The Worldwide LHC Computing Grid

Nils Høimyr
IT Department

LHC accelerator and detectors

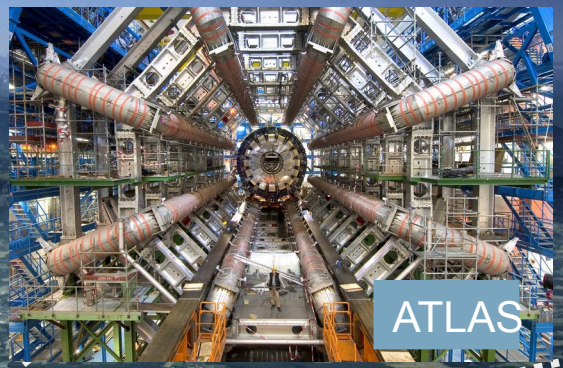

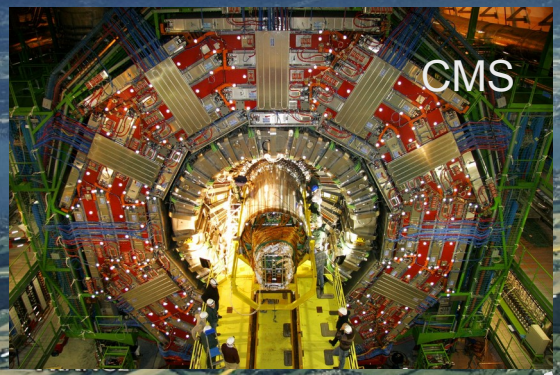

Ultra high vacuum, colder than outer space

The ATLAS experiment

CERN IT Department

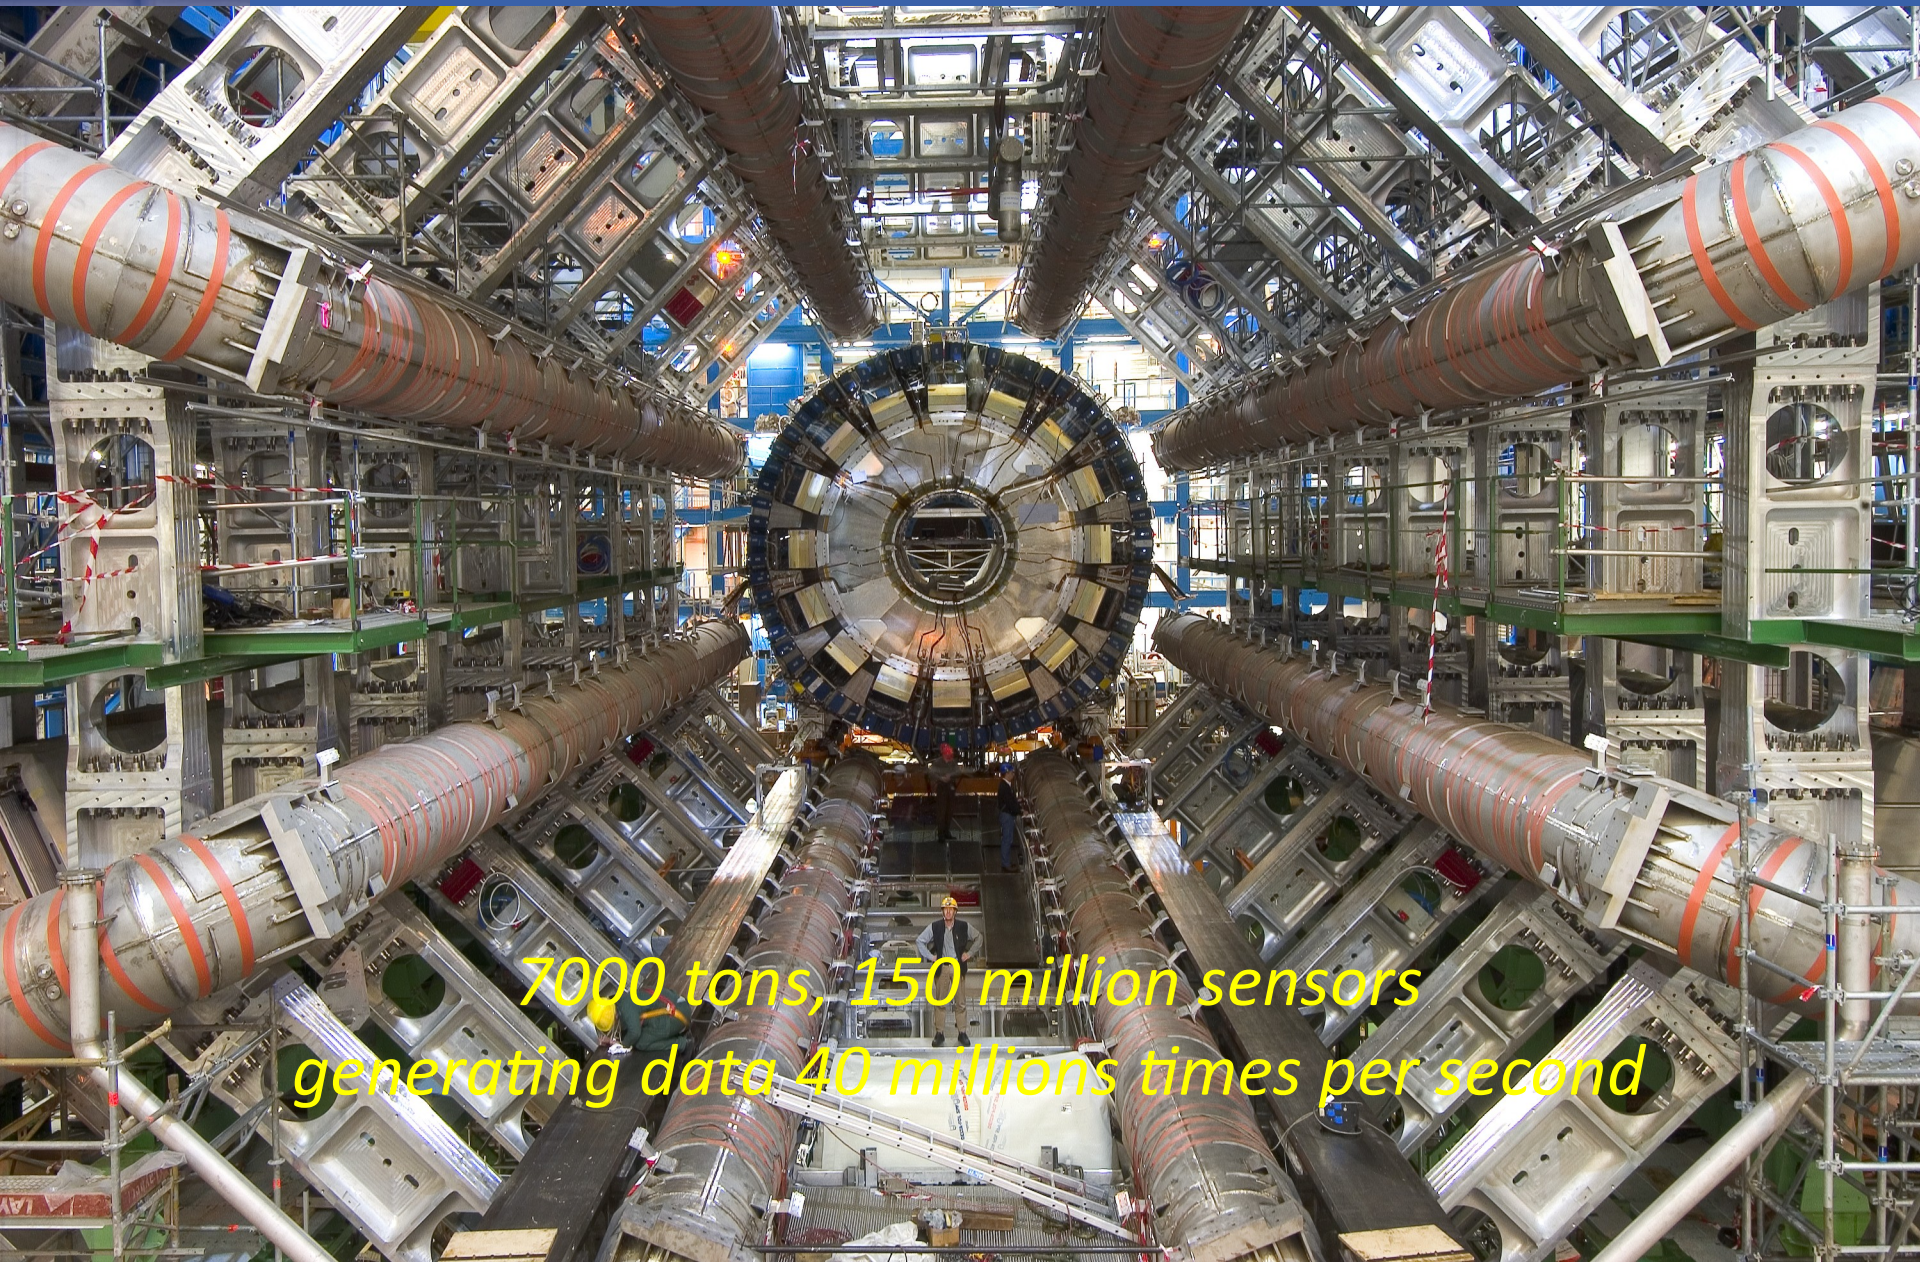

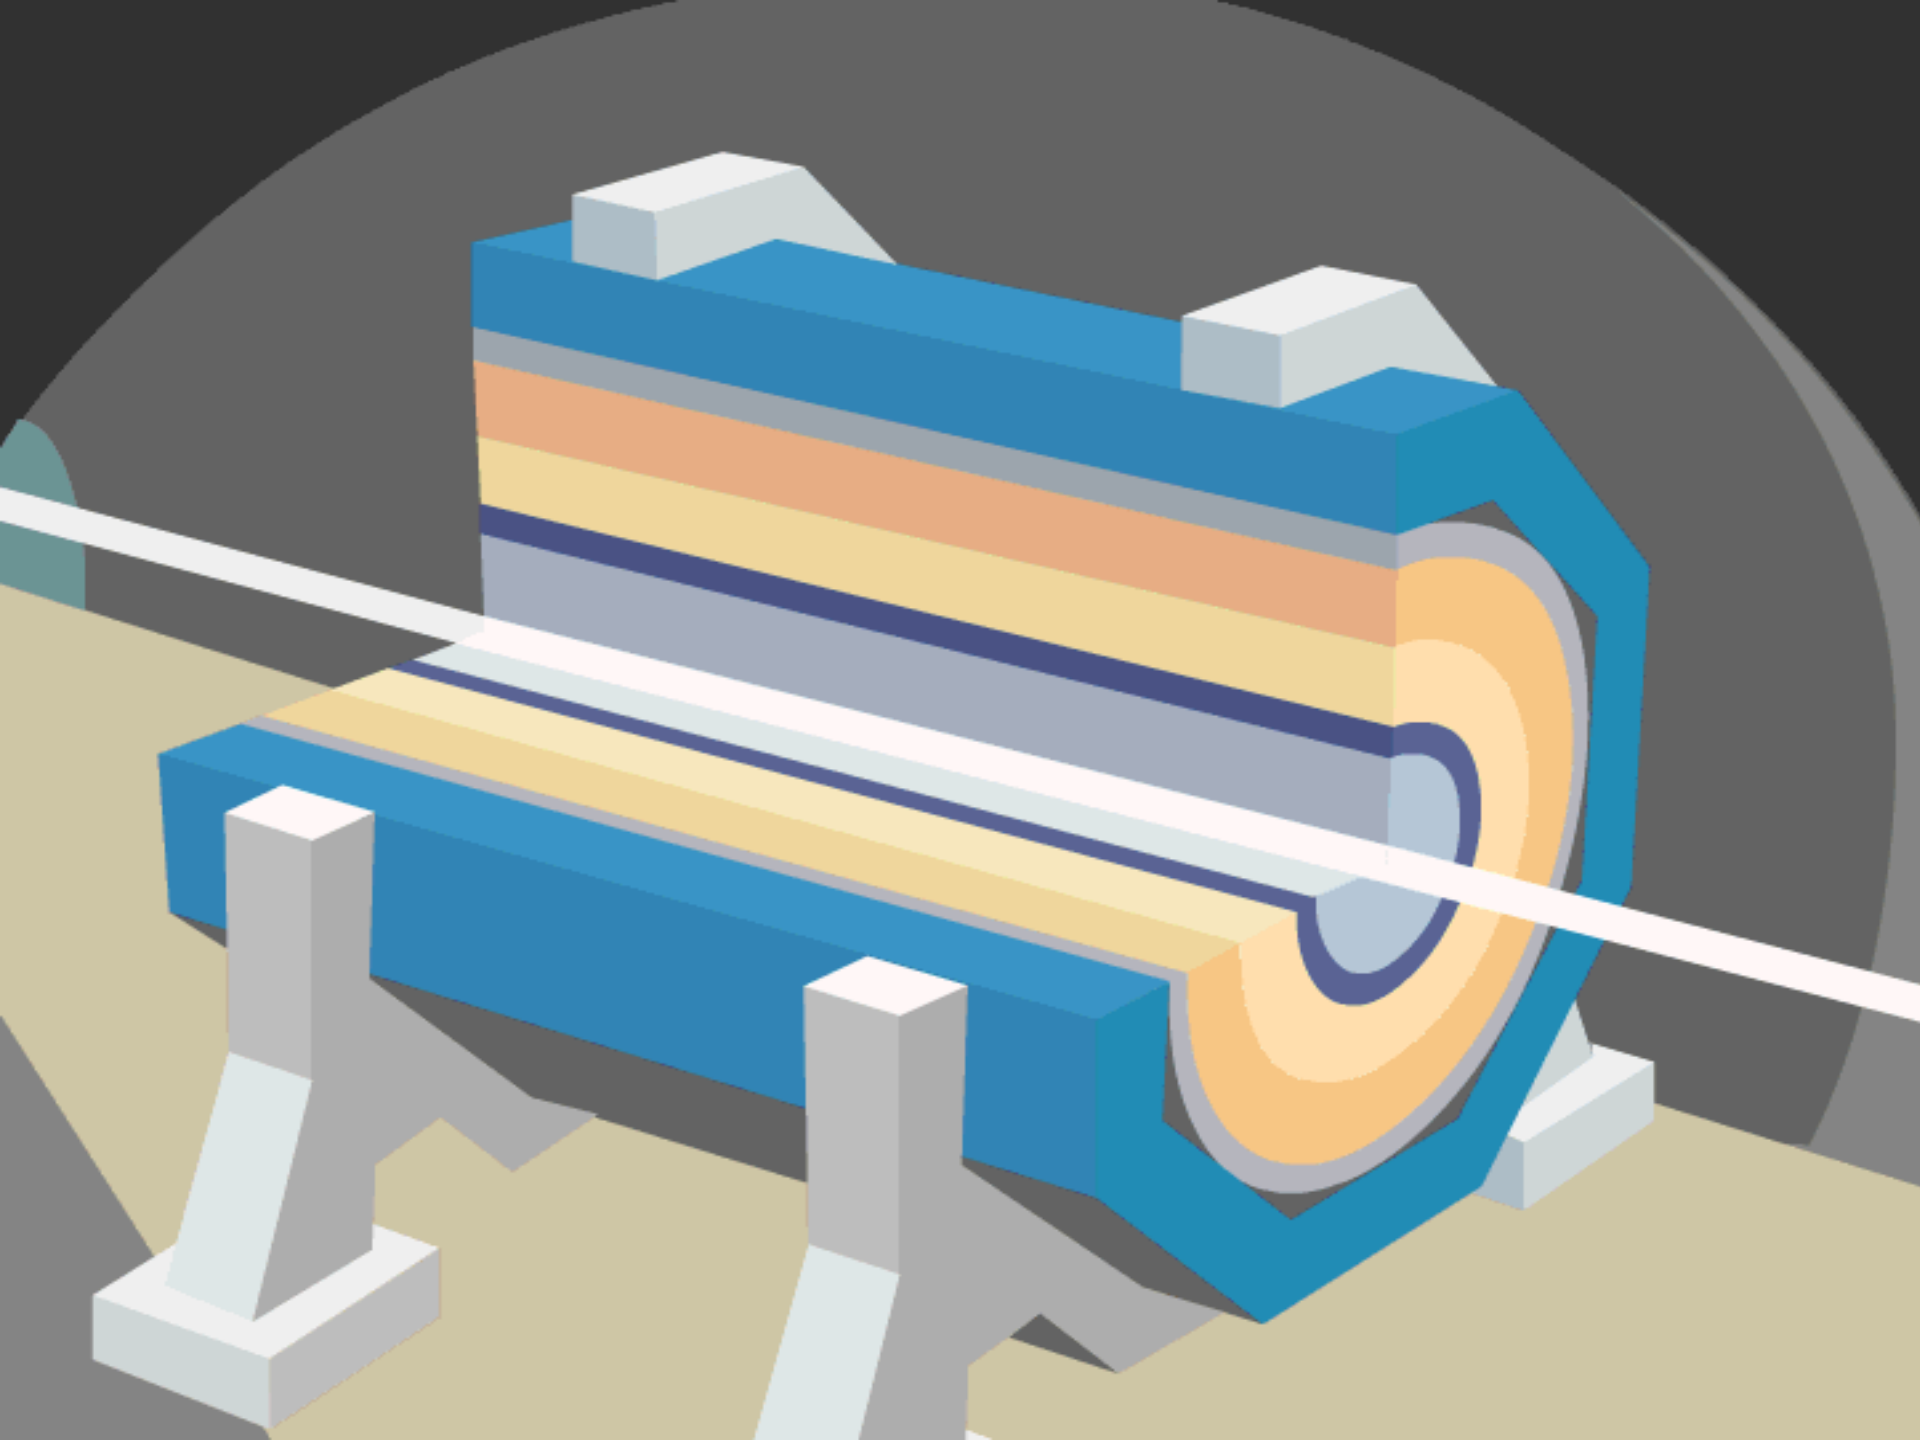

Collisions Produce 1PB/s

