EIROforum

October 2020

WLCG:
an infrastructure for HEP computing

Simone Campana (CERN)

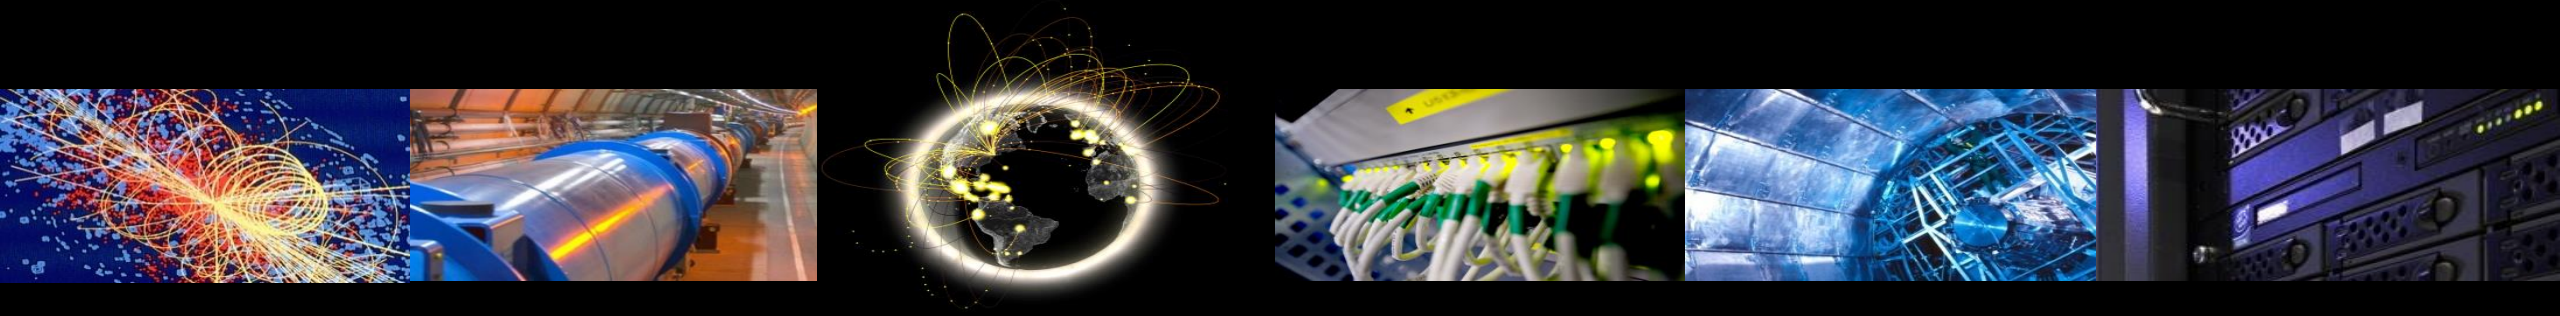

WLCG
Worldwide LHC Computing Grid

CERN

Evolution of the Universe

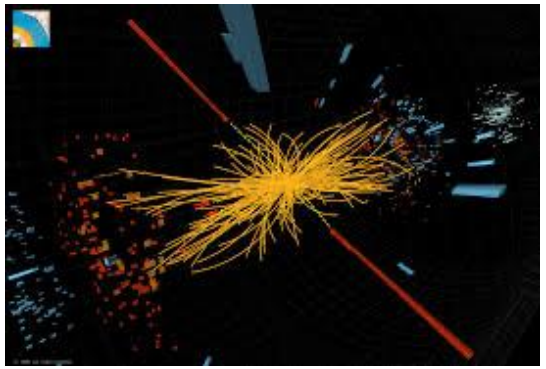

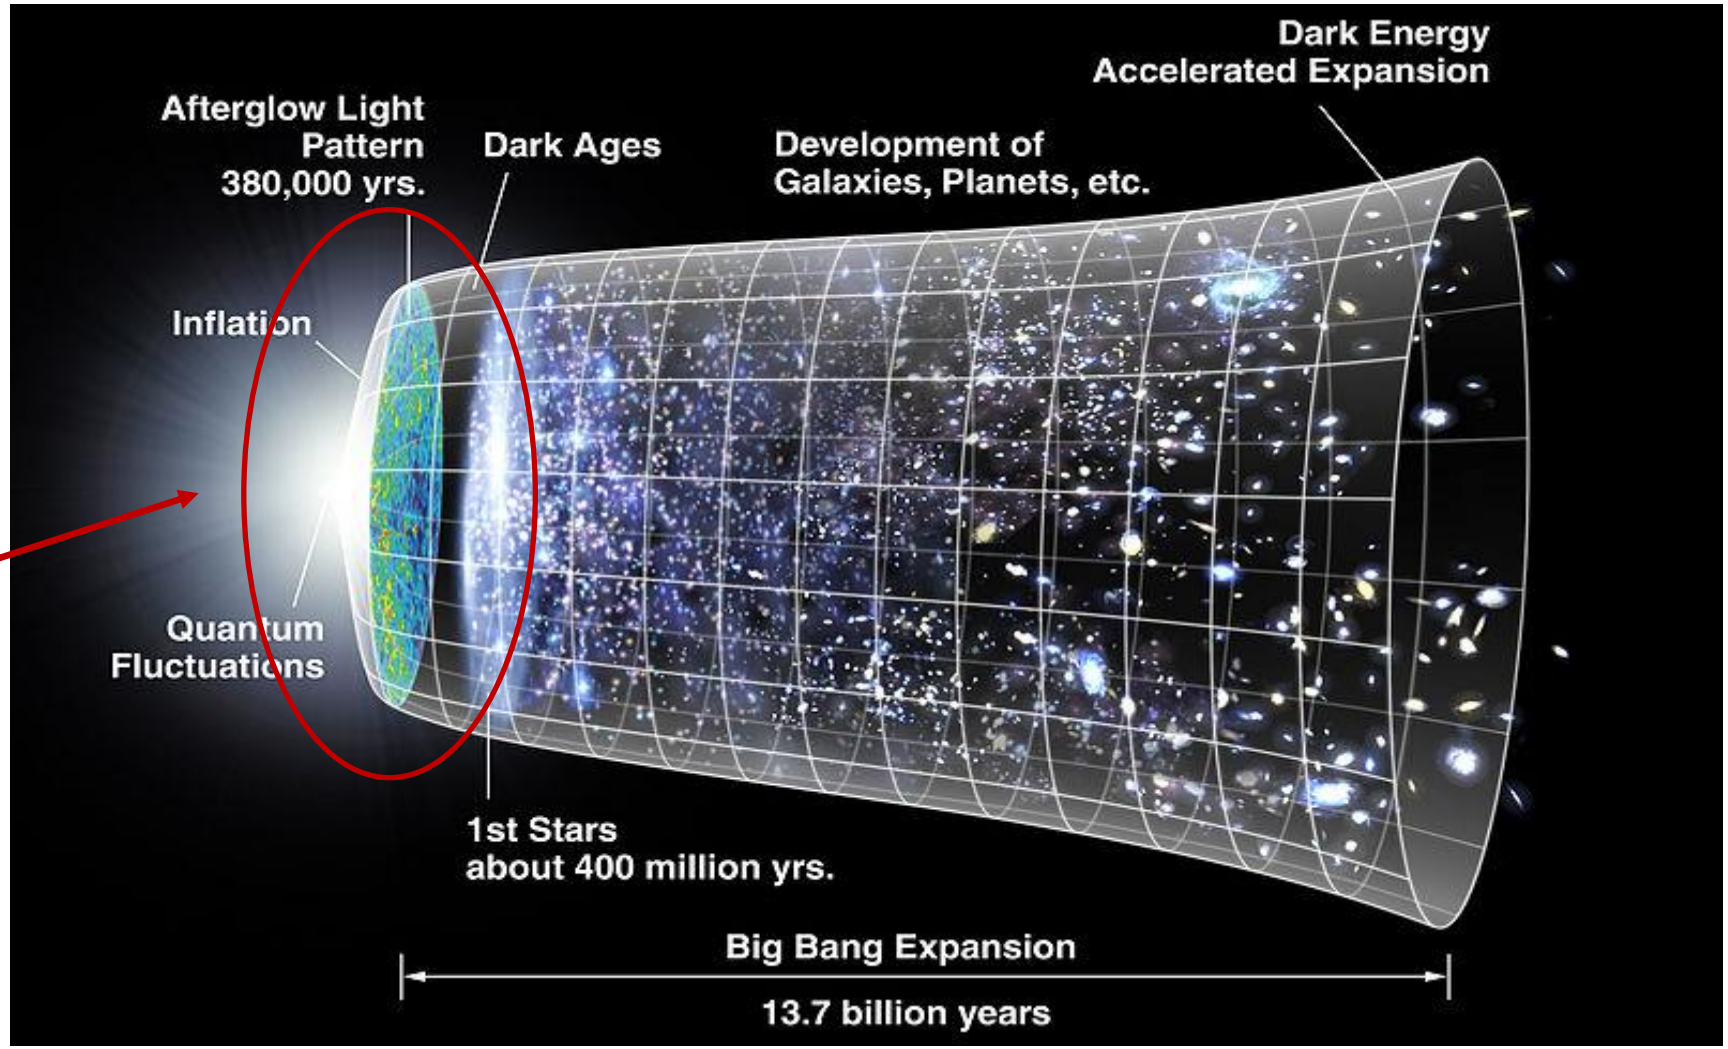

The Large Hadron Collider @ CERN

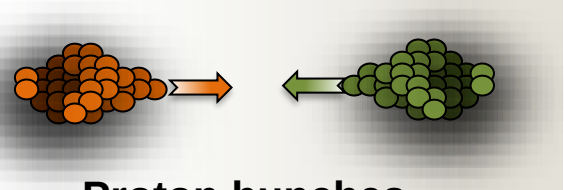

Proton bunches
>10^{11} protons/bunch
(colliding at ~**40MHz)**

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

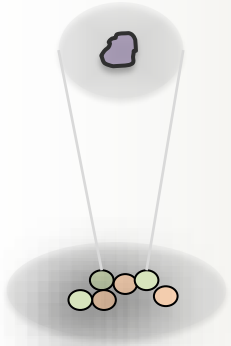

New Particle!
(<<**mHz?**)

