TIFR-Mumbai

September 2024

HEP Computing: successes and future challenges

Simone Campana (CERN)

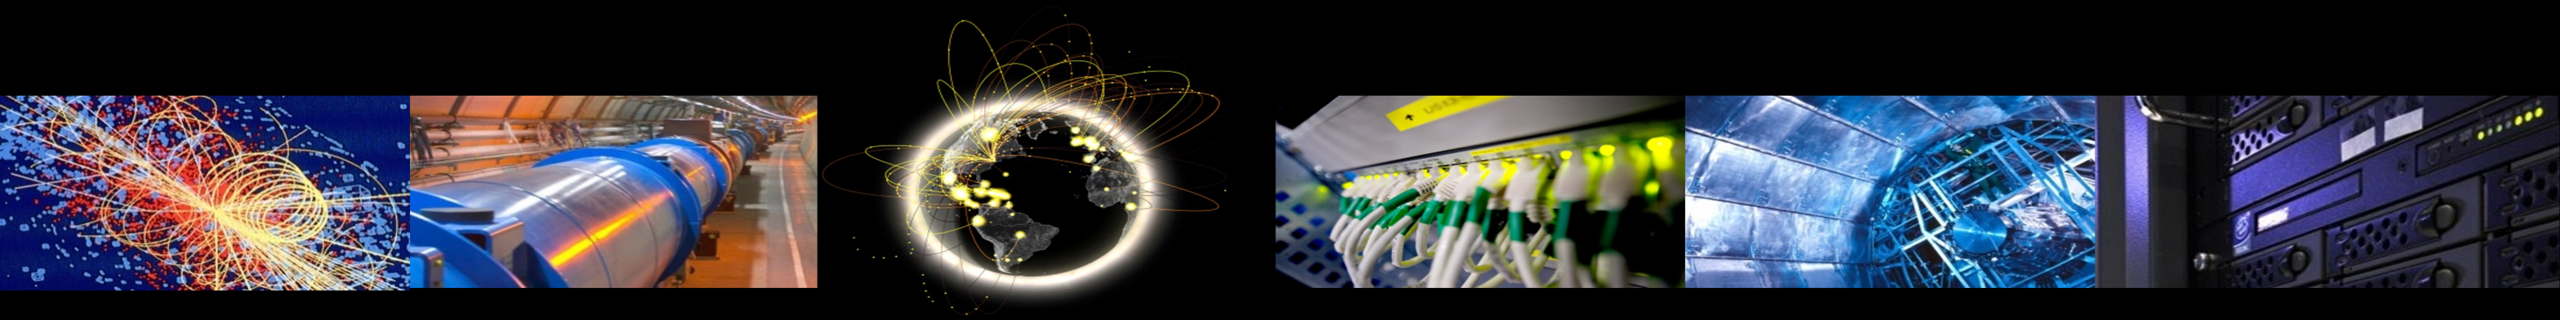

High Energy Physics

"The branch of physics dealing with the constitution, properties, and interactions of elementary particles especially as revealed in experiments using particle accelerators" (Webster)

Recreate the conditions of the early universe

– as "early" as possible –

by colliding particles as very high energy

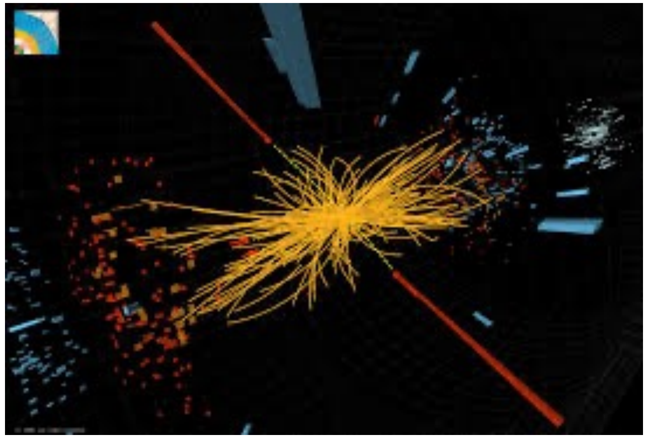

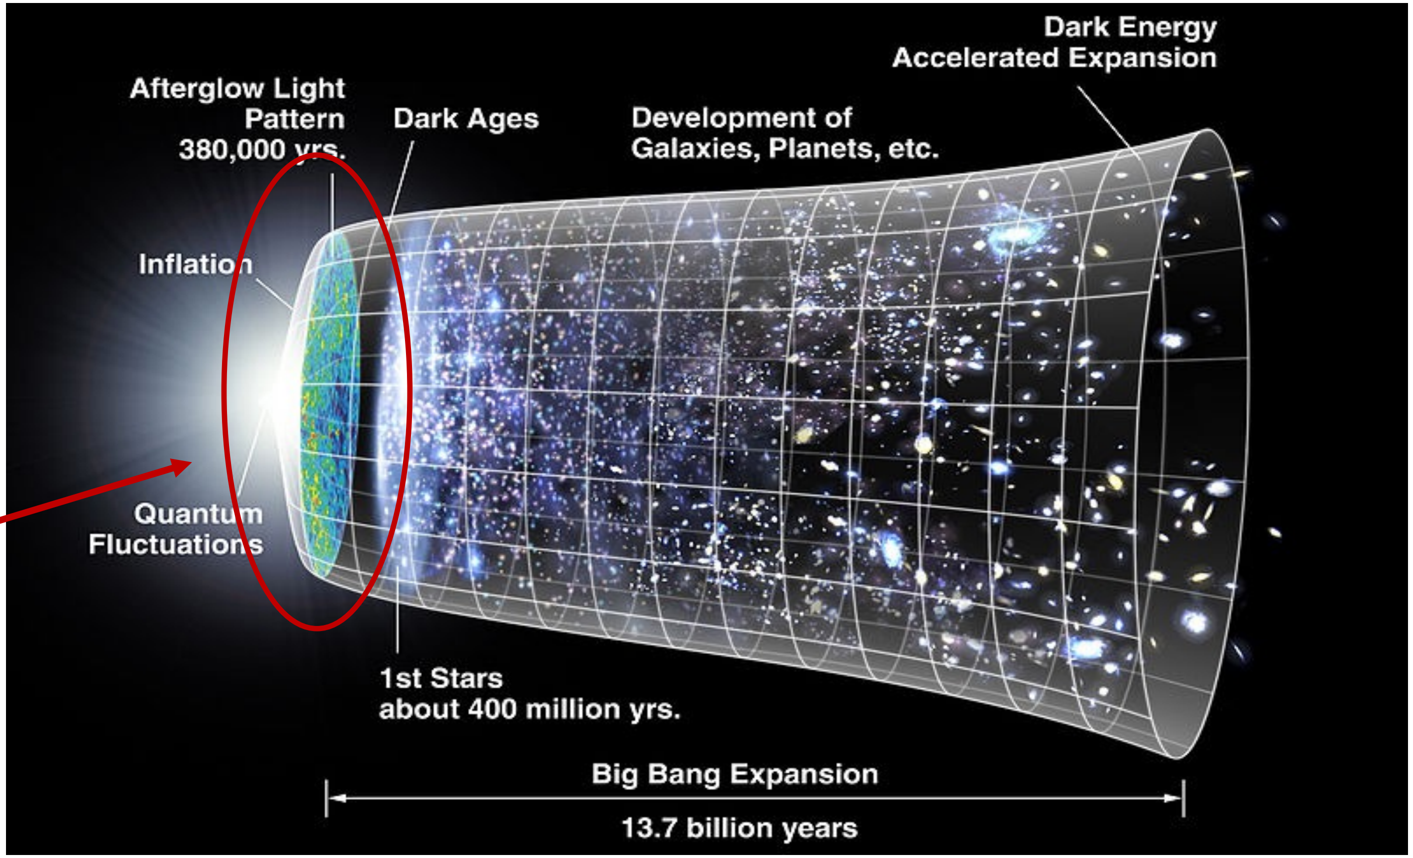

The Large Hadron Collider @ CERN

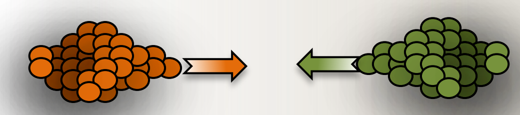

Proton bunches
$>10^{11}$ protons/bunch
(colliding at **~40MHz)**

p 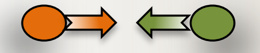 p

p-p collisions with interesting parton interactions **(<kHz)**

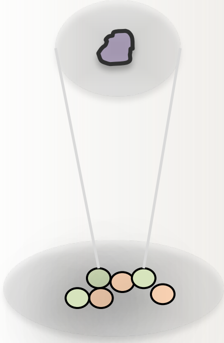

New Particle!
(<<**mHz?**)

27 km ring of superconducting magnets and accelerating cavities

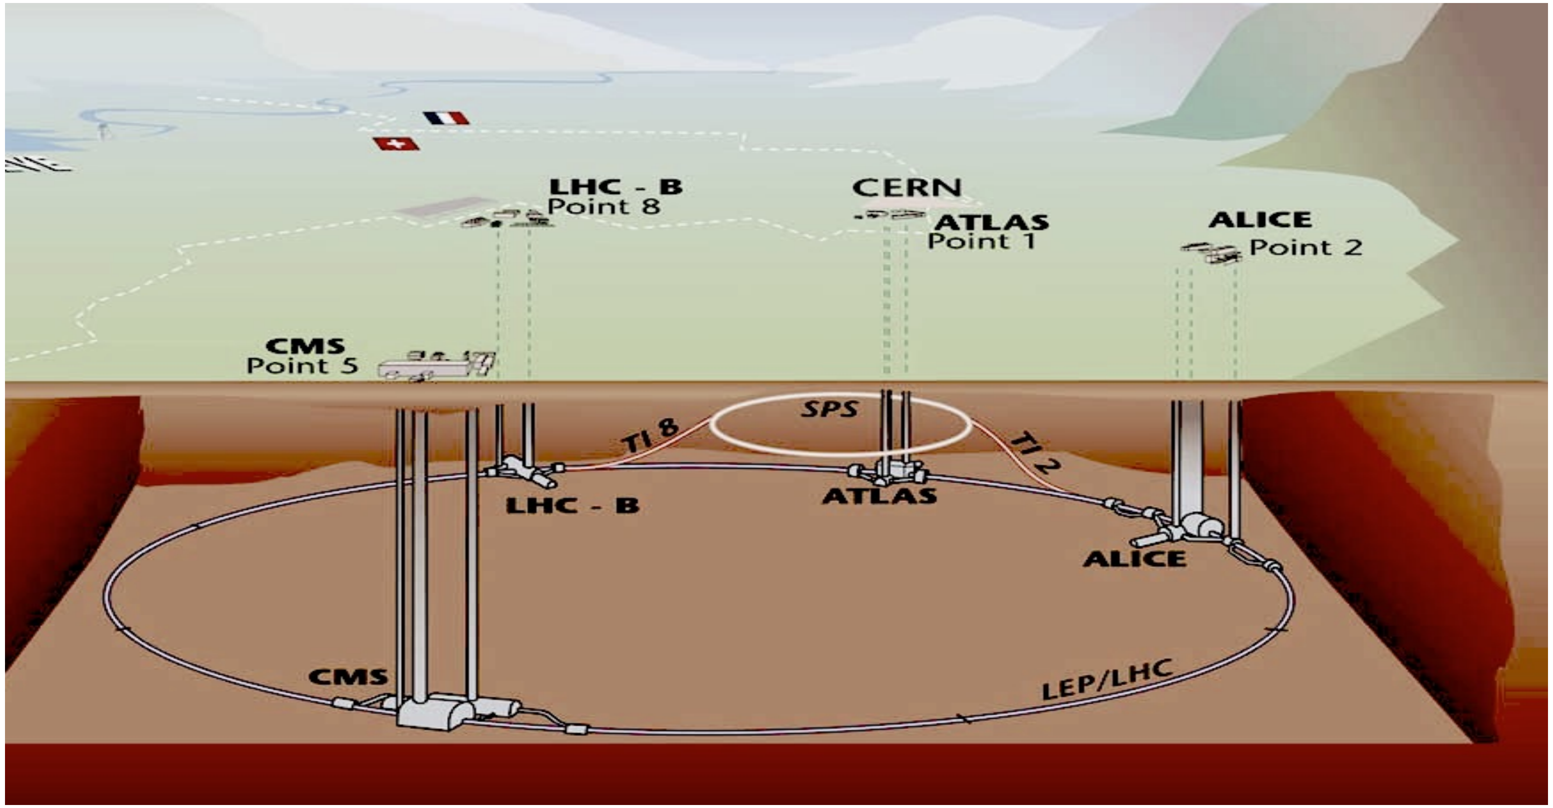

Four large detectors – ALICE, ATLAS, CMS, LHCb

Computing @ High Energy Physics

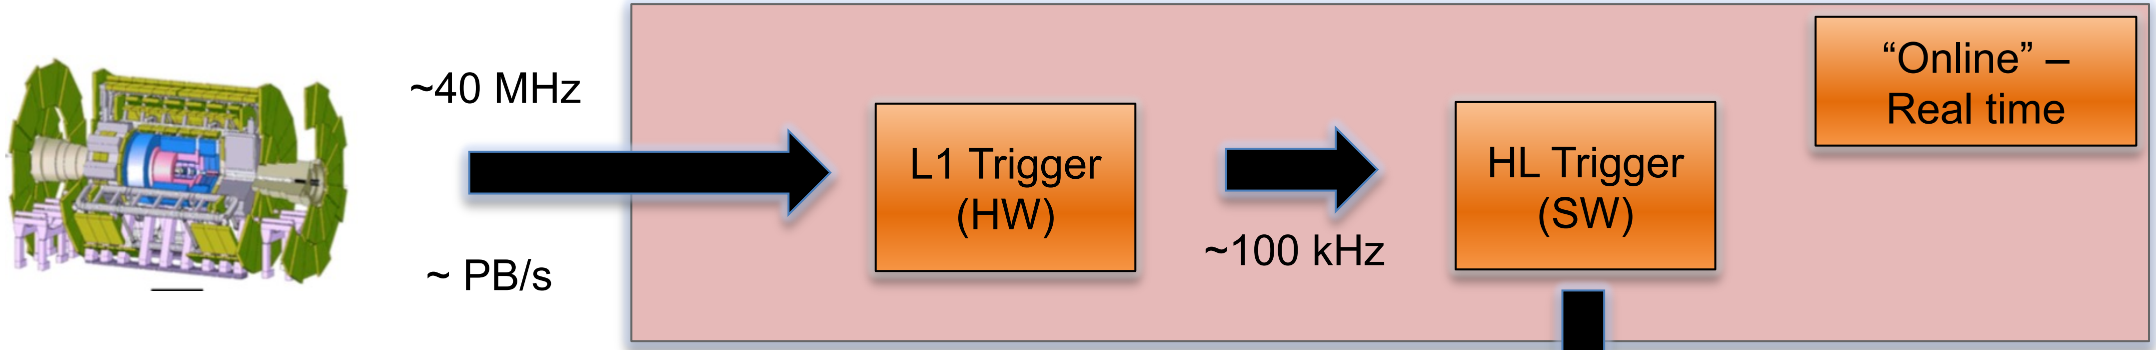

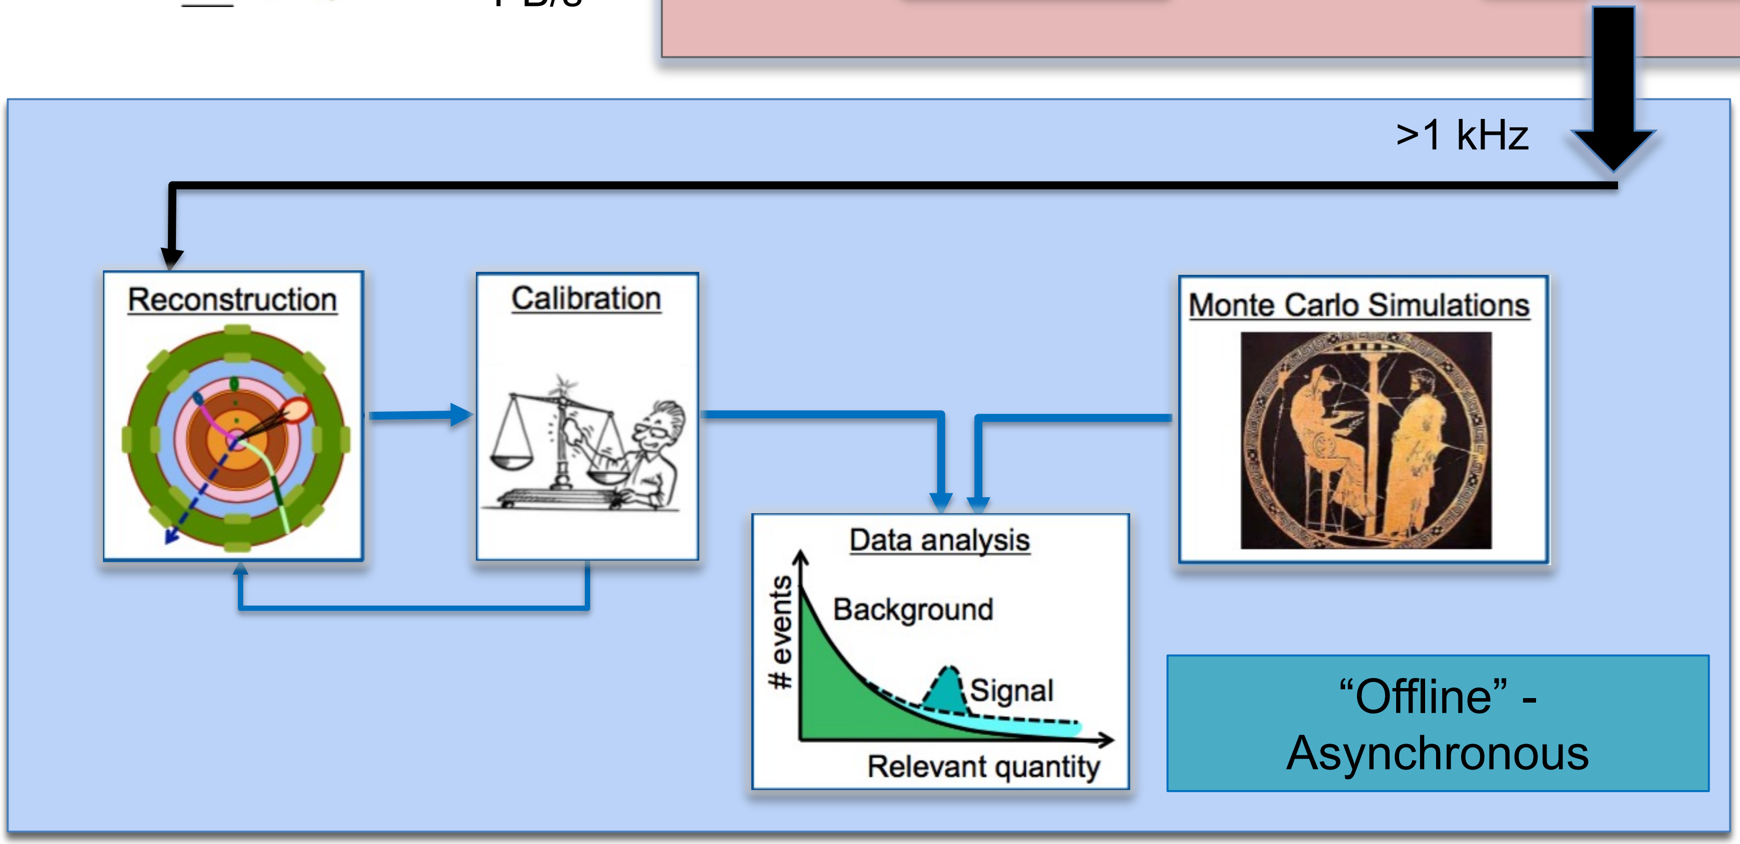

From Big Data (40 Million collisions/s

To Useful Data (>1000 events/s) for analysis

>1MB/event

100 PB/year

Distributed vs Centralised Computing

Large scale HEP computing embraced a distributed model since early 2000s

Based on network service technologies, federating national and international initiatives

The distributed grid model fits:

- **the national strategies of the Funding Agencies:** develop the cyberinfrastructure for scientific computing in the country – data centres, networks and the required expertise

- **the distributed nature of the HEP community:** hundreds of institutes contributing as part of the HEP experiments and international collaborations

WLCG Worldwide LHC Computing Grid

The WLCG Collaboration provides the distributed computing infrastructure for the needs of the CERN Large Hadron Collider experiments