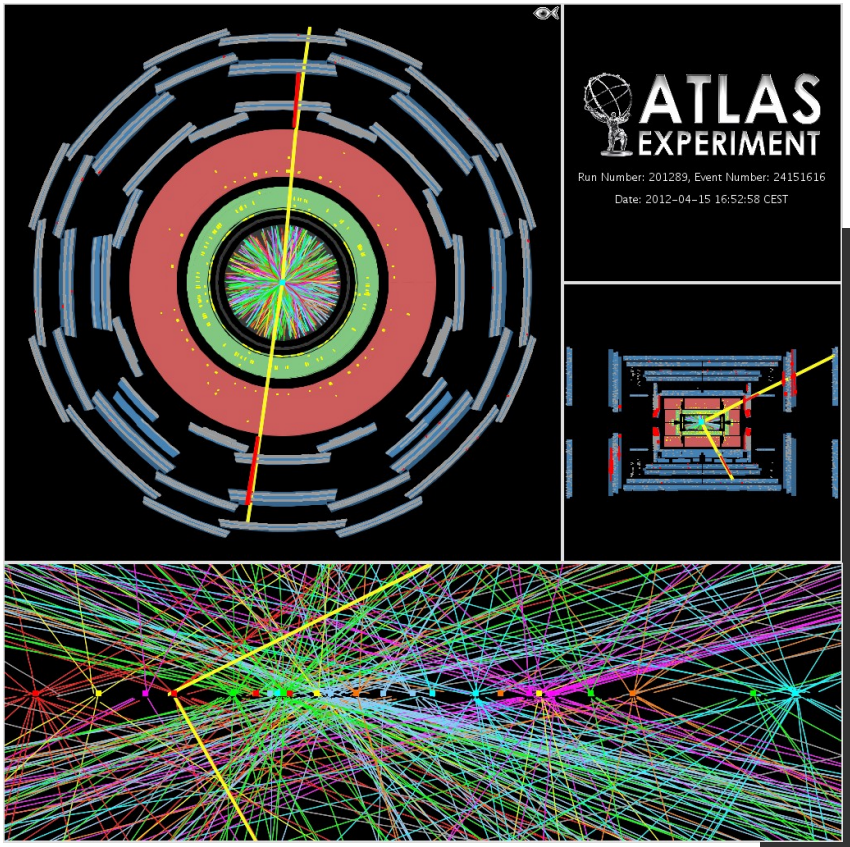

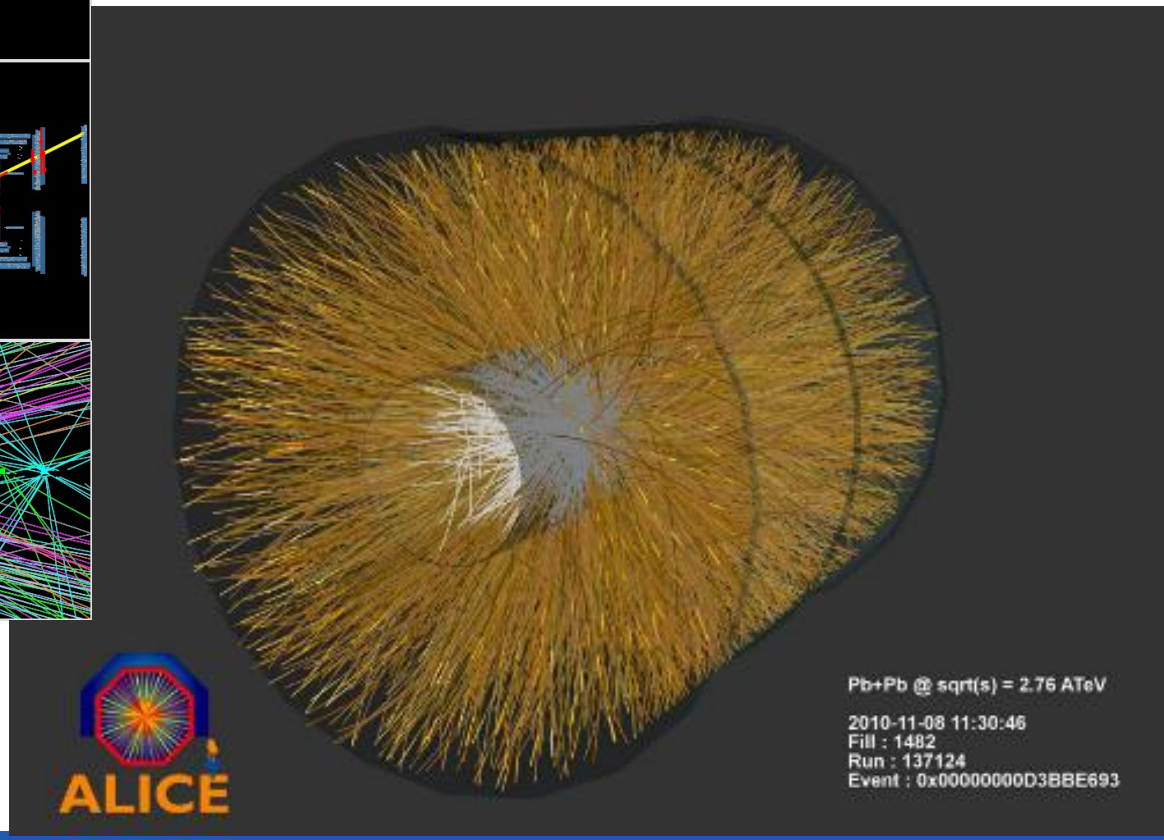

Pick the interesting events

- 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

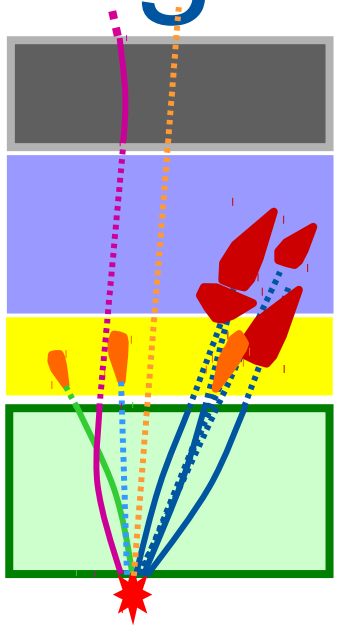

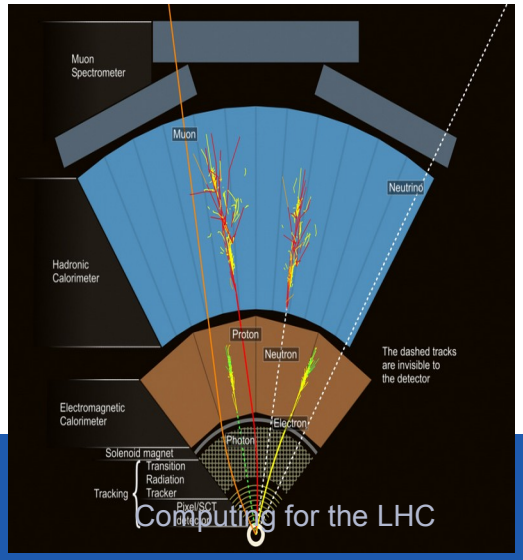

Pick the interesting events: Data size

- 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

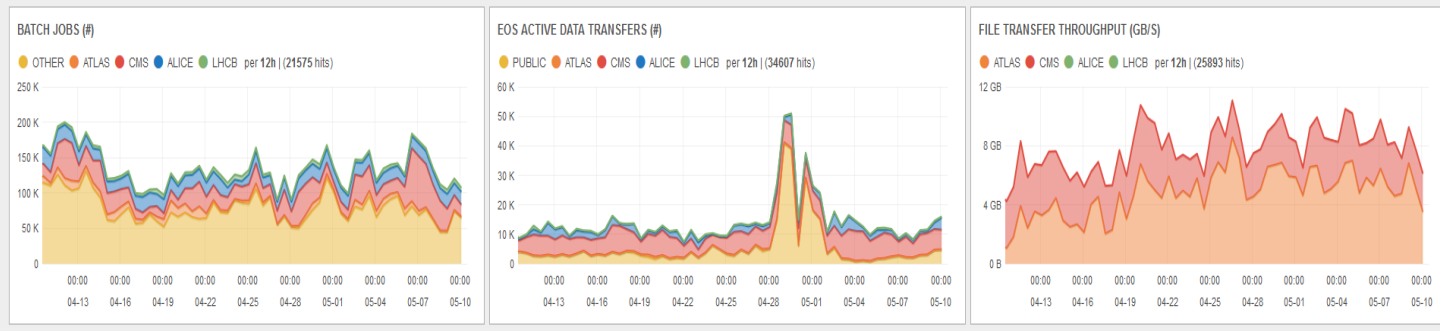

- Built in the 70s on the CERN site (Meyrin-Geneva), 3.5 MW for equipment
- Extension located at Wigner (Budapest), 2.7 MW for equipment
- Connected to the Geneva CC with 3x100Gb links (24 ms RTT)
- Hardware generally based on commodity
- 15,000 servers, providing 190,000 processor cores
- 80,000 disk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