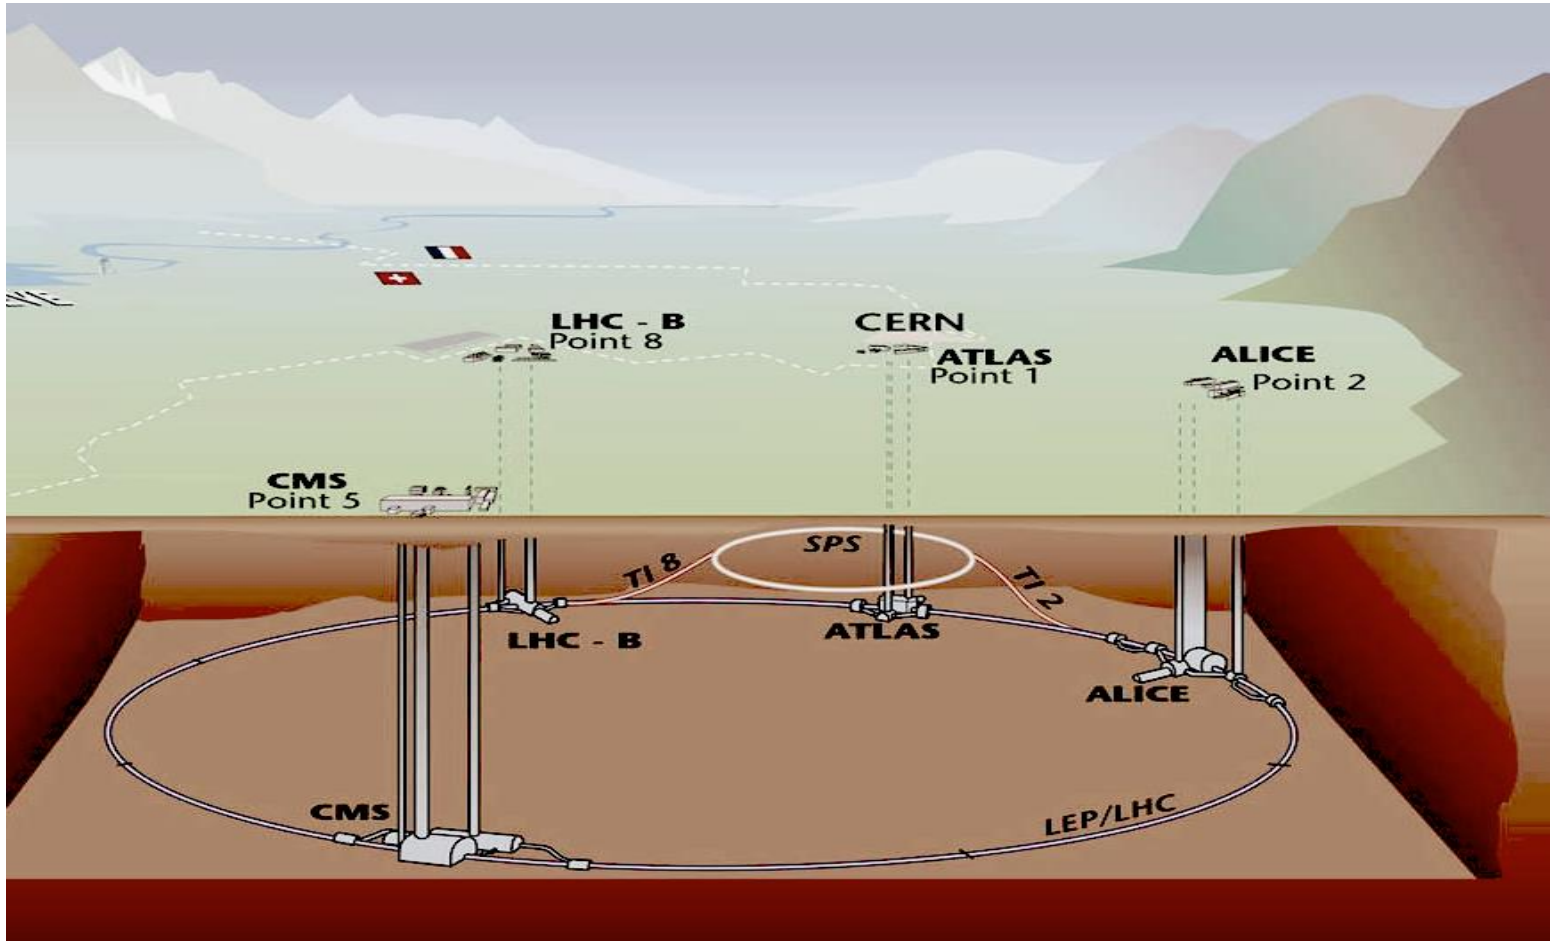

27 km ring of superconducting magnets and accelerating cavities

Some pictures …

The Large Hadron Collider

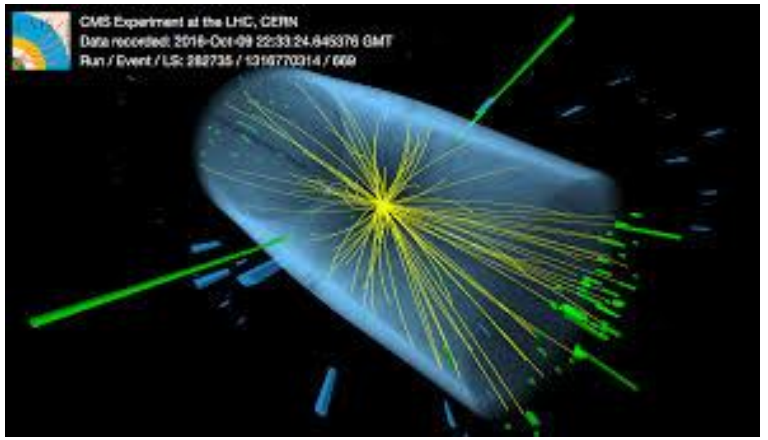

The Higgs particle

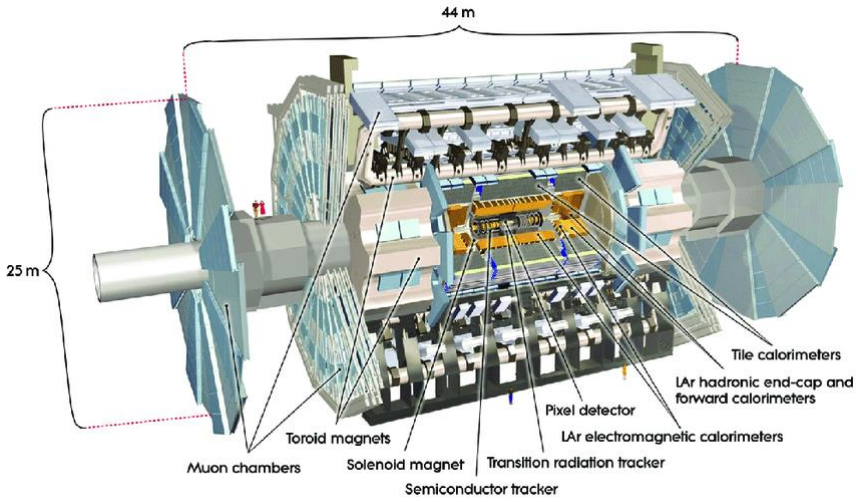

The ATLAS detector

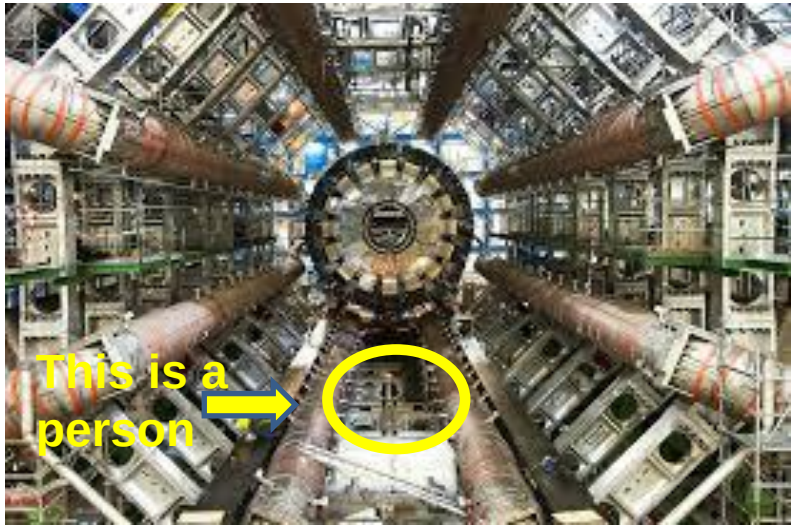

Computing @ High Energy Physics

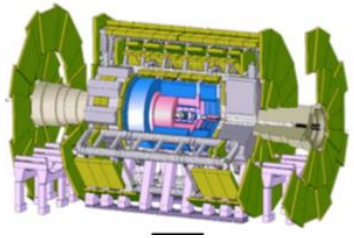

~40 MHz

~ PB/s

"Online" – Real time

L1 Trigger (HW)

~100 kHz

HL Trigger (SW)

~1 kHz

Reconstruction

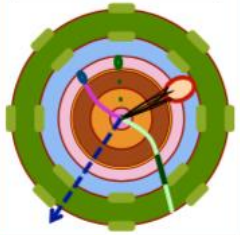

Calibration

Monte Carlo Simulations

Data analysis

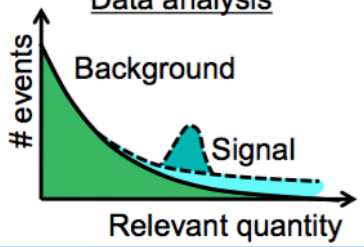

"Offline" - Asynchronous

From Big Data (40 Million collisions/s

To Useful Data (1000 events/s) for analysis

1MB/event

50 PB/year

Large scale HEP computing embraced a distributed model since early 2000s
Based on network service technologies, federating national and international initiatives

Tier-0 (CERN):
data recording, reconstruction and distribution

Tier-1: permanent storage, re-processing, analysis

Tier-2: Simulation, end-user analysis

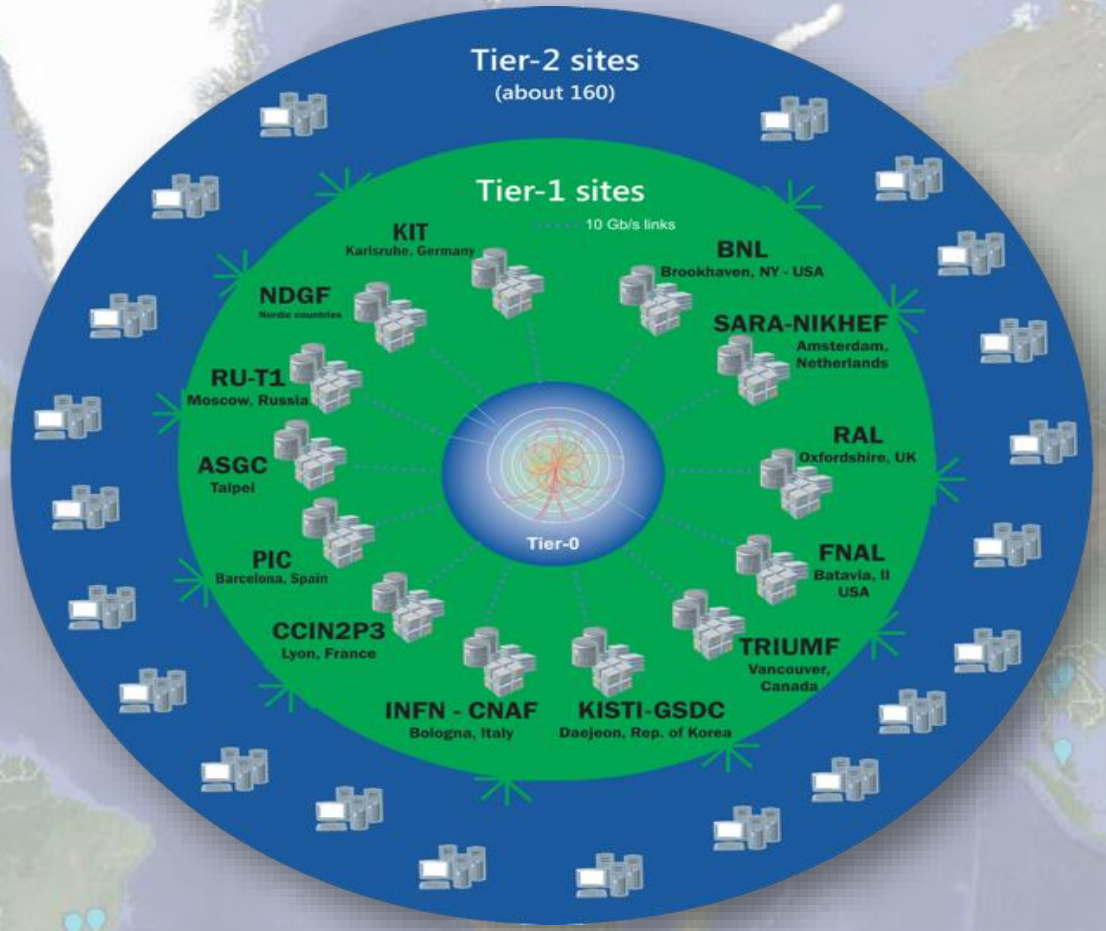

161 sites,
42 countries

~1M CPU cores

~1 EB of storage

> 2 million jobs/day

10-100 Gb links

WLCG – Worldwide LHC Computing Grid: an international collaboration to distribute and analyse LHC data

Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists.

