The Worldwide LHC Computing Grid

Nils Høimyr
IT Department

LHC accelerator and detectors

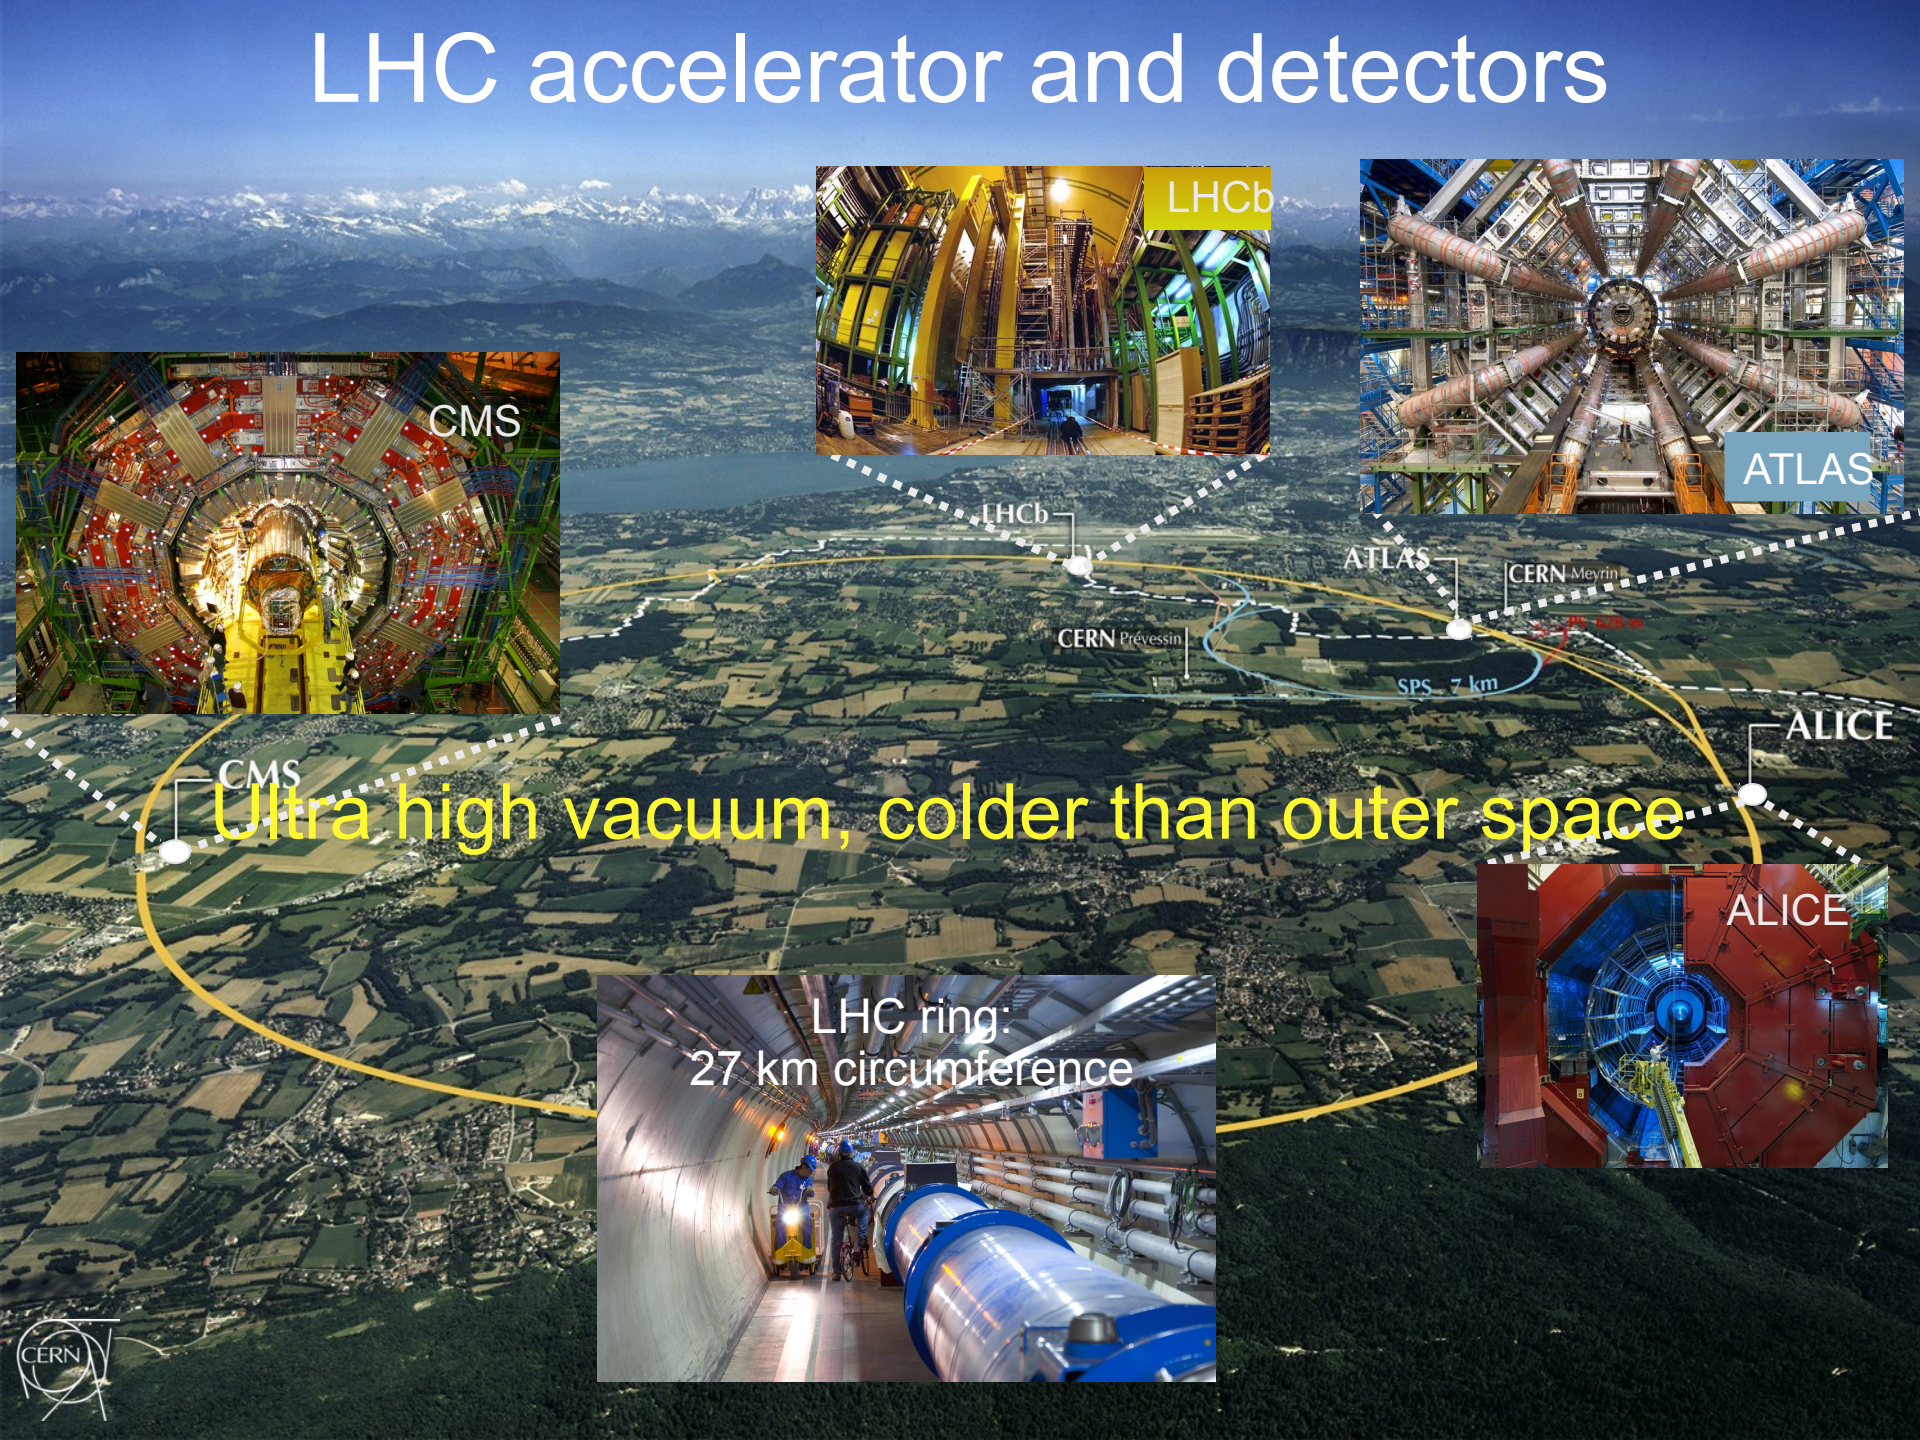

Ultra high vacuum, colder than outer space

The ATLAS experiment

CERN IT Department