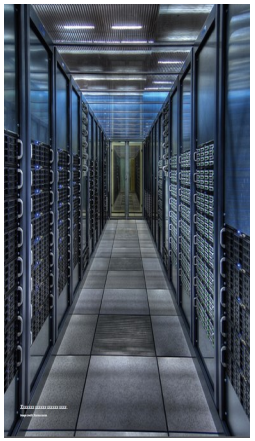

WLCG Collaboration

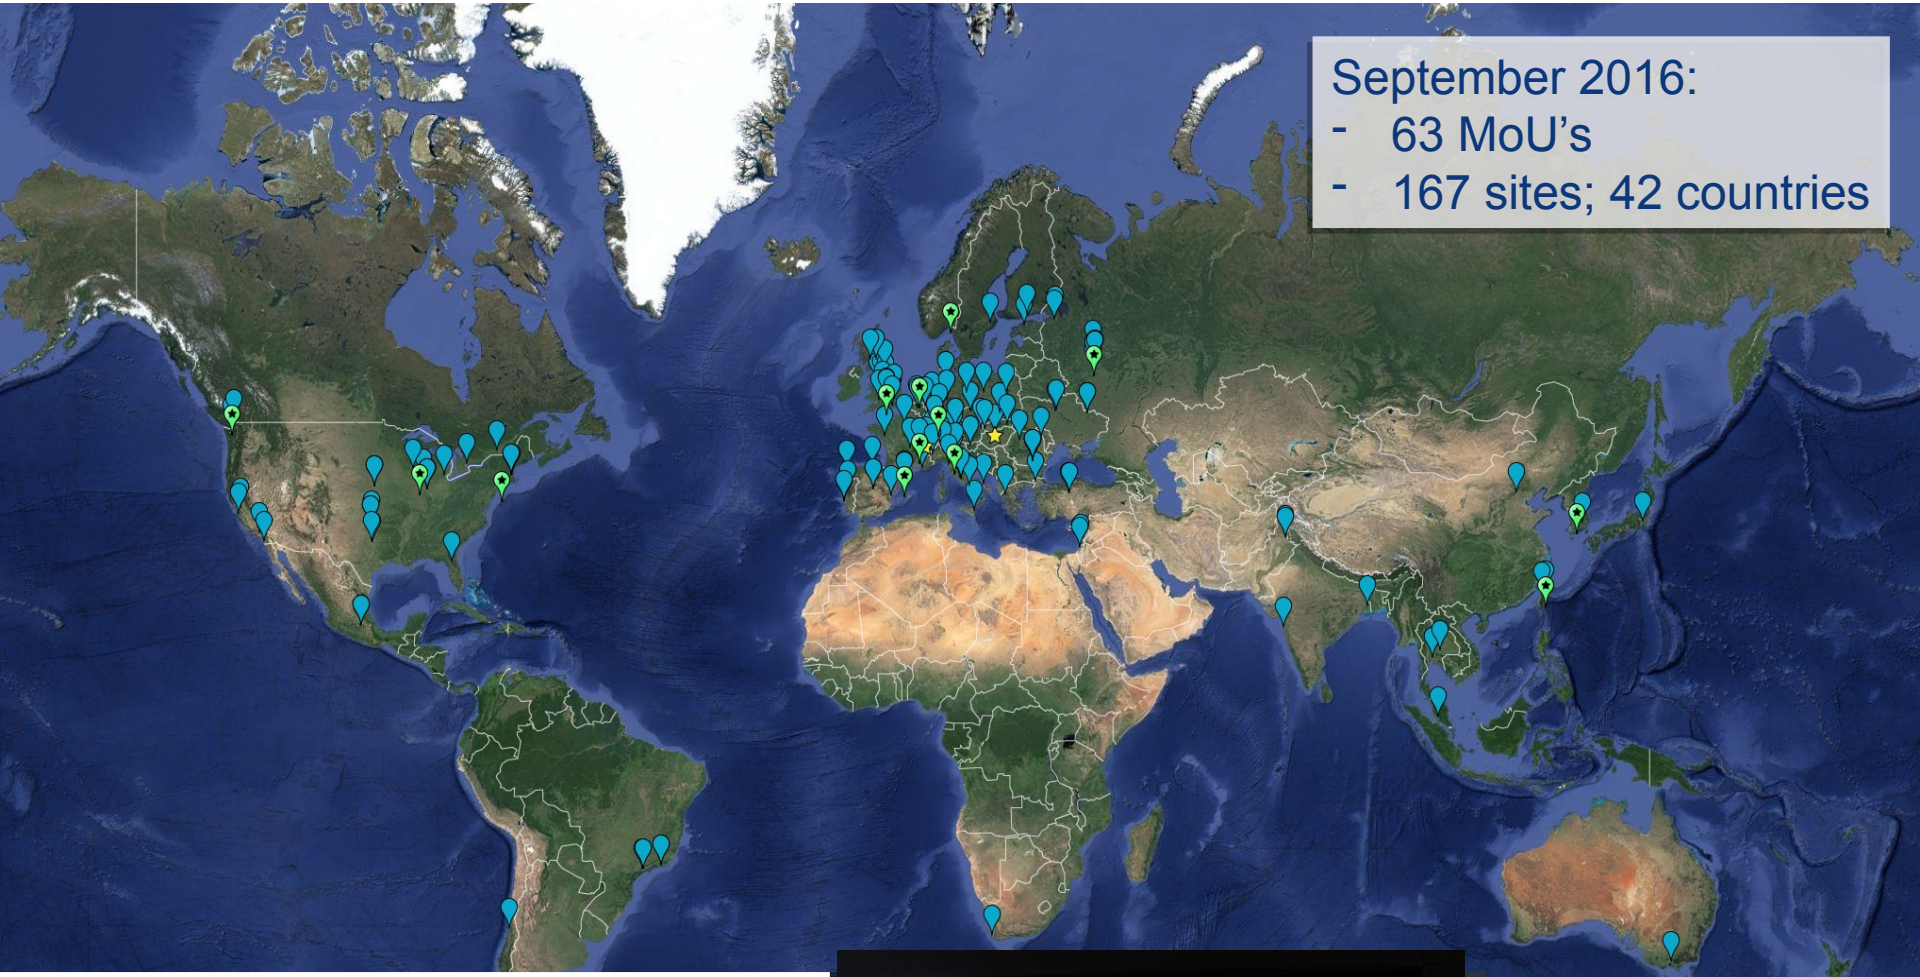

- CPU: 3.8 M HepSpec06
  - If today's fastest cores: ~ 350,000 cores
  - Actually many more (up to 5 yr old cores)
- Disk 310 PB
- Tape 390 PB

