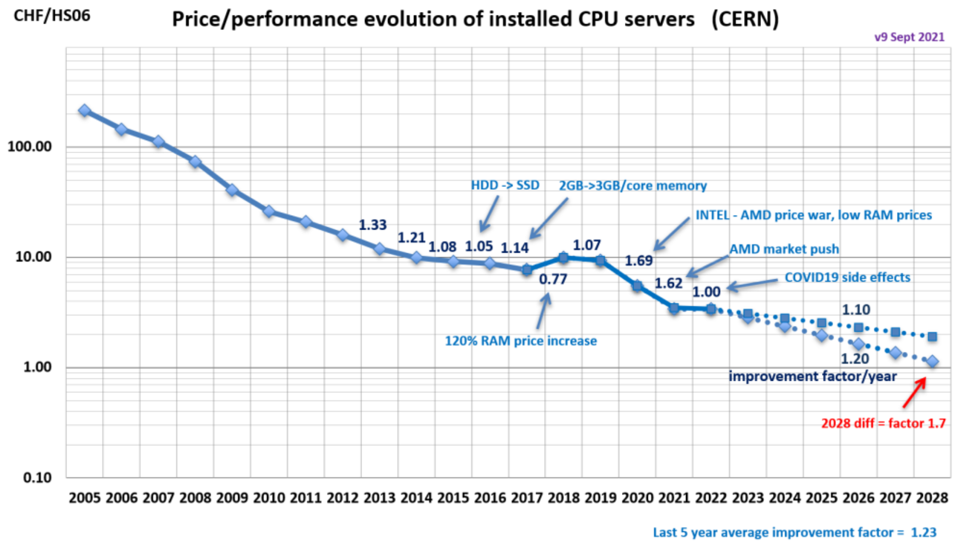

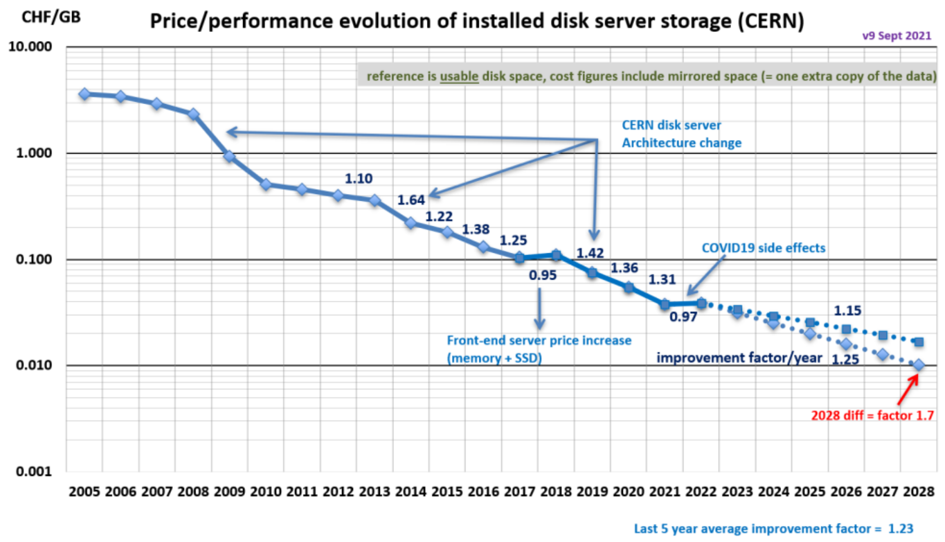

Opportunistic resources

A mitigation for the gap in resources comes from opportunistic CPUs. HPC centers offer a unique opportunity

Accessing and using resources at HPC centers comes with different challenges:

- Diversity in access and usage policies, edge services, system architectures
- Heterogeneous computing architectures: non x86 CPUs and GPUs

It is important to be selective and aim at the cases with maximum benefit/cost. Software portability and the success in integrating accelerators will play an important role