The WLCG Collaboration - October 2020

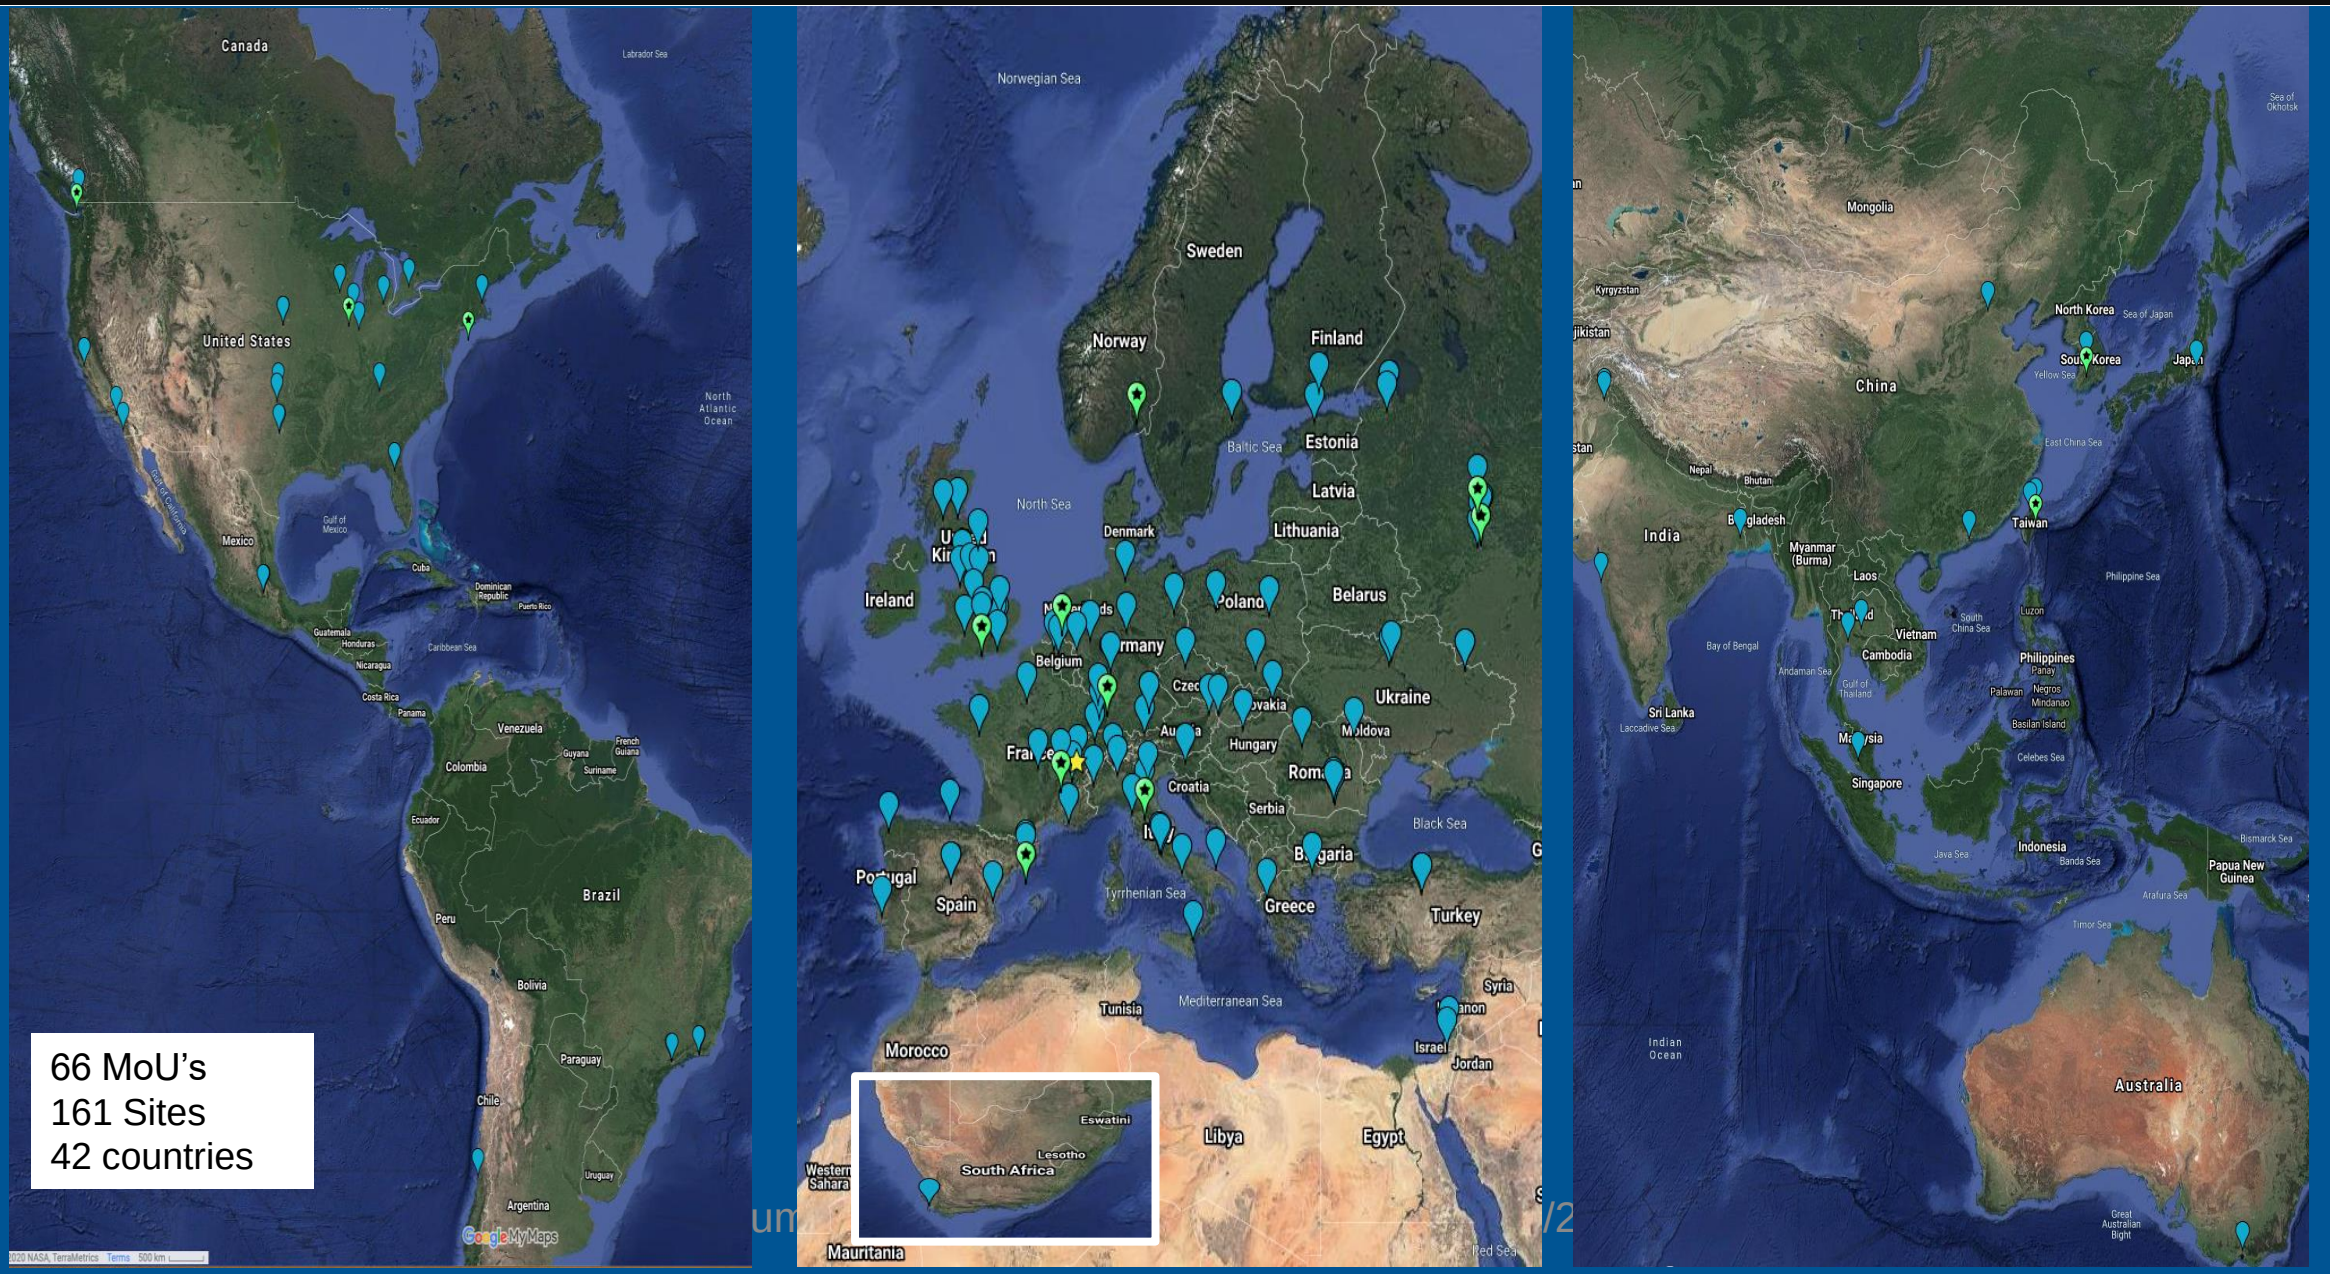

66 MoU's
161 Sites
42 countries

WLCG
Worldwide LHC Computing Grid

CERN role in WLCG as the Tier-0

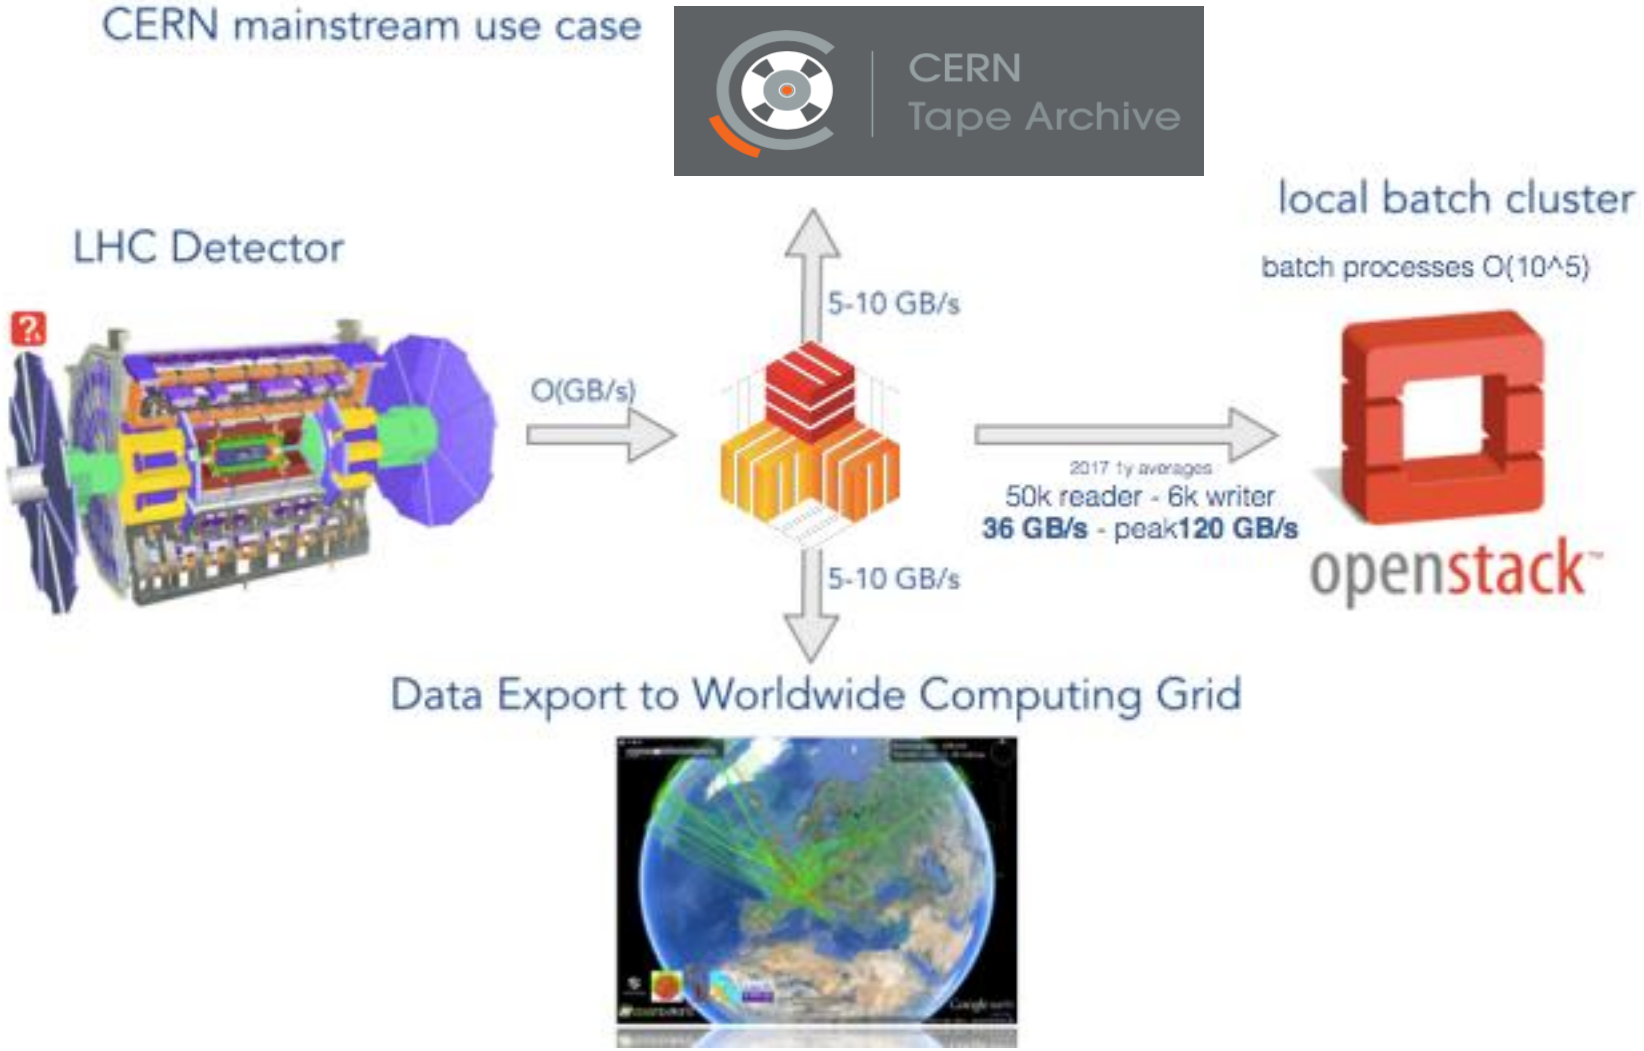

- Receive the RAW data from the experimental areas and archive them for long term preservation

- 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.

The LHC RAW data archive

The LHC physics program is organized in ”runs” (generally ~3 years), interleaved by “Long - Shutdown periods (order of ~2 years)

Tape writing at CERN consists mostly of RAW data for the LHC experiments

One can clearly see the run structure in tape writing but not tape reading

Tape is not “cold storage” for LHC

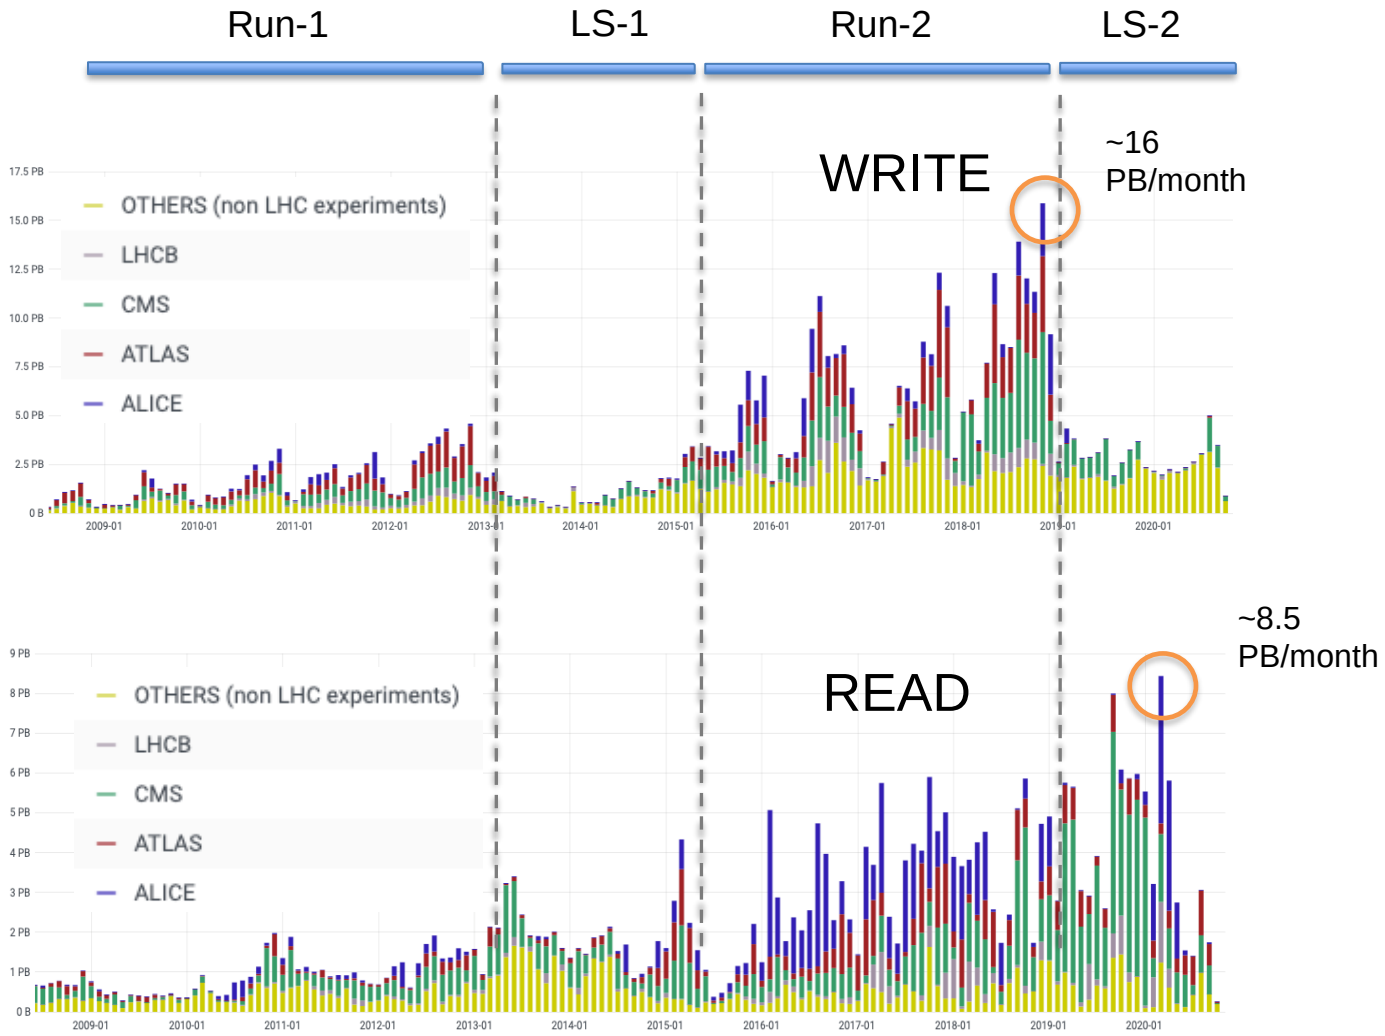

The initial Computing Model

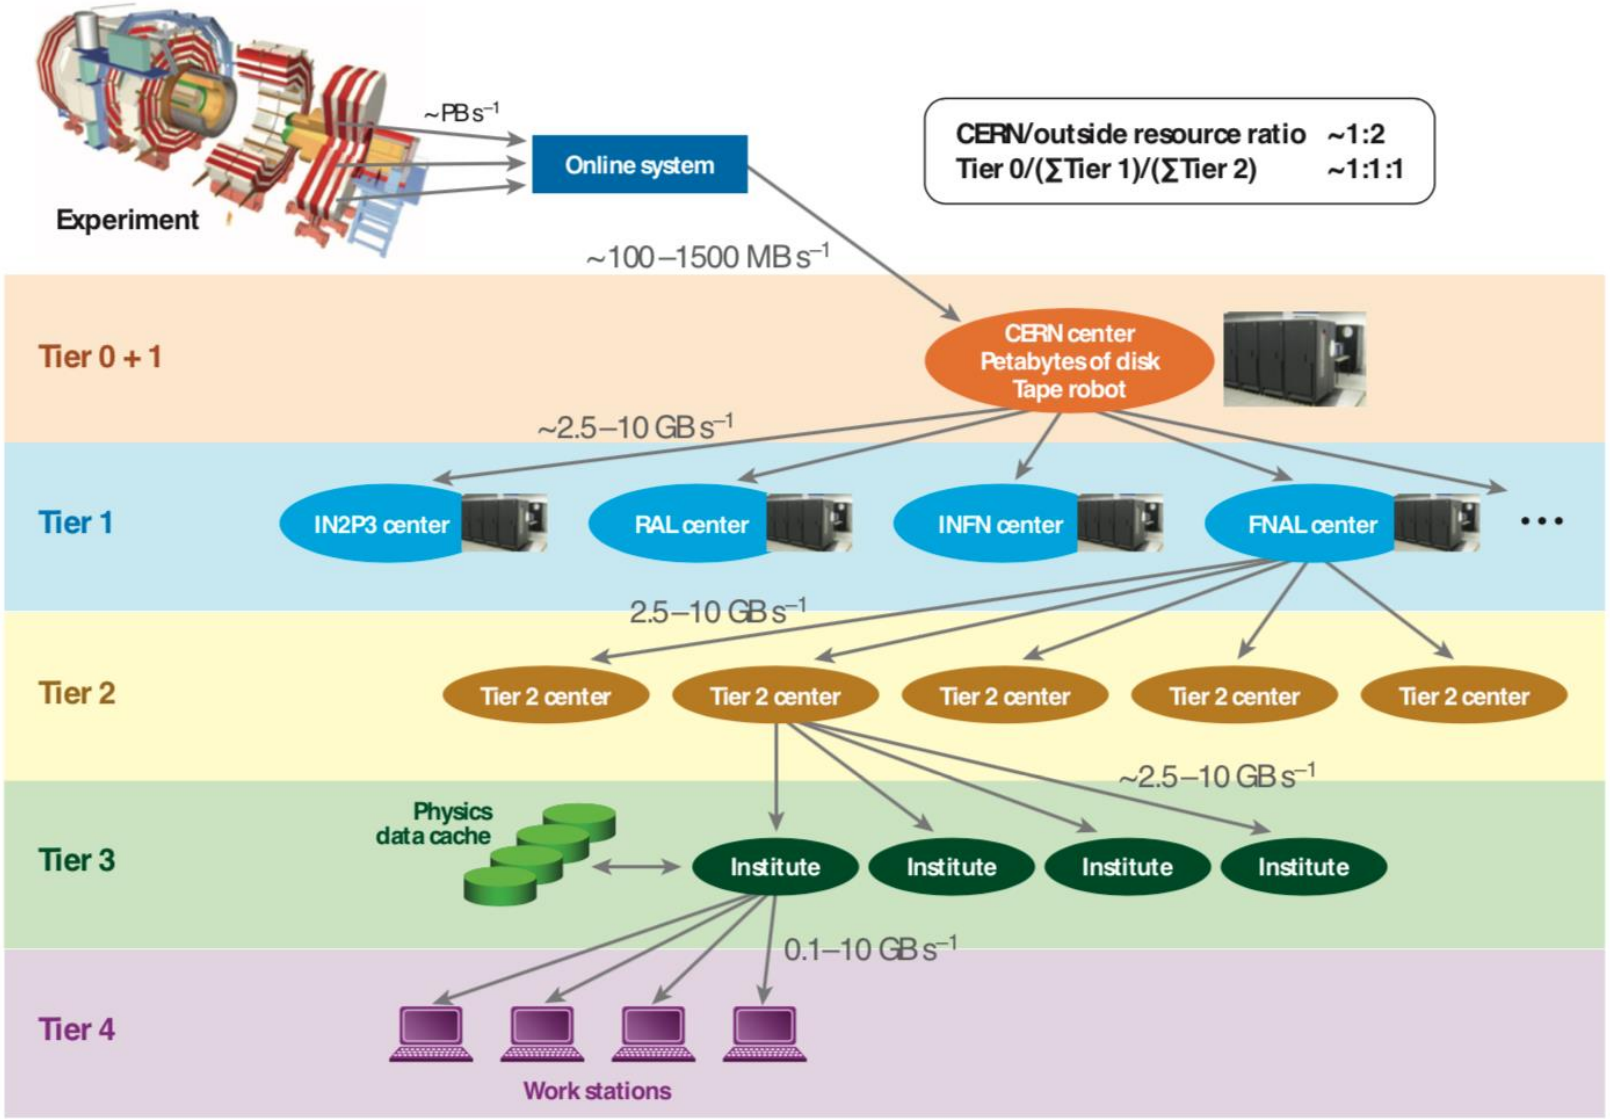

This was the initial computing model (1999)