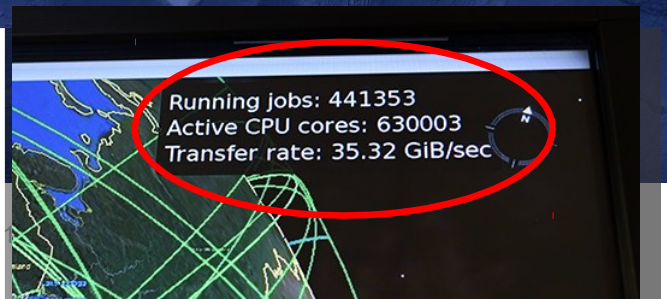

Final update - Data - 2018

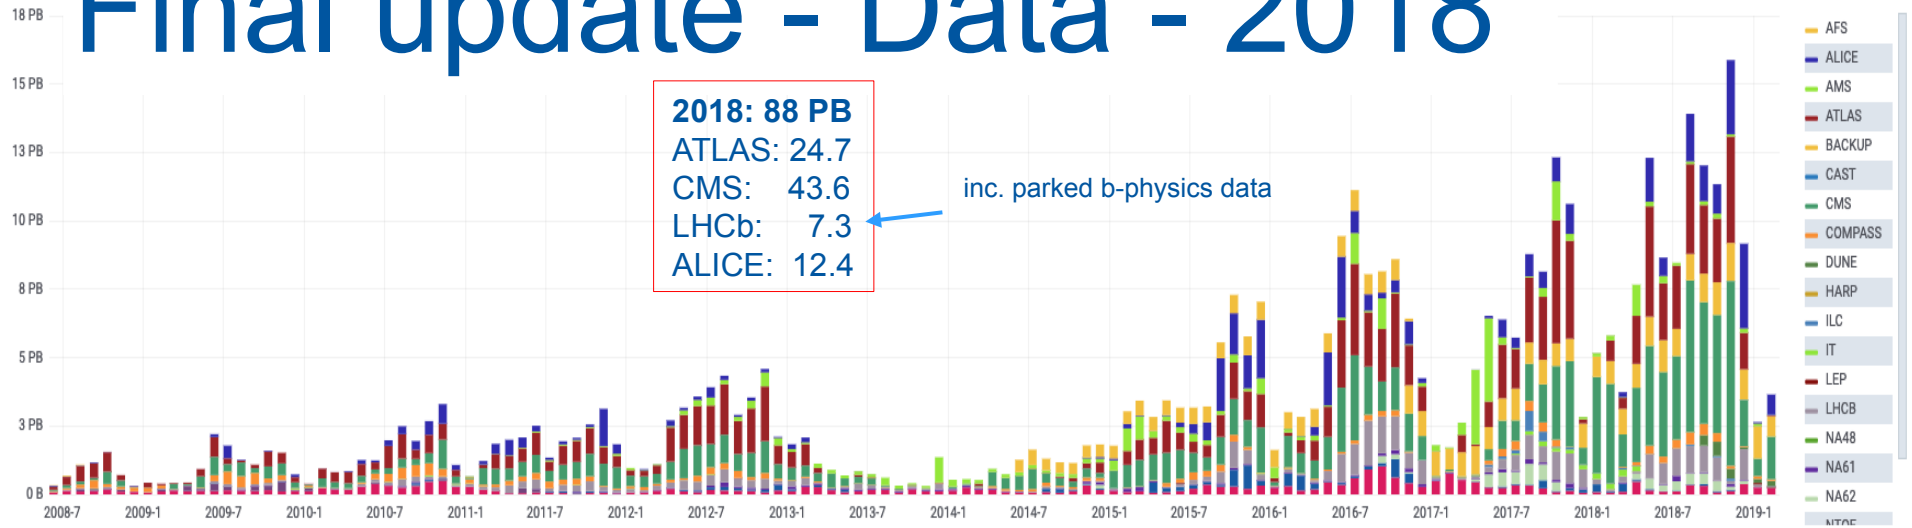

2018: 88 PB
ATLAS: 24.7
CMS: 43.6
LHCb: 7.3
ALICE: 12.4

inc. parked b-physics data

Data transfers

HI Run

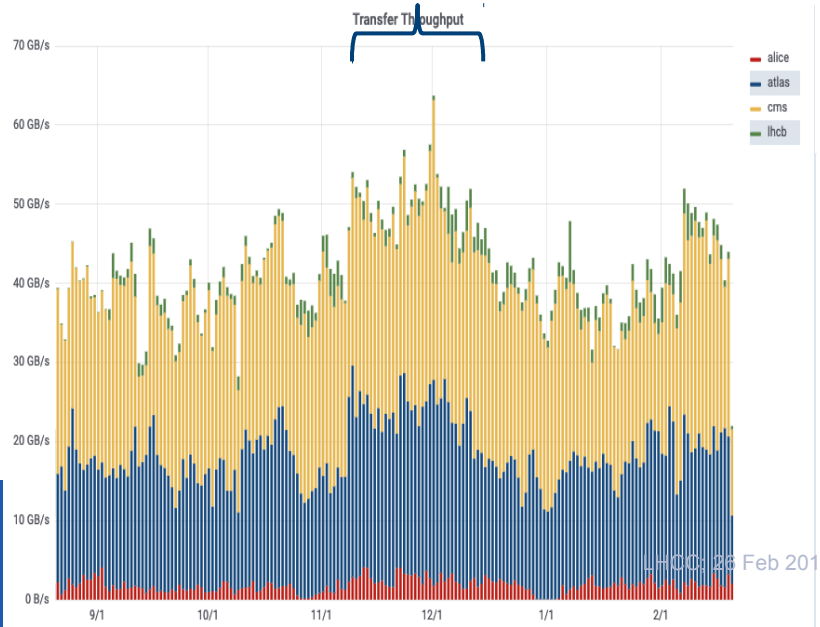

2018: 19.8 PB
ATLAS: 5.2
CMS: 7.7
LHCb: 1.2
ALICE: 5.7

HI Run

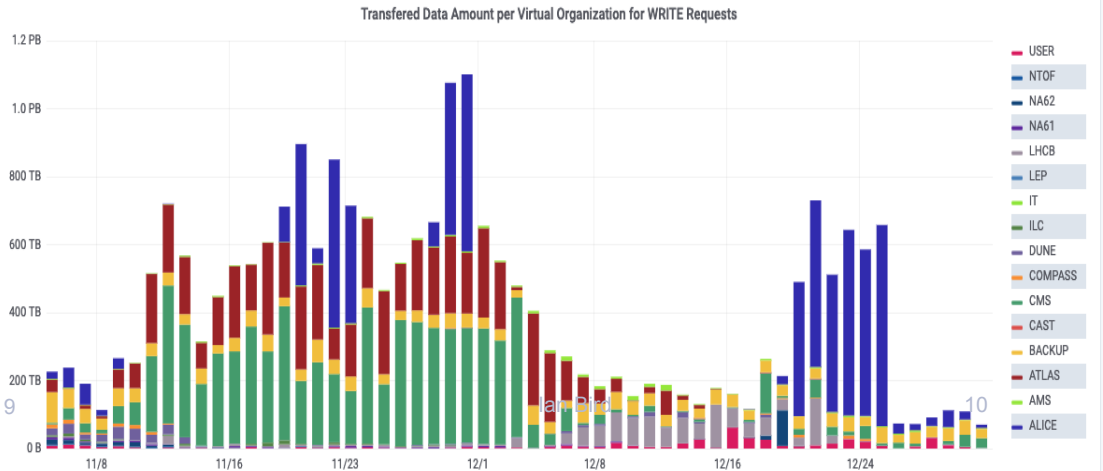

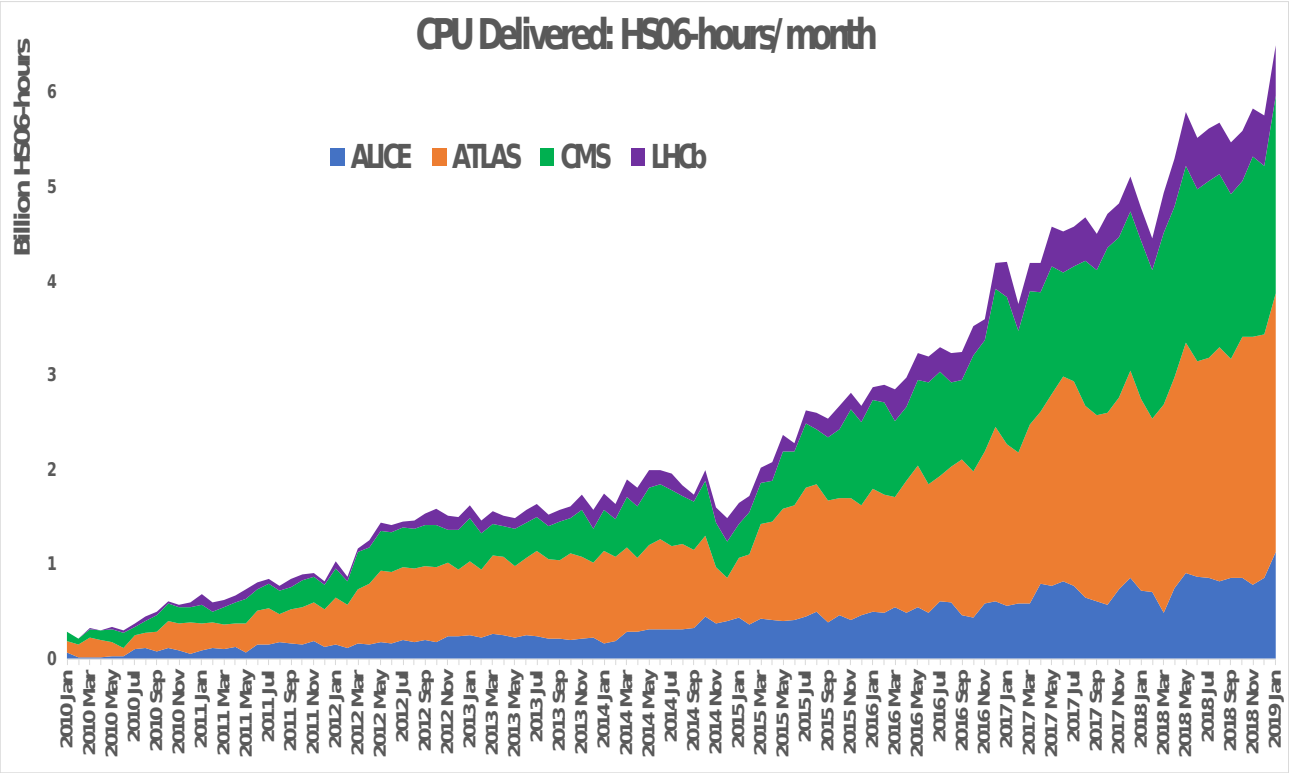

CPU Delivered

New peak: ~271 M HS06-days/month

~ 875 k cores continuous

(From sites that pledge)

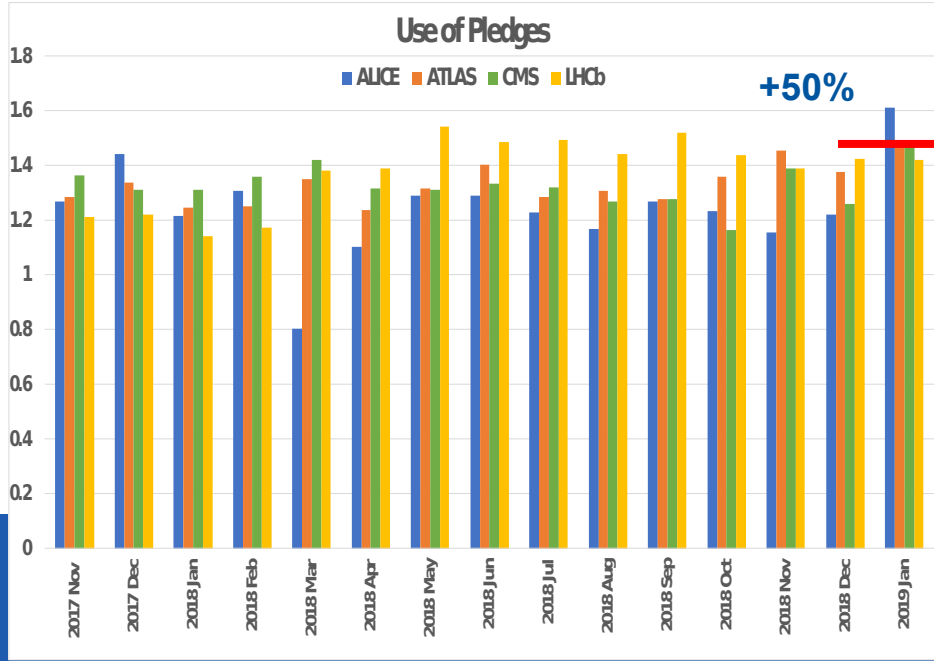

Grid vs Cloud

CERN IT Department

- "Cloud computing" has become the standard
  - 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
  - "Volunteer Cloud" - LHC@home 2.0