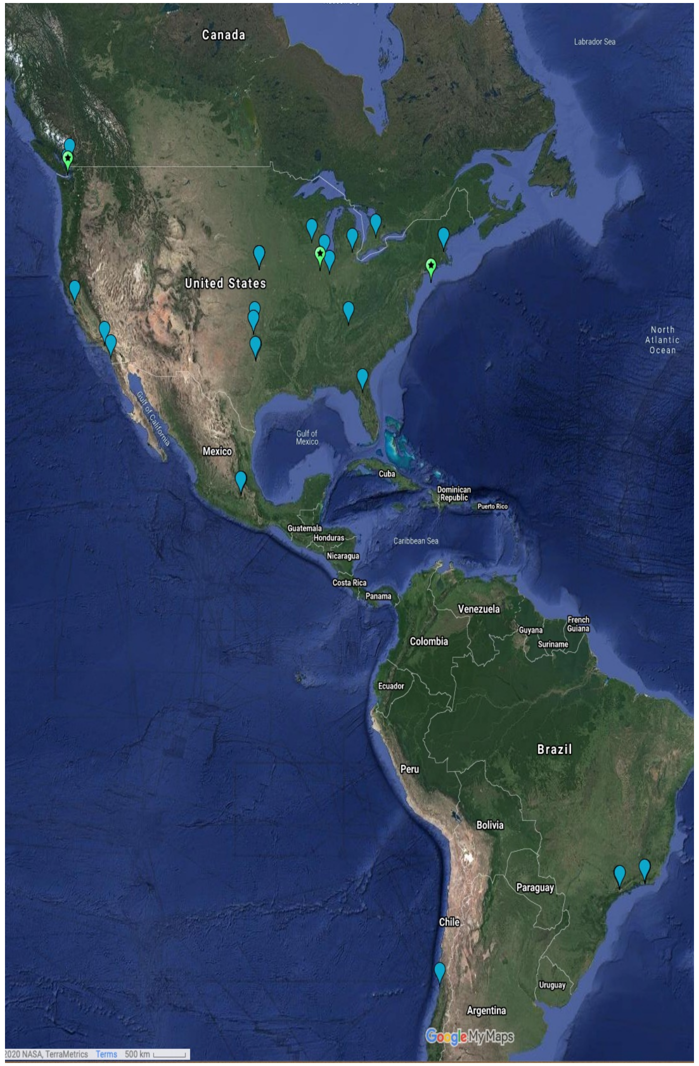

WLCG
Worldwide LHC Computing Grid

WLCG in Numbers…

65 MoUs, 40+ counties, 159 sites:

Tier-0 (CERN): data recording, reconstruction, distribution, archive. > 99% target availability

14 Tier-1s: archive storage, re-processing, analysis. > 98% target availability

144 Tier-2: Simulation, analysis. > 95% target availability

WLCG Resources 2024	
CPU (kHS23*)	12,315
Disk (PB)	1,255
Tape (PB)	2,461

*1 core ~ 10 - 15 HS23

Hierarchical Structure

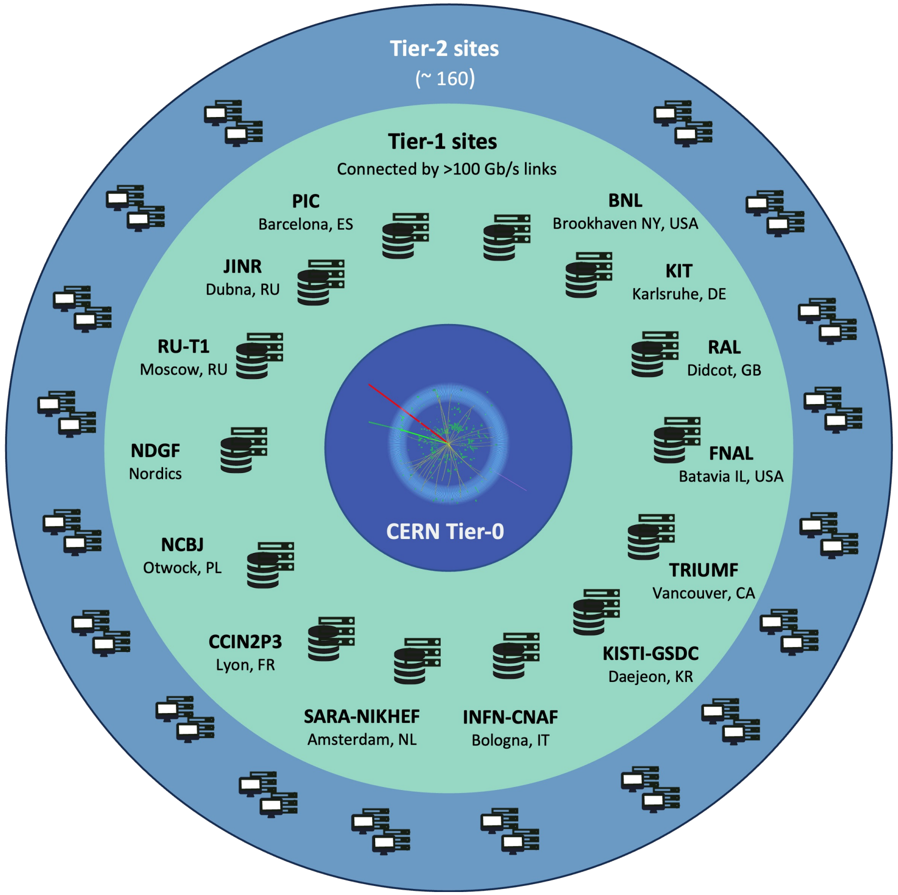

WLCG
Worldwide LHC Computing Grid

India in WLCG

Two Tier-2 sites in Kolkata (ALICE) and Mumbai (CMS). MoU signed by DAE in June 2006

6% of the ALICE T2 resources and 8% of CMS T2 resources

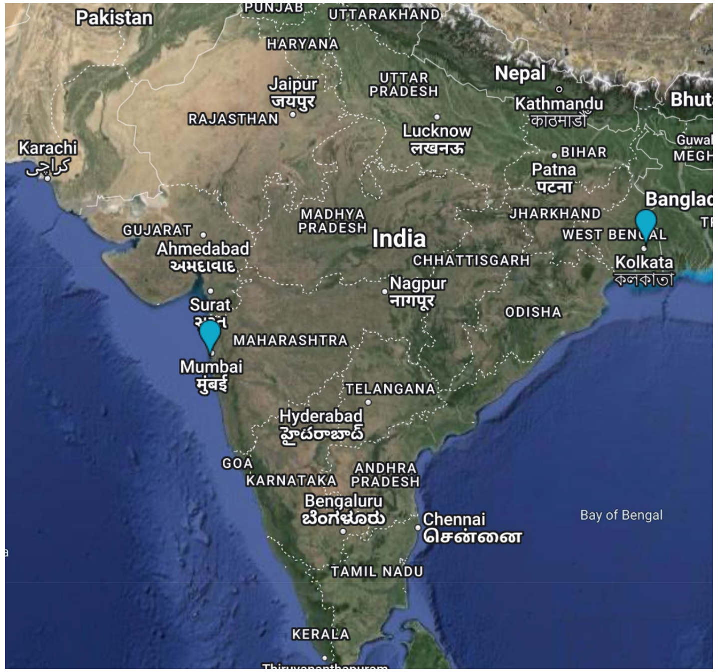

Federation	Tier	VO	Country	Year	Type	Pledge
IN-DAE-KOLKATA-TIER2	2	ALICE	India	2024	CPU	60000 HEPscore23
IN-DAE-KOLKATA-TIER2	2	ALICE	India	2024	Disk	3000 TBytes
IN-INDIACMS-TIFR	2	CMS	India	2024	CPU	140000 HEPscore23
IN-INDIACMS-TIFR	2	CMS	India	2024	Disk	11000 TBytes

Tier-0 workflow in 2024

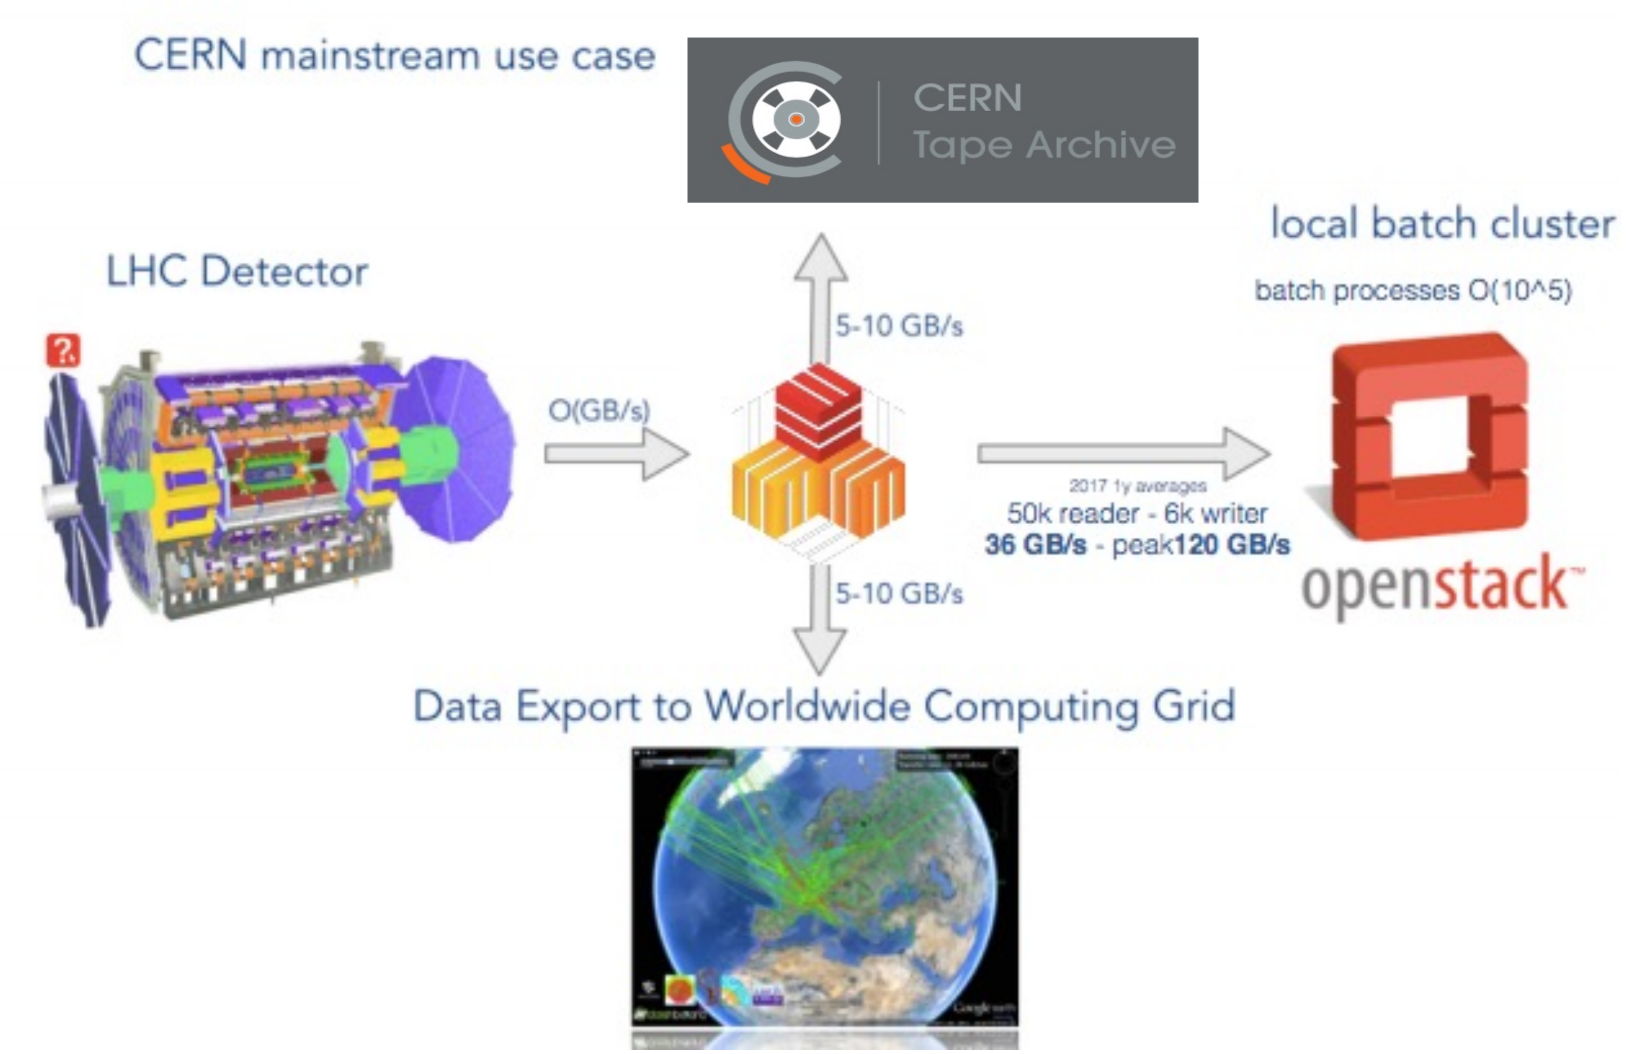

- Receive the RAW data from the experimental areas and archive them for long term preservation
- Perform data quality assessment and reconstruct the events into formats suitable for analysis.
- Export RAW and reconstructed data data to WLCG centers. RAW data is archive at Tier-1s. Reconstructed data is further distributed for analysis.

WLCG Worldwide LHC Computing Grid

The LHC RAW data archive

The LHC physics program is organized in ”runs” interleaved by “Long -Shutdown periods

