The LHC computing model and its evolution

Dr Bob Jones

CERN

Bob.Jones <at> CERN.ch

CERN was founded 1954: 12 European States
Today: 20 Member States

~ 2300 staff
~ 790 other paid personnel
> 10000 users

Budget (2011) ~1000 MCHF

20 Member States: Austria, Belgium, Bulgaria, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Netherlands, Norway, Poland, Portugal, Slovakia, Spain, Sweden, Switzerland and the United Kingdom

1 Candidate for Accession: Romania

8 Observers to Council: India, Israel, Japan, the Russian Federation, the United States of America, Turkey, the European Commission and UNESCO

Answering fundamental questions…

- How to explain particles have mass?

  We have theories but need experimental evidence

- What is 96% of the universe made of ?

  We can only see 4% of its estimated mass!

- Why isn't there anti-matter in the universe?

  Nature should be symmetric…

- What was the state of matter just after the « Big Bang » ?

  Travelling back to the earliest instants of the universe would help…

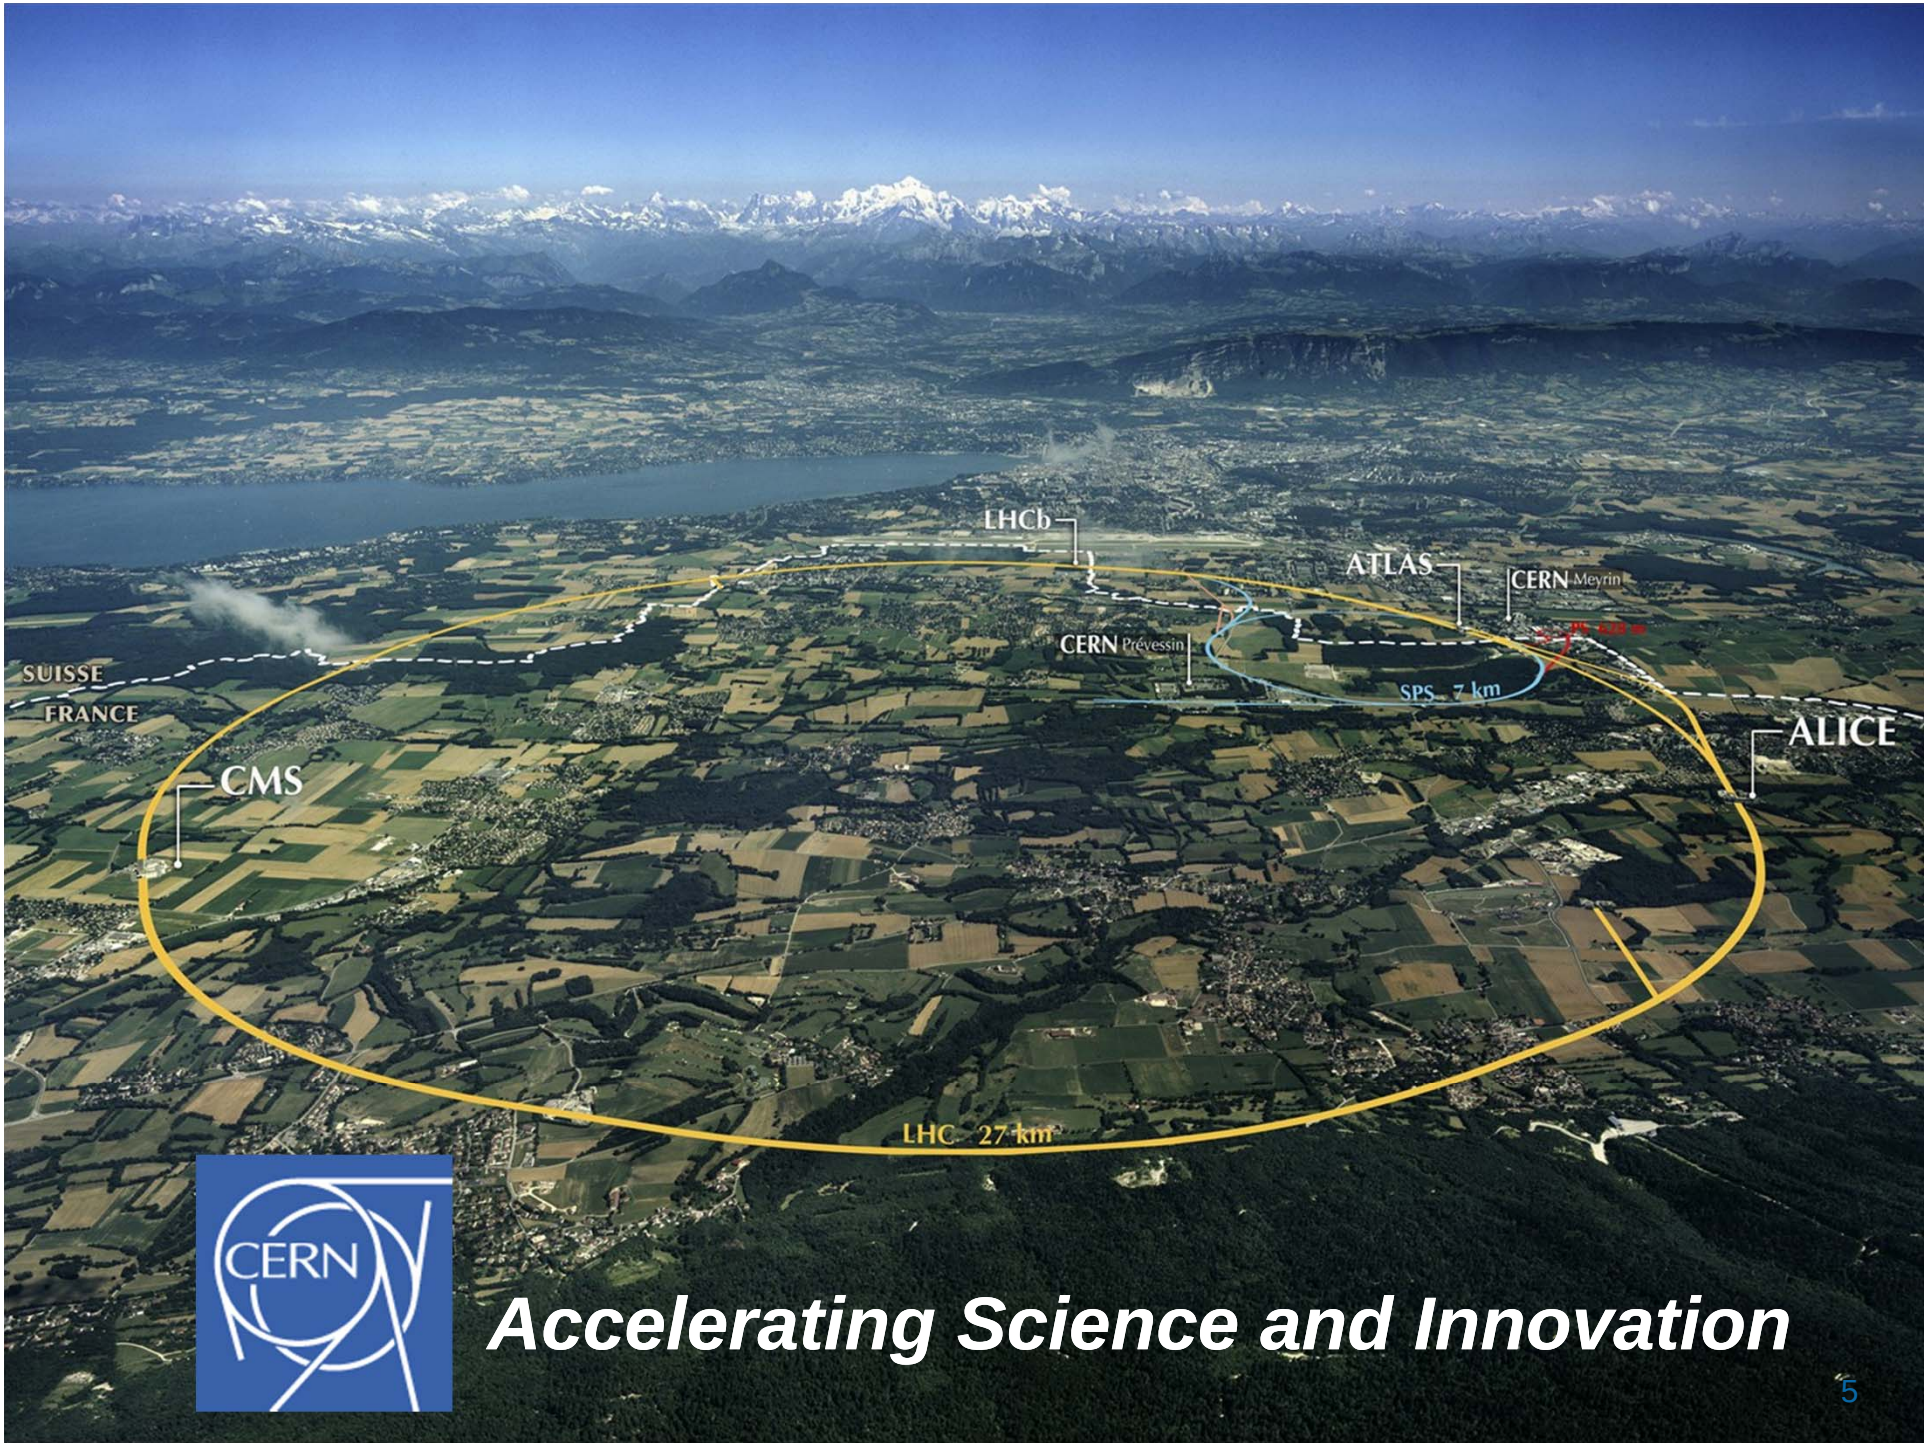

CERN

Accelerating Science and Innovation

Inside the Large Hadron Collider

Enter a New Era in Fundamental Science

Start-up of the Large Hadron Collider (LHC), one of the largest and truly global scientific projects ever, is an exciting turning point in particle physics.