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

Cloud Infrastructure

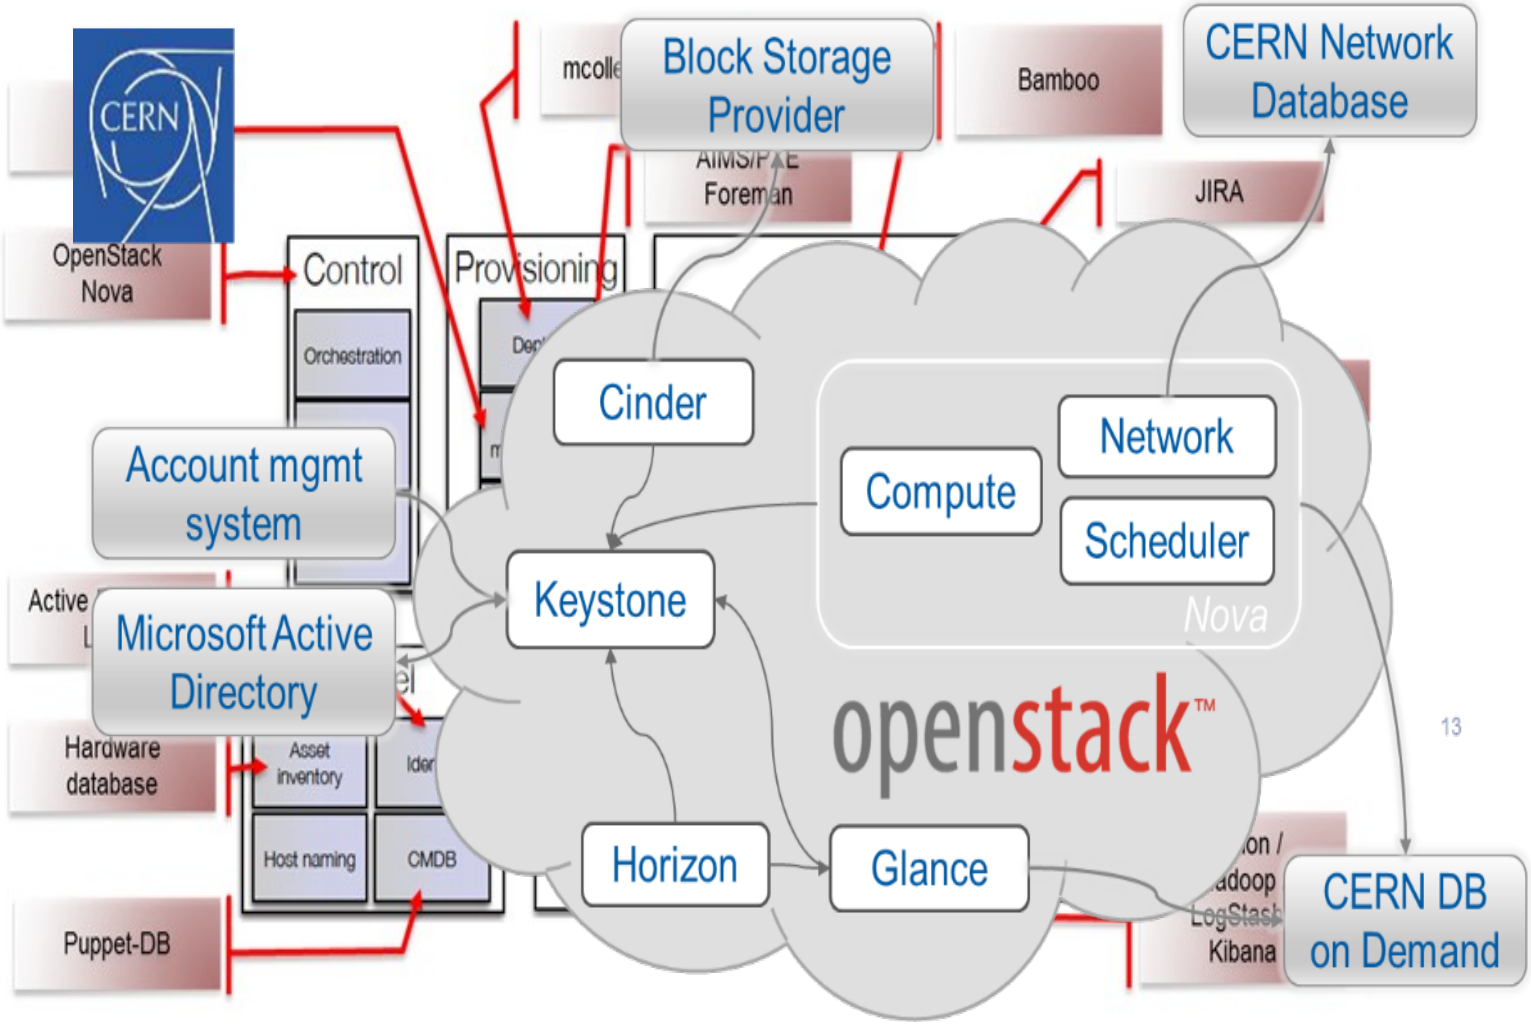

Volunteer grid - LHC@home

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**

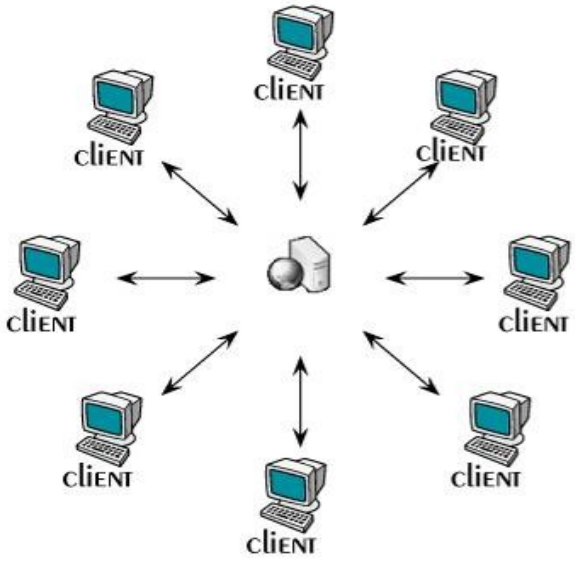

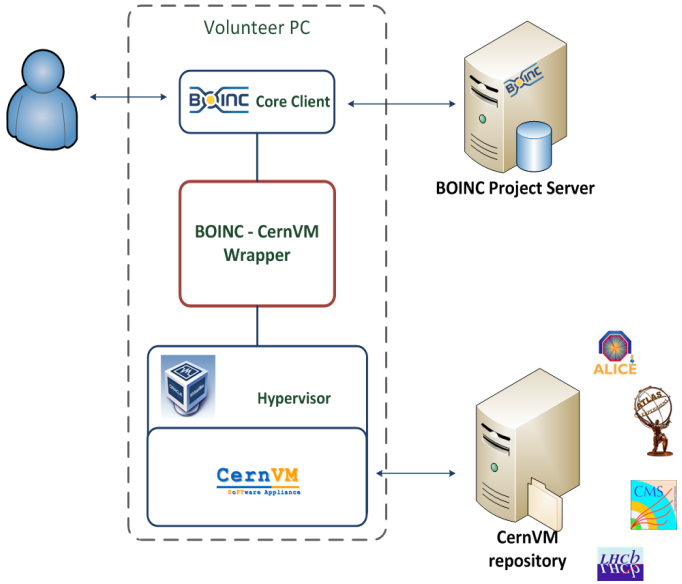

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

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

Exciting, but challenging, times ahead

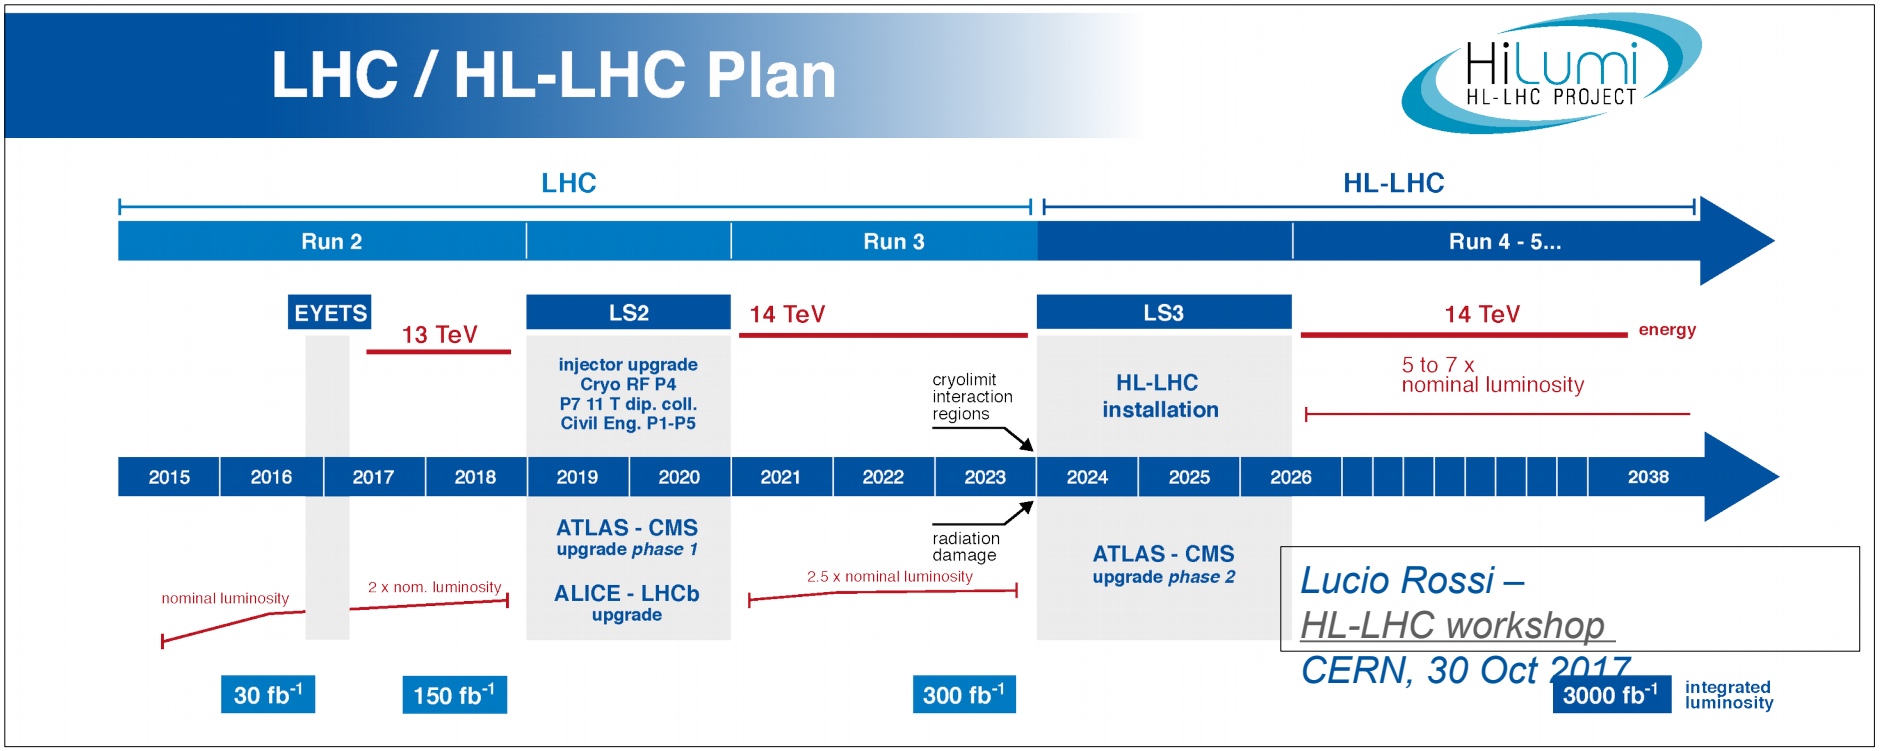

- Run3 (2021): major LHCb/ALICE upgrades