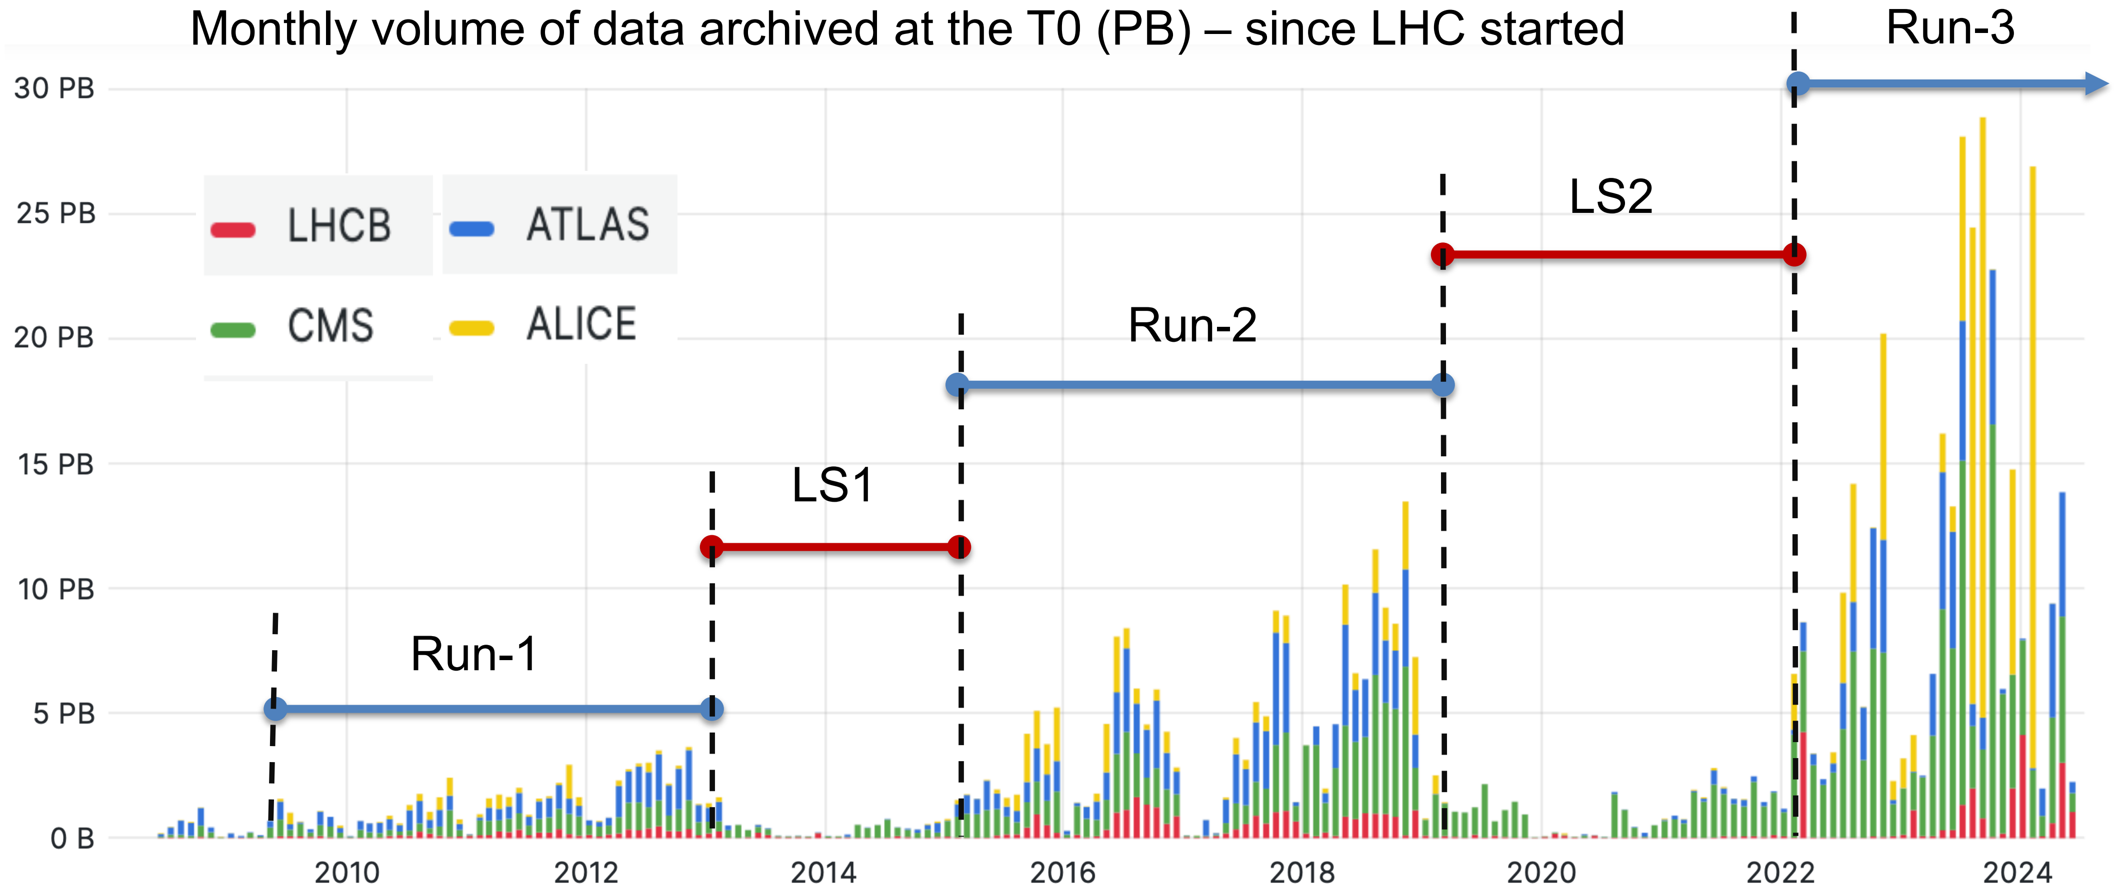

The initial Computing Model

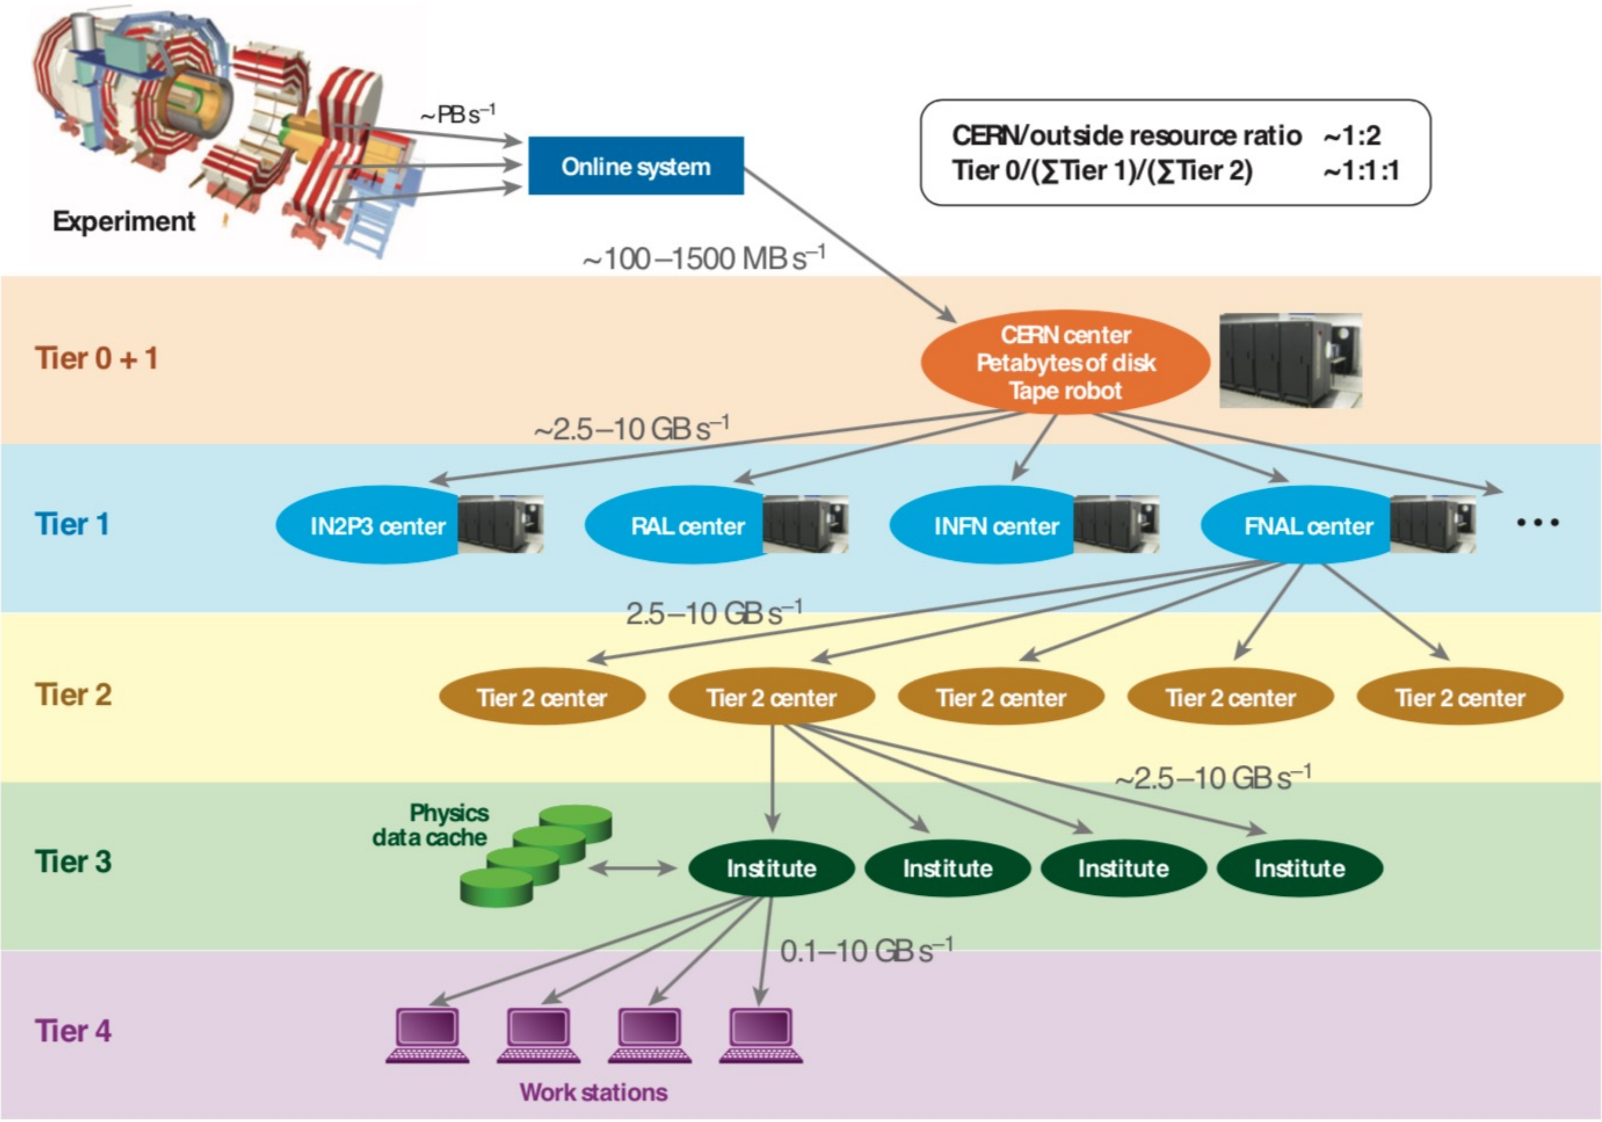

This was the initial computing model (1999)

Uncertainty over network performance, reliability

Focus on distributing data *globally* to compute resources in a hierarchical structure

Every center runs a compute and a storage service

No concept of data remote from compute

Predefined roles of Tiers: T2s for simulation and analysis, T1s for reconstruction and archive

Quickly evolved …

WLCG
Worldwide LHC Computing Grid

WLCG networks

WAN connectivity increased x10 in the last 10 years

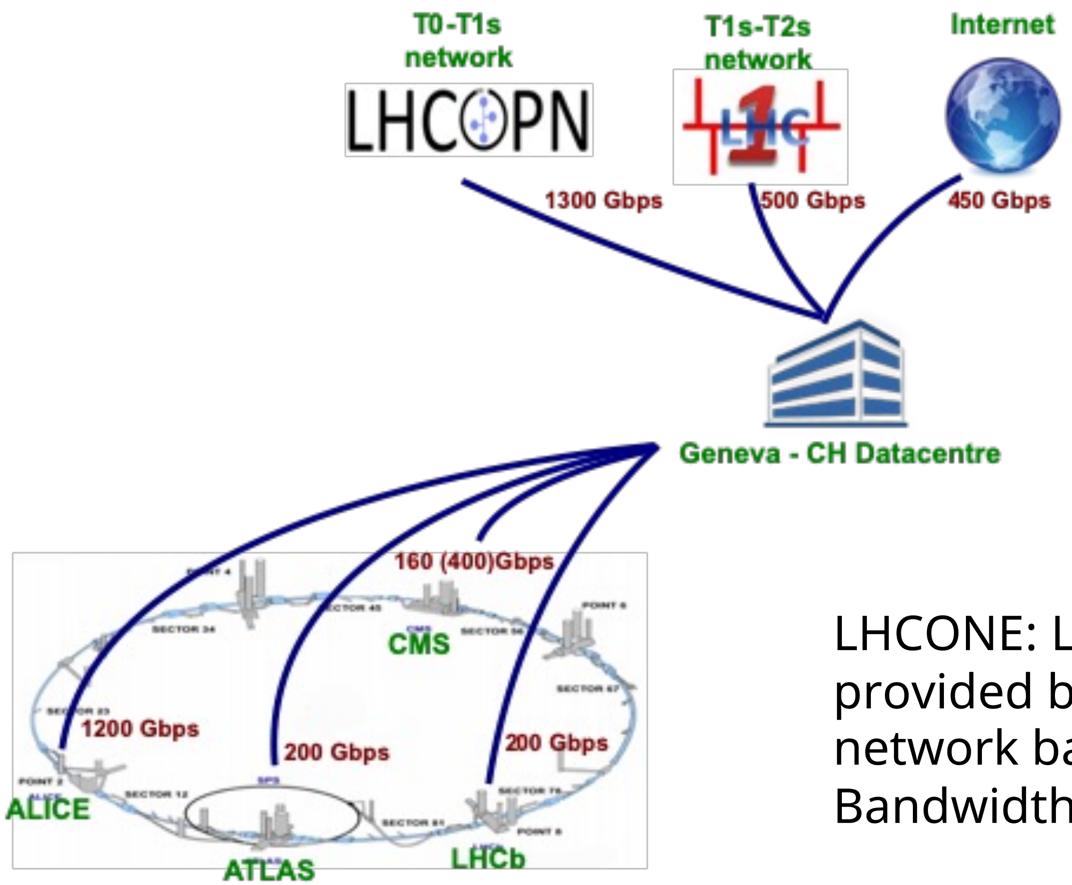

WLCG networks are an opportunity, not a limitation as initially expected

LHCOPN: Private network connecting Tier0 and Tier1s. Dedicated to LHC data transfers and analysis

LHCONE: Layer3 (routed) Virtual Private Network provided by the R&E network providers Worldwide network backbone connecting Tier1s and Tier2s. Bandwidth dedicated to High Energy Physics

WLCG Worldwide LHC Computing Grid

WLCG traffic (GB/s) in the last 10 years

WLCG network used for data transfers but also Monte Carlo simulations

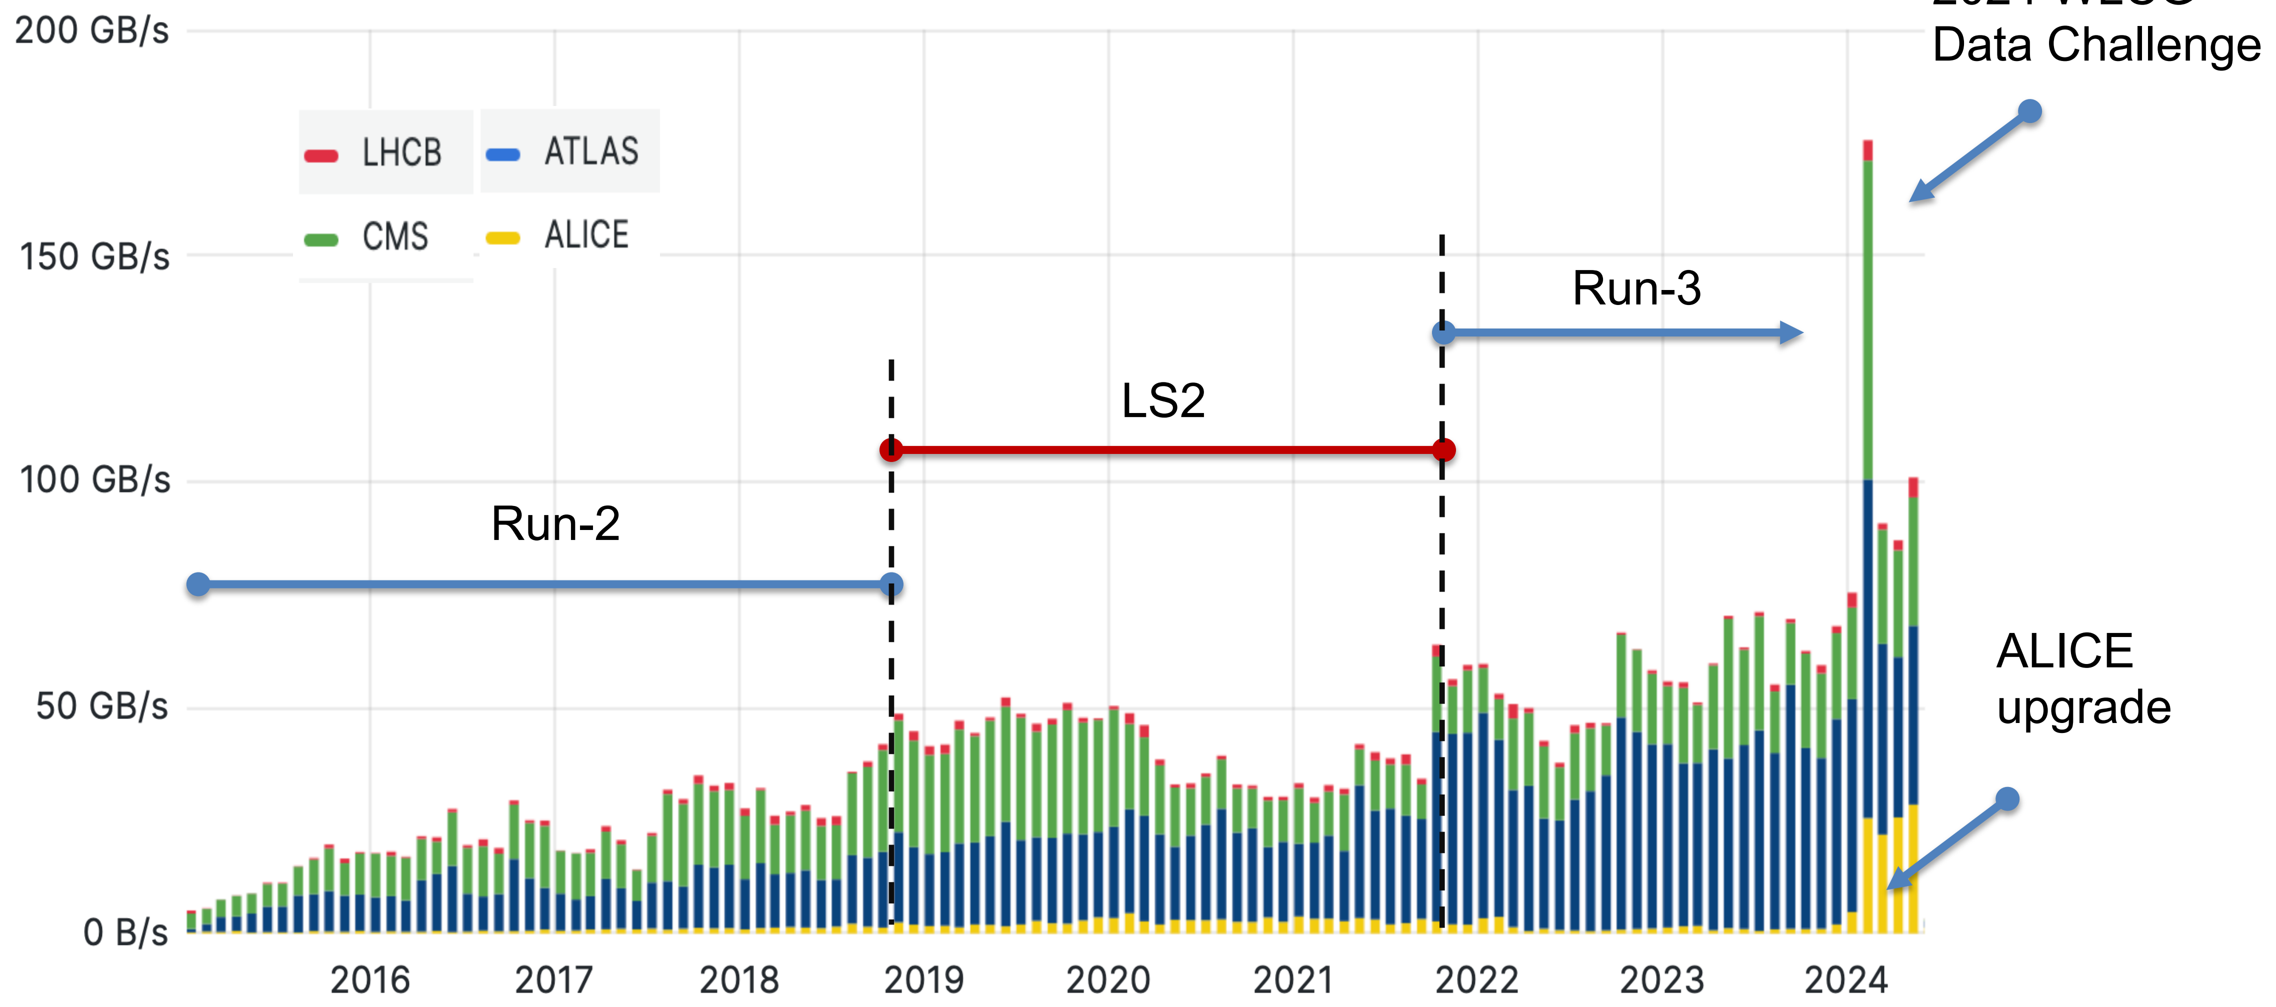

Data processing, simulation, analysis

CPU use in WLCG continuously increases despite shutdown periods

- Simulations require a lot of CPU
- A full analysis requires many years

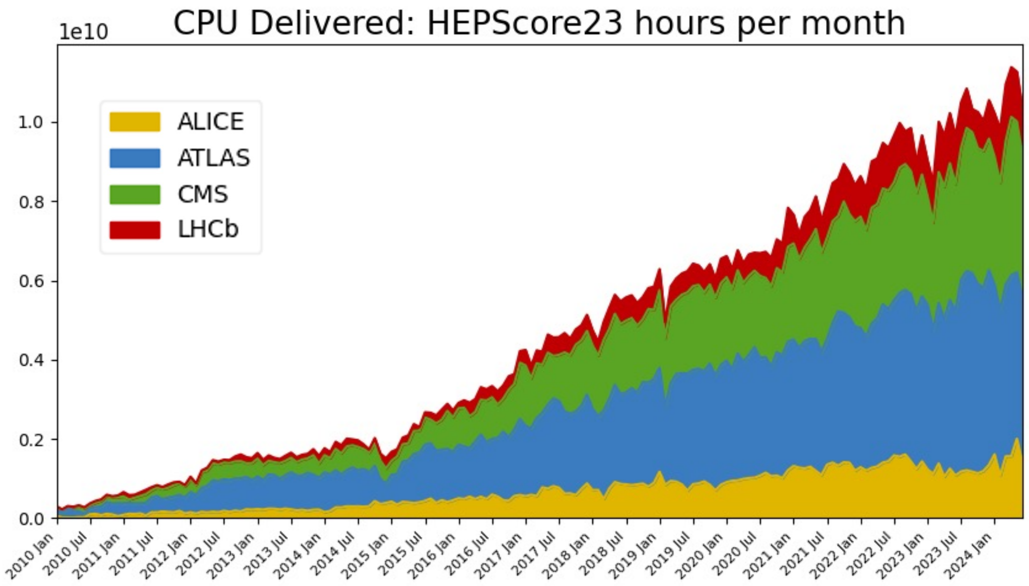

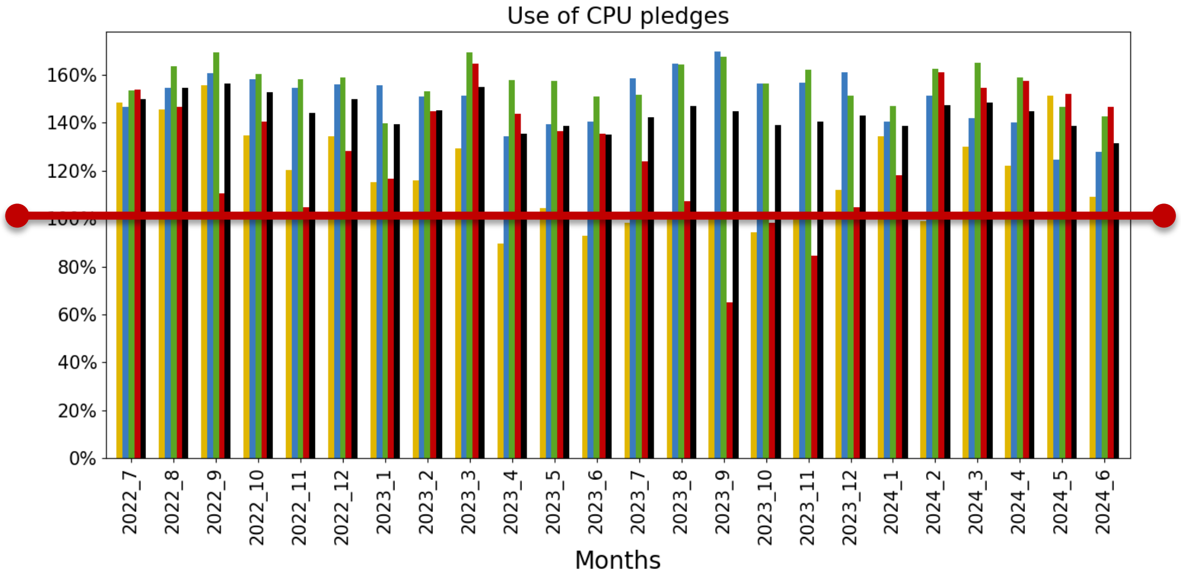

WLCG sites provide ~40% additional capacity on top of their commitments

- The grid model favors opportunistic provisioning of resources

WLCG performance and operations

WLCG Operations and Security Coordination teams ensure the functioning of WLCG services leveraging international initiatives

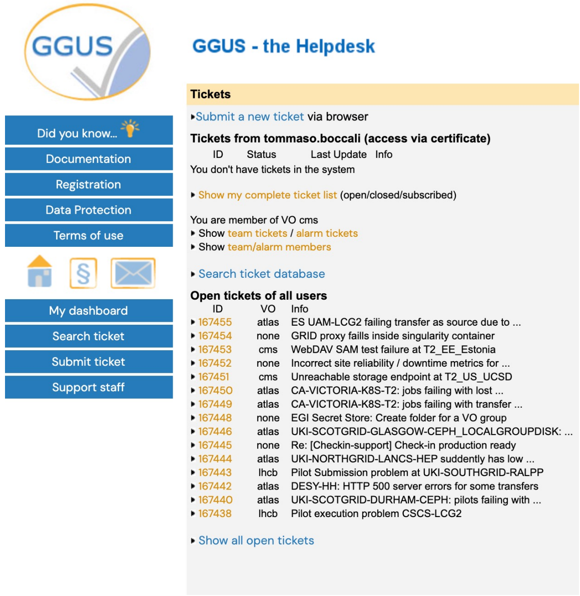

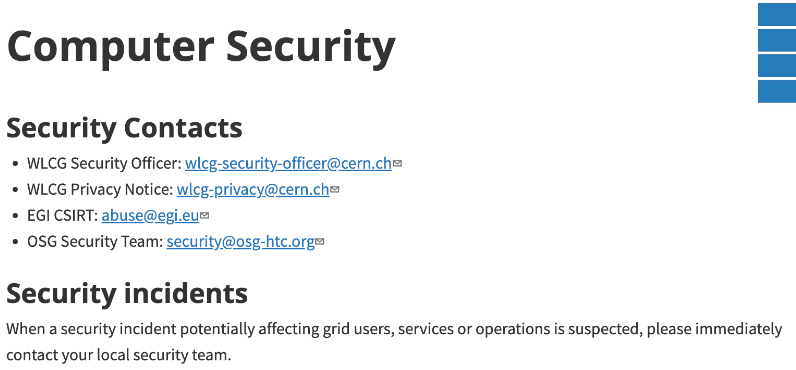

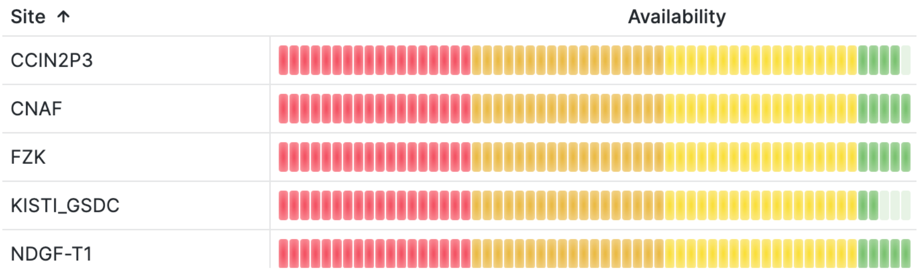

We monitor continuously and review regularly the performance of the sites against MoU targets