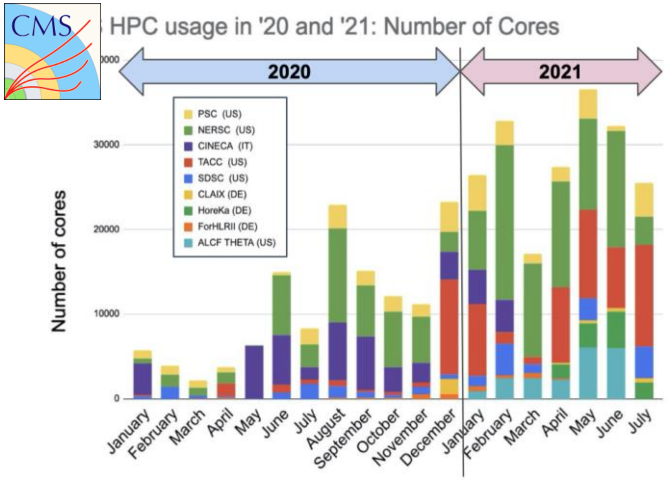

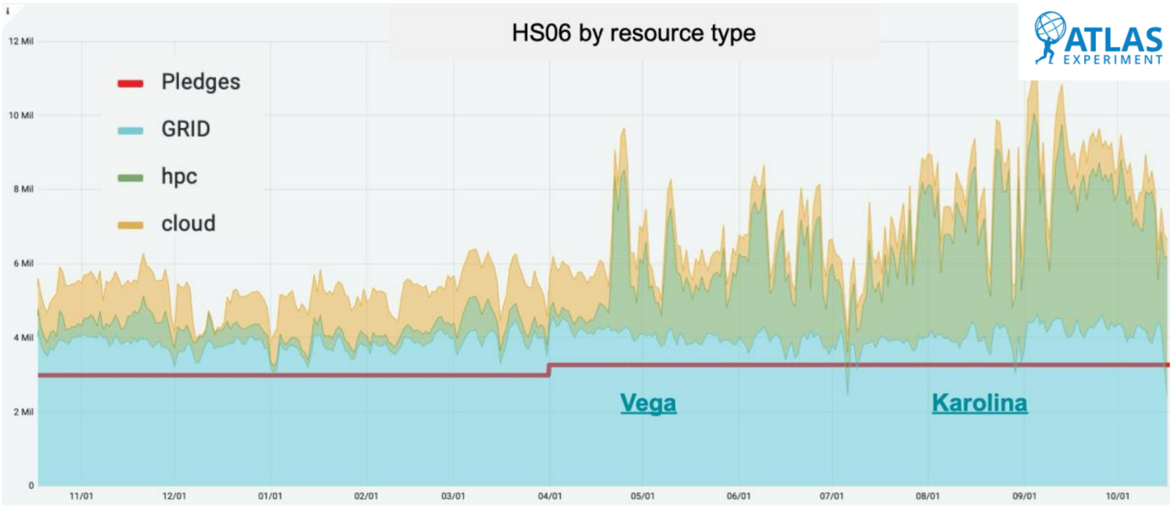

Heterogeneous Architectures

Heterogeneous architectures: complementing CPU capacity with accelerators (e.g. GPUs)

Playing a fundamental role in Run-3 already, in most online systems. Non exhaustive examples:

LHCb: exploitation of heterogeneous architectures, thanks to Allen framework:

- for partial reconstruction in Run-3
- for full reconstruction in Run-4?
- for fast reconstruction in Run-4 or Run-5 ?

Ongoing work in all experiments to leverage heterogeneous architectures in offline computing

Alice O2: Speed up from GPU usage + from algorithmic improvements + tuning on CPUs

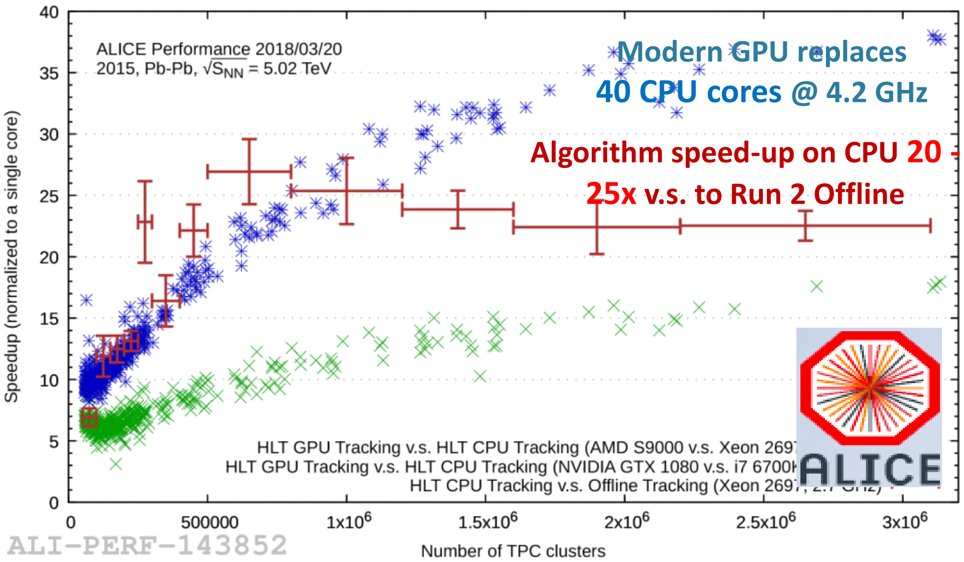

Infrastructure and services sustainability

The HL-LHC challenge is not just about resources. The sustainability of the infrastructure over the next 15 years is also challenging

- Focus on open-source community solutions, not specific to HEP
- Progress towards a more flexible model integrating a heterogeneous landscape of facilities