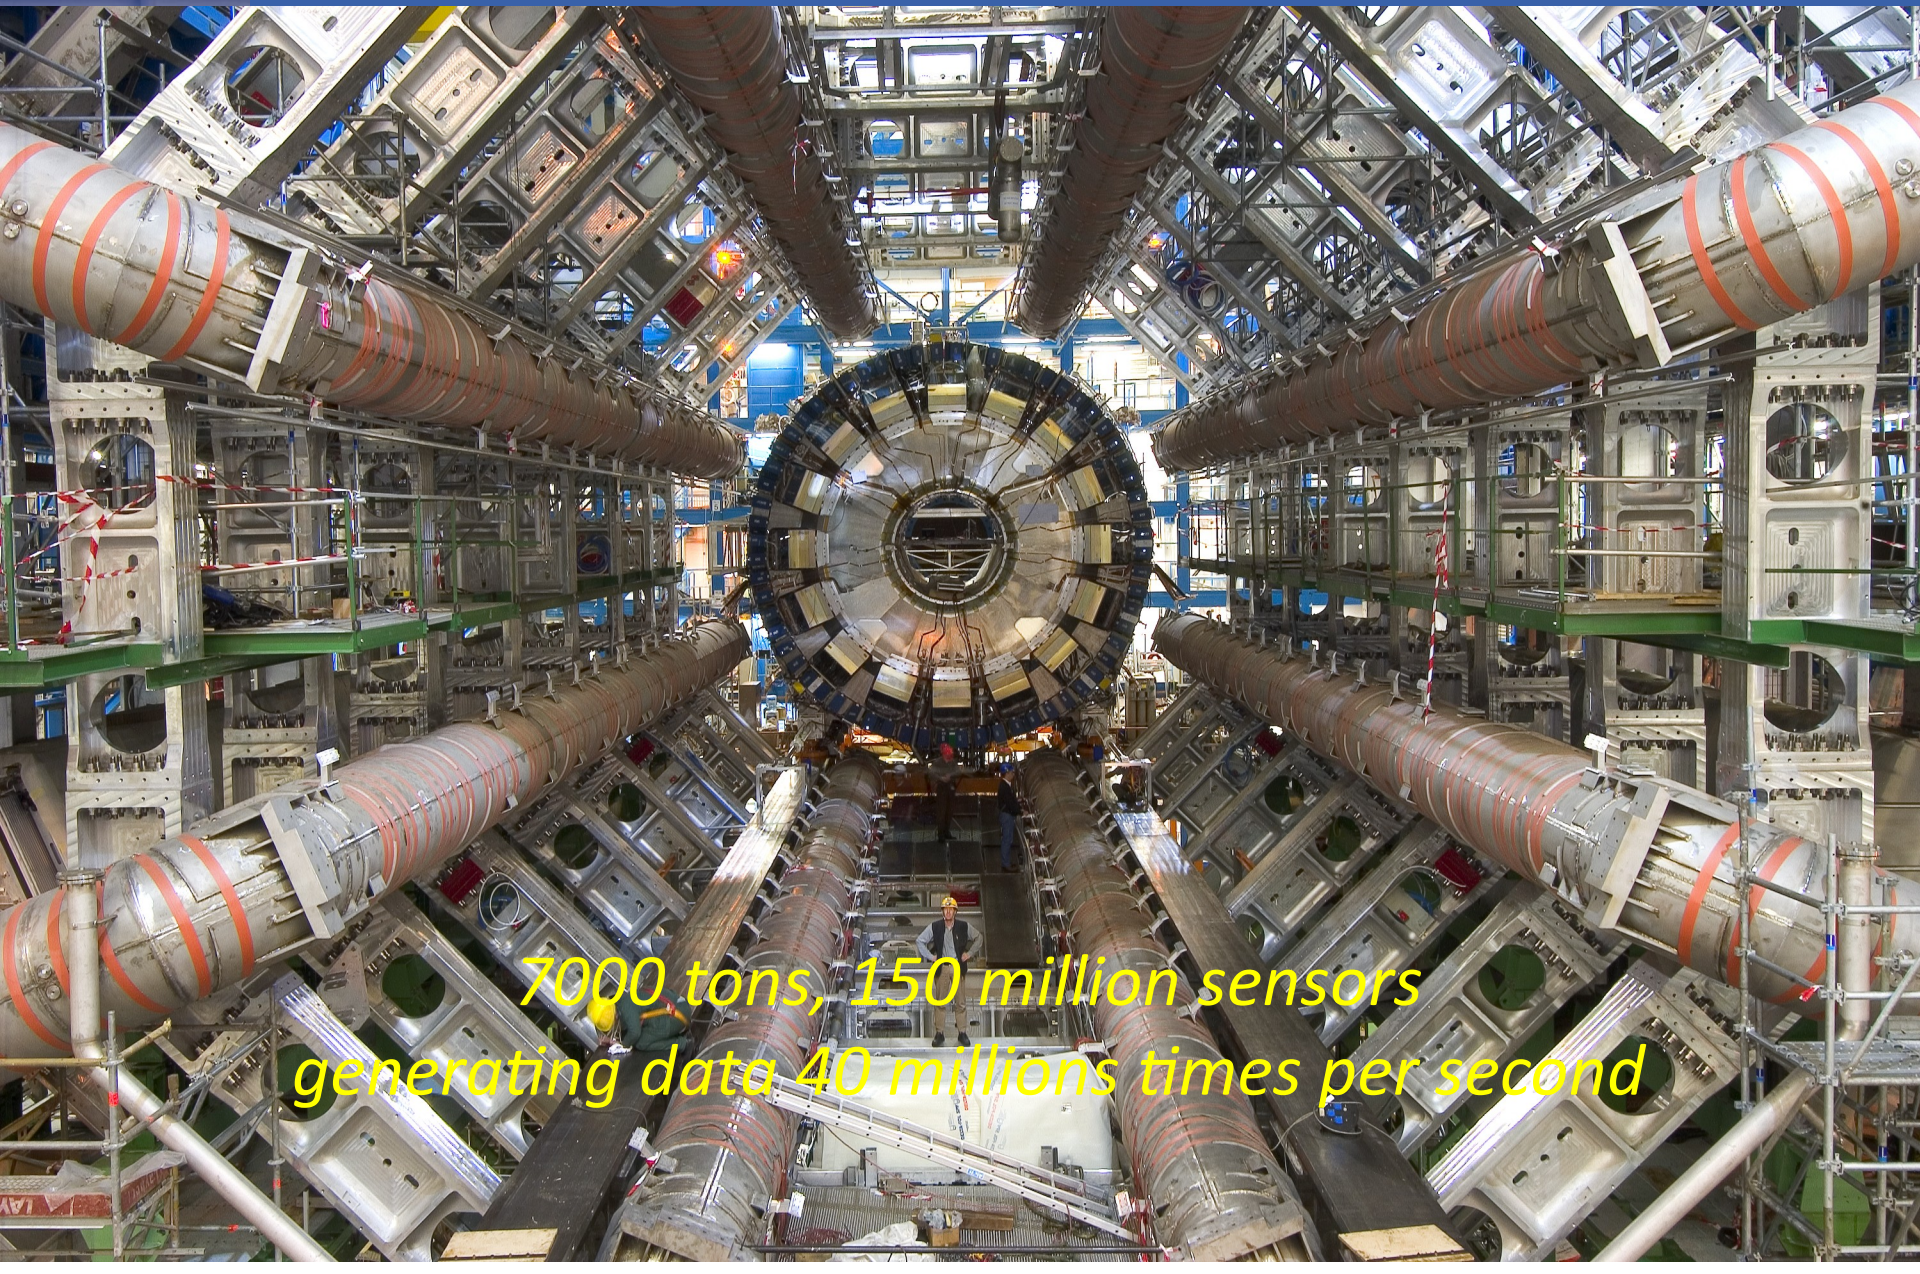

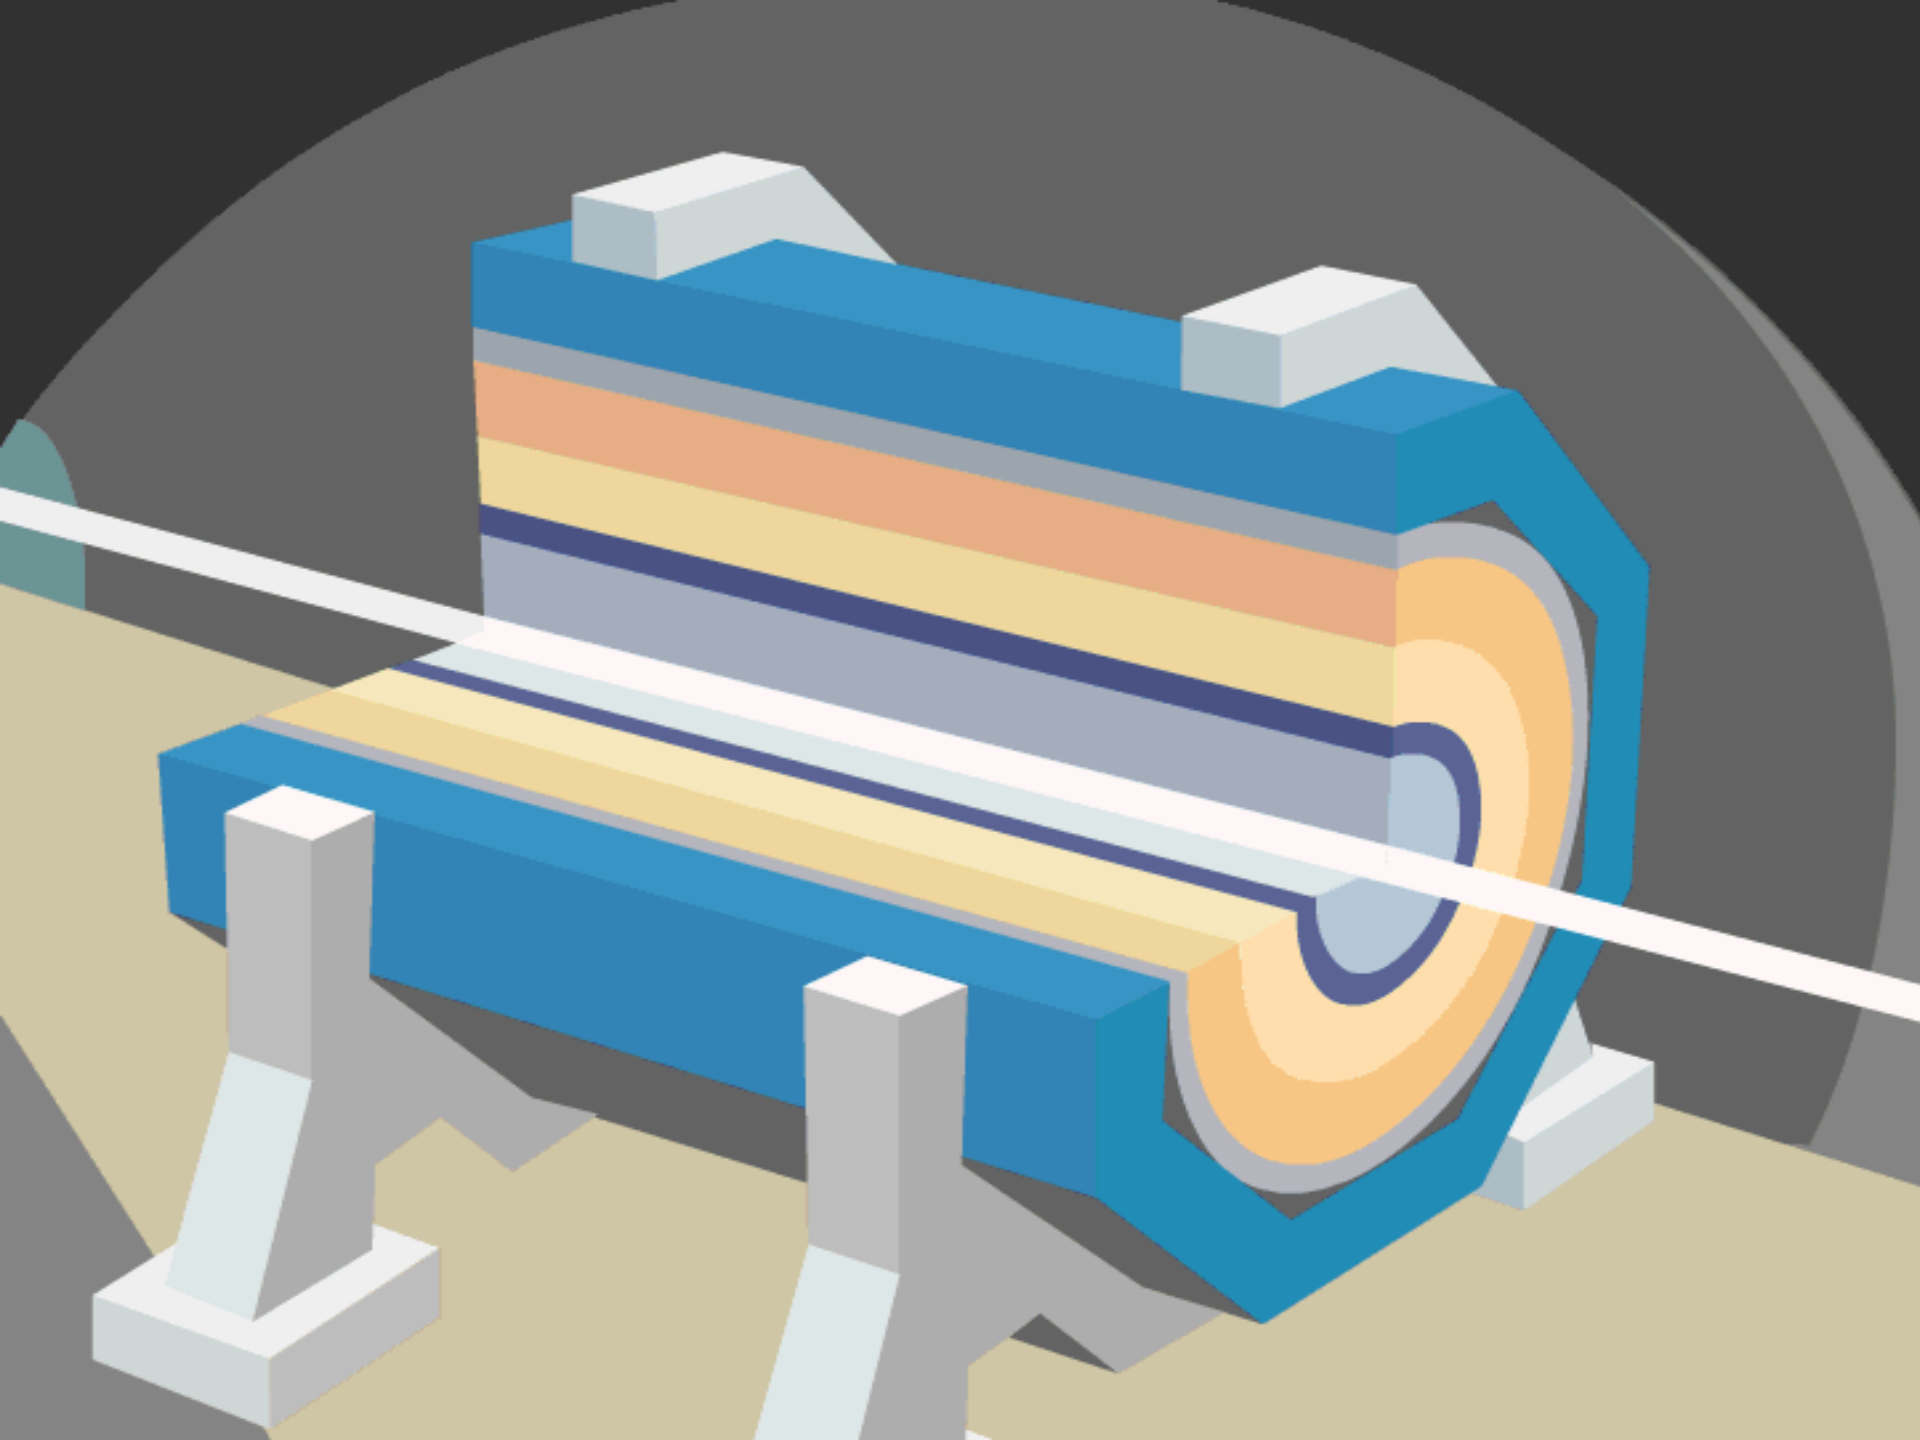

Collisions Produce 1PB/s

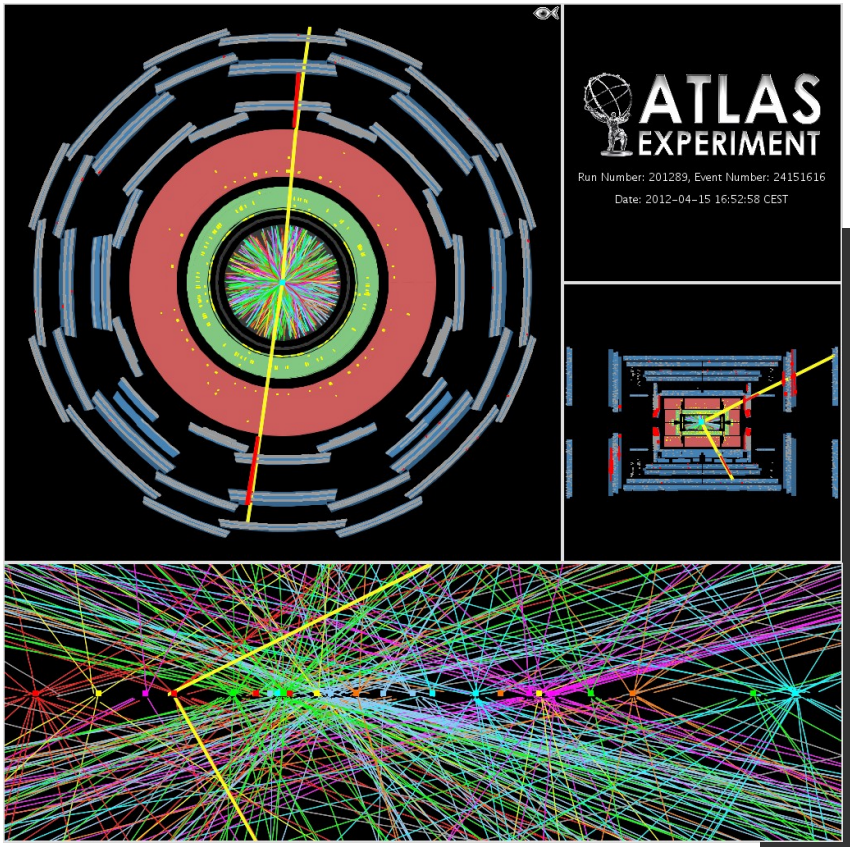

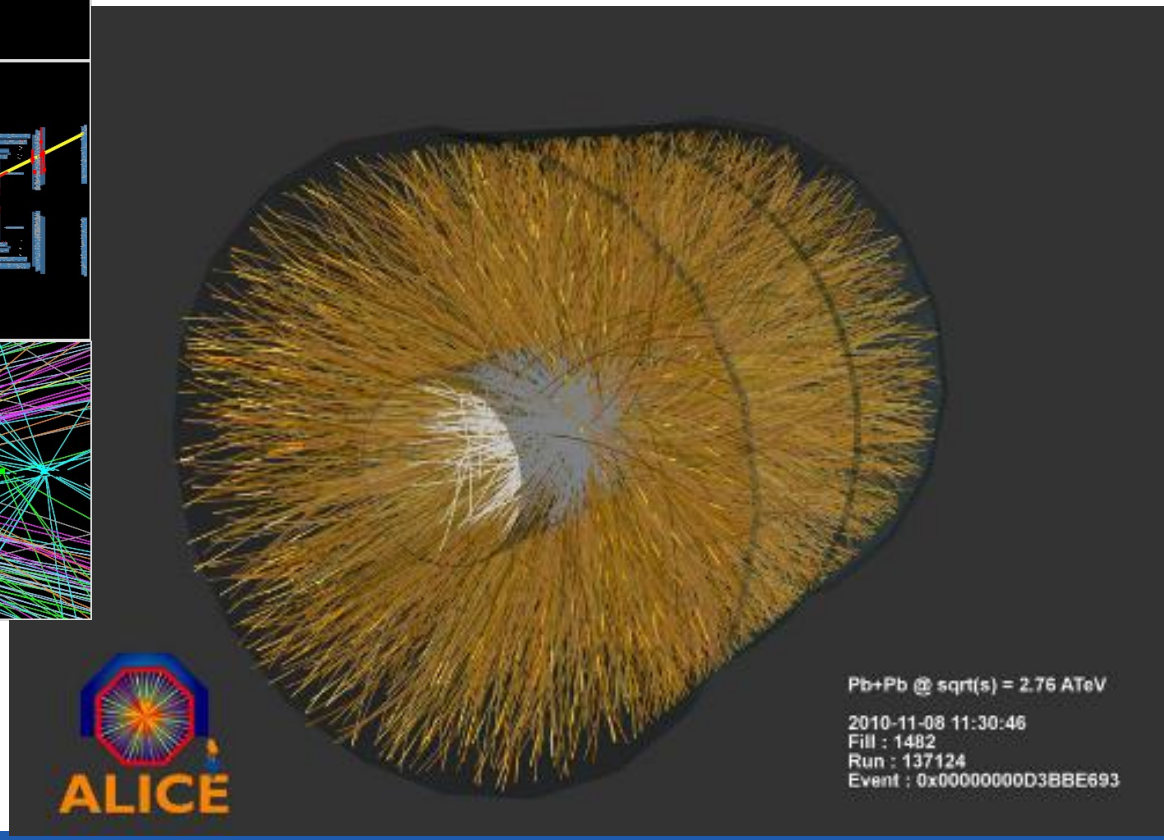

Pick the interesting events

- 40 million per second
  - Fast, simple information
  - Hardware trigger in a few micro seconds
- 100 thousand per second
  - Fast algorithms in local computer farm
  - Software trigger in <1 second
- Few 100 per second
  - Recorded for study