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 …

LHCOPN

Private network connecting Tier0 and Tier1s:

- Dedicated to LHC data transfers and analysis
- Secured: only declared IP prefixes can exchange traffic
- Advanced routing: communities for traffic engineering, load balancing
- Star topology around the Tier-0 (CERN)

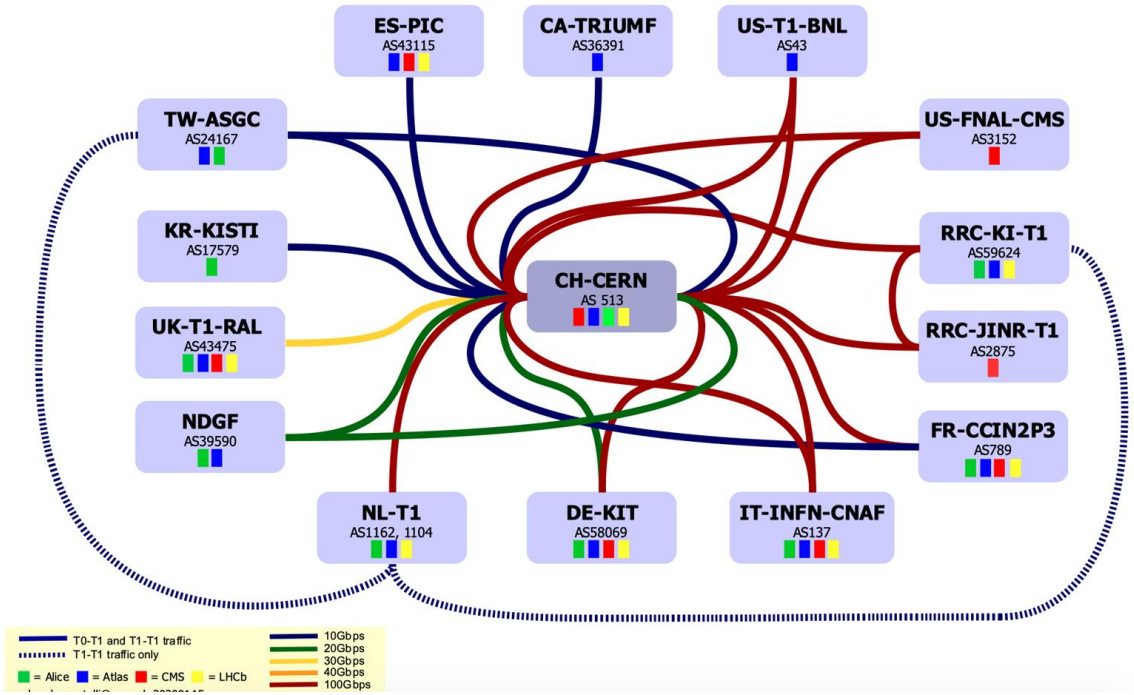

Numbers

- 14 Tier1s + 1 Tier0
- 12 countries in 3 continents
- Dual stack IPv4-IPv6
- 1.1Tbps to the Tier0
- Moved ~288 PB in 12 months

WLCG Worldwide LHC Computing Grid

LHCONE

Layer3 (routed) Virtual Private Network:
- Worldwide network backbone connecting Tier1s and Tier2s at high bandwidth
- Provided by Research and Education Network providers
- Bandwidth dedicated to High Energy Physics data transfers
- Trusted traffic that can be allowed to bypass slow perimeter firewalls

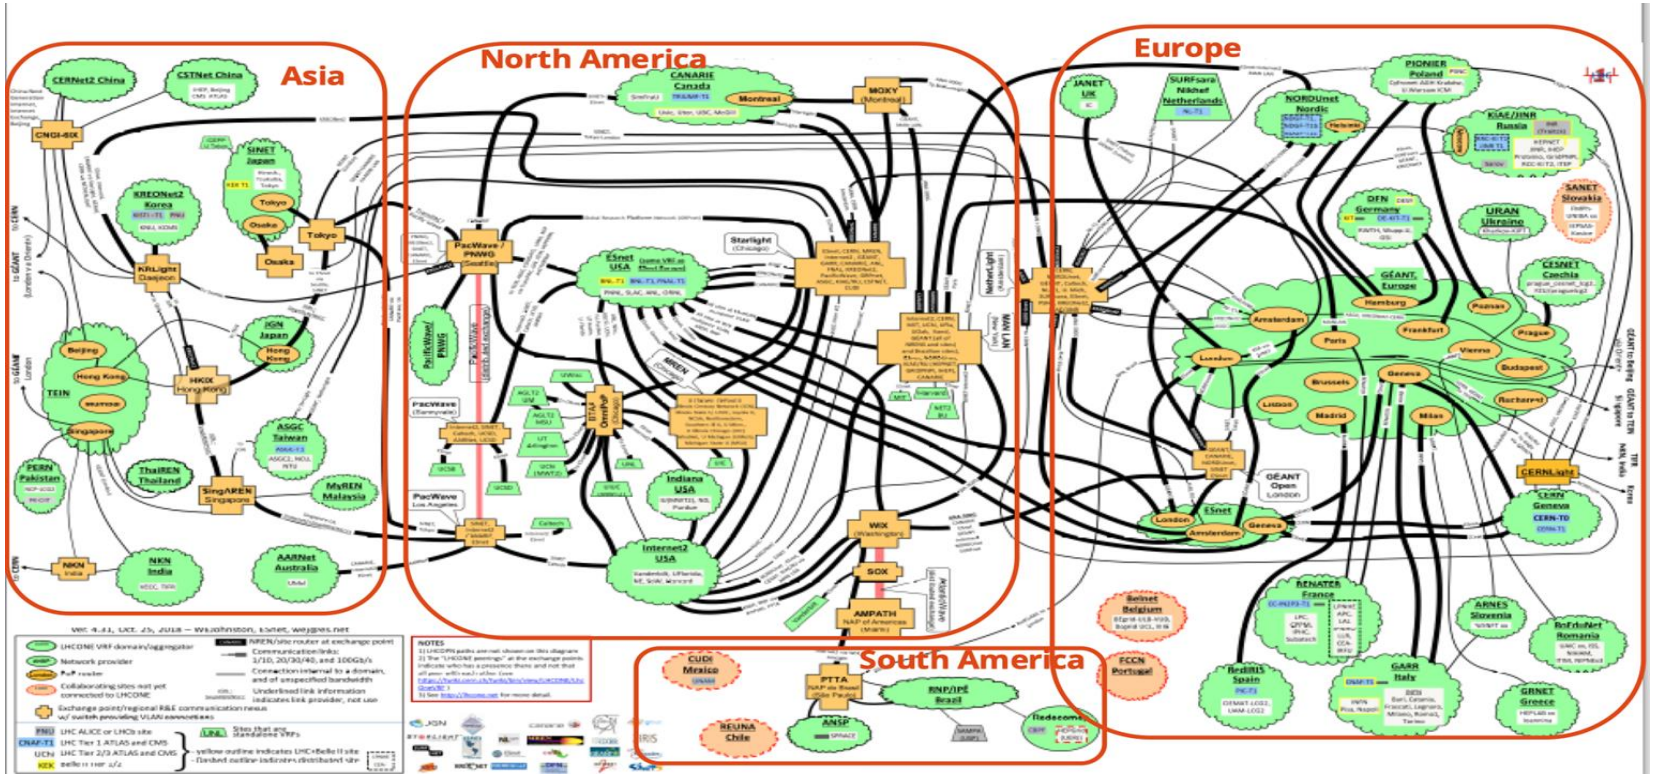

Numbers
- 28 R&E networks
- 14 Tier1s and ~90 Tier2s in 5 continents
- ~250 perfSONAR instances