- Run4 (2026): major ATLAS/CMS upgrades, **high-luminosity LHC**
  - **more data** (higher luminosity) – and correspondingly more MC to generate/simulate
  - **more complex events** (higher pileup)

Computing resource challenges

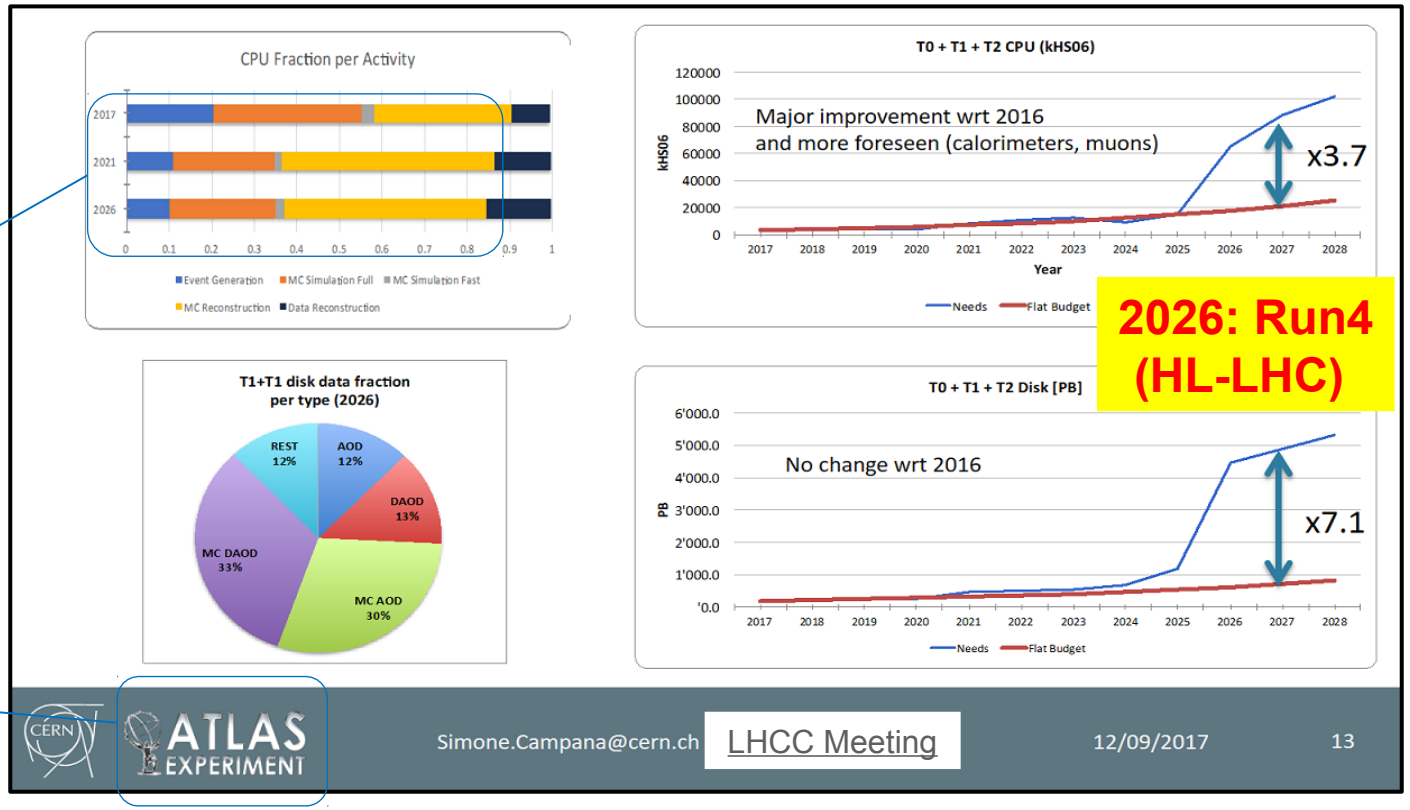

(80-90% of CPU for MC, not data)

(similar expectations for CMS)

Shortfall of resources for HL-LHC: factor ~7 for disk, factor ~4 for CPU

"The amount of data that experiments can collect and process in the future will be limited by affordable software and computing, not by physics"