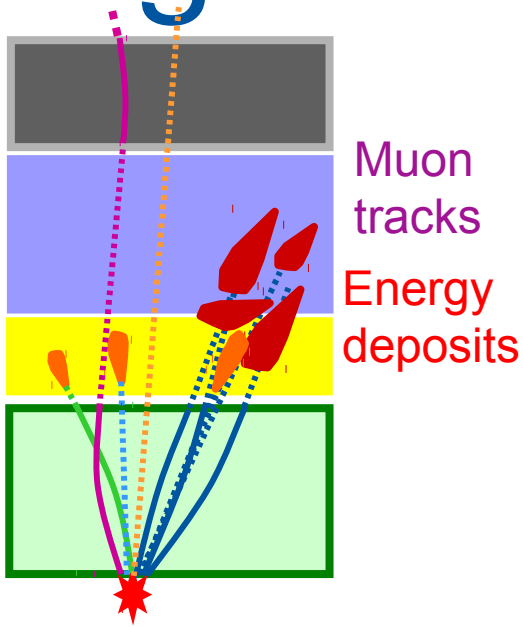

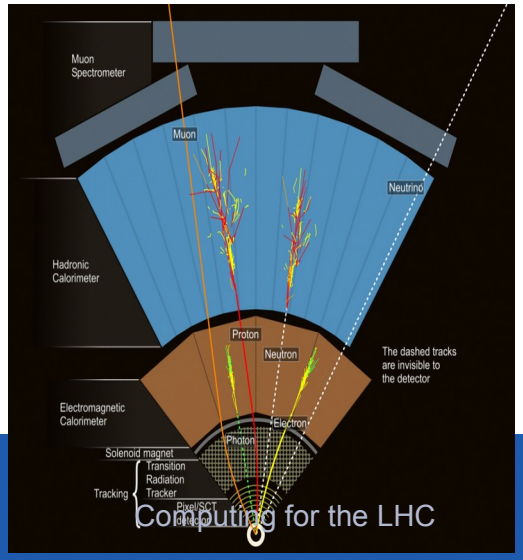

Pick the interesting events: Data size

- 40 million per second
  - Fast, simple information
  - Hardware trigger in a few micro seconds
- 100 thousand per second
  - Fast algorithms in computers
  - Software trigger
- Few 100 per second
  - Recorded for study

Ⓟ ~1 Petabyte per second?

- Cannot afford to store it
  - 1 year's worth of LHC data at 1 PB/s would cost few hundred *trillion* dollars/euros
- Have to filter in real time to keep only "interesting" data
- We keep 1 event in a million
  - Yes, 99.9999% is thrown away

Ⓟ >>6 Gigabytes per second

CERN Data Centre

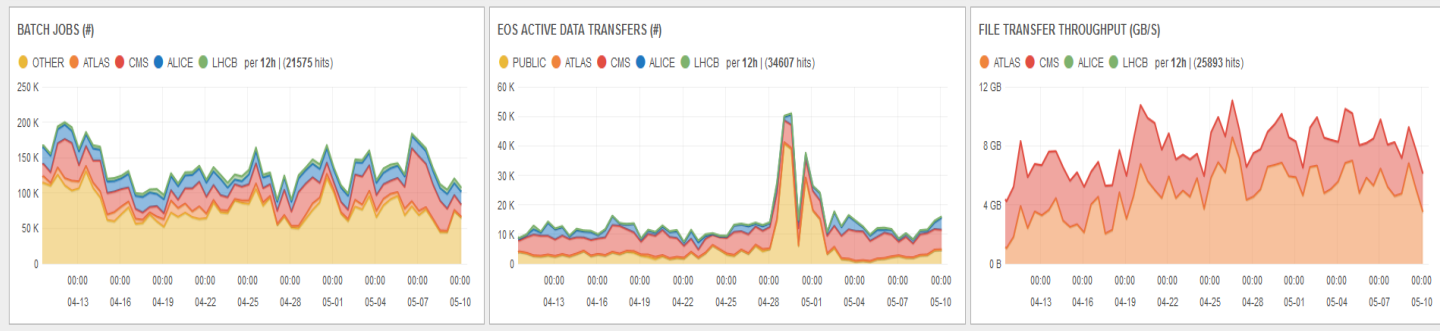

- Built in the 70s on the CERN site (Meyrin-Geneva), 3.5 MW for equipment
- Extension located at Wigner (Budapest), 2.7 MW for equipment
- Connected to the Geneva CC with 3x100Gb links (24 ms RTT)
- Hardware generally based on commodity
- 15,000 servers, providing 190,000 processor cores
- 80,000 disk drives providing 250 PB disk space
- 104 tape drives, providing 140 PB

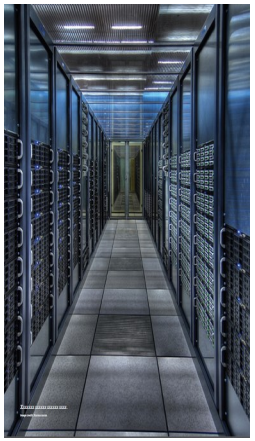

Worldwide computing

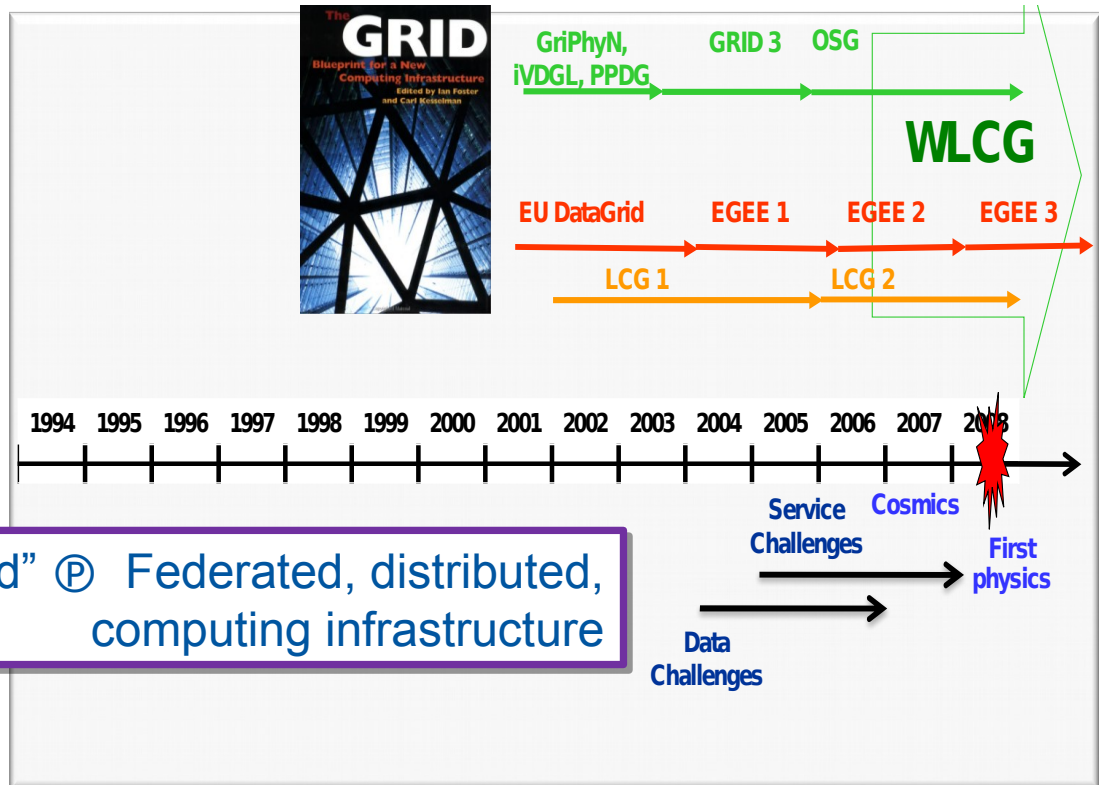

"Grid" ® Federated, distributed, computing infrastructure

WLCG for LHC

Public national and international grid infrastructures

- When we started LHC computing (~2001)
  - There were no internet companies, **no cloud computing** – Google was a search engine, Amazon, etc. did not exist
- We had to invent all of the tools from scratch
  - At CERN we had no tools to manage a data centre at the scale we thought was needed (no commercial or OS tools existed)
    - Initial tools developed through EU Data Grid – **Open Source** from the beginning
- Grid ideas from computer science did not work in the real world at any reasonable scale
  - We (EU, US, LHC grid projects) had to make them work at scale
  - We had to invent trust networks to convince funding agencies to open their resources to federated users
- Our users were not convinced that any of this was needed ;-)