Evolution to a more flexible Computing Model

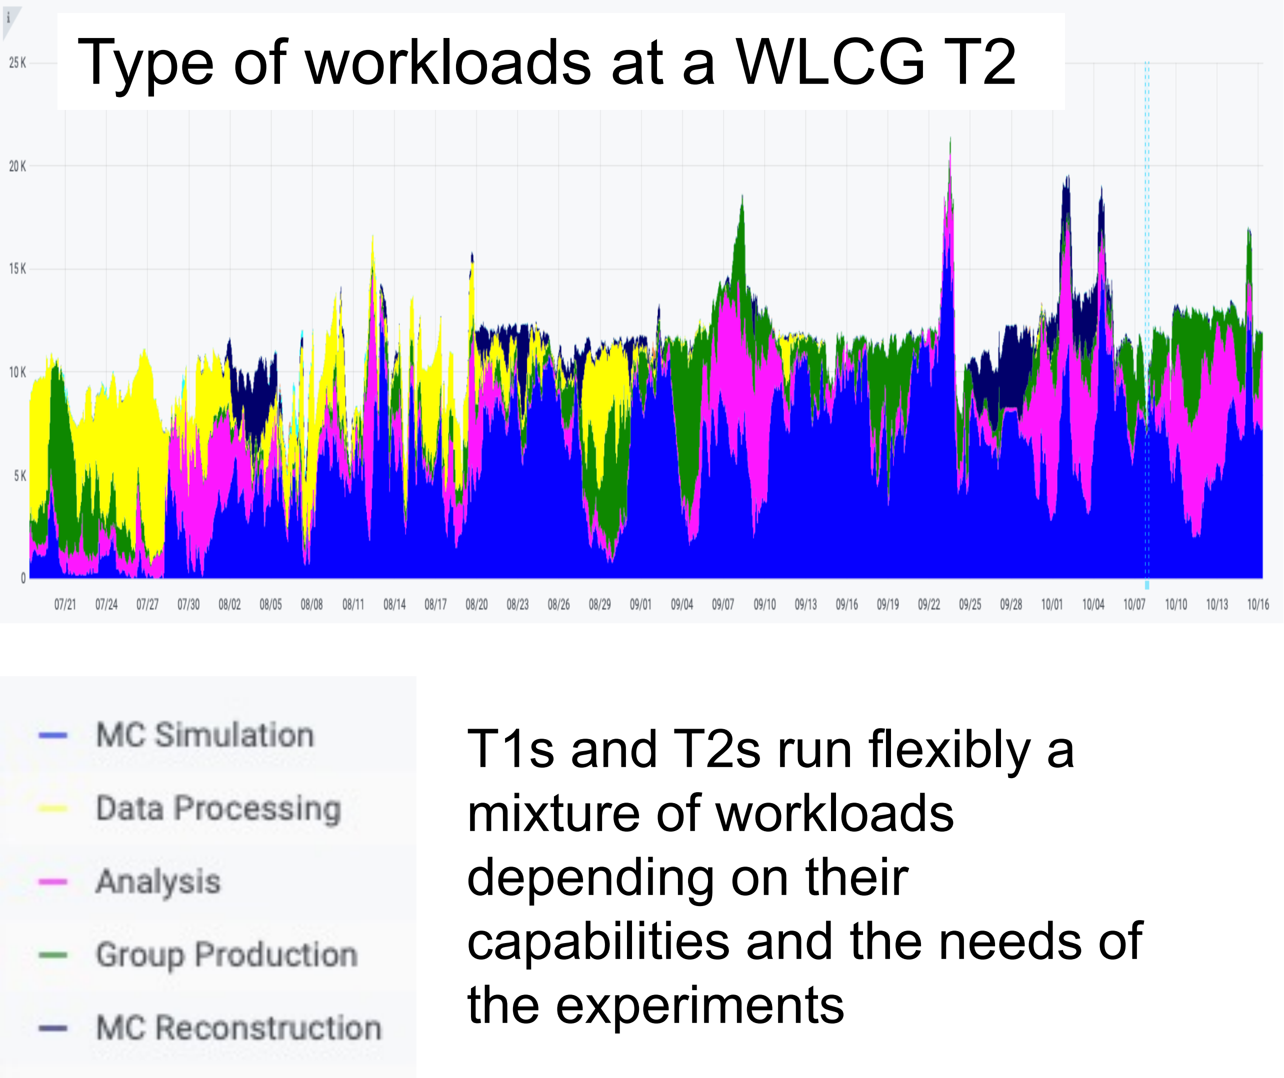

T1s and T2s run flexibly a mixture of workloads depending on their capabilities and the needs of the experiments

Data is transferred in a full mesh rather than hierarchically

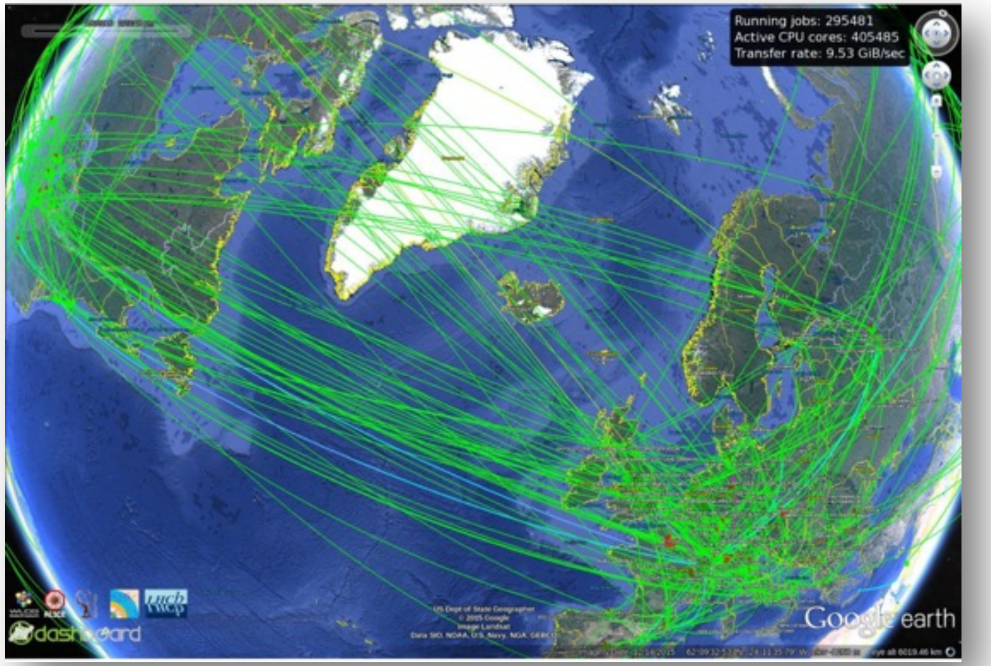

WLCG
Worldwide LHC Computing Grid

Worldwide data processing and data transfers 24/7/365, since 2006, increasing over time. Very high availability and reliability. Solid funding model.

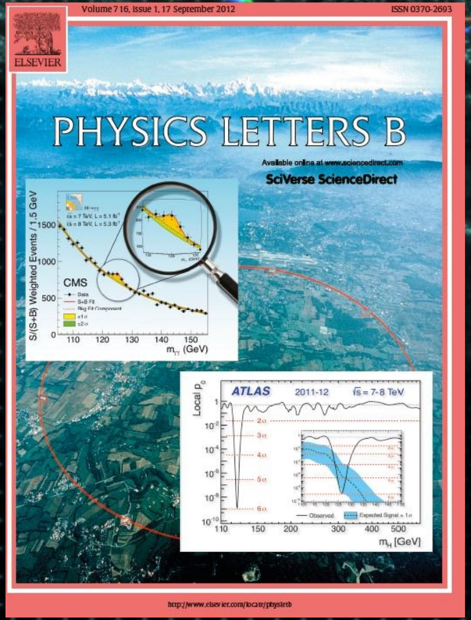

Never a showstopper to physics. Fast turnaround.

It simply works … so what?

WLCG
Worldwide LHC Computing Grid

The HL-LHC challenge – data volume and complexity

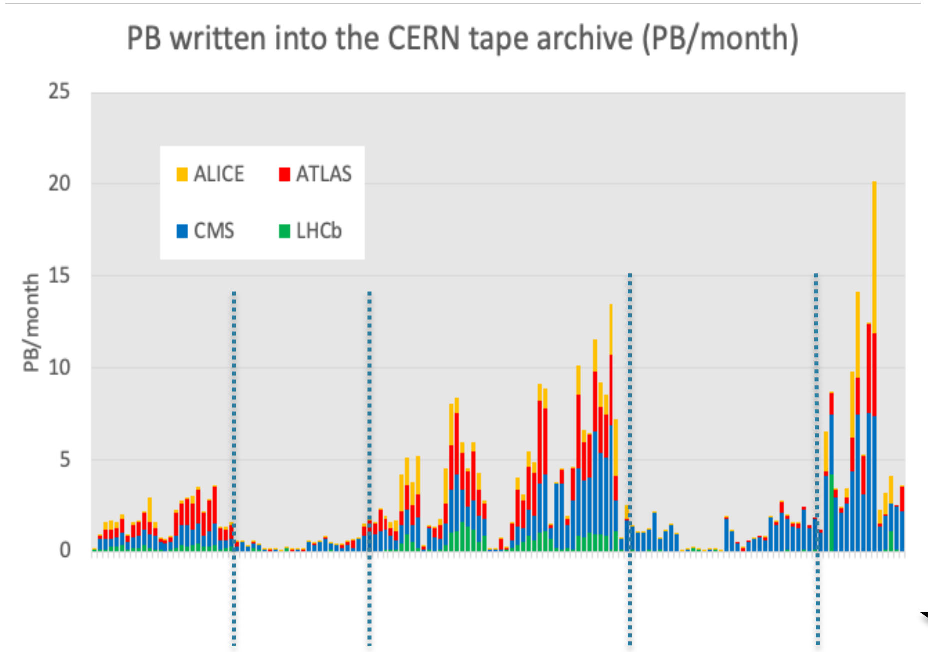

2011 Higgs candidate

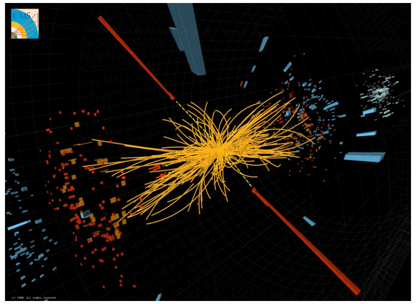

HL-LHC simulated event

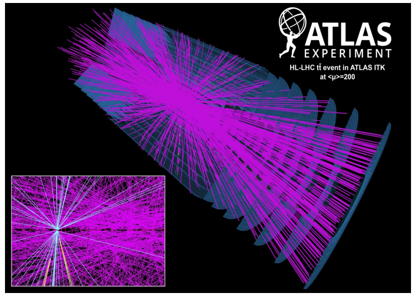

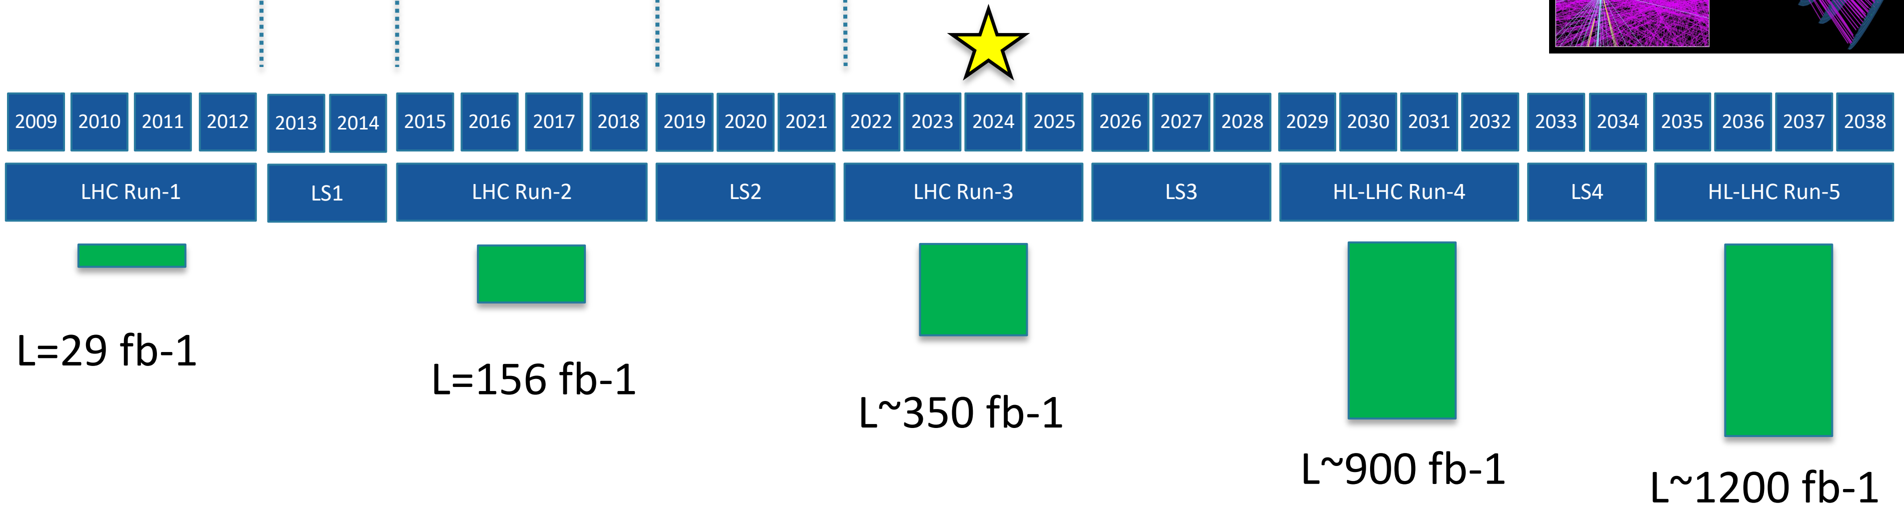

WLCG
Worldwide LHC Computing Grid

Events at HL-LHC

CMS: a very large event from 2017 with 78 reconstructed vertices

ATLAS simulation for HL-LHC with 200 reconstructed vertices

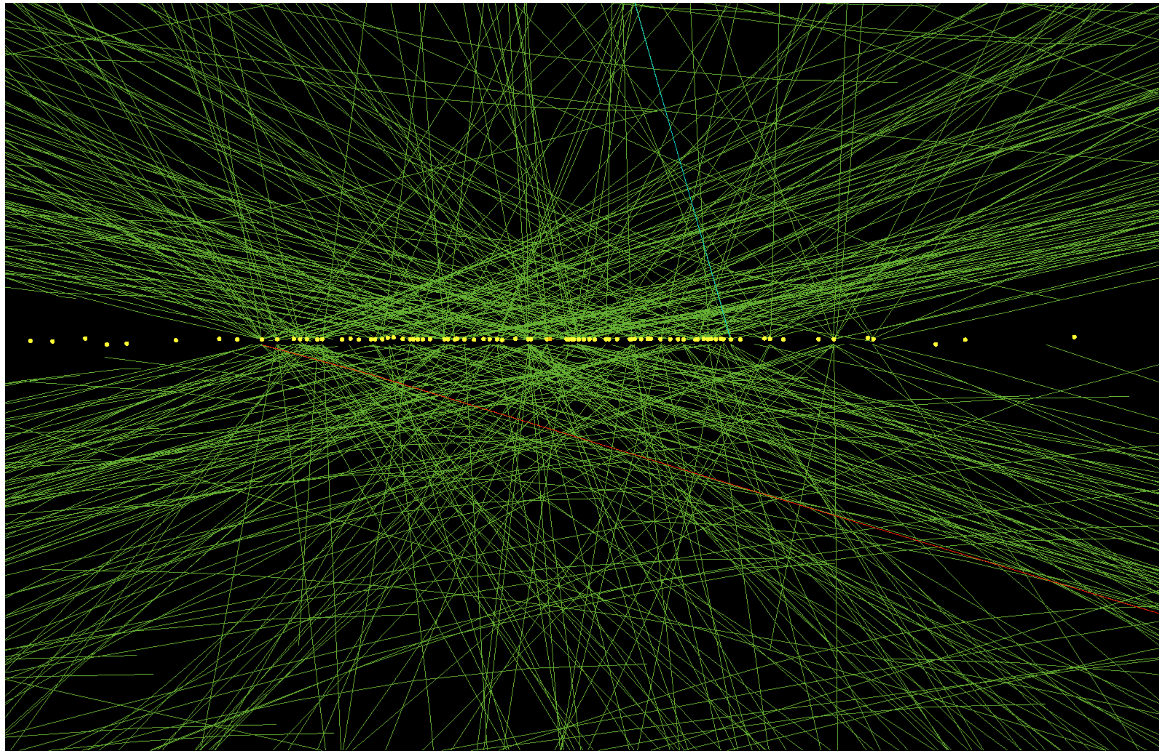

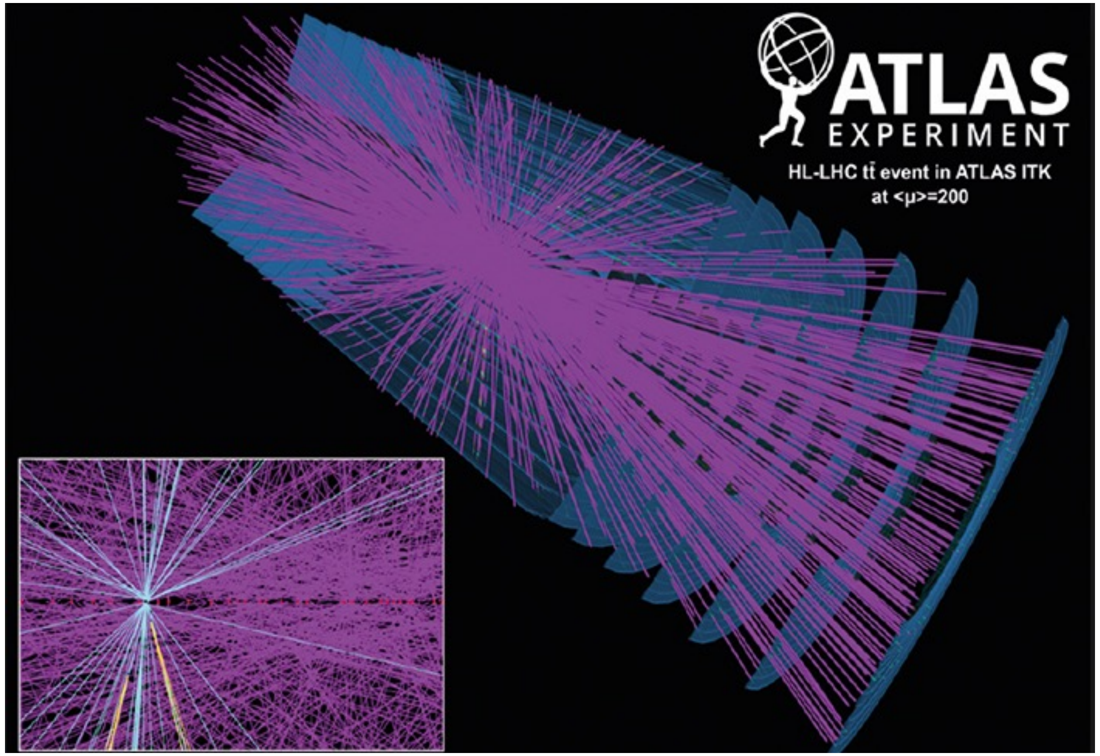

The HL-LHC challenge –Hardware Trends

- Cost of hardware decreasing vs time 😄 … but not as steeply as before 😳
- In general, trends driven by market (revenues) rather than technology