WLCG Worldwide LHC Computing Grid

WLCG traffic (GB/s) in the last 5 years

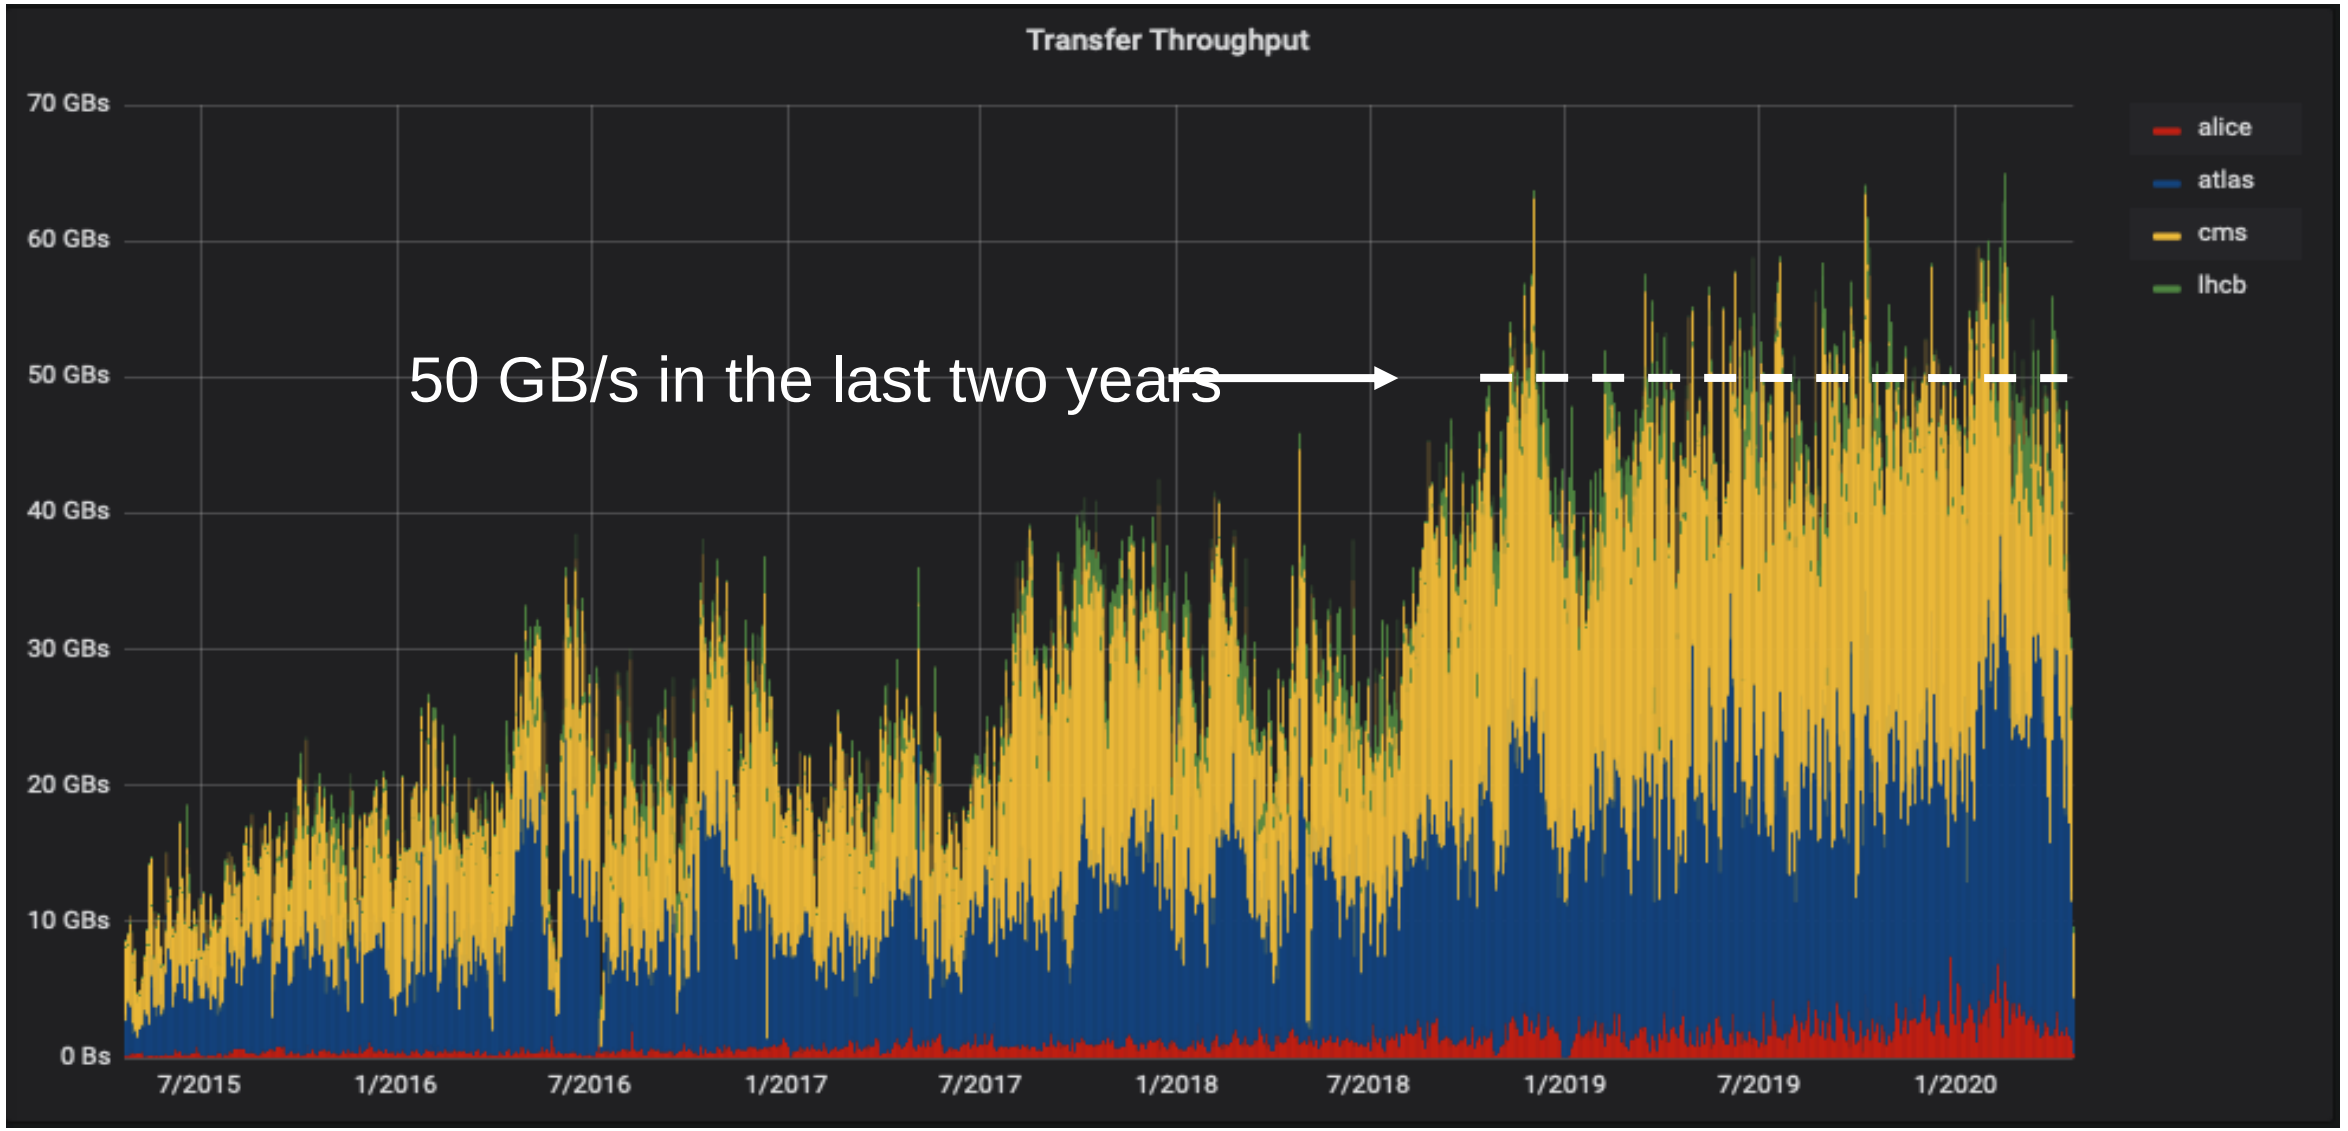

Evolution to a more flexible Computing Model

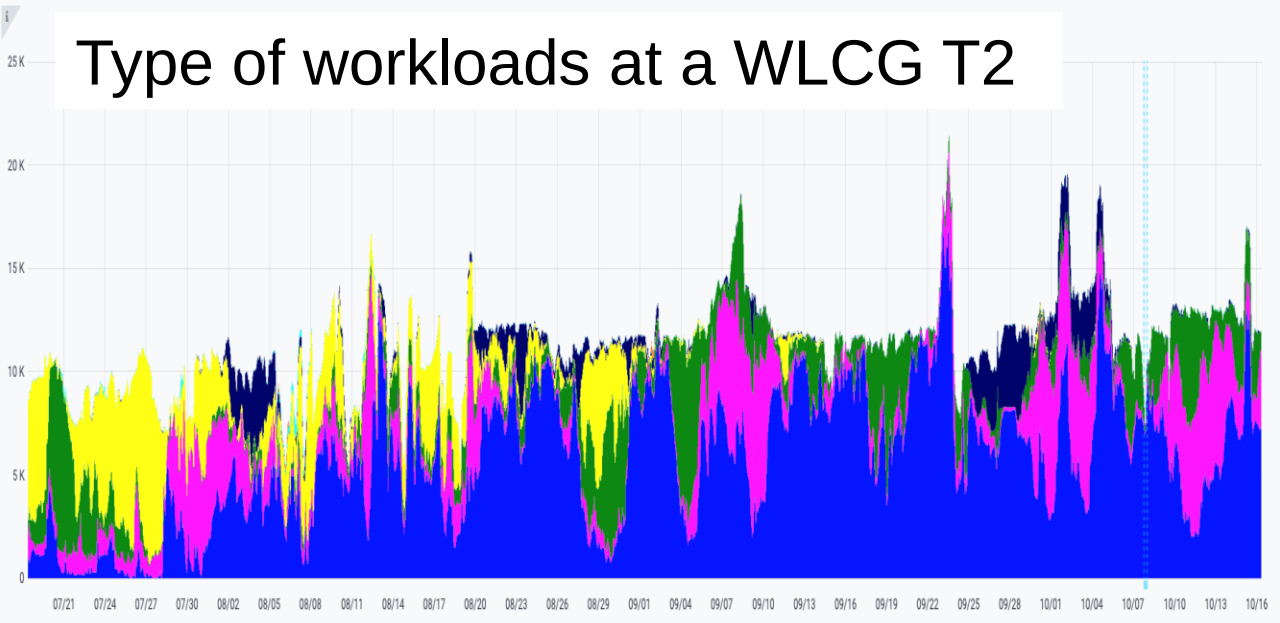

T1s and T2s run flexibly a mixture of workloads depending on their capabilities and the needs of the experiments

Data is transferred in a full mesh rather than hierarchically

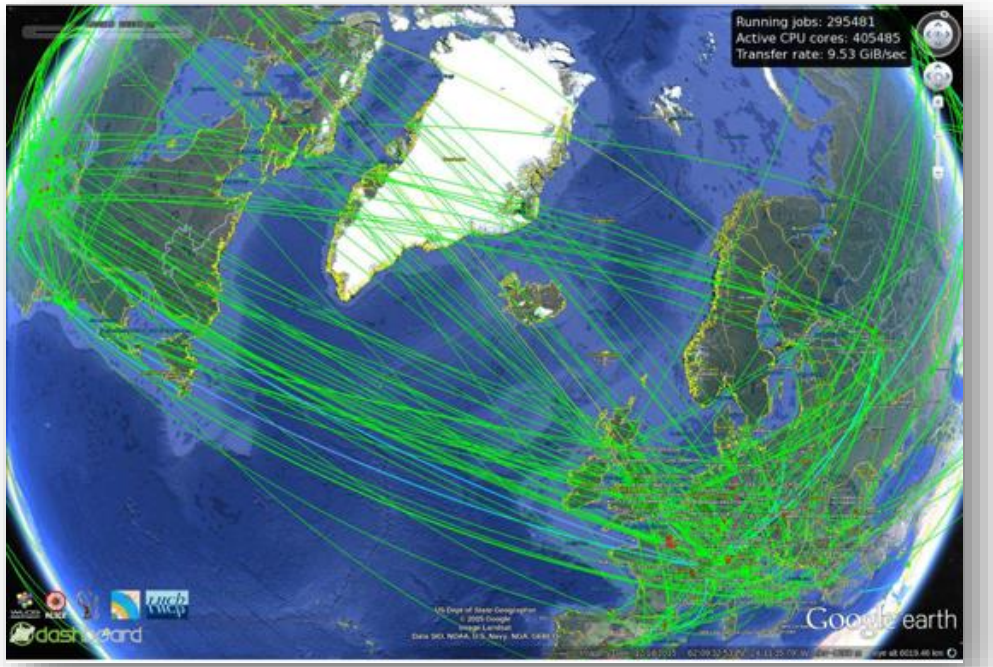

WLCG usage and performance

CPU use in WLCG continuously increases despite shutdown periods

We operate 24/7/365

1M cores concurrently in use

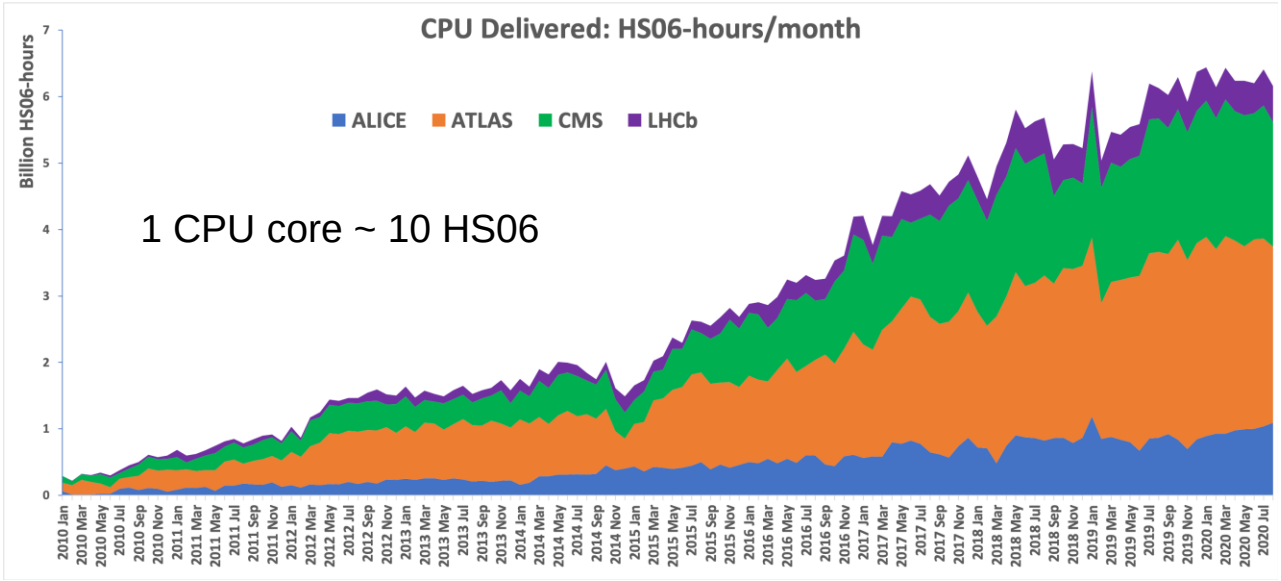

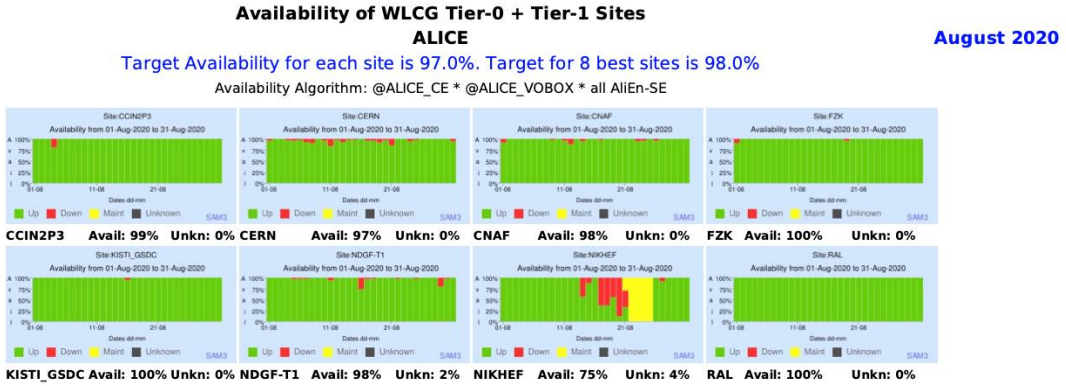

We monitor continuously and and review regularly the performance of the sites against the MoU targets