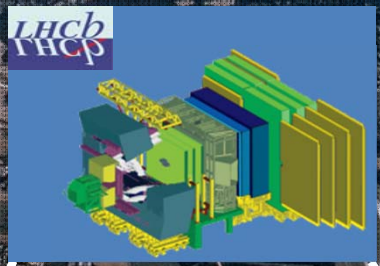

Exploration of a new energy frontier
Proton-proton collisions at E_{CM} = 14 TeV

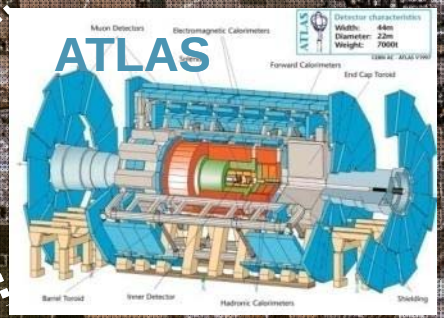

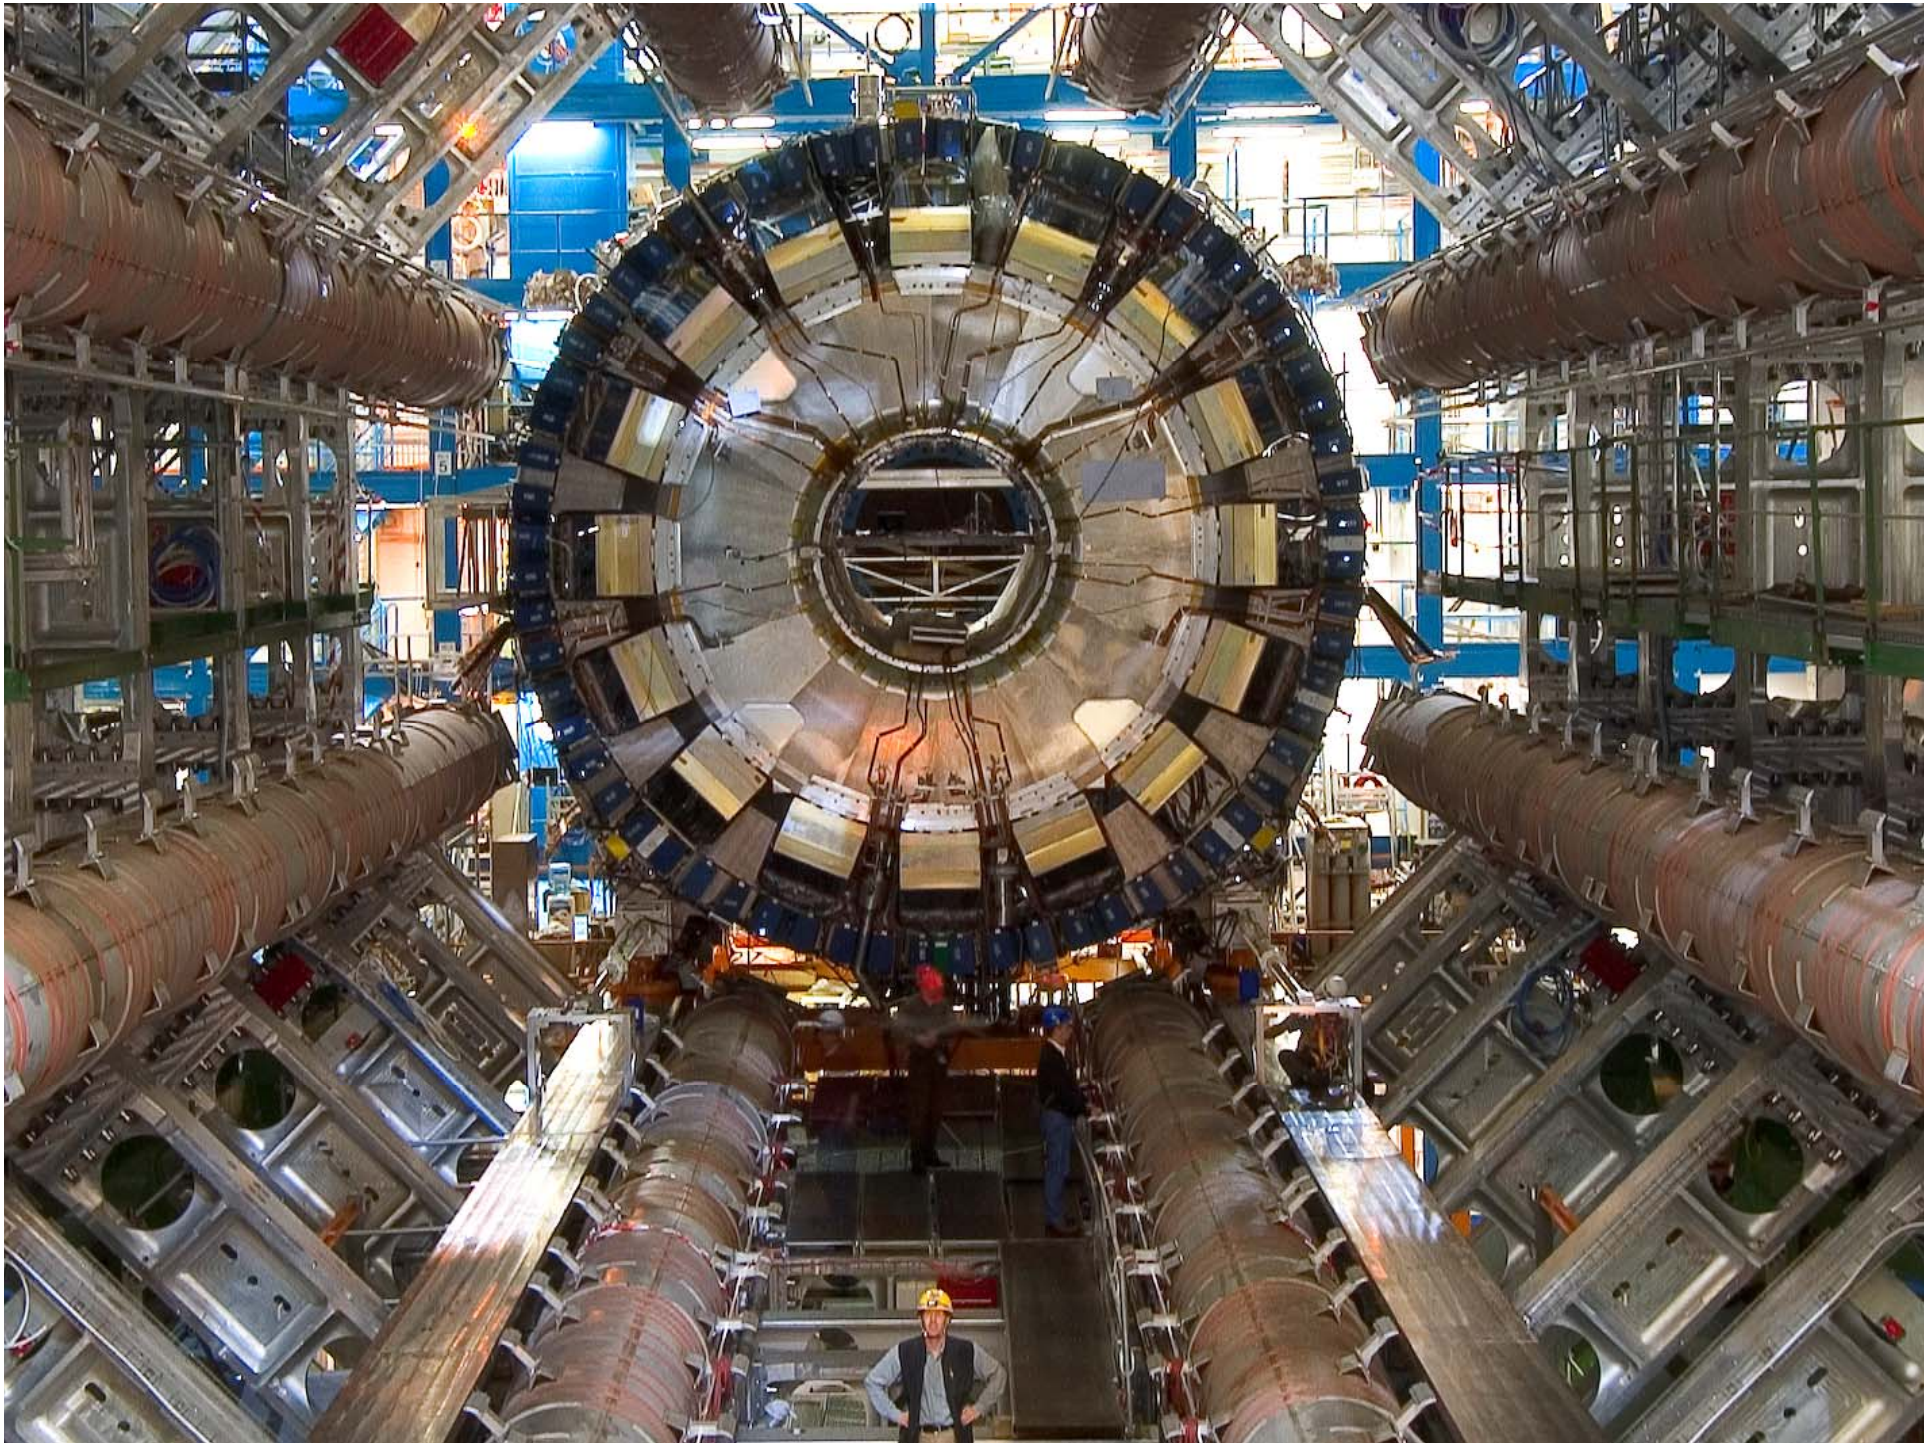

A collision @ LHC