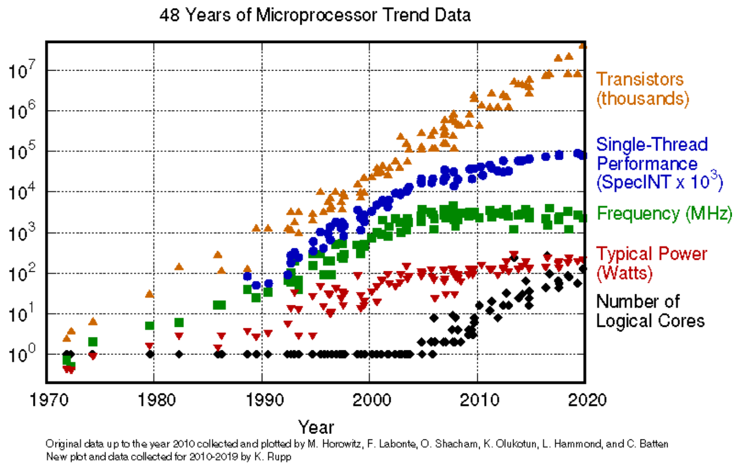

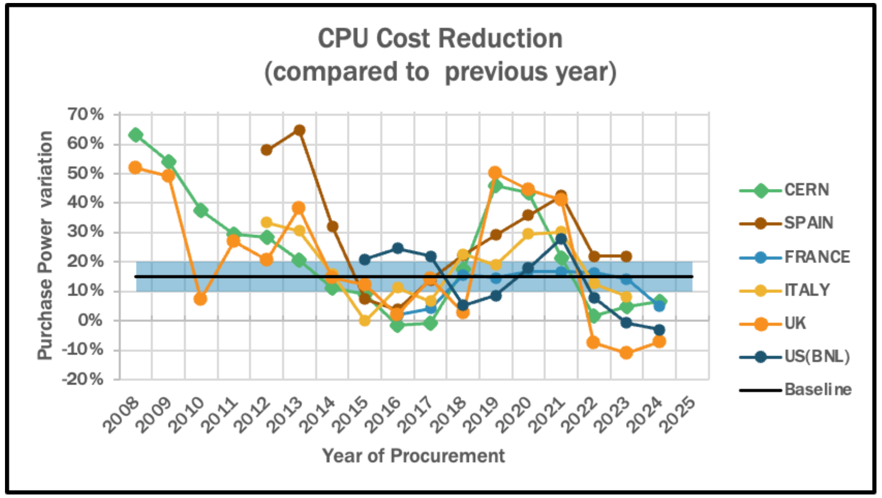

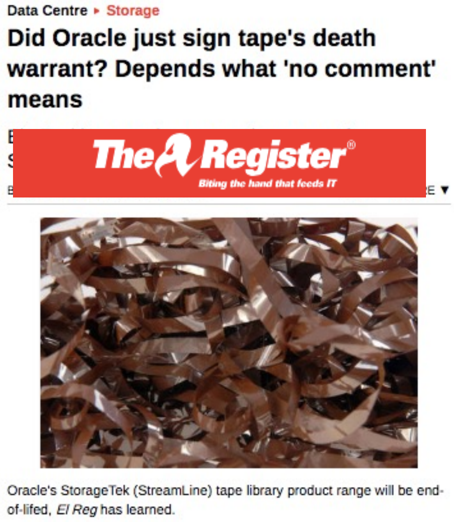

Less benefits from technology improvements + general loss of long-term predictability

ATLAS and CMS computing needs for HL-LHC

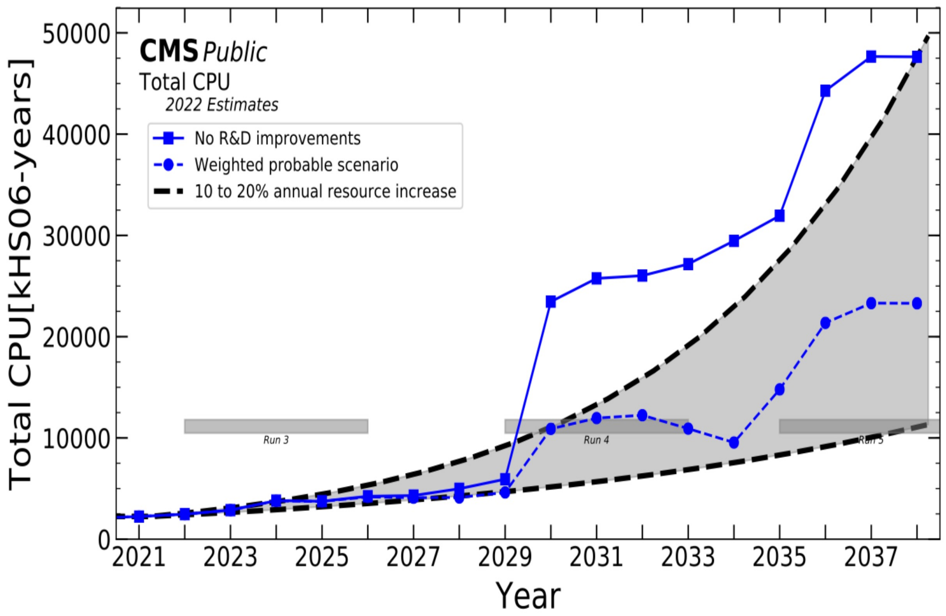

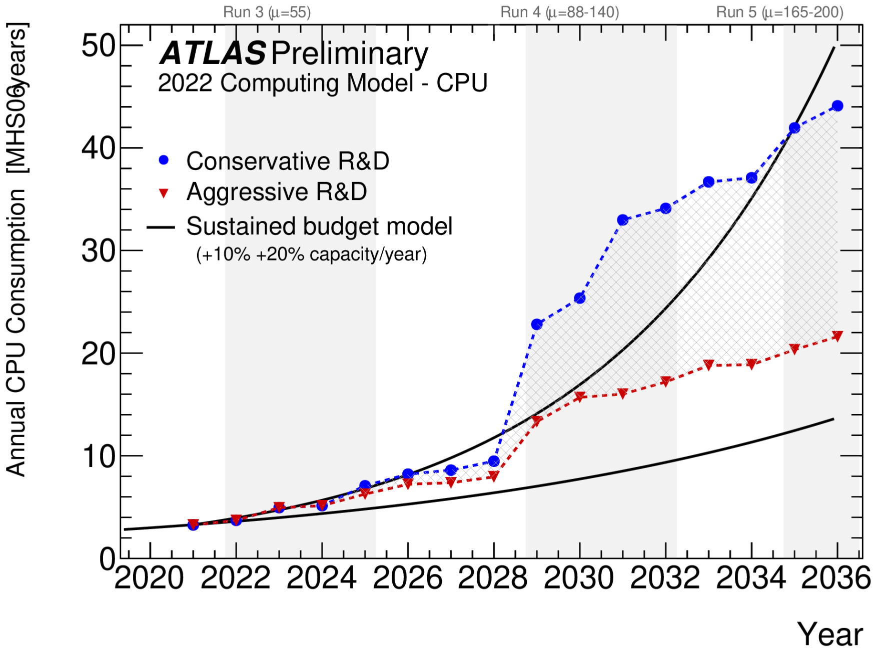

The projected CPU needs of ATLAS and CMS at HL-LHC. Estimates produced for the 2021 HL-LHC computing review and updated in 2022 to reflect the changes in the LHC schedule

WLCG Worldwide LHC Computing Grid

ALICE and LHCb computing needs for HL-LHC

ALICE and LHCb will require considerably less in Run-4, while no firm estimates are available for Run-5. For LHCb, some early thinking shows a challenging scenario in Run-5

Naive extrapolation of the computing model parameters of the current upgrade to the conditions foreseen for Run-5

Today: 10GB/s from online to offline
Run-5: 75GB/s from online to offline

The needs are many factors off the budget for offline

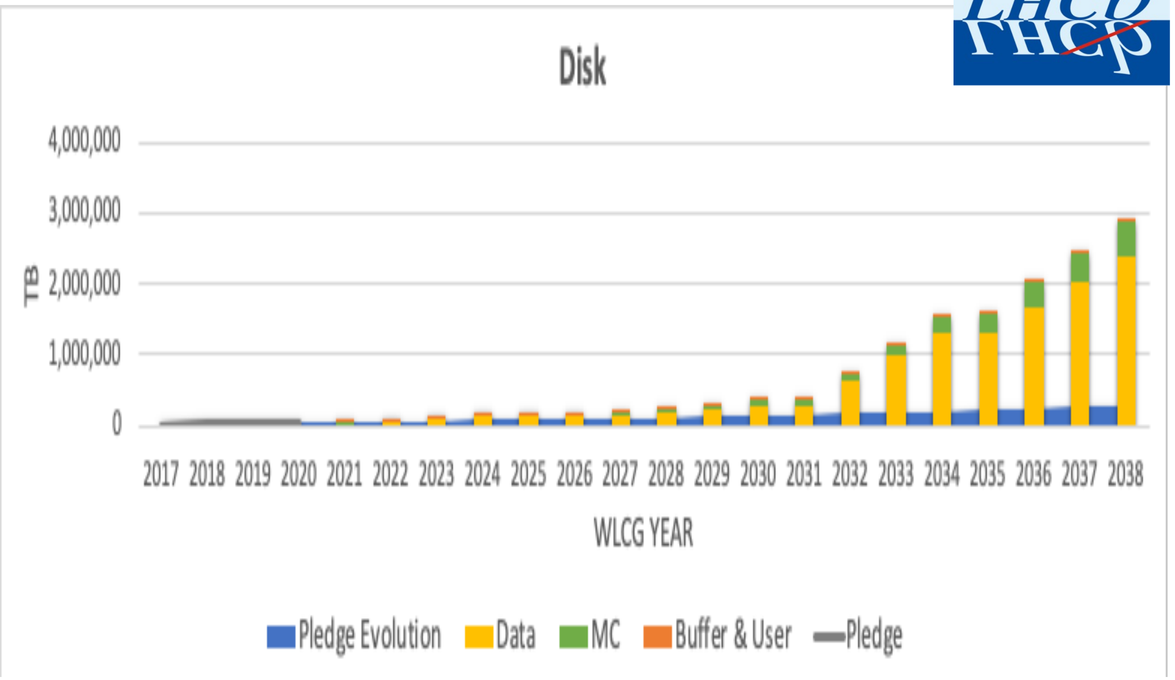

HL-LHC computing needs evolution

2015

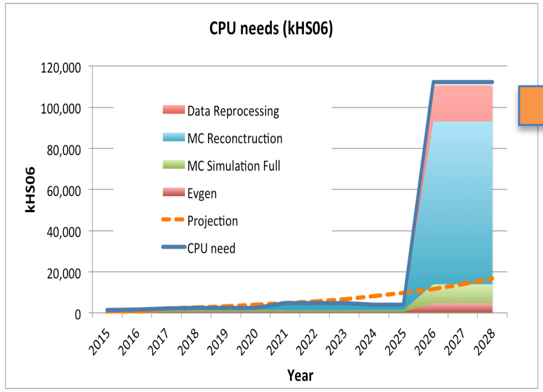

2017

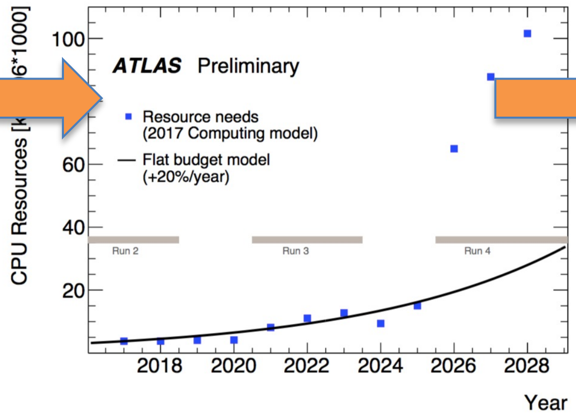

2018

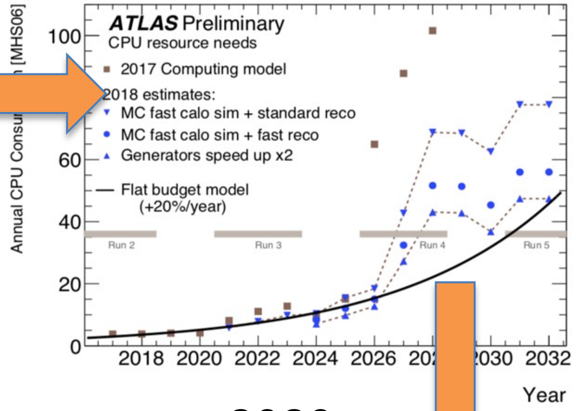

2022

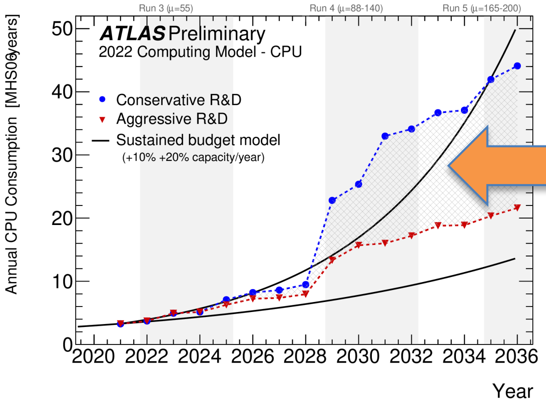

2020

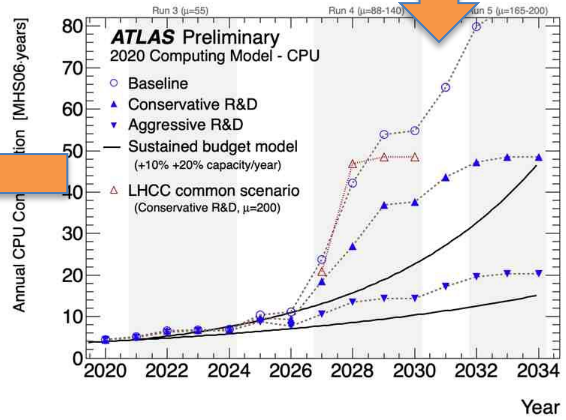

2015 projections: resource needs = 10x more than budget allows

2022 projections: resource needs compatible with budget (optimistic scenario)

WLCG
Worldwide LHC Computing Grid

Not just LHC .. more HEP experiments in the 2020s

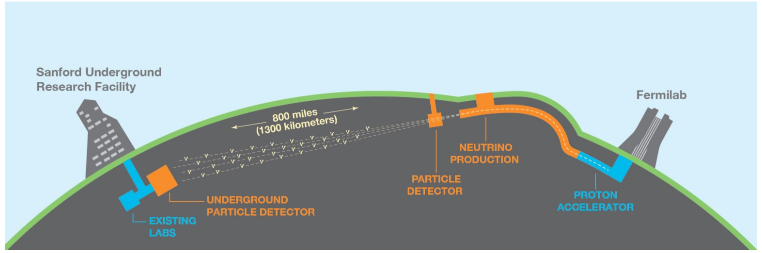

DEEP UNDERGROUND NEUTRINO EXPERIMENT

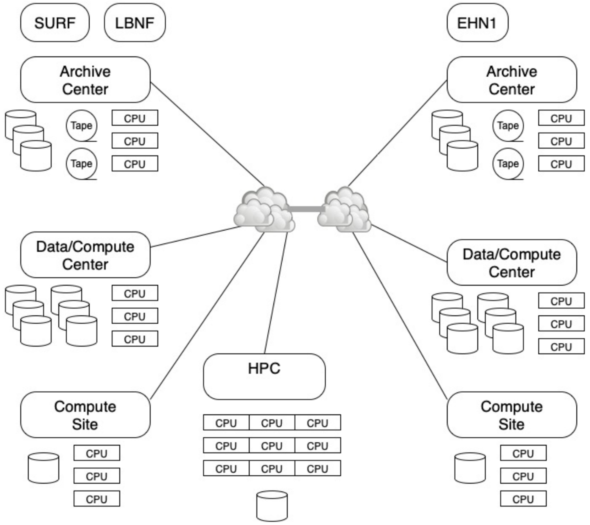

Computing needs:
~10% of HL-LHC

Belle II

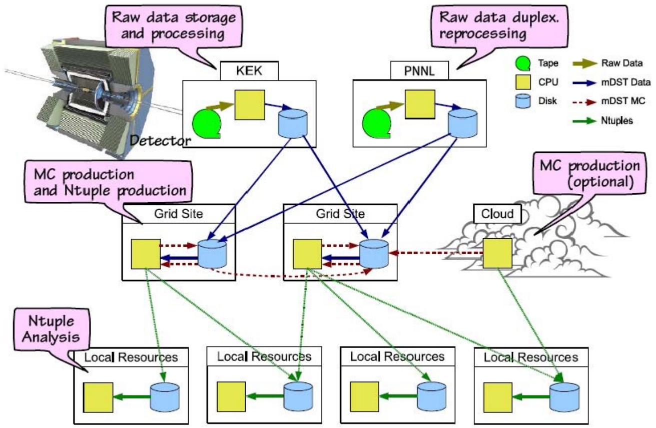

JUNO

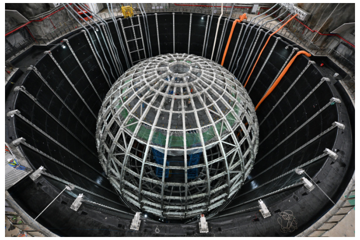

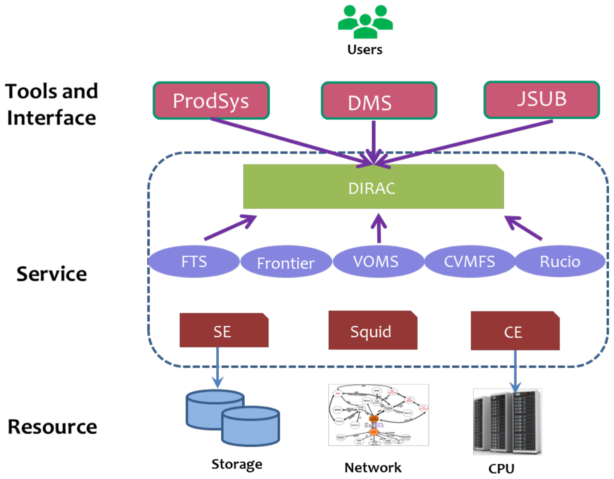

WLCG Worldwide LHC Computing Grid

And more sciences …

Other sciences will have data needs like LHC in the same time horizon.

SKA is the main example

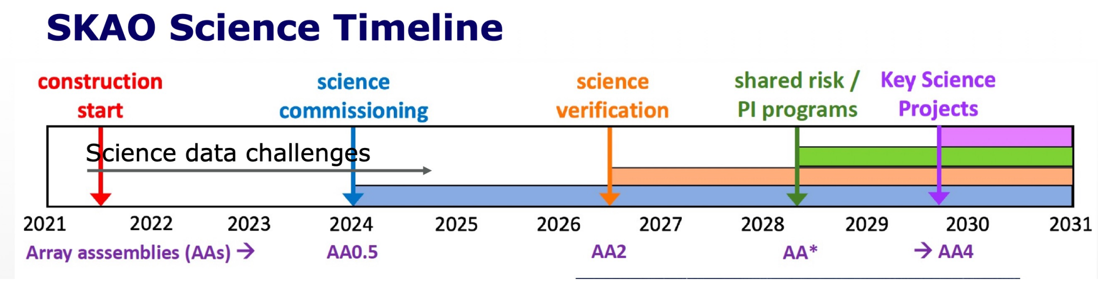

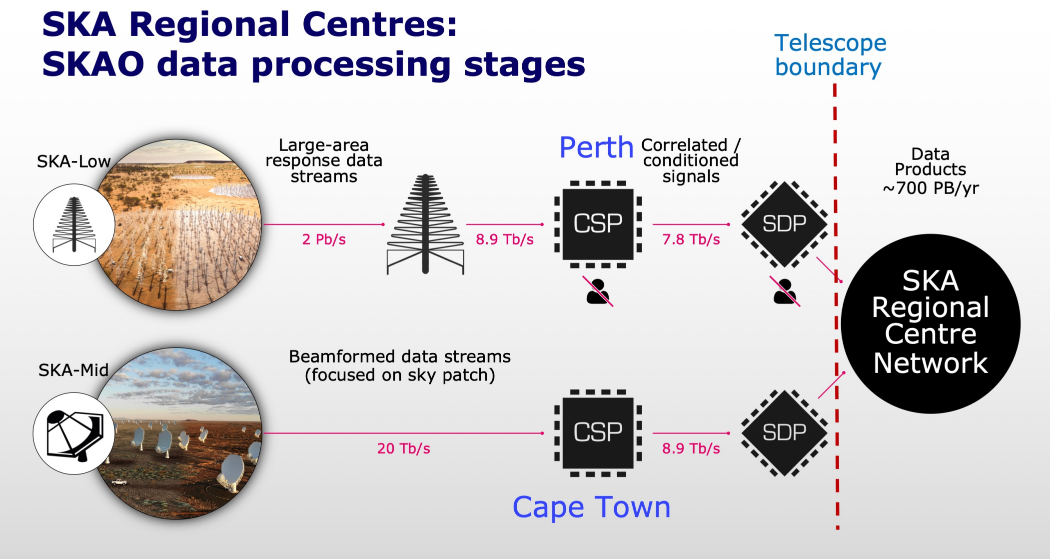

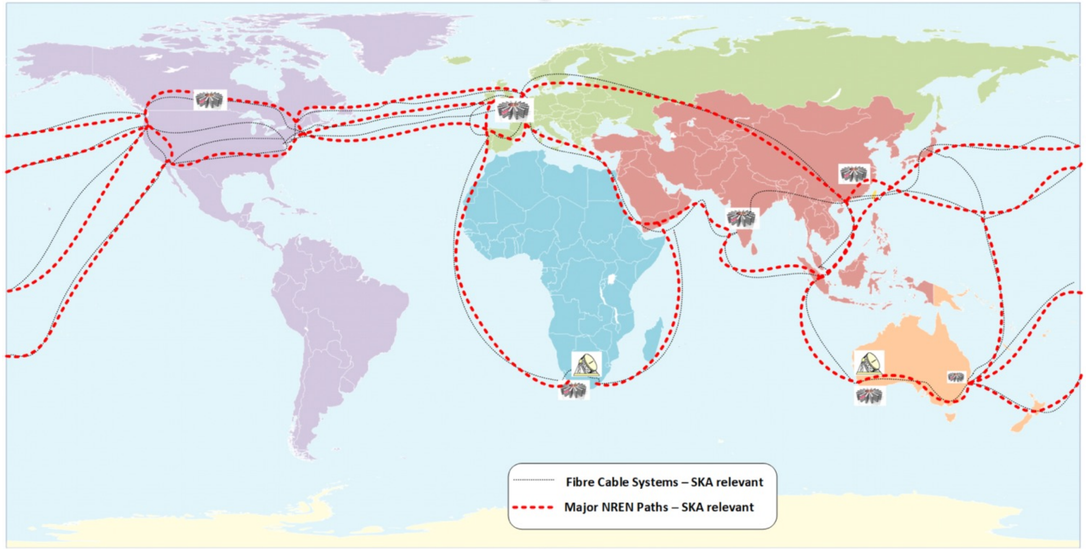

WLCG Worldwide LHC Computing Grid

The future HEP computing challenge in 6 points

1) 10x more data/year and 5x more complex. HL-LHC computing is the main driver

2) No considerable increase in funding for HEP computing expected in the coming years

3) Trends in hardware technology and costs not favorable to HEP

4) Geopolitical situation very unstable, impacting a distributed infrastructure and collaboration

5) Many HEP computing services were conceived in a different millennium and the Information and Communication Technology landscape changed considerably

6) The Funding Agencies want to see their investments to benefit multiple sciences

WLCG strategic vision: Innovation and Collaboration

WLCG presented:

- its vision about a common scientific computing infrastructure at the European Strategy for Particle Physics in 2019
- a joint paper with DUNE and Belle-2 to the Snowmass 2021 process, which detailed the strategic directions to address the computing challenges of the experiments over the next decade

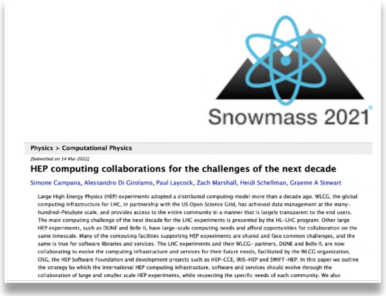

European Strategy Update

The WLCG strategy for 2024-2027 has a strong focus on innovation and collaboration

- Innovation: modernise software and services to leverage the most modern technologies and architectures
- Collaboration: leverage synergies between HEP experiments and other sciences

Strike the right balance between common/custom, standard/specific

WLCG Worldwide LHC Computing Grid

Leveraging Heterogeneous Facilities - HPCs

HPCs have been integrated as WLCG resources since ~ 20 years

Some HPC are very challenging to use: limited connectivity, special harware (GPUs, non-X86 CPUs) and configs (RAM, local disk), different AAI. The benefits are considerable in many cases.