The next challenge

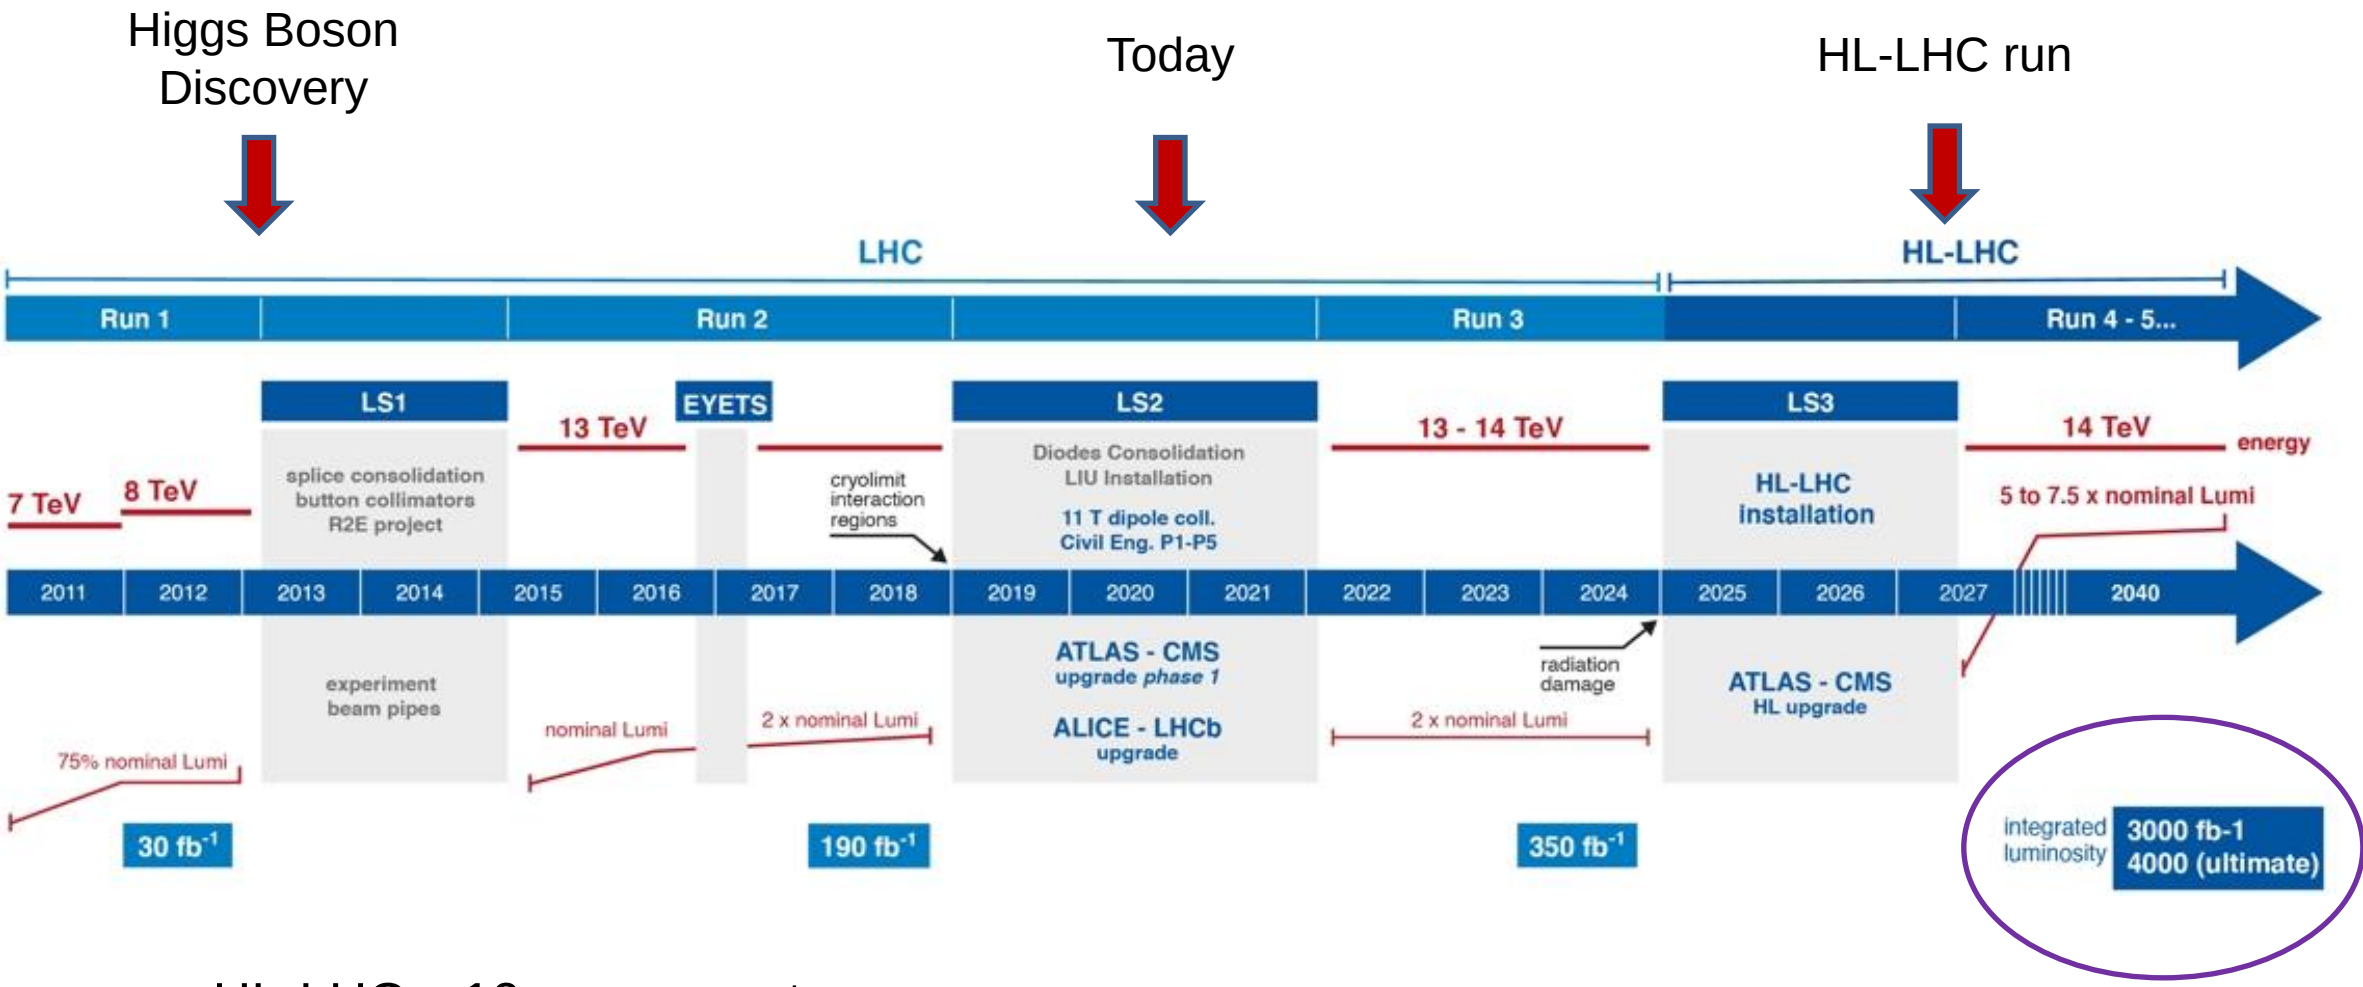

HL-LHC: x10 more events.
Bigger and more complex events (x5)

An unprecedented challenge!

Events at HL-LHC

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

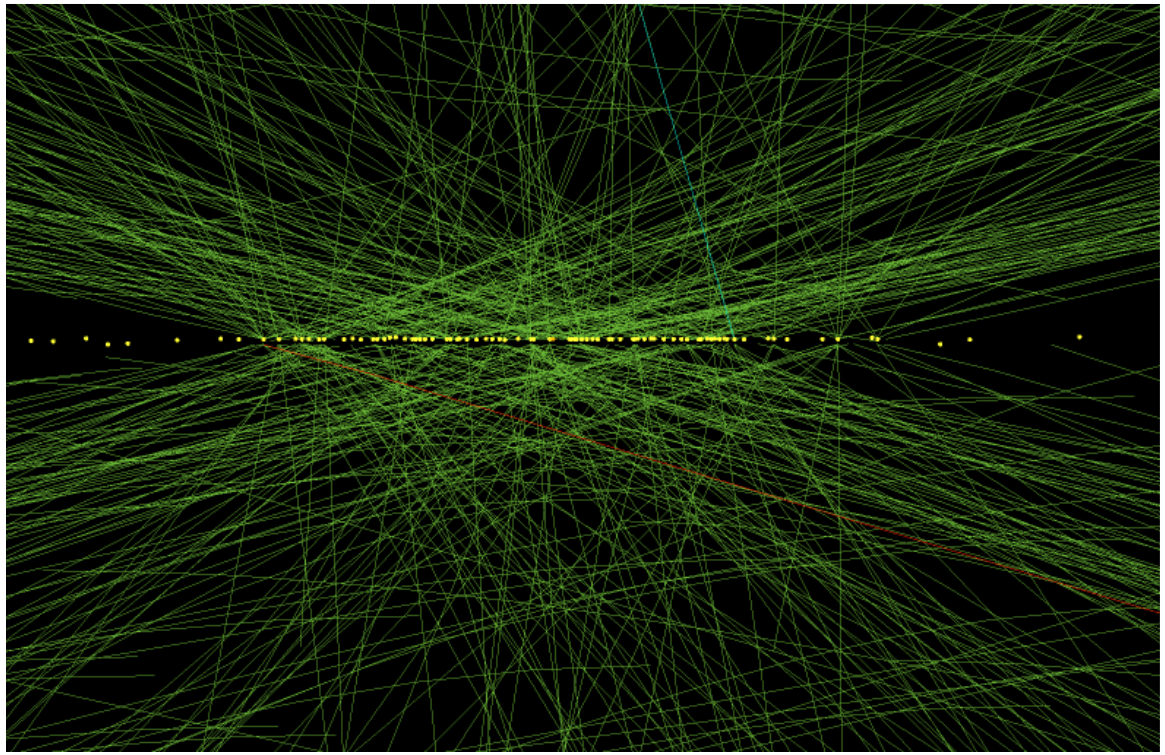

ATLAS simulation for HL-LHC with 200 reconstructed vertices

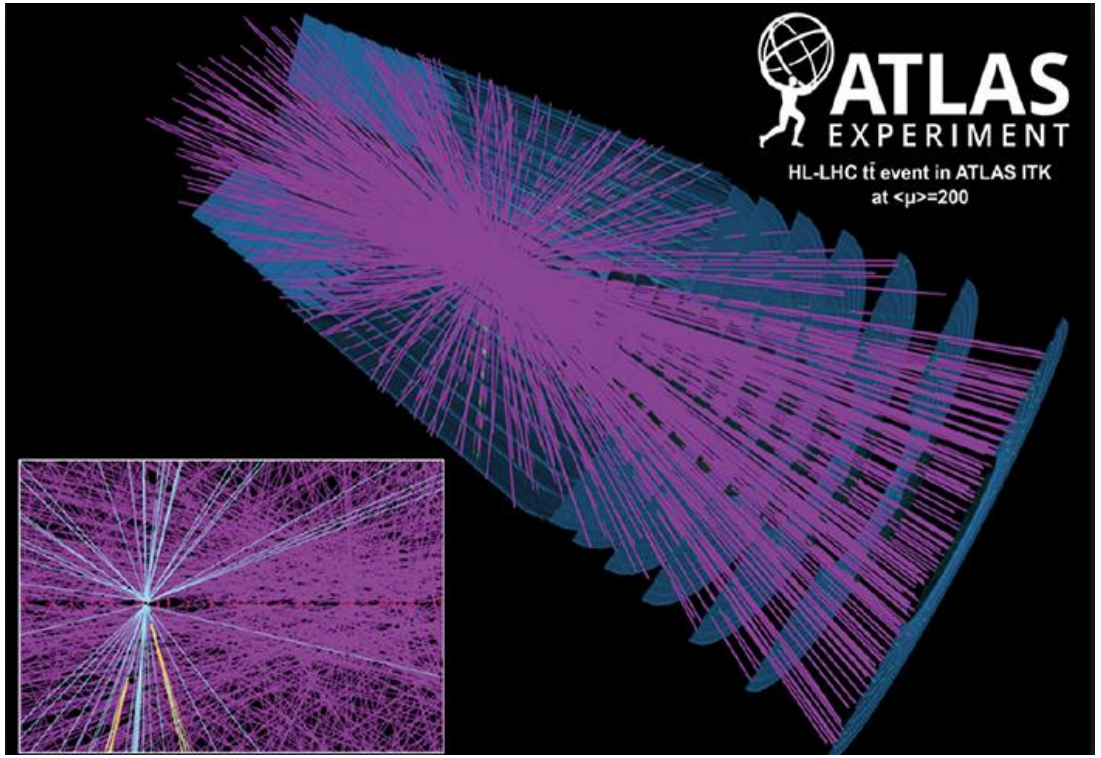

Hardware Trends

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

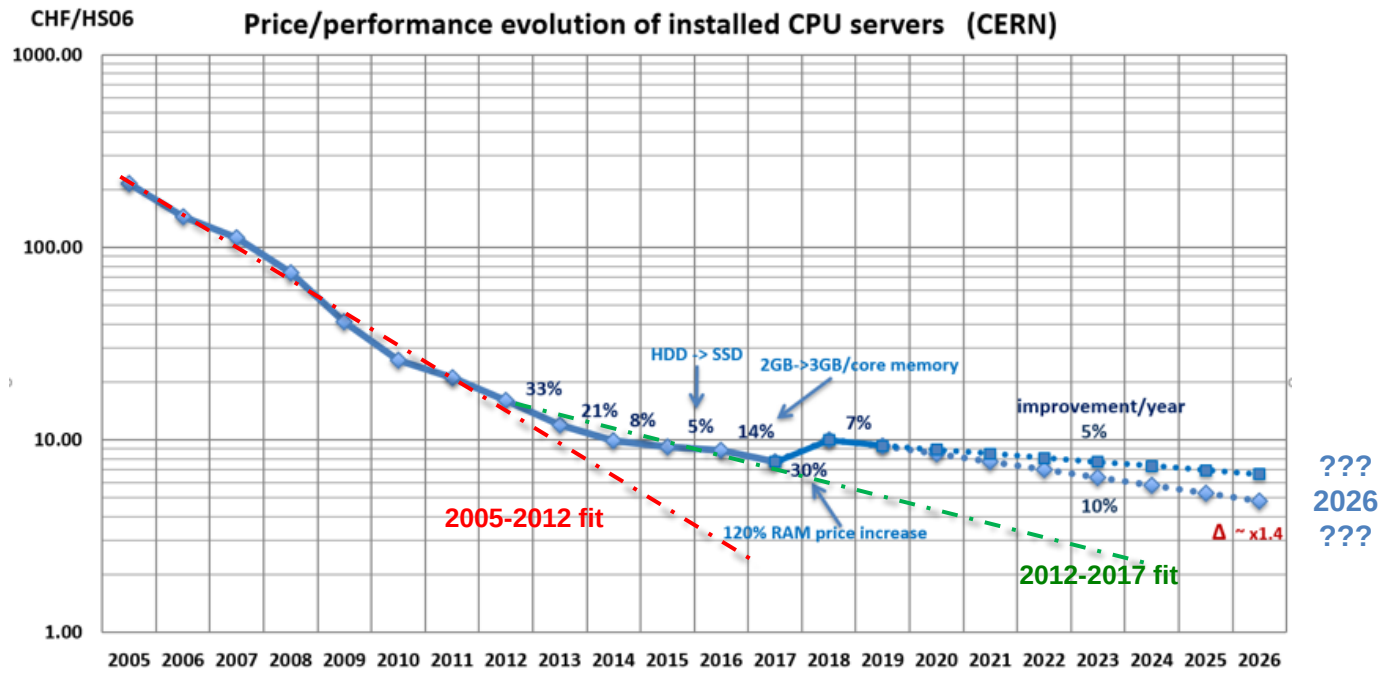

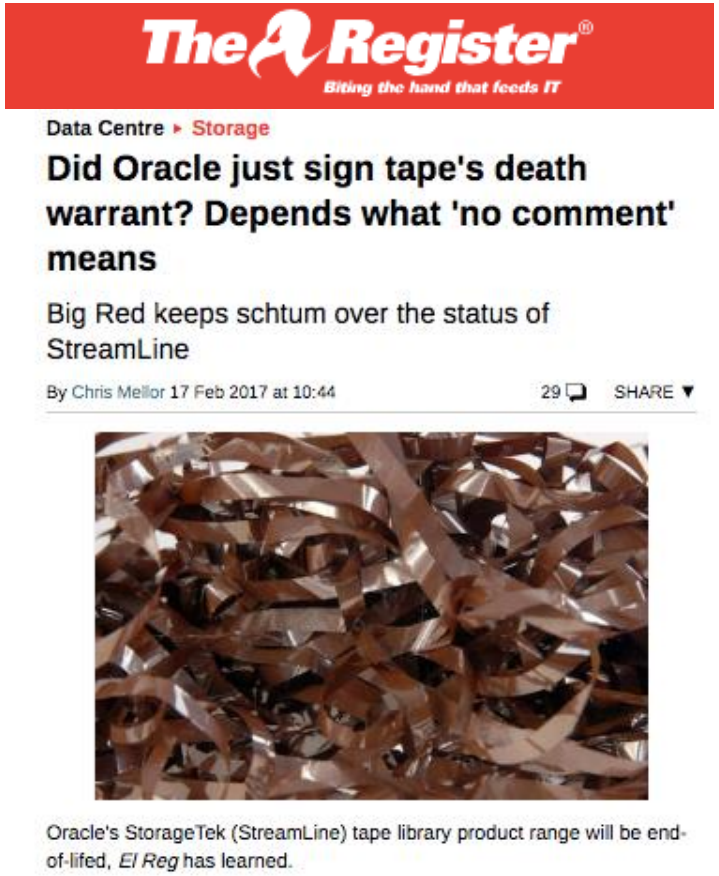

General loss of long term predictability

The HL-LHC computing challenge

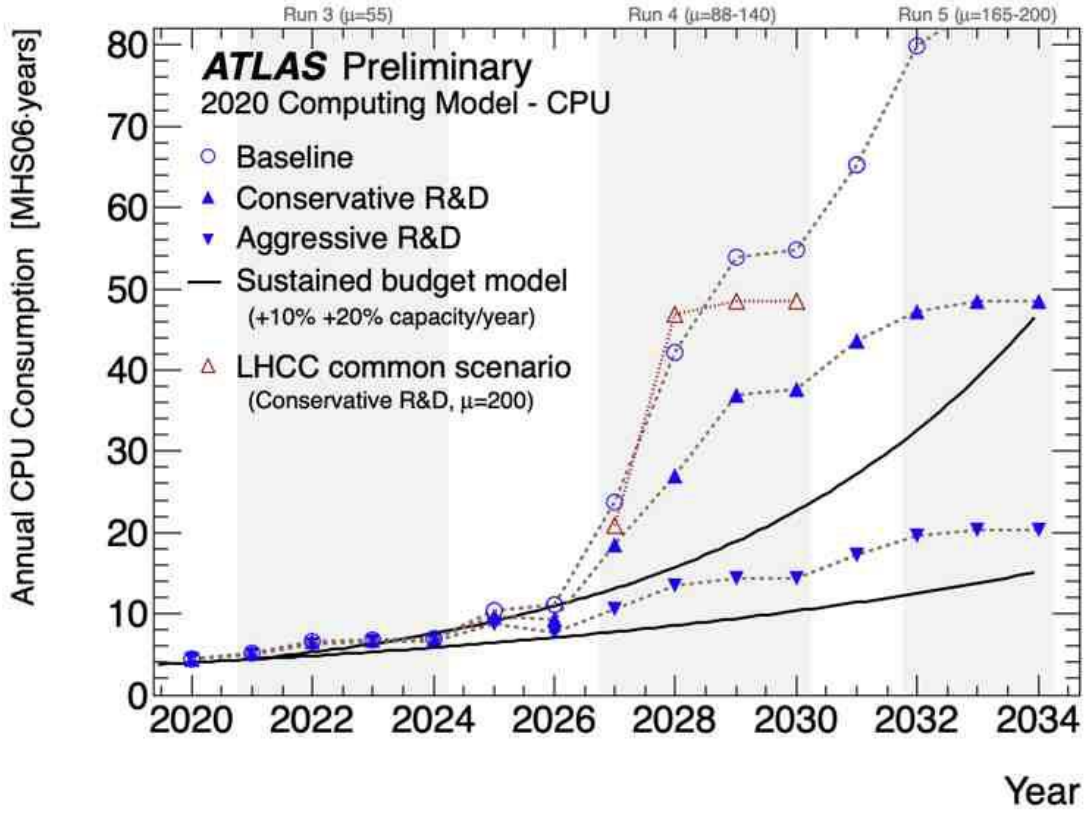

We do not foresee an increase in funding for LHC Computing in the coming years

An aggressive investment in R&D is needed to address the resource challenge

The infrastructure will need to evolve as well, not only in terms of technology but also policies

Long term sustainability of the infrastructure is at least as important as filling the gap in the resource need

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

The HL-LHC computing strategy

Strengths

- WLCG established, trust network among stakeholders. Reputation of a success story.
- Lightweight decision making processes
- Consolidated infrastructure – common building blocks, policies, federation, resource sharing, etc.