Data - 2018

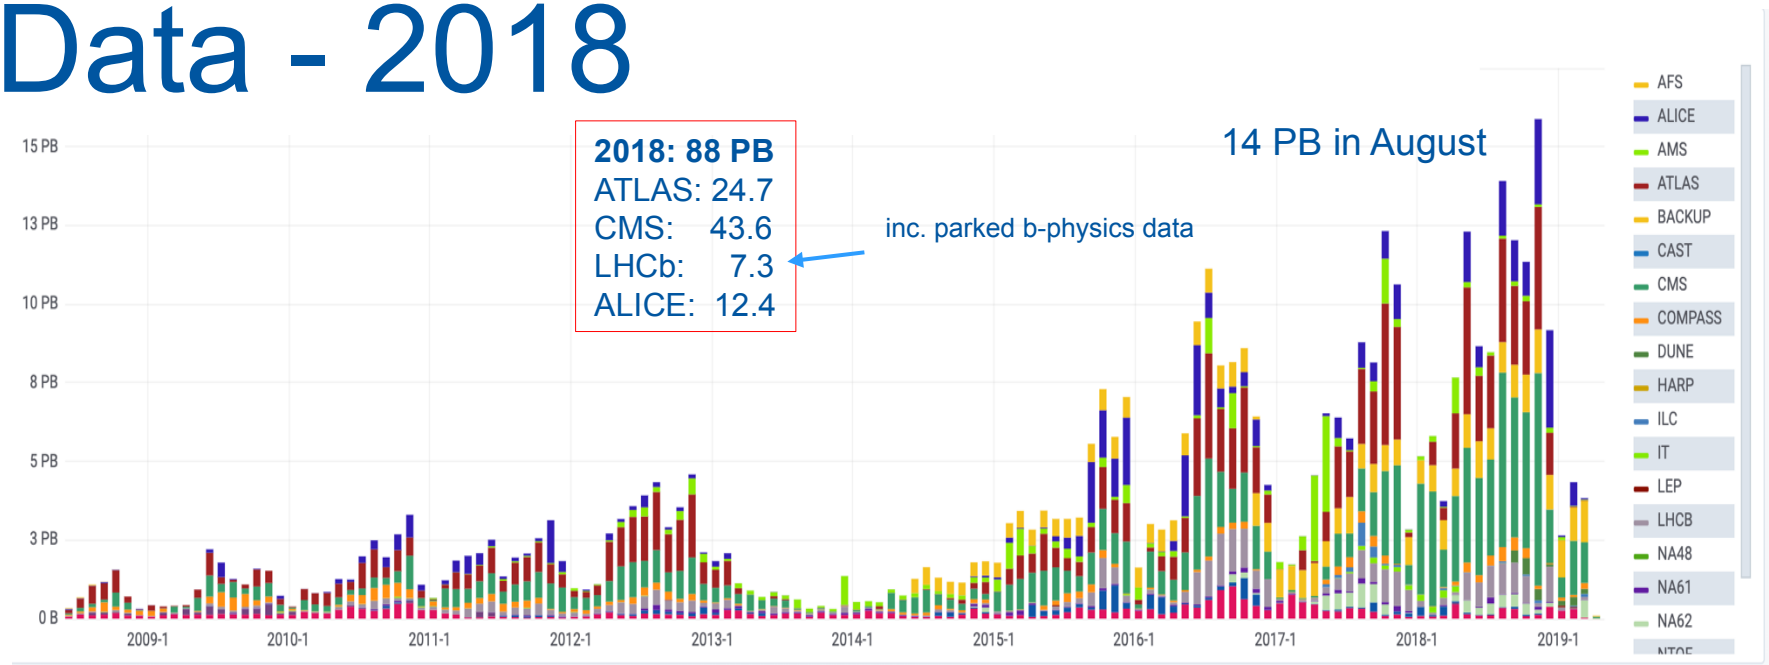

Data transfers

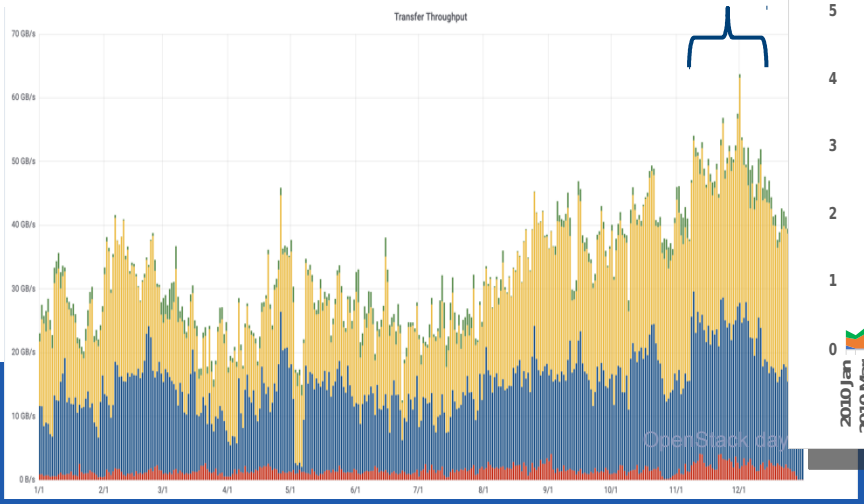

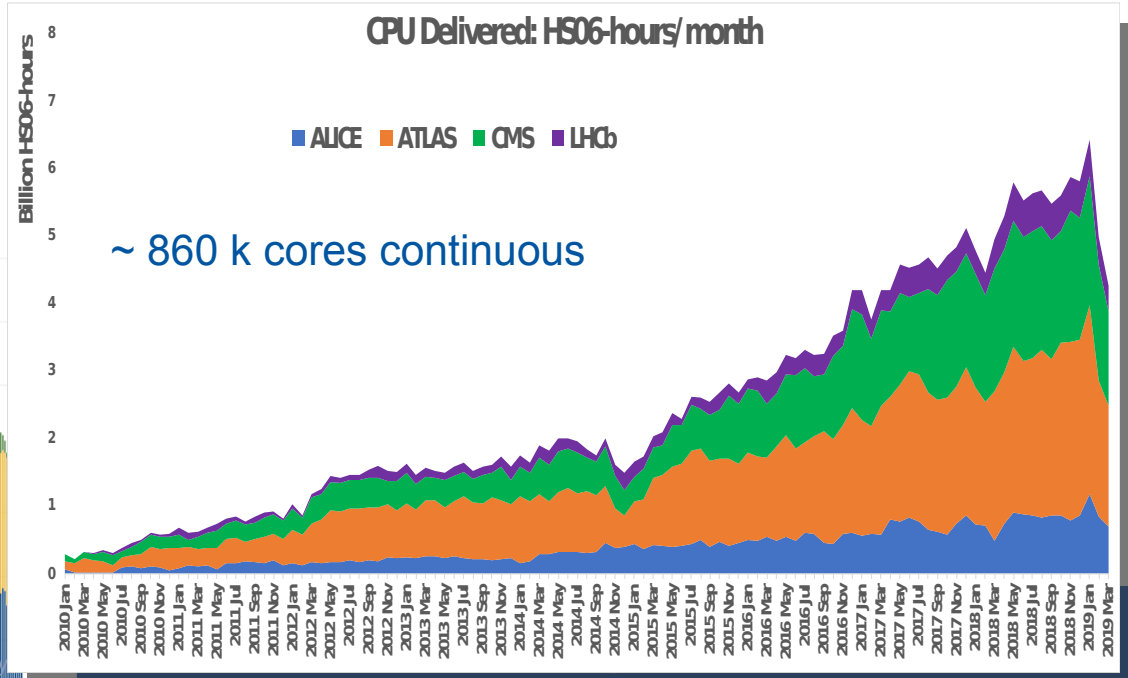

Heterogenous computing

- The majority of today's HEP processing is performed on dedicated clusters of commodity processors ("x86")
- Recently: opportunistic use of many types of compute, in particular HPC systems, and HLT
- In future, this heterogeneity will expand; we must be able to make use of all types:
  - Non-x86 (esp GPU), HPC, clouds, HLT farms (inc FPGA?)

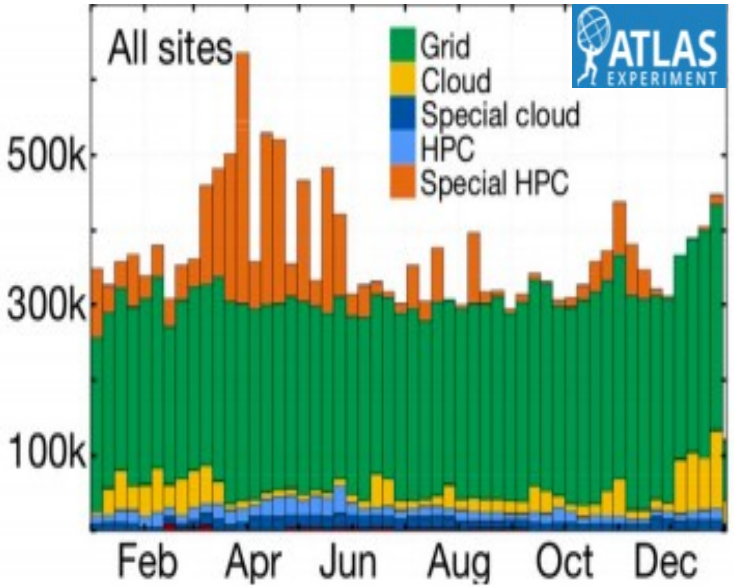

HPC Use is challenging
HEP engagement with DOE & NSF in USA and (together with SKA) with PRACE and EuroHPC in Europe and participating in BDEC2 workshops

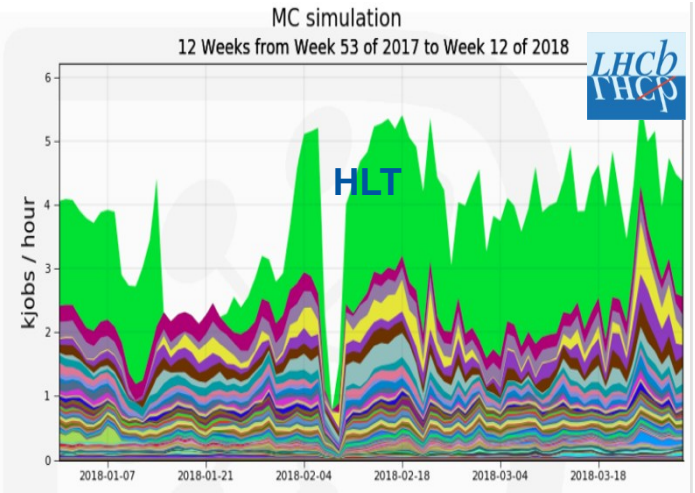

Heterogenous compute

- Requires:
  - Common provisioning mechanisms, transparent to users
  - Facilities able to control access (cost), appropriate use, etc
- HPC, Clouds, HLT will not have (affordable) local storage service (in the way we assume today)
  - Must be able to deliver data to them when they are in active use