The benefits are considerable in many cases. Attempts to move the collaboration with HPC centers into consolidated partnerships

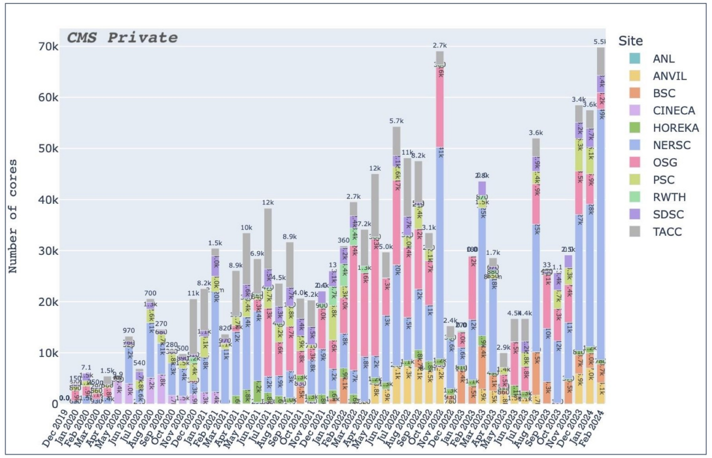

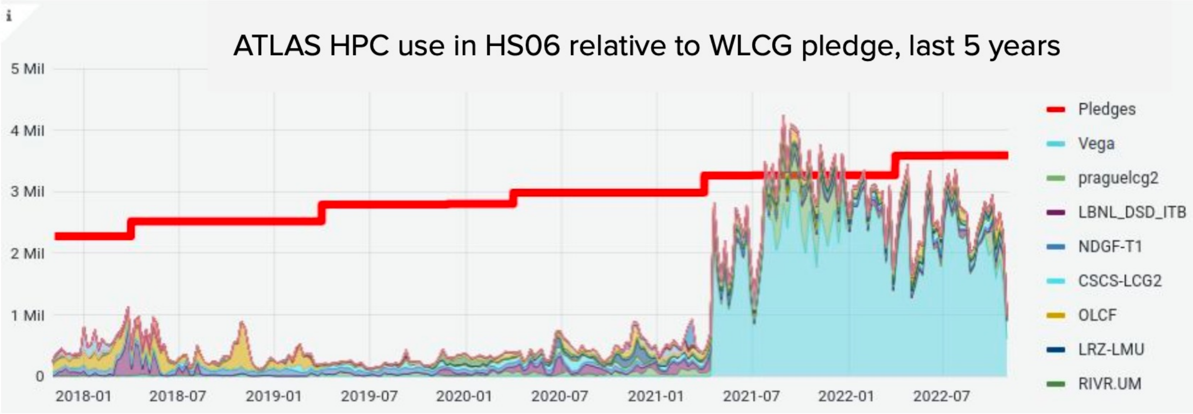

Leveraging Heterogeneous Facilities - Clouds

Cloud technologies were integrated in WLCG in ~2012. Academic clouds in use since then

Many initiatives integrating commercial clouds. Large variation in cost depending on choices being made

Commercial cloud TCO vs on premise at WLCG facilities is not obvious: see a recent and some previous studies

Not generally financially advantageous for large scale productions. Many benefits in flexibility (capacity, resource type).

WLCG needs to continue being prepared using commercial cloud resources from different vendors

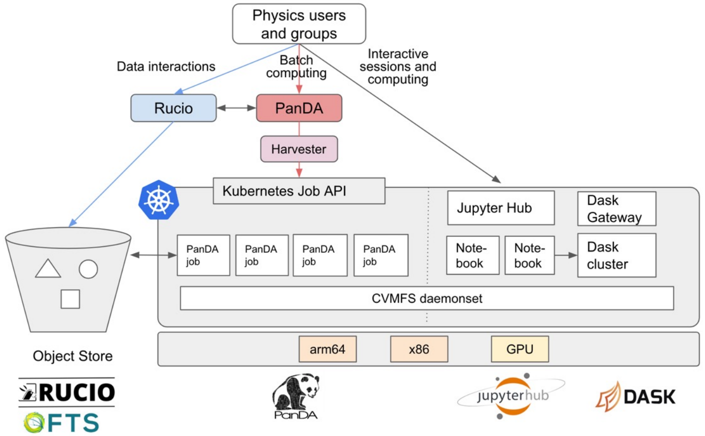

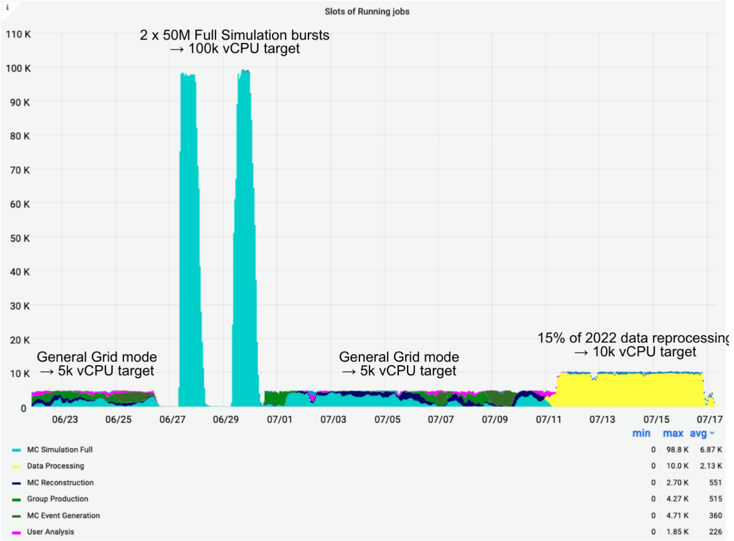

WLCG
Worldwide LHC Computing Grid

Software Performance and Portability

Software performance, modernization, innovation and portability is one of the key areas to address the future HEP computing challenge. It allows to:

- access available resources otherwise not useful (e.g GPUs, non-X86 CPUs)
- optimize procurement and deployment to guarantee the best value for money
- support different experiments and sciences on the same resources
- leverage the most modern hardware features (e.g vectorization)
- reduce the resource needs/event (compute and storage)

It needs upfront large investments and forward-looking solutions

- E.g. LHC experiments have a 20+ years codebase and rewriting from scratch is not an option

Table 6. SLOCCount measured lines of source code for ATLAS and CMS.

Experiment Type	Source Lines of code (SLOC)	Development effort (person-years)	Total estimated cost to develop
ATLAS	5.5M	1630	220 M$
CMS	4.8M	1490	200 M$

Linux Kernel is: 15M sloc, 4800 FTEy, 650M$ (3x CMS);

Geant4 is: 1.2M sloc, 330 FTEy, 45 M$ (1/4x CMS)

WLCG Worldwide LHC Computing Grid

Non-X86 CPUs

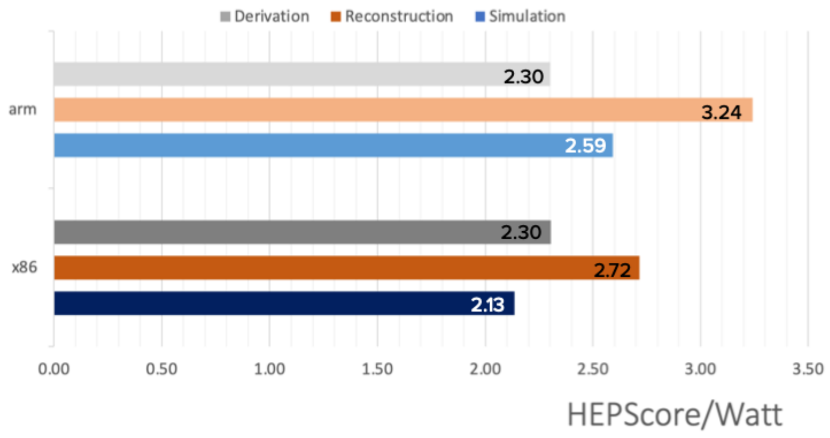

ARM CPUs process more events/Watt wrt X86

- This is the case for the HW in Glasgow Tier-2
- Differences between workflows
- Modern AMDs perform similarly to ARM - again based on the HW in Glasgow Tier-2

See this presentation for details

All experiments ported main workflows to ARM

- ATLAS, ALICE: physics validation successful
- CMS, LHCb: in progress

ARM resources available at various WLCG sites

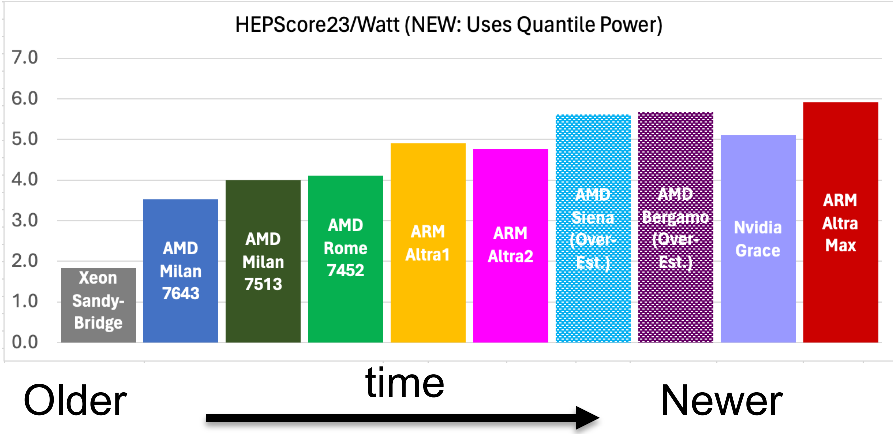

Heterogeneous architectures in reconstruction

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

Playing a fundamental role in Run-3 already, in most online systems.

LHCb: exploitation of heterogeneous architectures in online system, thanks to Allen framework

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

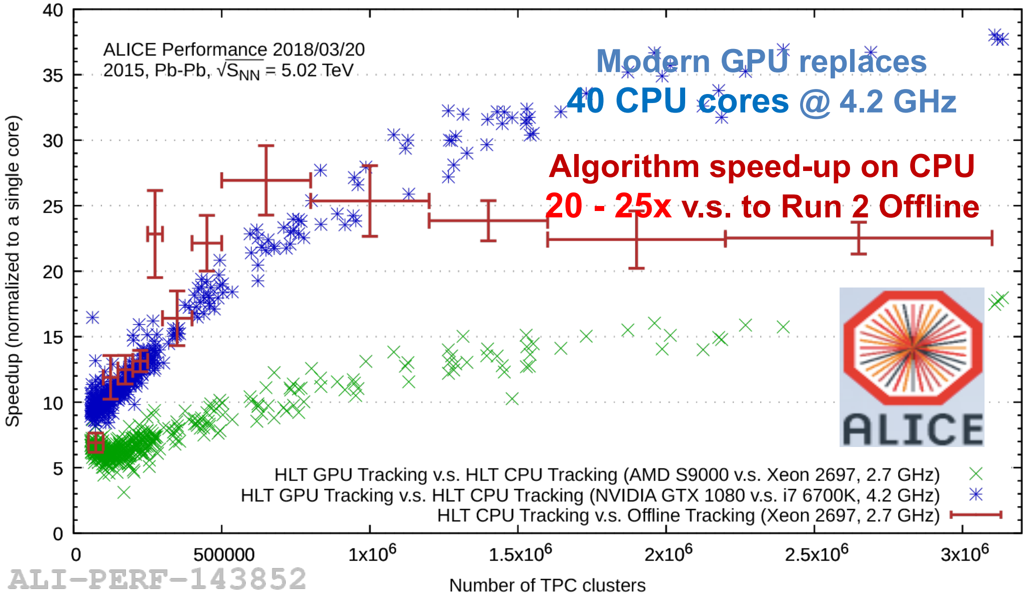

The utilization of AI is not new in HEP. What is game changing is:

- the availability of modern models (Transformers, LLMs, Normalizing Flows, …)
- the availability of large (optimized) processing systems

AI expected to become a major tool by HL-LHC

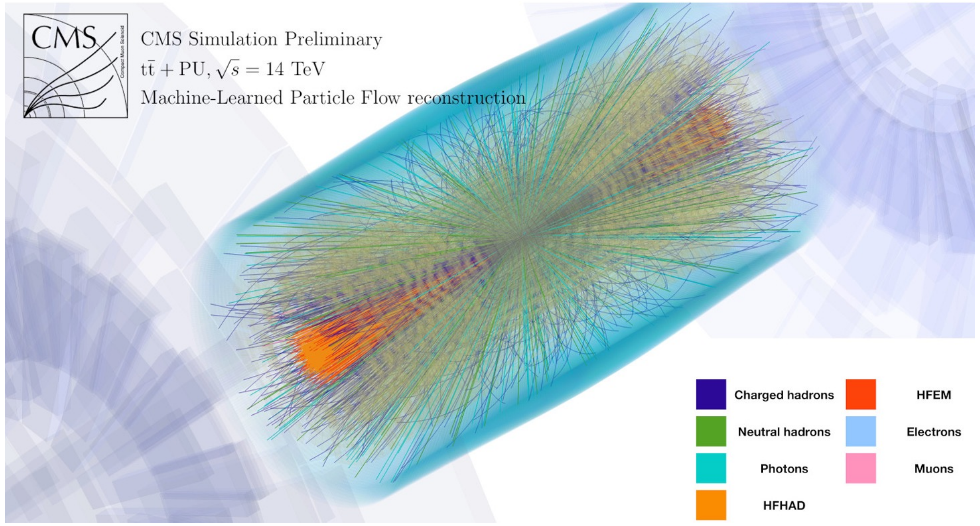

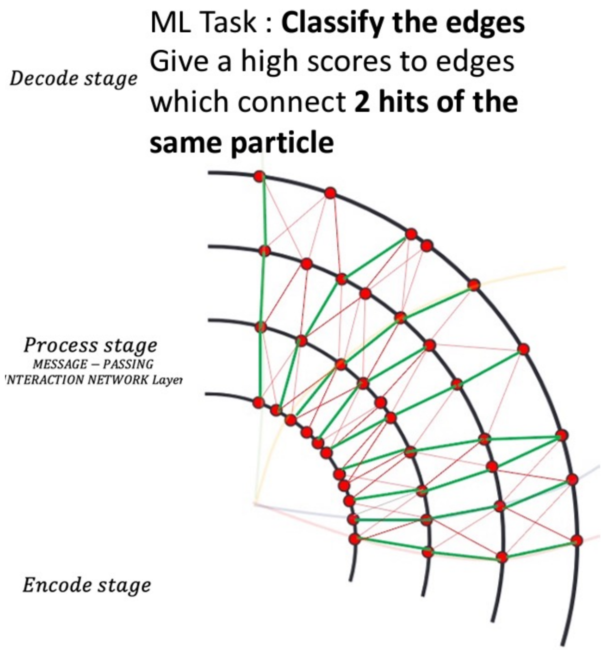

Software-related activities and project

Software performance, modernization, innovation and portability is one of the key areas in order to address the HL-LHC challenge. It needs upfront large investments

The HEP Software Foundation: facilitates the cooperation and common efforts in HEP software and computing

The Institute for Research and Innovation in Software for High Energy Physics (IRIS-HEP): R&D for the software for acquiring, managing, processing and analyzing HL-LHC data.

CERN Openlab: a public-private partnership that works to accelerate the development of cutting-edge ICT solutions for the worldwide LHC community

WLCG Worldwide LHC Computing Grid

Custom or common or ... ?

We want to use as much commodity software products as possible, to minimize cost and manpower:

- In some cases, using commercial libraries, e.g. for state-of-the-art Machine Learning (pyTorch, TensorFlow etc.) is a straightforward choice for using ML tools in HEP.