Weaknesses

- Lack of investment in software for the future; initial investment in WLCG applications has levelled off; future need more weight on software
- WLCG does not have development effort – reliance on external providers – some developments and efforts were of marginal use

Threats

- Flat funding but expectation of more and more computing and data
- Other sciences now compete with WLCG for resources at similar scales; WLCG may lose influence
- Need to use facilities not particularly suited to our problems – HPC
- No financial support for software nor career recognition

Opportunities

- Leverage experience and be central to evolution of scientific computing and become the reference, for the benefit other sciences, resource providers, funding agencies
- Leverage extra resources (HPCs) by investing in software and building relationships
- Retain significant leadership role in setting direction

Heterogeneous Computing

Today we get opportunistic use of many type of compute in particular HPC systems and the experiment online farms presented as Cloud resources

In the future this heterogeneity will expand and we need to be able to make use of the resources provided to us

A lot of work going on at the level of policies, infrastructure/services and software

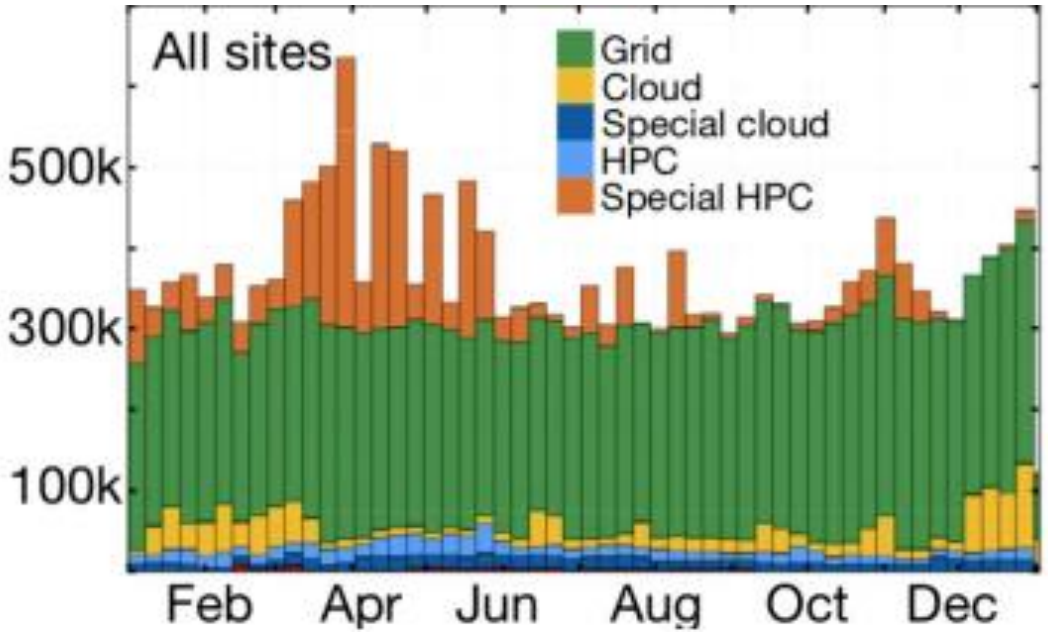

New Summit HPC @ ORNL:
9.96 Pflops in CPU (**non X86_64**)
215.7 Pflops in GPUs

Moving to a network-centric model

Datalake model:

Fewer number of facilities operating storage services, less data replication

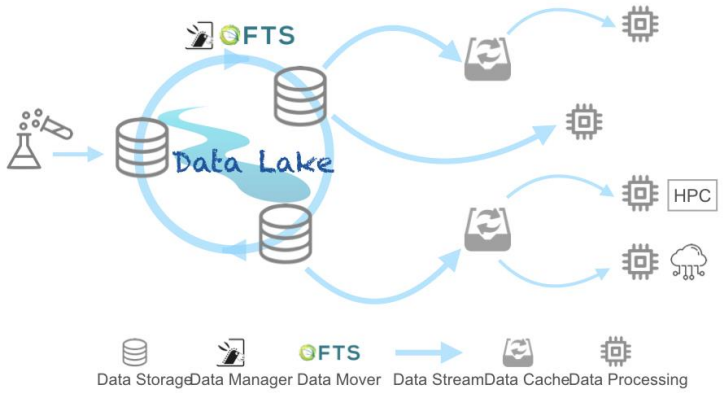

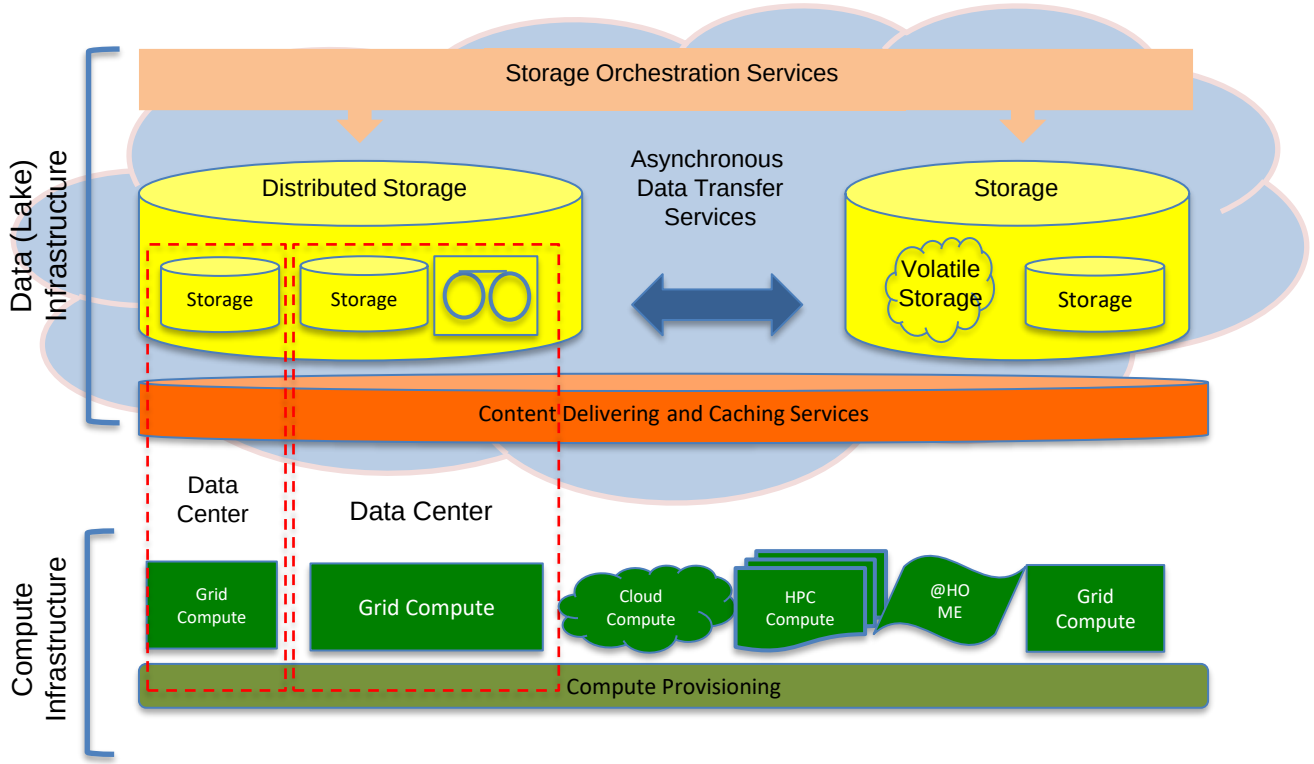

CPUs and storage not necessarily co-located: need to deliver the content over the WAN and/or cache it

Toward a sustainable, open and shared infrastructure

WLCG launched a set of R&D projects to prototype such a data management infrastructure – and associated tools

Aims:

- Reduce the global cost of storage (hw and operations)
- Enable a more effective use of existing storage
- Efficiently and scalably deliver data to large, remote, heterogenous, compute resources

Build a common set of DM tools based on open and standard protocols that can be used by a broad set of scientific experiments

Evolution of the Authentication/Authorisation towards federated identities and token-based systems in line with most modern network services