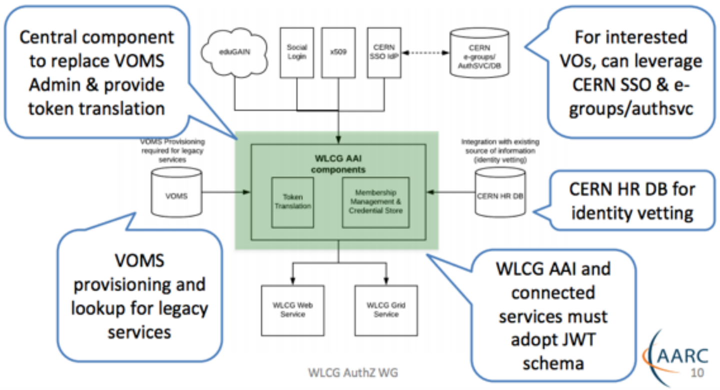

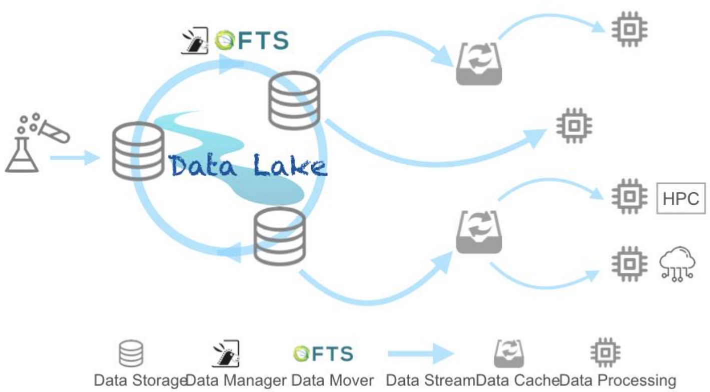

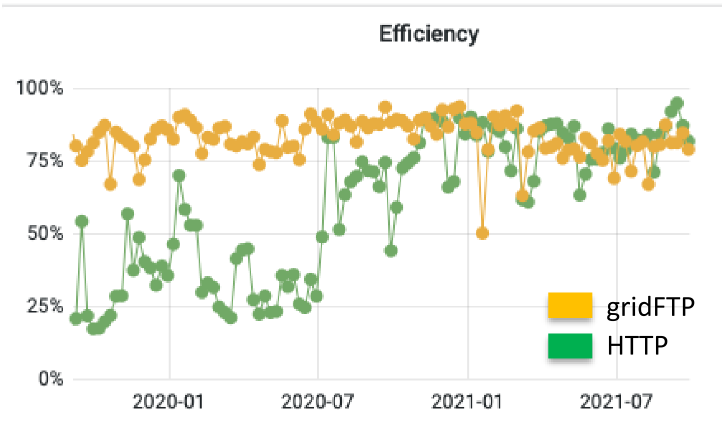

Sustainability through collaboration

The Belle-2 and DUNE HEP experiments leverage the same infrastructure as WLCG (and they are "associate members")

DUNE DEEP UNDERGROUND NEUTRINO EXPERIMENT

Belle II

Evolution of Scientific Computing in the next decade: HEP and beyond

WLCG Overview Board
17th December 2018

Contact: Ian Bird (Ian.Bird@cern.ch),
Simone Campana (Simone.Campana@cern.ch)

Astroparticle Physics European Consortium (APPEC)

APPEC

APPEC Contribution to the European Particle Physics Strategy

December 17, 2018

Editorial Board:
S. Katsanevas, A. Masiero, T. Montaruli, J. de Kleuver, A. Haungs

Contact Person:
T. Montaruli (APPEC Chair from Jan. 1, 2019)
Email: teresa.montaruli@unige.ch

Website: http://www.appec.org

European Science Clusters

The Science Clusters are EU project launched in 2019 to link the world-class Research Infrastructures with the European Open Science Cloud (and influence its evolution)

We are part of the ESCAPE cluster (HENP, Astronomy, Astro-Particle)

The European Science Clusters produced a joint Position Statement on expectations and long-term commitment (role) in open science

Many opportunities to potentially leverage:

- Influence the direction of EOSC, EuroHPC
- Create synergies with other sciences in other clusters
- A step forward in our vision to share technologies and services with other experiments and sciences, with the same needs, and in many cases sharing the same infrastructure, for long term sustainability

ENVRI FAIR

EOSC-Life

ESCAPE

panosc

SSHOC

Conclusions

- We went through the process of defining the strategy to address the HL-LHC challenge. We are now implementing the plan following that strategy
- The HL-LHC is not a solved problem, however it is converging to a solution
- There is still a gap between needs and affordable resources with the currently available effort. More effort would contribute closing that gap, as the concrete priorities have been identified
- There are risks that could materialize and for which we need to be prepared for with mitigation strategies
- The HL-LHC challenge is not just resources. The infrastructure sustainability is at least as important