...but there is lot still specific to HEP (and/or science):

- specific tools for simulation (Geant4), data persistency and analysis (ROOT),
- Experiment and WLCG services for distributed computing

... and these all need long-term continuity of support ! To share the load and optimize resources, a lot of effort has been made (e.g. HSF)

User Analysis

What should also not be neglected are the ways ’users’ want to use the computing resources and software environment:

- Yesterday: program via ROOT/C++macros
- Today: jupyter notebooks/browser frontend is a desired and even an expected approach

Analysis Facilities: infrastructures tailored to provide an analysis ecosystem - tools and services

- Possibly also centers for large scale AI training

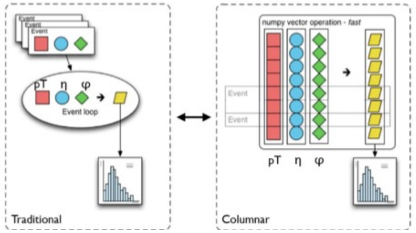

Coffea Analysis Framework

ROOT RDataFrame

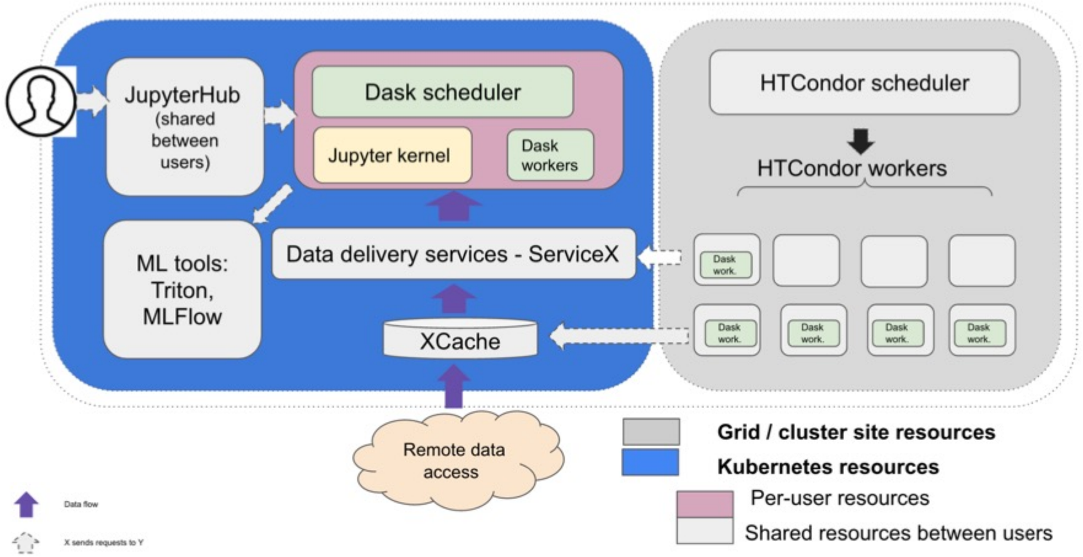

WLCG Worldwide LHC Computing Grid

WLCG partners and collaborators

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

DUNE, Belle-2, JUNO and VIRGO are now WLCG **partners**

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

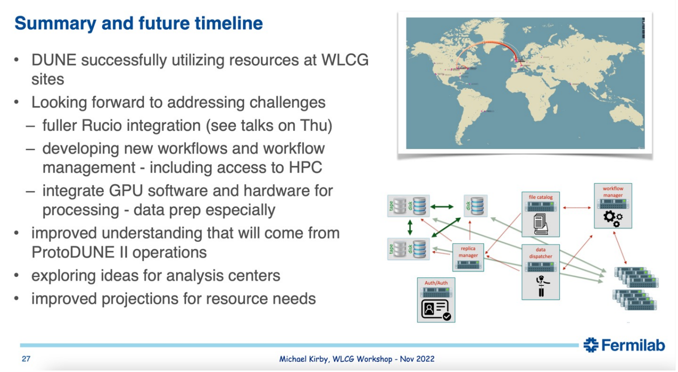

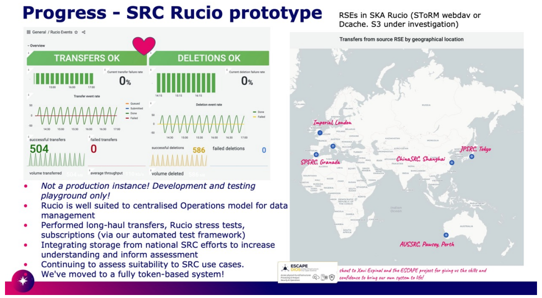

Snowmass 2021

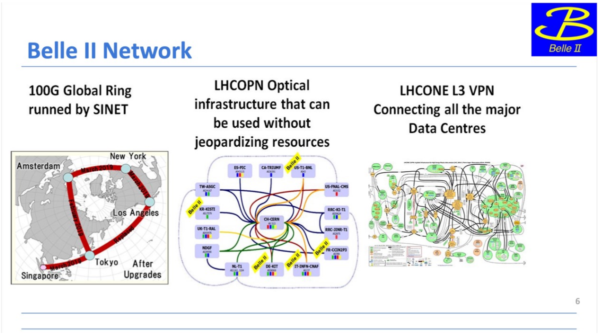

Fostering this collaboration is essential for the sustainability of the WLCG infrastructure: optimise investments and effort

WLCG
Worldwide LHC Computing Grid

Collaboration across sciences

https://zenodo.org/record/4889503

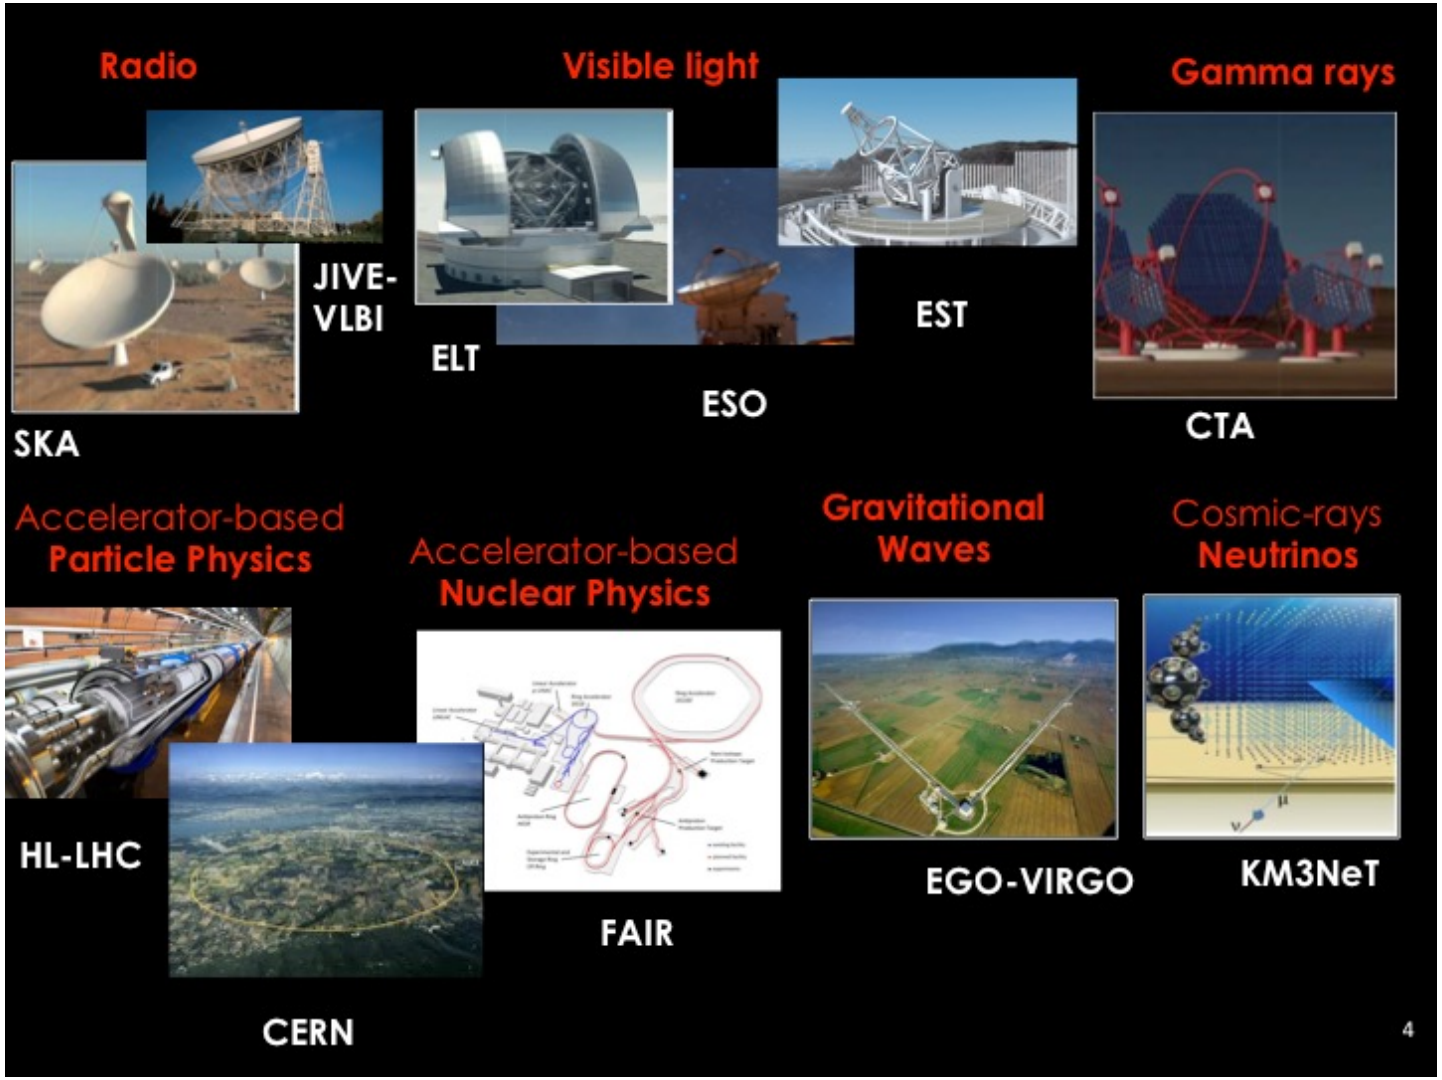

Scientific Data Management

Scientific Data Management is a clear example of successful collaboration in HEP and beyond

Community experiences

- Rucio has become the de-facto standard for open scientific data management
  - Used by CERN-based experiments AMS, ATLAS, CMS
  - And non-CERN experiments Belle II, SKA, CTAO, LBNF/DUNE, SBN/ICARUS, KIS Solar, LIGO/VIRGO/KAGRA, Vera Rubin Observatory, XENON, …
  - Under evaluation by many others EIC/ePIC, KM3NeT, …
  - Long tail of science Fermilab, RAL, CERN run Rucio as a service for small experiments
  - Used by several EU projects ESCAPE, InterTwin, DaFab

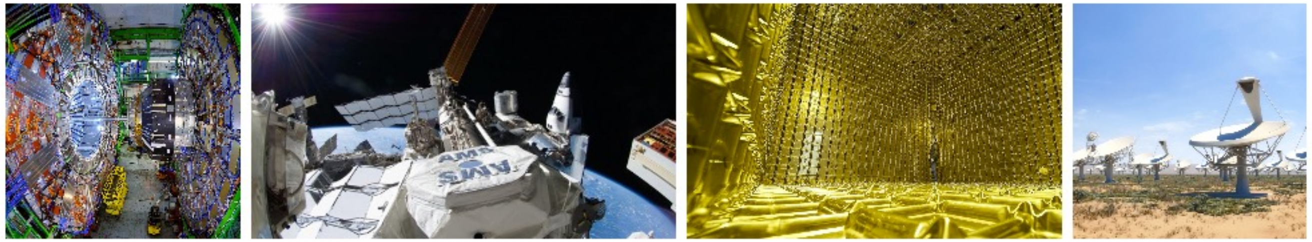

2024-03-14 Rucio @ CERN Storage Day - M. Barisits 9

SCIENTIFIC DATA MANAGEMENT RUCIO

Rucio: initially developed for/by the LHC ATLAS experiment

Now a standard de-facto in HEP and beyond

WLCG Worldwide LHC Computing Grid

Scientific Data Management evolution

The WLCG DOMA - Data Organization, Management, Access - initiative drives the evolution of Data Management tools, services and infrastructure for the needs of the experiments at HL-LHC

- Modernize the data management services to leverage modern, non-HEP-specific technologies
- Commission the WLCG data management infrastructure for the HL-LHC scale

Do you something about SRM and GridFTP?

How about I tell you HTTP?

Ever used X509?

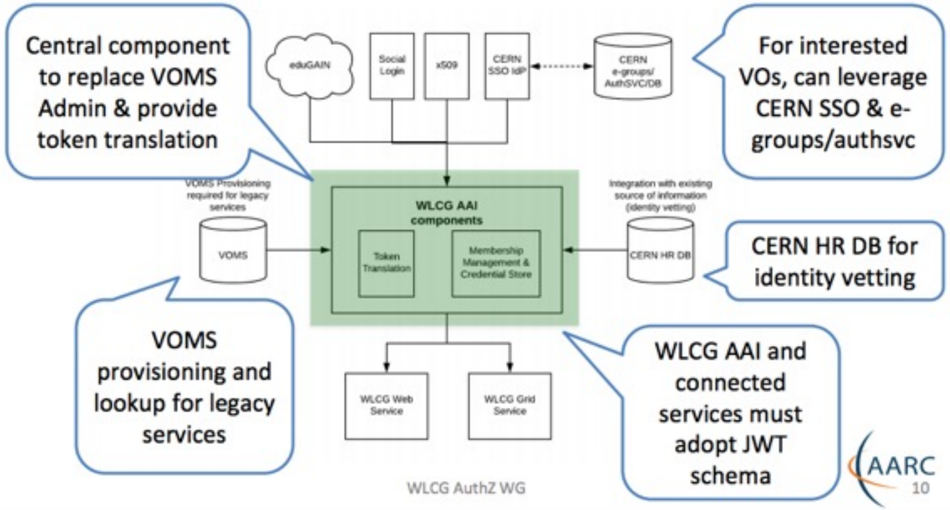

WLCG Worldwide LHC Computing Grid

WLCG Data Challenges program

The WLCG Data Challenges program was initiated to

- Increasingly commission the WLCG data management infrastructure to the HL-LHC scale
- Progressively evolve the service technology and introduce innovative solutions

Started in 2021, run every 2 to 3 years. DC21 (10% of HL-LHC) lessons documented here.

DC24 had 3 goals:

- Measure the end-to-end data transfer capabilities at WLCG sites (target is 25% of HL-LHC needs)
- Assess the progress integrating new technologies (e.g. tokens and monitoring)
- Assess the status of different R&D initiatives

DC24: from Feb 12 to Feb 23 in 2024

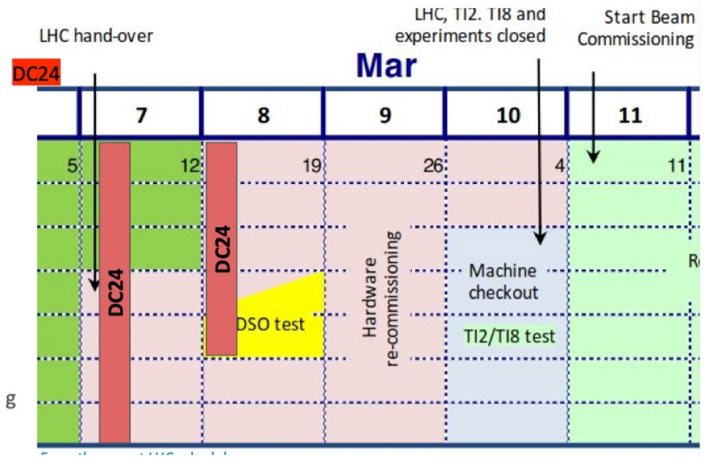

Data Challenge 2024 - Highlights

DC24 WLCG data transfers (Gbps) – 15 days: **all targets achieved**

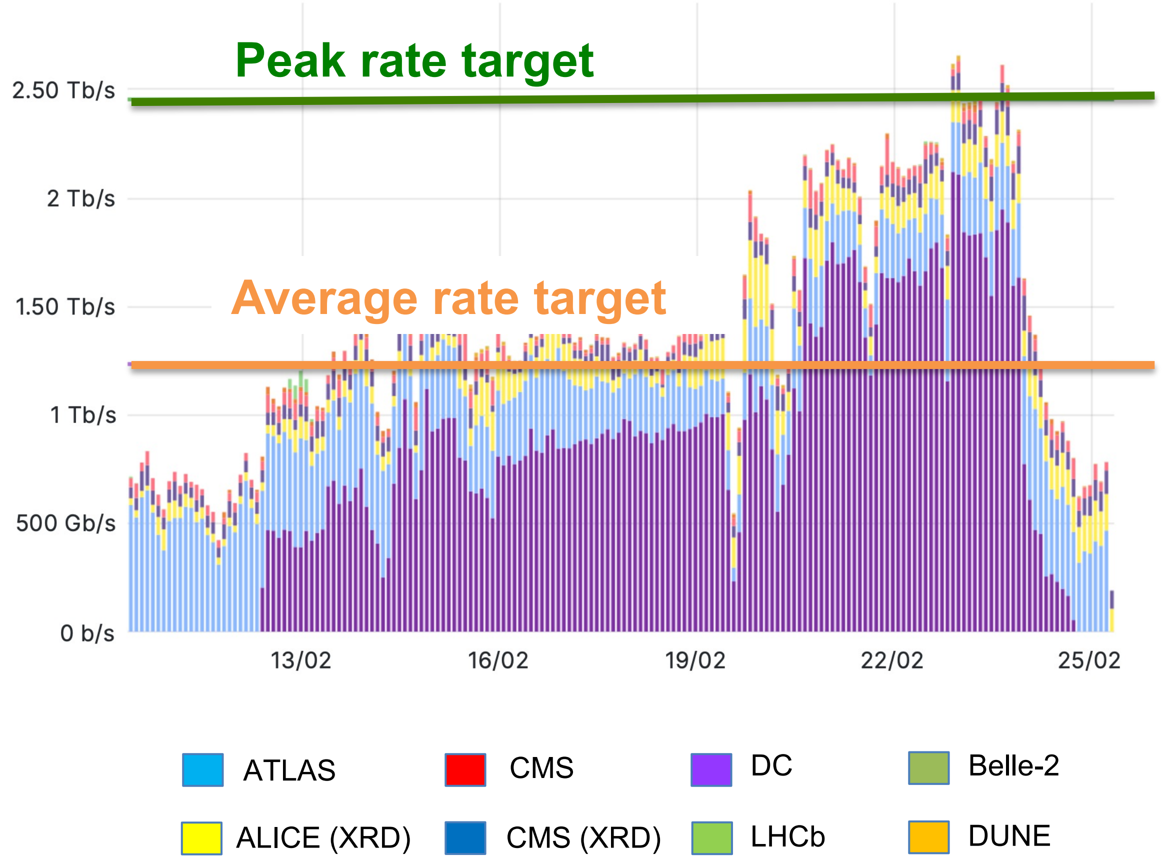

New technologies (e.g. authentication tokens) introduced and validated

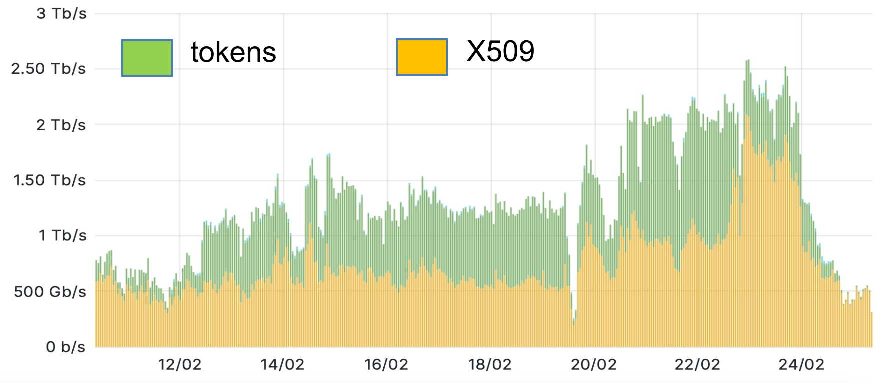

WLCG services successfully supports DUNE and Belle-2 computing models

WLCG Worldwide LHC Computing Grid

Evolve towards a network-centric model

Fewer and larger facilities operating storage services

CPUs and storage not necessarily co-located: to deliver the content over the WAN and/or cache it

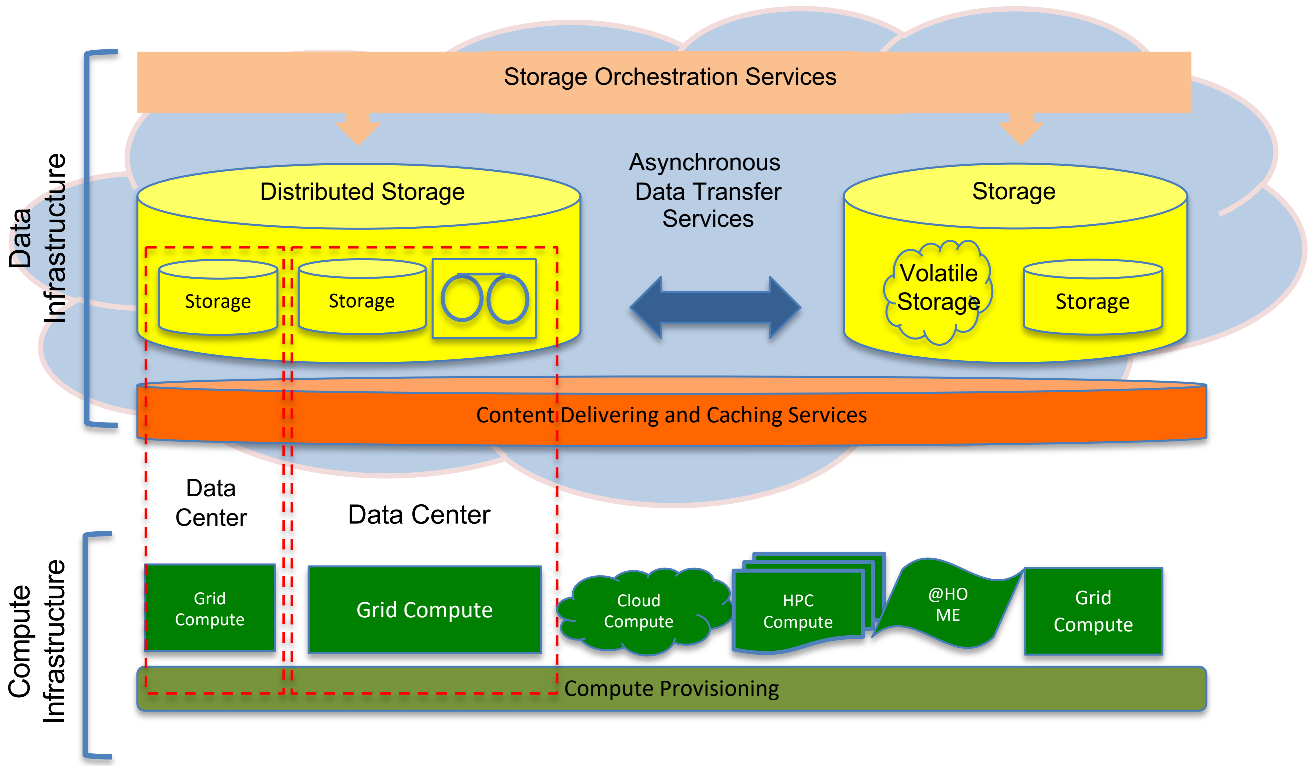

A model being pursued in different countries. Allows optimizing the operational costs and leveraging the existing expertise in challenging and impactful tasks

WLCG Worldwide LHC Computing Grid

From Data Management to a Research Environment

ESCAPE prototyped a distributed data infrastructure across Europe for HEP, Astronomy and Astro-particle physics. Based on many of the WLCG building blocks and on top of many WLCG facilities