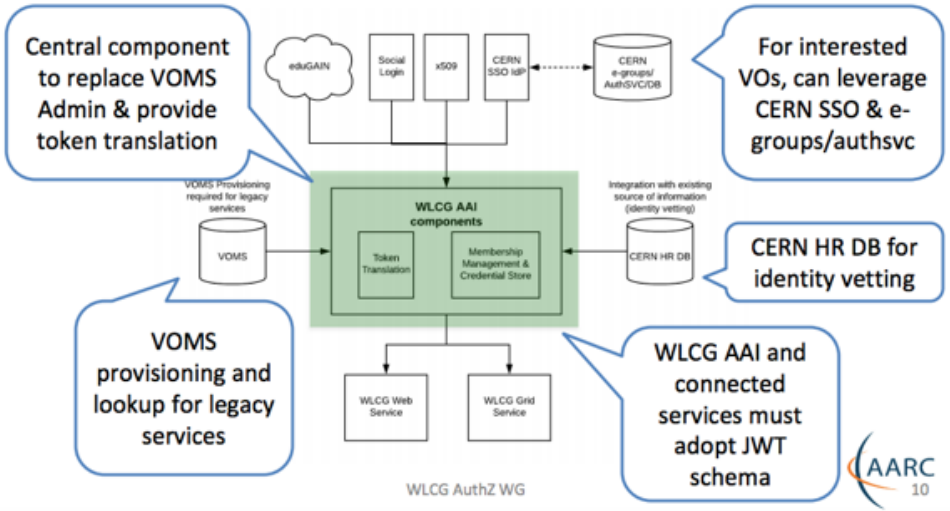

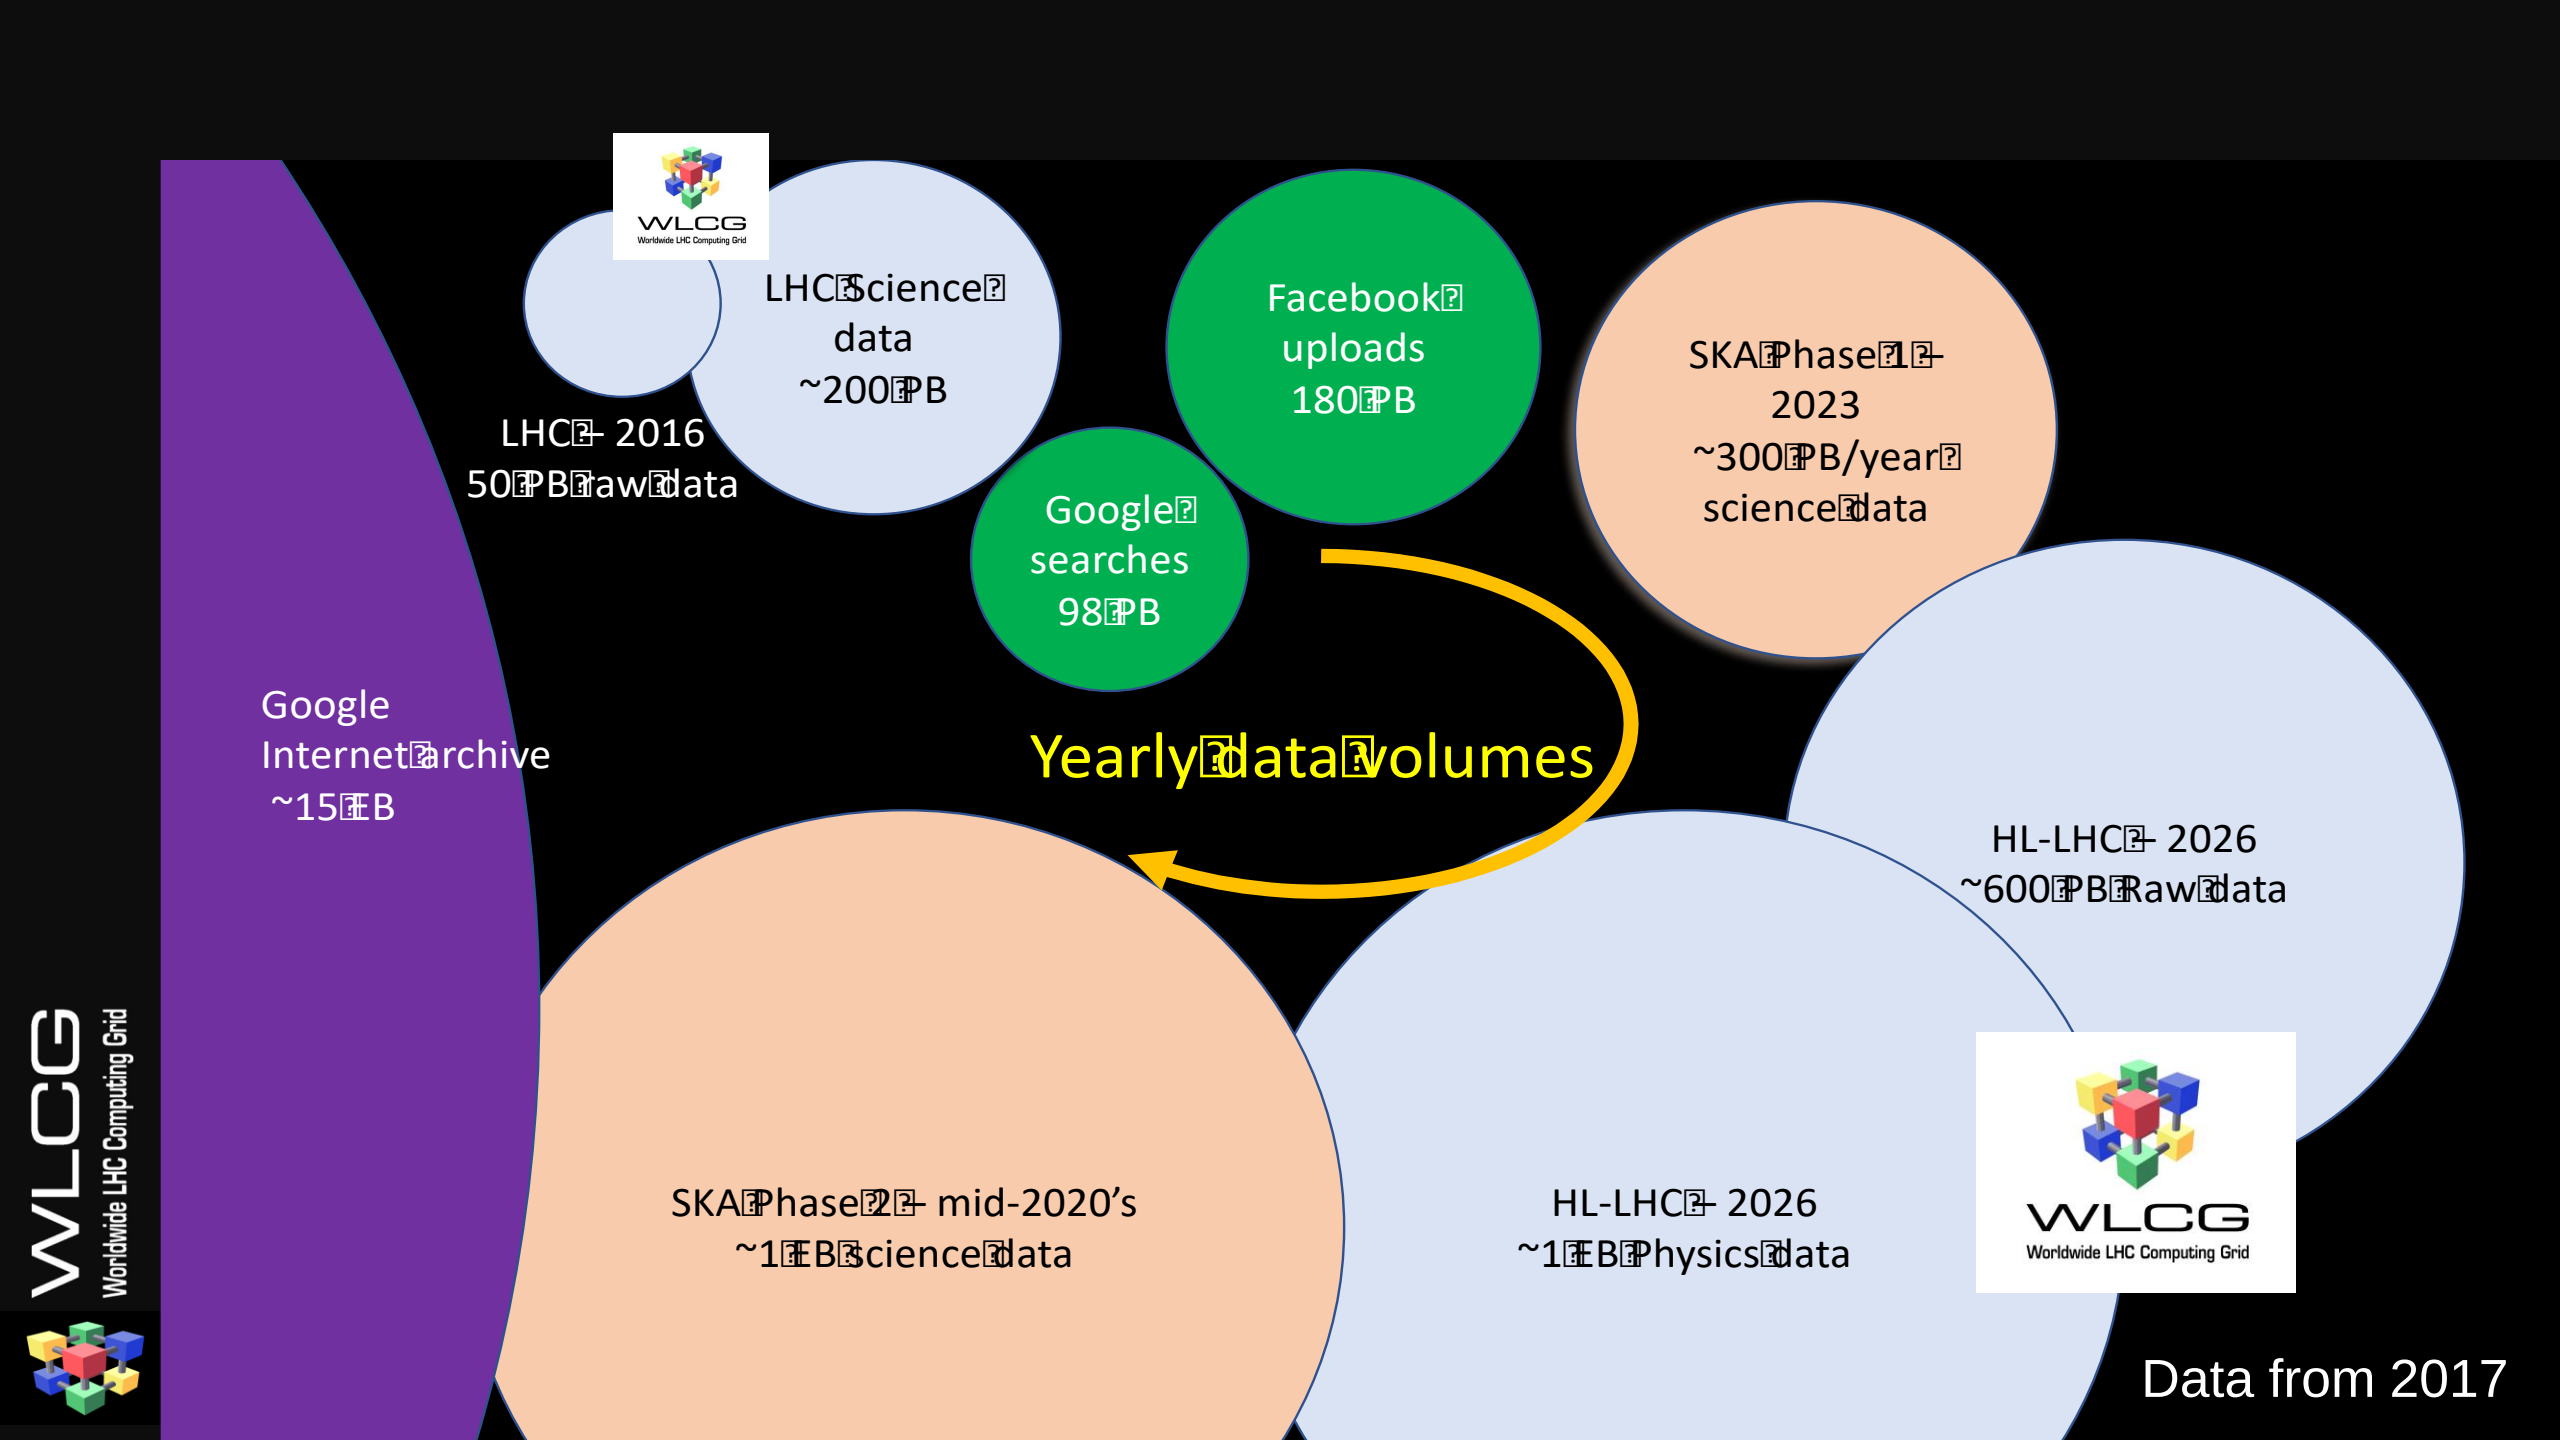
LHC - 2016
50 PB raw data
LHC Science data
~200 PB
Facebook uploads
180 PB
Google searches
98 PB
SKA Phase 1 - 2023
~300 PB/year science data
Yearly data volumes
Google Internet archive
~15 EB
HL-LHC - 2026
~600 PB Raw data
SKA Phase 2 - mid-2020's
~1 EB science data
HL-LHC - 2026
~1 EB Physics data
WLCG
Worldwide LHC Computing Grid
Data from 2017

Collaboration with data intensive sciences

CERN COURIER

Aug 11, 2017

SKA and CERN co-operate on extreme computing

Big-data co-operation agreement

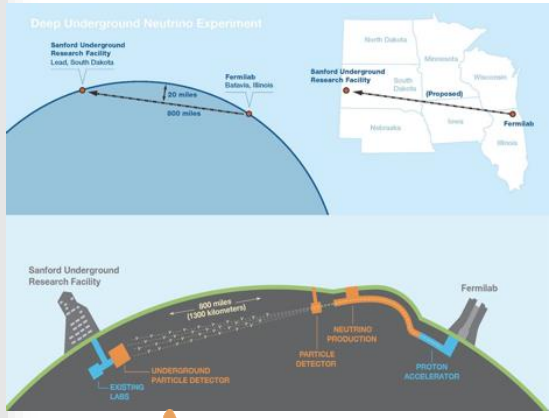

DUNE DEEP UNDERGROUND NEUTRINO EXPERIMENT

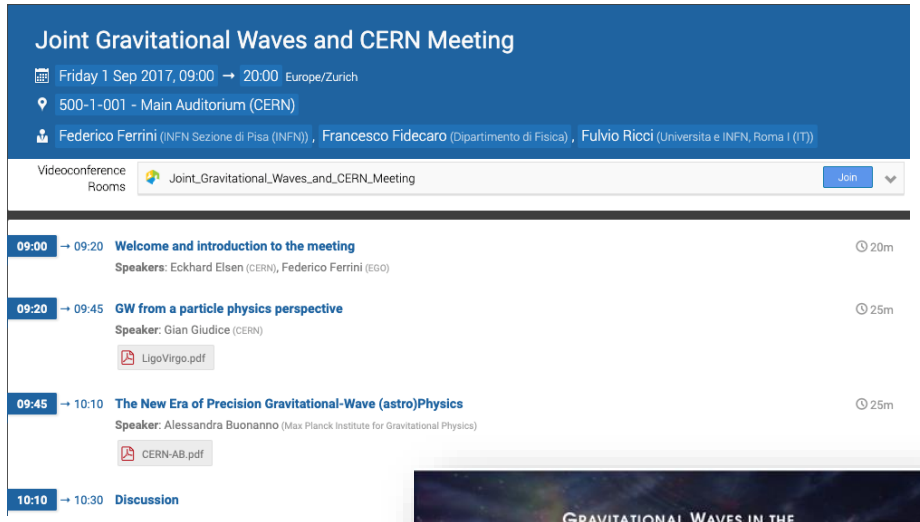
Joint Gravitational Waves and CERN Meeting

Friday 1 Sep 2017, 09:00 → 20:00 Europe/Zurich

500-1-001 - Main Auditorium (CERN)

Federico Ferrini (INFN Sezione di Pisa (INFN)), Francesco Fidecaro (Dipartimento di Fisica), Fulvio Ricci (Universita e INFN, Roma I (IT))

Videoconference Rooms: Joint_Gravitational_Waves_and_CERN_Meeting — Join

09:00 → 09:20 **Welcome and introduction to the meeting** 20m
Speakers: Eckhard Elsen (CERN), Federico Ferrini (EGO)

09:20 → 09:45 **GW from a particle physics perspective** 25m
Speaker: Gian Giudice (CERN)
LigoVirgo.pdf

09:45 → 10:10 **The New Era of Precision Gravitational-Wave (astro)Physics** 25m
Speaker: Alessandra Buonanno (Max Planck Institute for Gravitational Physics)
CERN-AB.pdf

10:10 → 10:30 **Discussion**

Belle II

GRAVITATIONAL WAVES IN THE EUROPEAN STRATEGY FOR PARTICLE PHYSICS

ET Steering Committee
http://www.et-gw.eu/

with the key contribution of the
'Gravitational Wave International Committee
3G subcommittee'
https://gwic.ligo.org/3Gsubcomm/

Contact Person: Dr. Michele Punturo
email: michele.punturo@pg.infn.it

Abstract

This document briefly describes some of the scientific and technological synergies that are possible between the nascent field of Gravitational Waves (GWs) and High Energy Particle Physics (HEPP). It is submitted by the ET steering committee under the supervision of GWIC-3G (a team of the Gravitational Wave International Committee (GWIC)) as contribution to the European Strategy for Particle Physics and in view of the submission of the Einstein telescope (ET) observatory project to the ESFRI Roadmap.

DECEMBER 2018

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

Belle-2 and DUNE leverage the same infrastructure as WLCG

First formal non-LHC “associate” members of WLCG

WLCG Worldwide LHC Computing Grid

Toward a sustainable, open and shared infrastructure

ESCAPE
European Science Cluster of Astronomy & Particle physics
ESFRI research infrastructures

Horizon 2020
funded project

EUROPEAN OPEN SCIENCE CLOUD

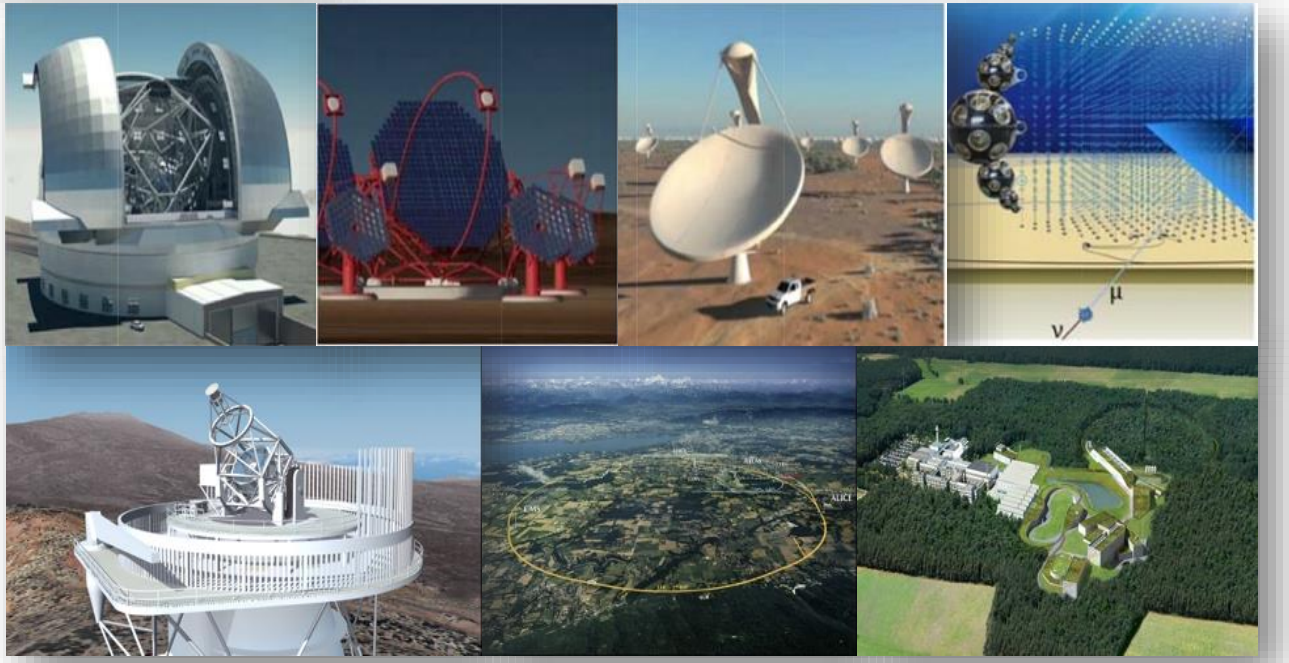

Goals:

Prototype an infrastructure adapted to the Exabyte-scale needs of the large science projects.

Driven by the sciences

Address *FAIR* data management

Science Projects

HL-LHC	SKA
FAIR	CTA
KM3Net	JIVE-ERIC
ELT	EST
EURO-VO (LSST)	EGO-VIRGO (CERN,ESO)

Data centres: CERN, INFN, DESY, GSI, Nikhef, SURFSara, RUG, CCIN2P3, PIC, LAPP, INAF

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 those sciences.