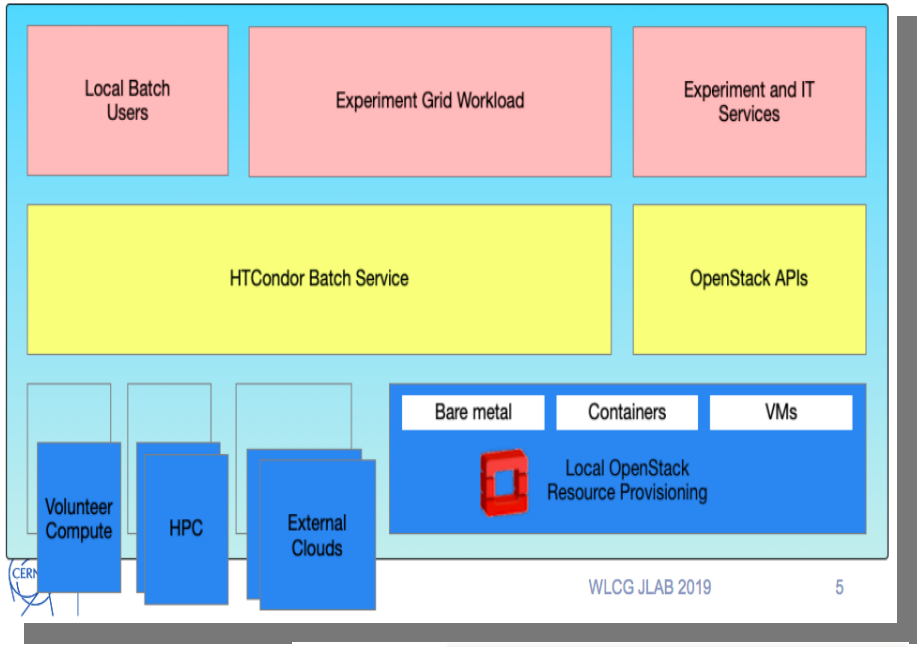

HELIX NEBULA THE SCIENCE CLOUD

Deployed in a hybrid cloud mode:
- Procurers' data centres
- commercial cloud service providers
- GEANT network and EduGAIN Federated Identity Management

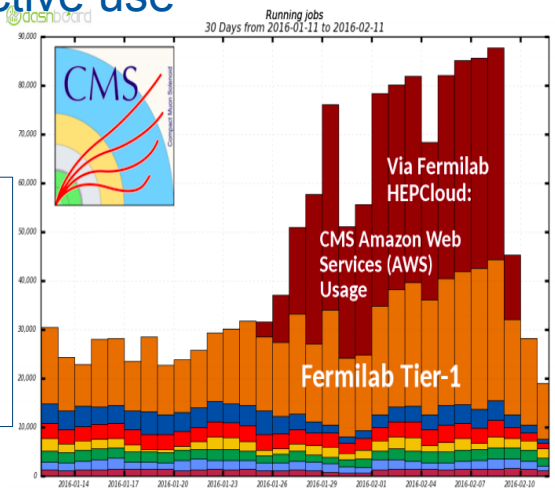

LHC Schedule

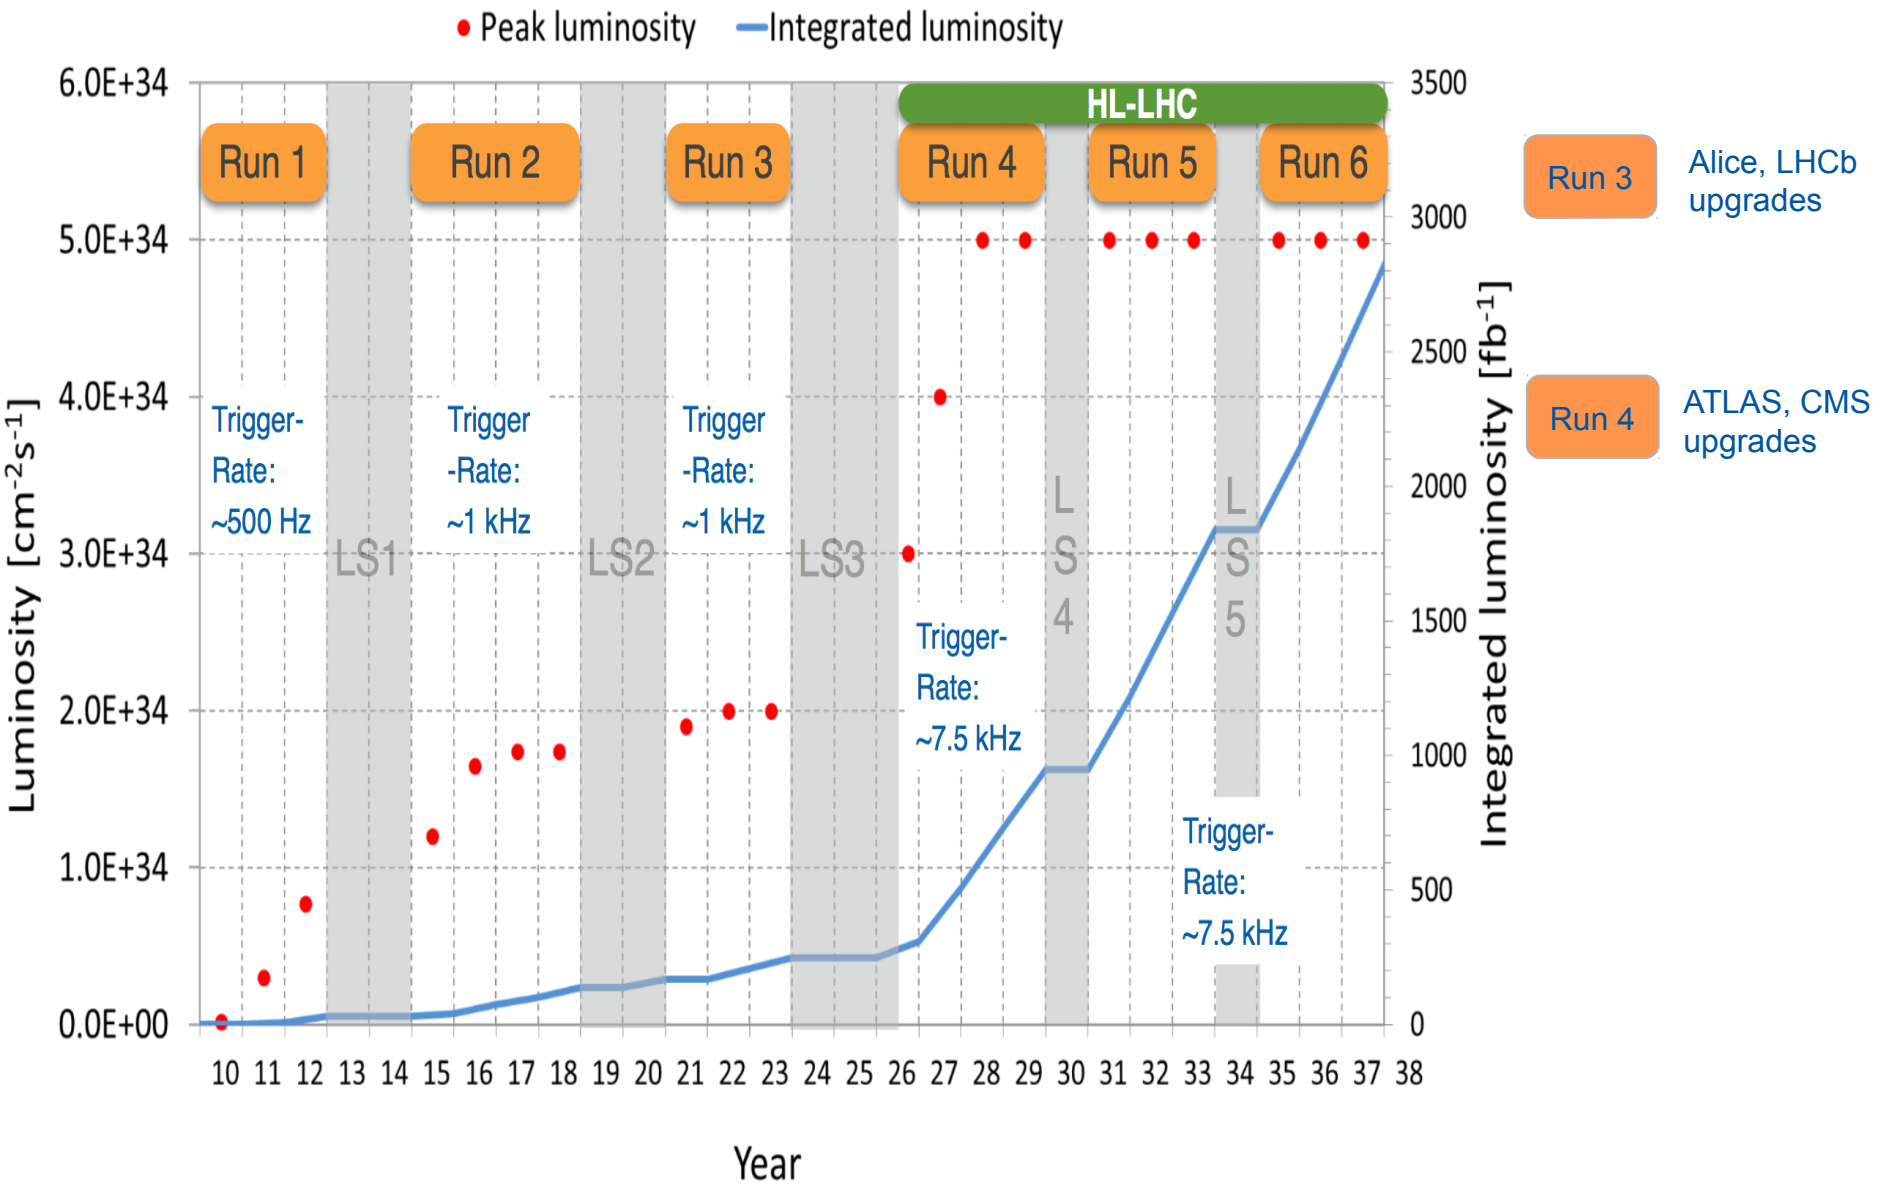

The HL-LHC computing challenge

- HL-LHC needs for ATLAS and CMS are above the expected hardware technology evolution (15% to 20%/yr) and funding (flat)
- The main challenge is storage, but computing requirements grow 20-50x