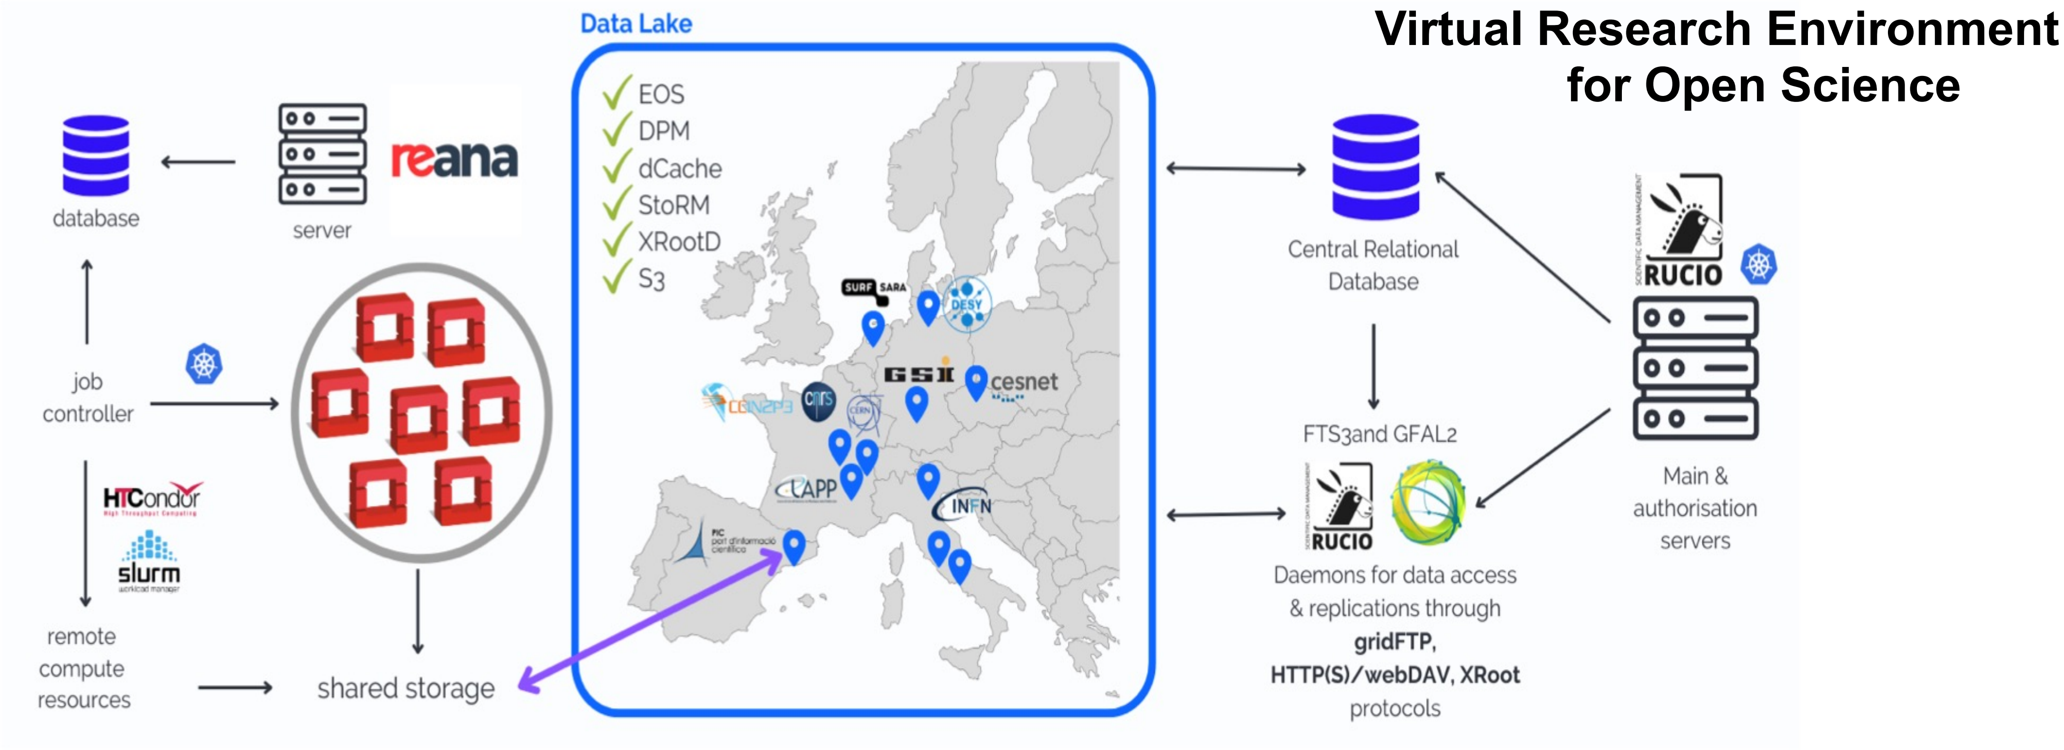

Open Science in a common Research Environment

SKA: data delivery from Perth and Cape Town to Europe and access through the data lake services

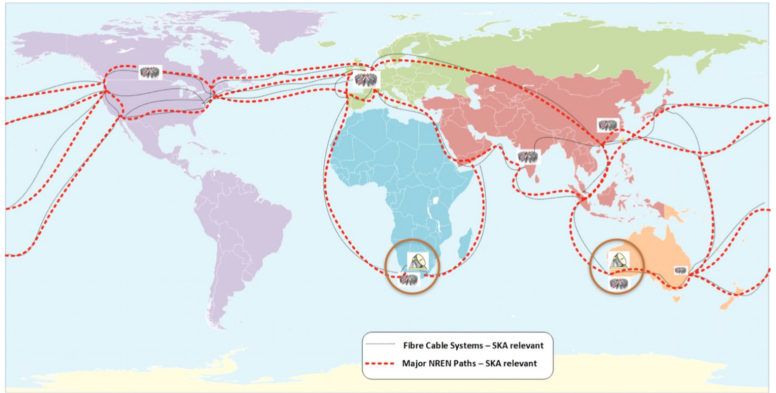

LOFAR: observations stored in the datalake, accessed and analyzed in a notebook service. Results and artifacts are all preserved.

I am told this is a Quasar …

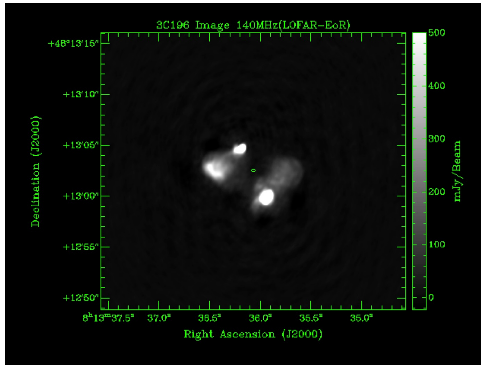

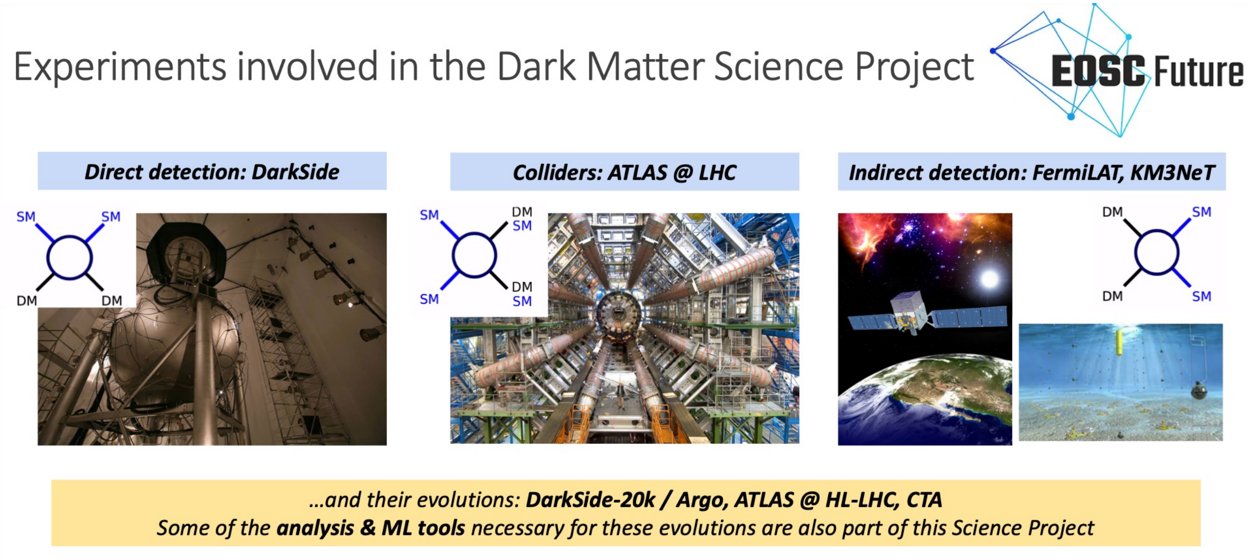

Energy efficiency

The electricity costs have been an unexpected development in the last couple of years. Environmental impact needs proper addressing!

What to do:

- Improve software performance
- Leverage modern architectures
- Invest in the facilities

There is no magic wand, however.

In a pessimistic scenario – little progress on all the above – the energy need is x2 than in an optimistic one

WLCG: the peak of energy need at the start of Run-5)

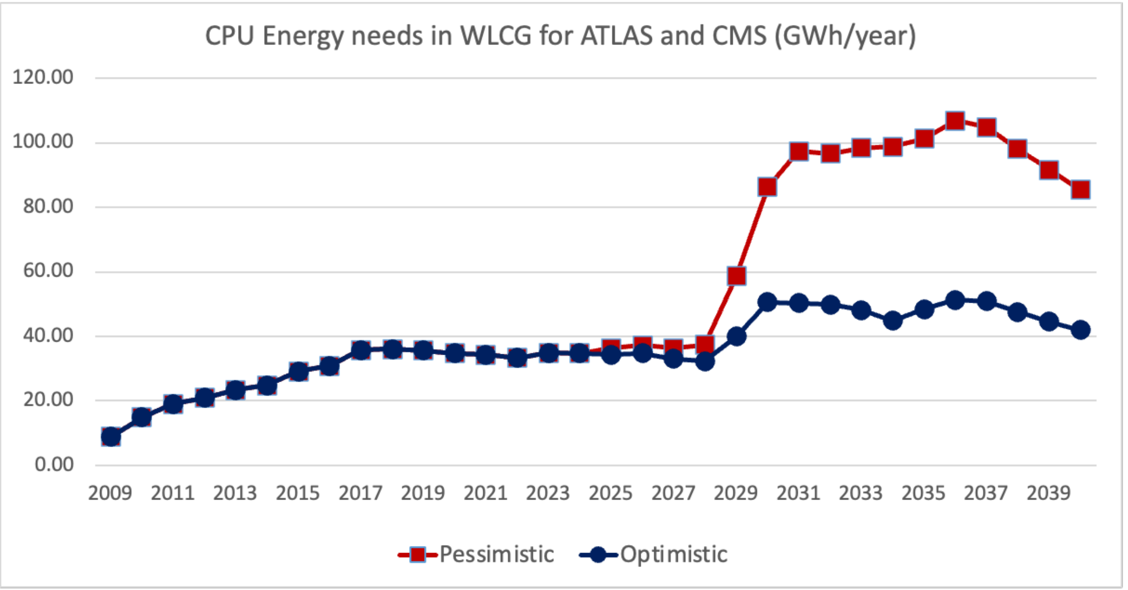

WLCG
Worldwide LHC Computing Grid

Energy Needs per unit of "science"

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

The scale on the right **(RED)** shows the energy and the scale on the left **(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

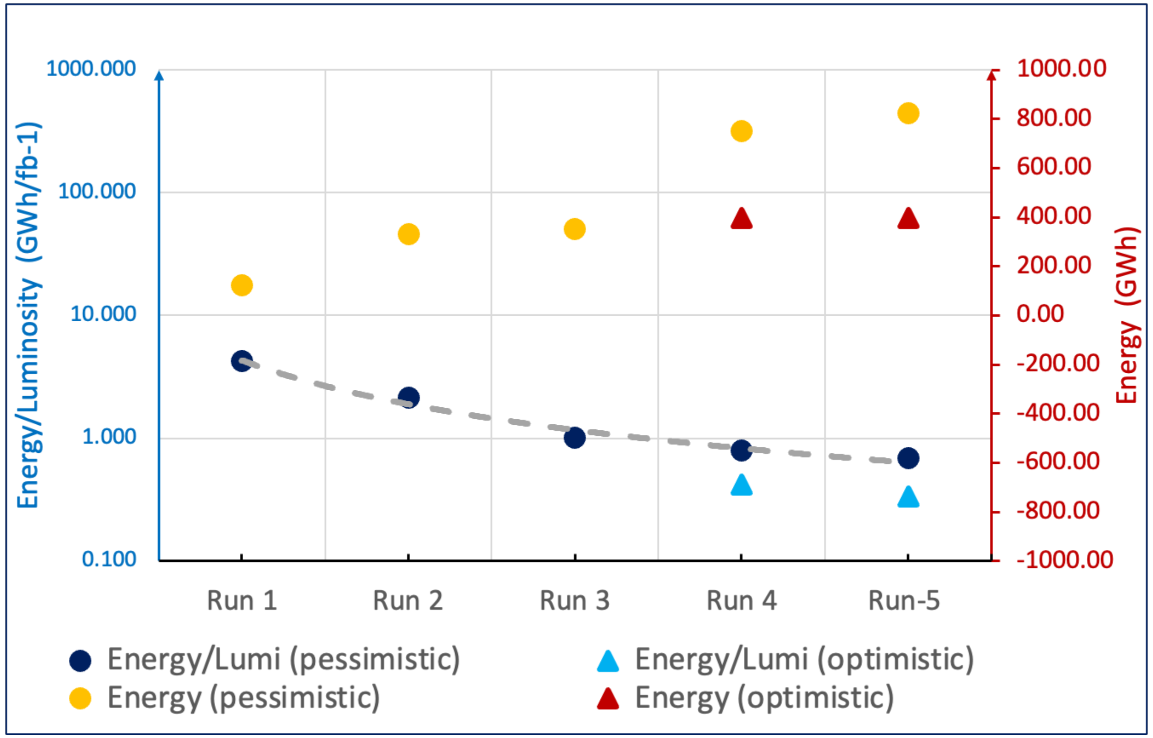

GWh/fb^{-1} decreases a factor 10 between Run-1 and Run-5 (exponential trend fits well)
In Run-5, GWh/fb^{-1} in the **optimistic** scenario is half compared to the **pessimistic** scenario

Invest in R&D

WLCG Worldwide LHC Computing Grid

Conclusions

The WLCG organization is providing a shared distributed infrastructure for the LHC experiments since 15 years.

The computing models together with the infrastructure and services changed with time adapting to the evolving landscape (experience and funding)

The HL-LHC will be an unprecedented challenge for us both in terms of scale and sustainability

Other very data intensive science projects will co-exist with LHC, with very similar use cases, on the same timescale and for a great majority on the same physical resources

There is an opportunity to share policies, tools, expertise, services and, when plausible, resources for the benefit of all our sciences

Backup

Event Generators on GPUs

A very good candidate for GPU acceleration with benefits for many experiments

CPU + GPU

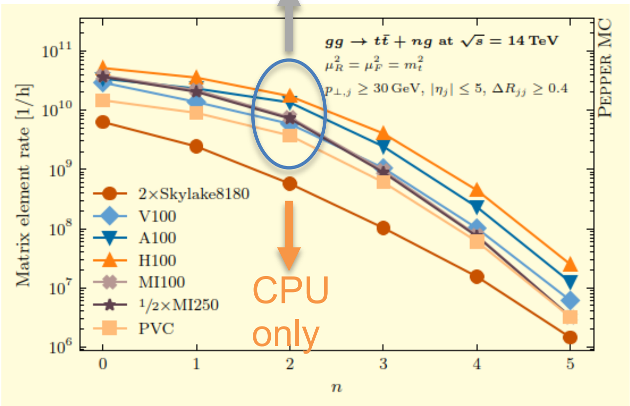

CPU only

Sherpa gg->tt+ng

Matrix Element event throughput: up to x10 gain when using GPUs

Available for production

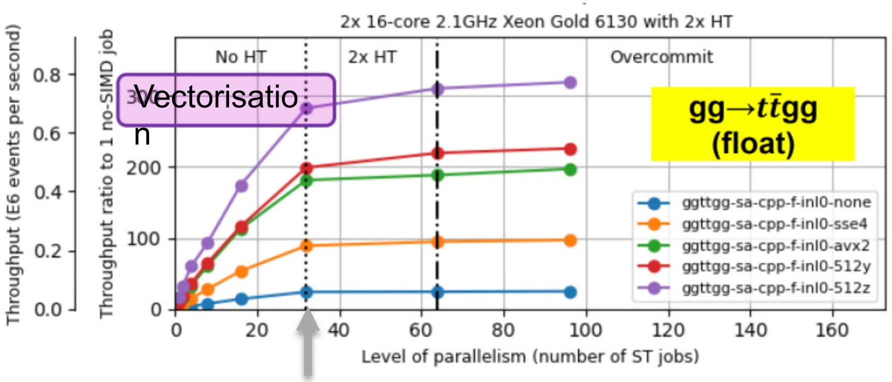

No Vectorisation

Madgraph gg->tt+ng (n=2)

GPU-enabled Leading Order: being released to production.
By-product: enabling of CPU vectorisation: up to x8 gain in ME event throughput (x6 global). Note: all CPUs in WLCG provide vectorisation

GPU-related work brings immediate benefits also on CPUs

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

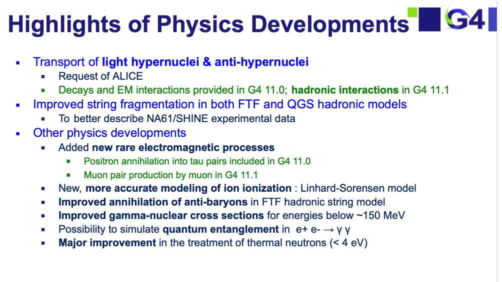

Highlights of Physics Developments

- 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- → γ γ
  - **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

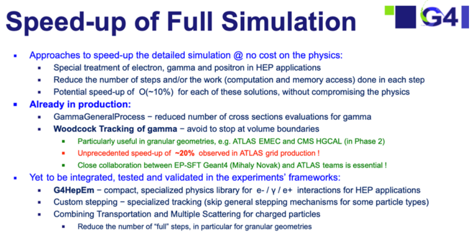

Speed-up of Full Simulation

- 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

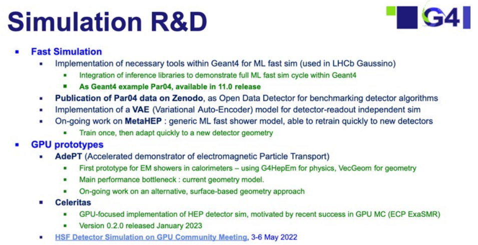

Simulation R&D

- **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

ROOT foundation layer for HL-LHC

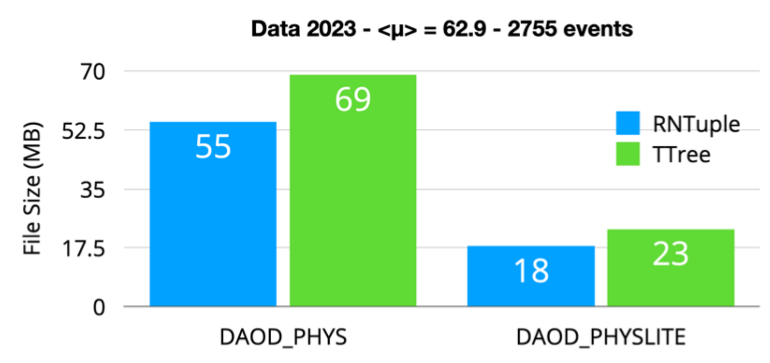

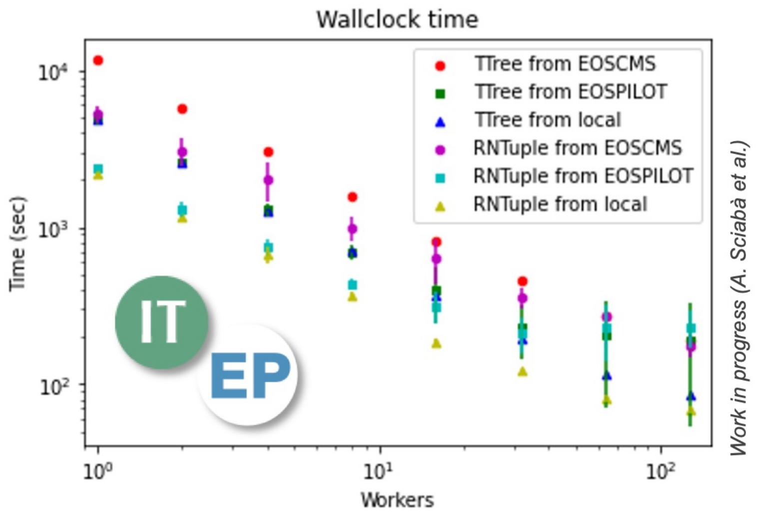

ROOT provides the common I/O layer for the LHC experiments.

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

RNTuple speed-up improvements measured in a real environment at CERN using a community standard analysis benchmark

RNTuple progress well on schedule for HL-LHC