Technology and market trends

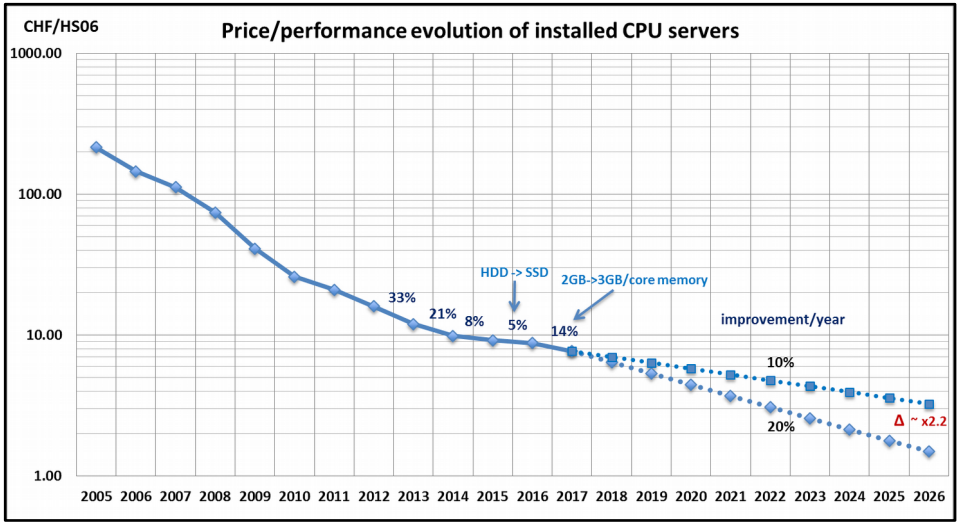

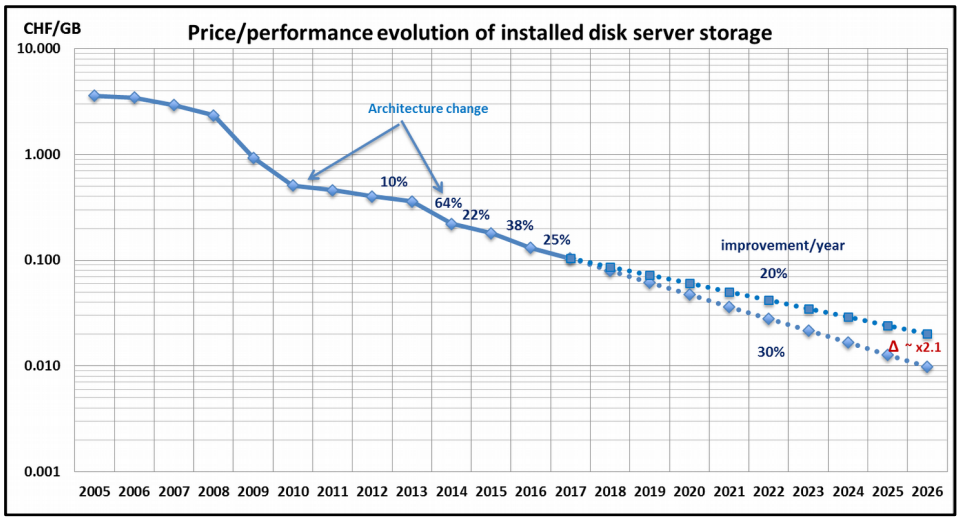

B. Panzer – Technology tracking twiki – May 2017

- Improvement per year: ~10-20% for CPU and ~20-30% for disk
  - already included in ATLAS and CMS projections at the LHCC
  - Moore's law and Kryder's law are slowing down
- More diverse landscape, too
  - multi- and many-core processors
  - wide vector registries
  - GPGPUs, FPGAs, ARM, HPCs...
  - memory bandwidth relatively low
  - *many programming models*
- *Technology alone will not solve the HL-LHC resource challenge*
- *Need a "software upgrade"!*

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, other architectures
- Data Analysis model and software changes
- Visualisation
- Storage and preservation

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

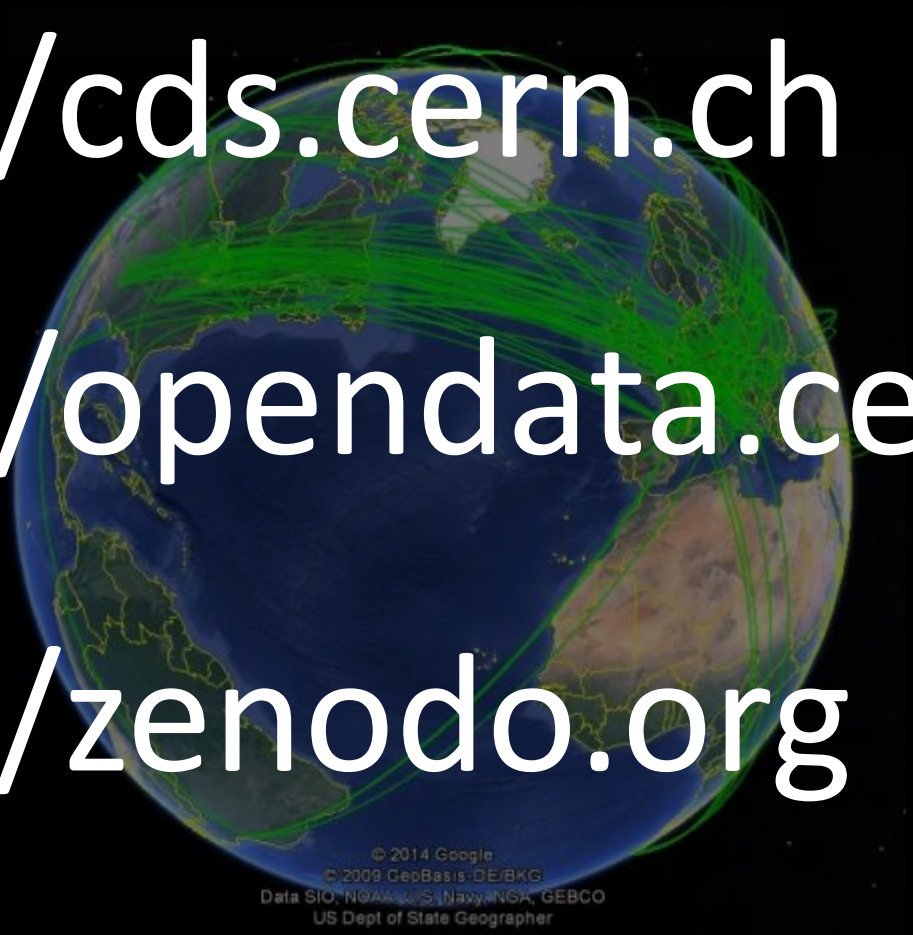

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/i:

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

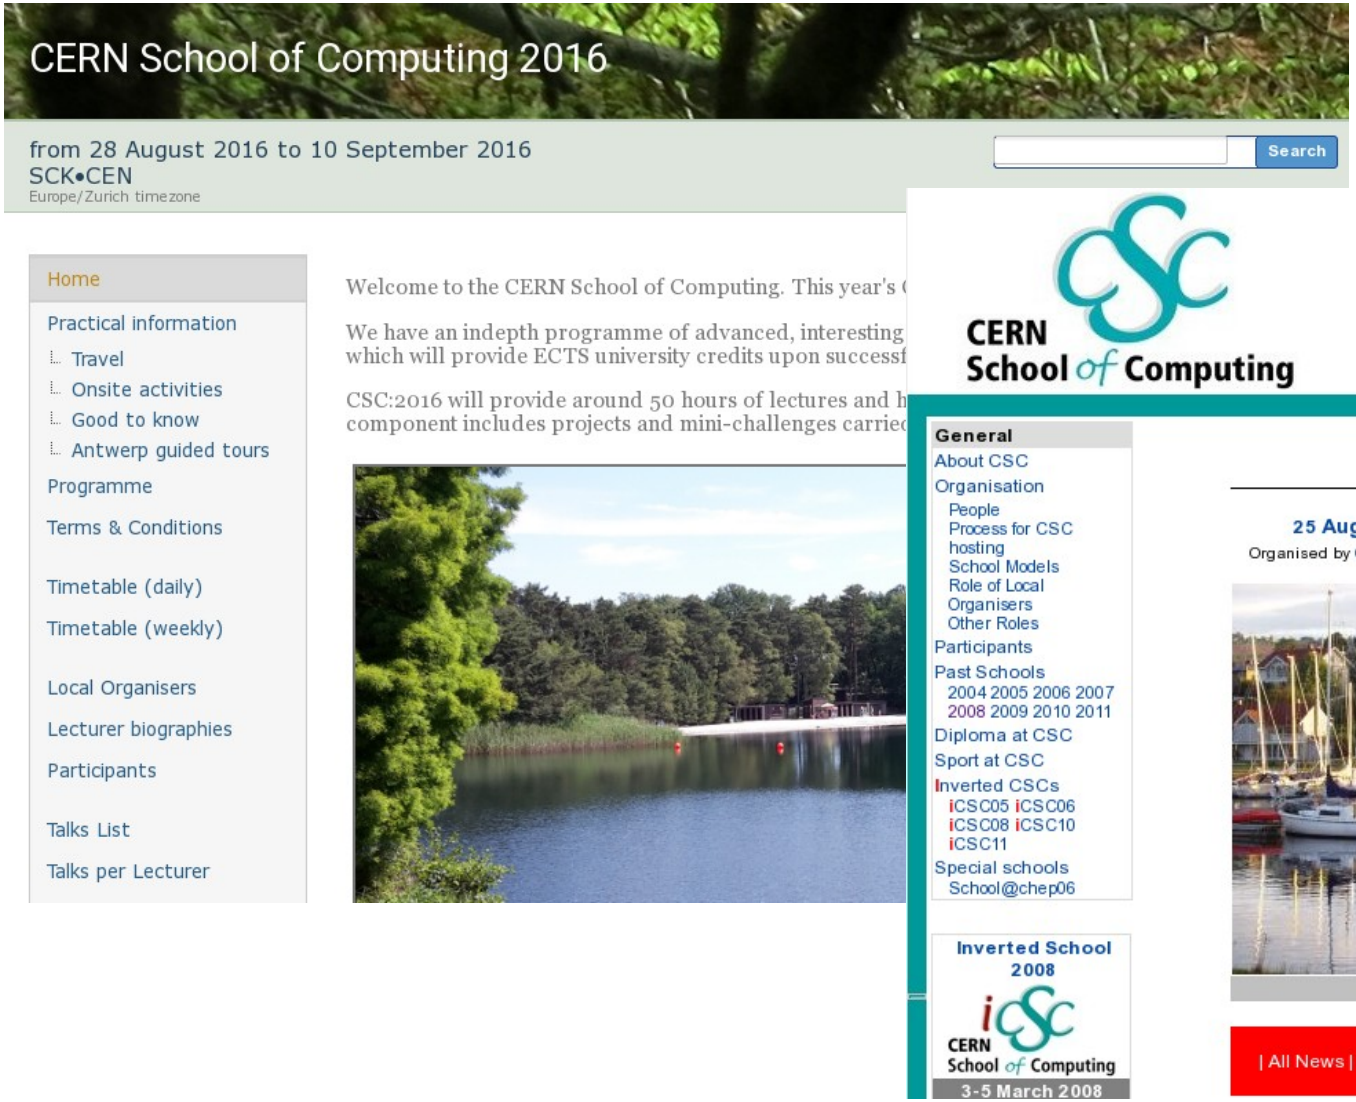

CERN openlab in a nutshell

- 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

PARTNERS

CONTRIBUTOR

ASSOCIATE

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, voice over IP...

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