CMS: event from 2017 with 78 reconstructed vertices

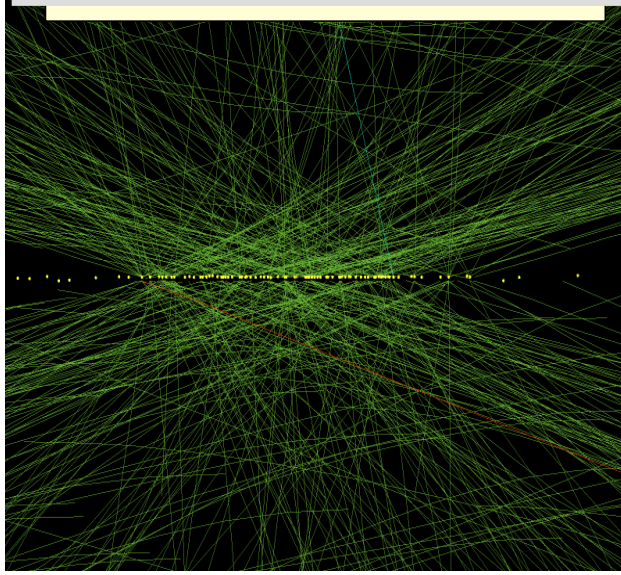

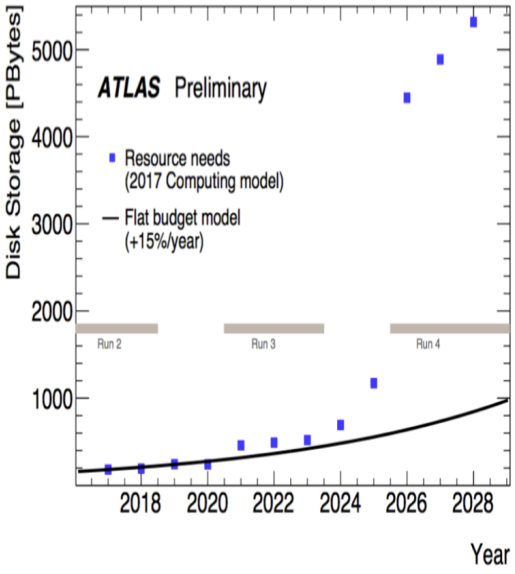

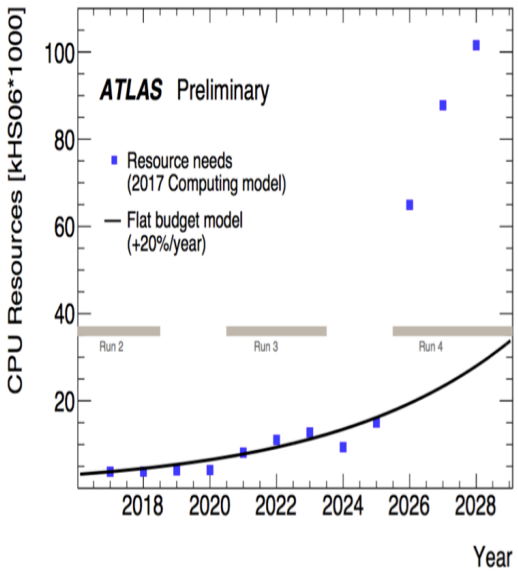

ATLAS: simulation for HL-LHC with 200 vertices

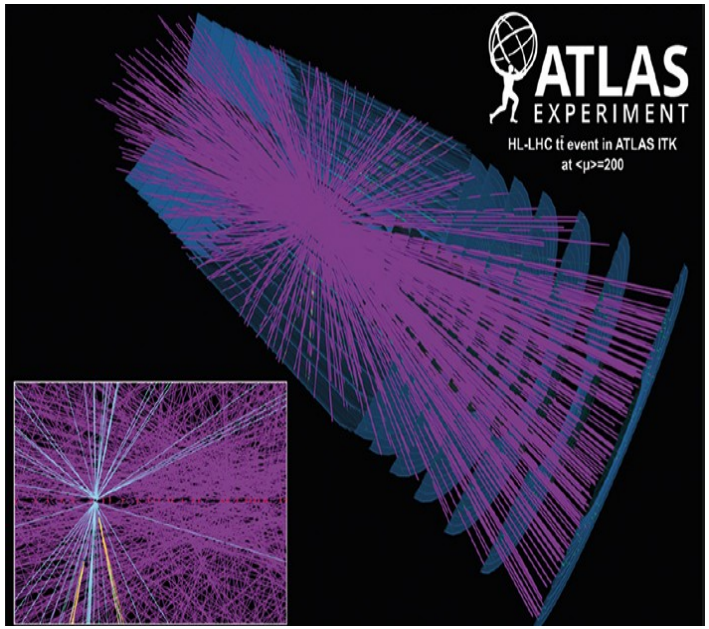

LHC – 2016
50 PB raw data

LHC Science data
~200 PB

Facebook uploads
180 PB

Google searches
98 PB

SKA Phase 1 – 2023
~300 PB/year science data

Google Internet archive
~15 EB

Yearly data volumes

HL-LHC – 2026
~600 PB Raw data

SKA Phase 2 – mid-2020's
~1 EB science data

HL-LHC – 2026
~1 EB Physics data

10 Billion of these

Many software challenges

- Improved algorithms, Machine Learning (ML)
  - “ML” as Neural Networks used for more than 20 years in HEP
  - A lot of development in the IT industry in this area, scope for re-use and improvements
- Vectorisation, GPUs, FPGAs, other processor architectures
- Data Analysis model and software changes
- Visualisation
- Storage and preservation

Grid vs Cloud

CERN IT Department

- "Cloud computing" everywhere
  - Web based solutions (http/https and RES)
  - Virtualisation, containers….
- GRID has mainly a scientific user base
  - Complex applications running across multiple sites, but works like a cluster batch system for the end user
  - Mainly suitable for parallel computing and massive data processing
- Technologies converging
  - "Internal Cloud" at CERN – OpenStack
  - Xbatch – extending to external cloud providers
  - CernVM – virtual machine running e.g. at Amazon
  - Google Cloud Kubernetes Higgs Challenge
  - "Volunteer Cloud" - LHC@home 2.0

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Volunteer grid - LHC@home

CERN IT Department

LHC@home

- **LHC volunteer computing**
  - Allows us to get additional computing resources for e.g. accelerator physics and theory simulations
- **Based on BOINC**
  - **"B**erkeley **O**pen **I**nfrastructure for **N**etwork **C**omputing"
  - Software platform for distributed computing using volunteered computer resources
  - Uses a volunteer PC's unused CPU cycles to analyse scientific data
  - Virtualization support - CernVM
  - Other well known projects
    - **SETI@Home**
    - **Climateprediction.net**
    - **Einstein@Home**

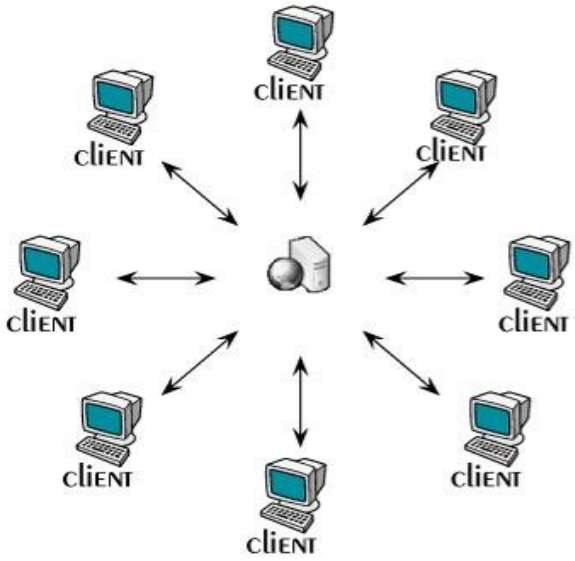

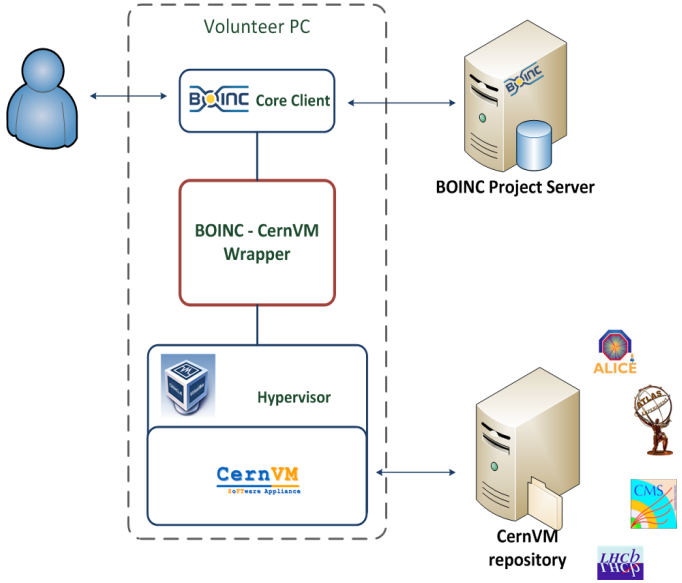

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

LHC@home
Volunteer computing for the LHC

You can help us!

- As a volunteer, you can help us by donating CPU when your computer is idle
- Connect with us on:
  - http://cern.ch/lhcathome

The Balance between Academic Freedom, Operations & Computer Security

http://cern.ch/security

CERN

Open Data

http://cds.cern.ch

http://opendata.cern.ch

http://zenodo.org

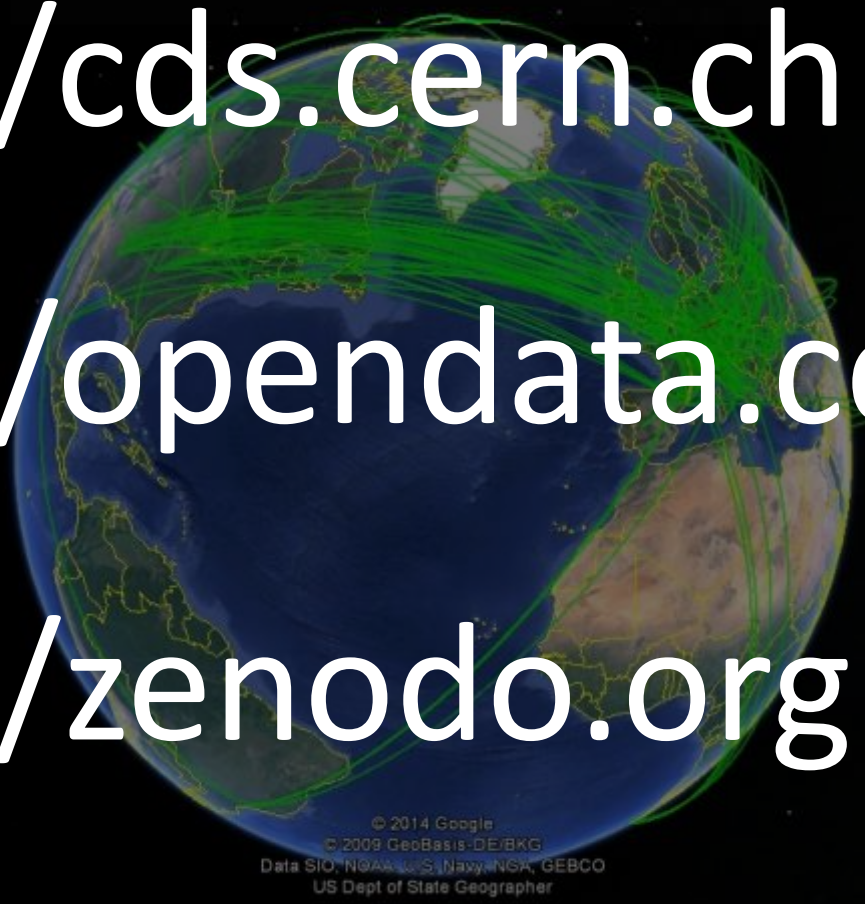

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Open Data – Open Knowledge

CERN & the LHC experiments have made the first steps towards Open Data (http://opendata.cern.ch/)

- Key drivers: Educational Outreach & Reproducibility
- Increasingly required by Funding Agencies
- Paving the way for Open Knowledge as envisioned by DPHEP (http://dphep.org)
  - ICFA Study Group on Data Preservation and Long Term Analysis in High Energy Physics

CERN has released Zenodo, a platform for Open Data as a Service (http://zenodo.org)[1]

- Building on experience of Digital Libraries & Extreme scale data management
- Targeted at the long tail of science
- Citable through DOIs, including the associated software
- Generated significant interest from open data publishers such as Wiley, Ubiquity, F1000, eLife, PLOS

[1]Initially cofunded by the EC FP7 OpenAire series of projects

Training

CERN School of Computing

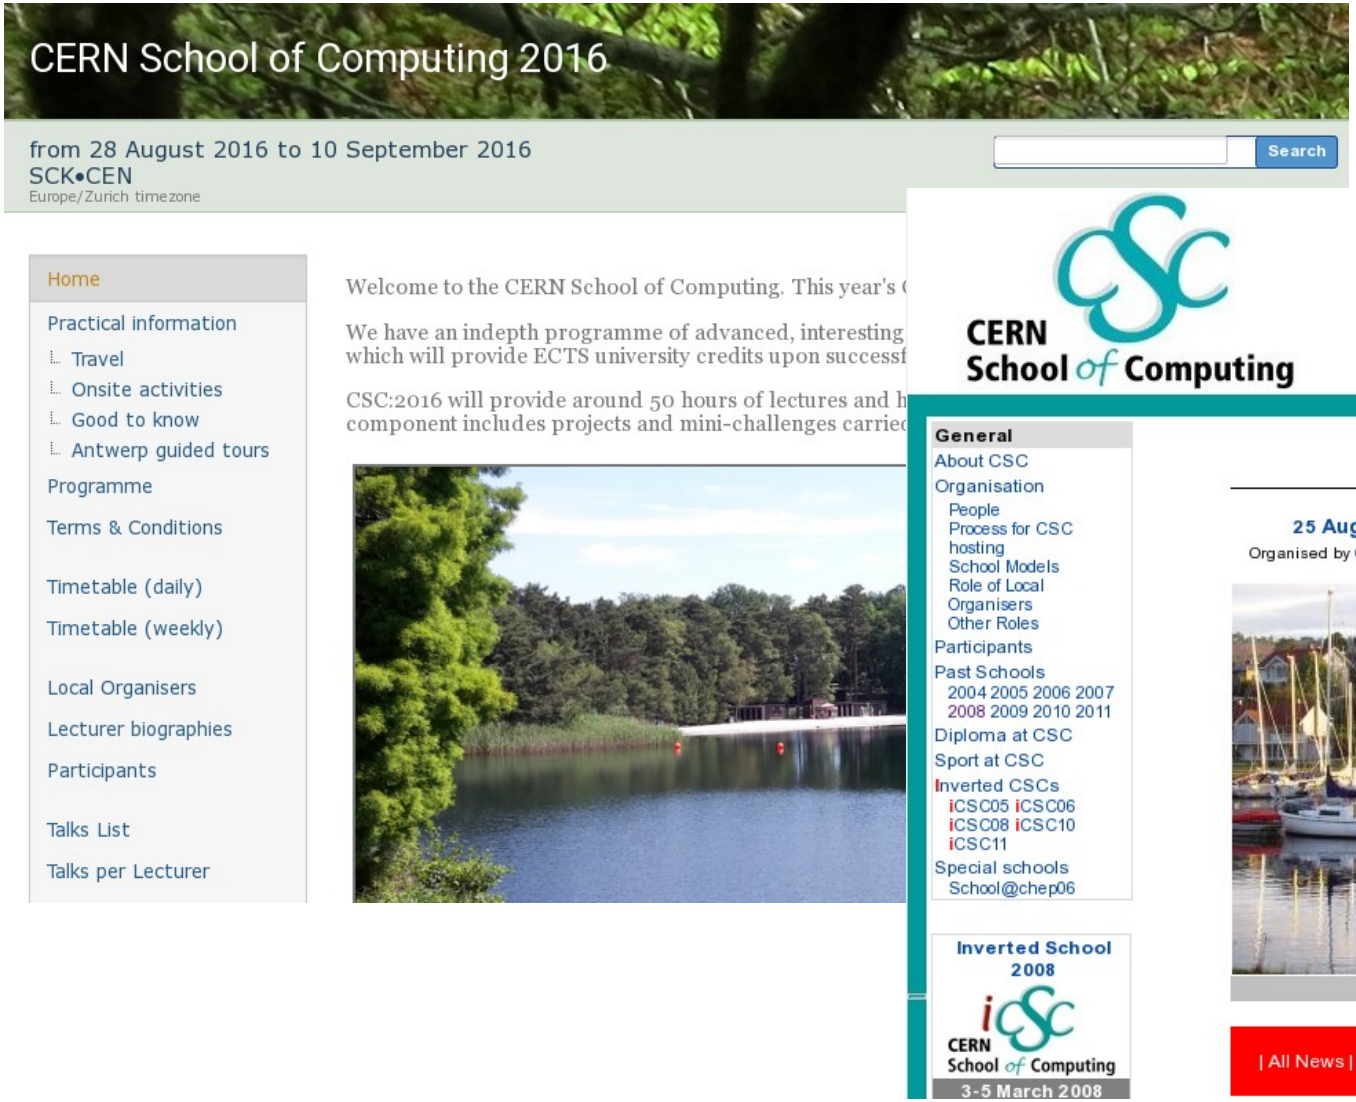

CERN openlab

- A science – industry partnership to drive R&D and innovation with over a decade of success
- Evaluate state-of-the-art technologies in a challenging environment and improve them
- Test in a research environment today what will be used in many business sectors tomorrow
- Train next generation of engineers/employees
- Disseminate results and outreach to new audiences

Our **collaborators**

Partners

Google, Micron, intel, SIEMENS, ORACLE

Contributors

be|studys, IBM, E4 COMPUTER ENGINEERING, EXTREME PARTNER NETWORK

Associates

COMTRADE, open systems

Research members

TU/e, Fermilab, EMBL-EBI, Newcastle University, King's College London, IVI INNOVATION VALUE INSTITUTE, INFN

IT at CERN – more than the Grid

CERN IT Department

- **Physics computing – Grids (this talk!)**
- **Administrative information systems**
  - Financial and administrative management systems, e-business...
- **Desktop and office computing**
  - Windows, Linux and Web infrastructure for day to day use
- **Engineering applications and databases**
  - CAD/CAM/CAE (Autocad, Catia, Cadence, Ansys etc)
  - A number of technical information systems based on Oracle, MySQL
- **Controls systems**
  - Process control of accelerators, experiments and infrastructure
- **Networks and telecom**
  - European IP hub, security, telephony software...

More information: *http://cern.ch/it*

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

LHCb
ATLAS
CERN Meyrin
CERN Prévessin
SPS 7 km
SUISSE
FRANCE
CMS
ALICE
LHC 27 km
Thank You!
CERN
Accelerating Science and Innovation
CERN/IT Nils Høimyr