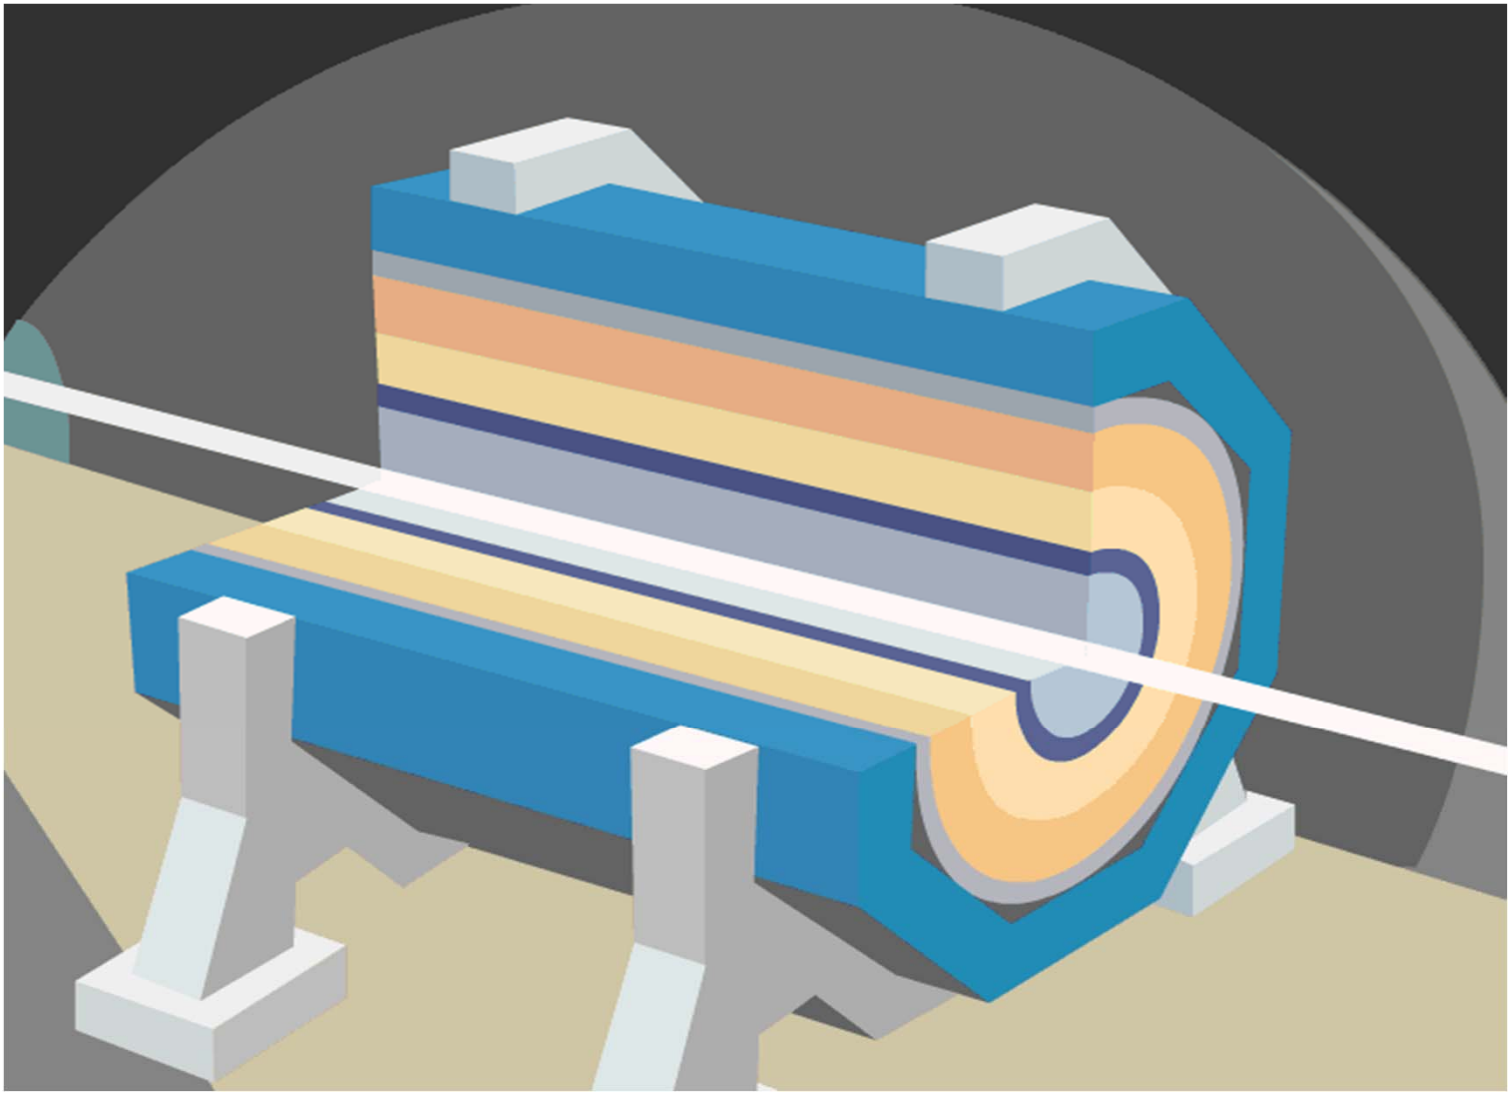

The Data Acquisition

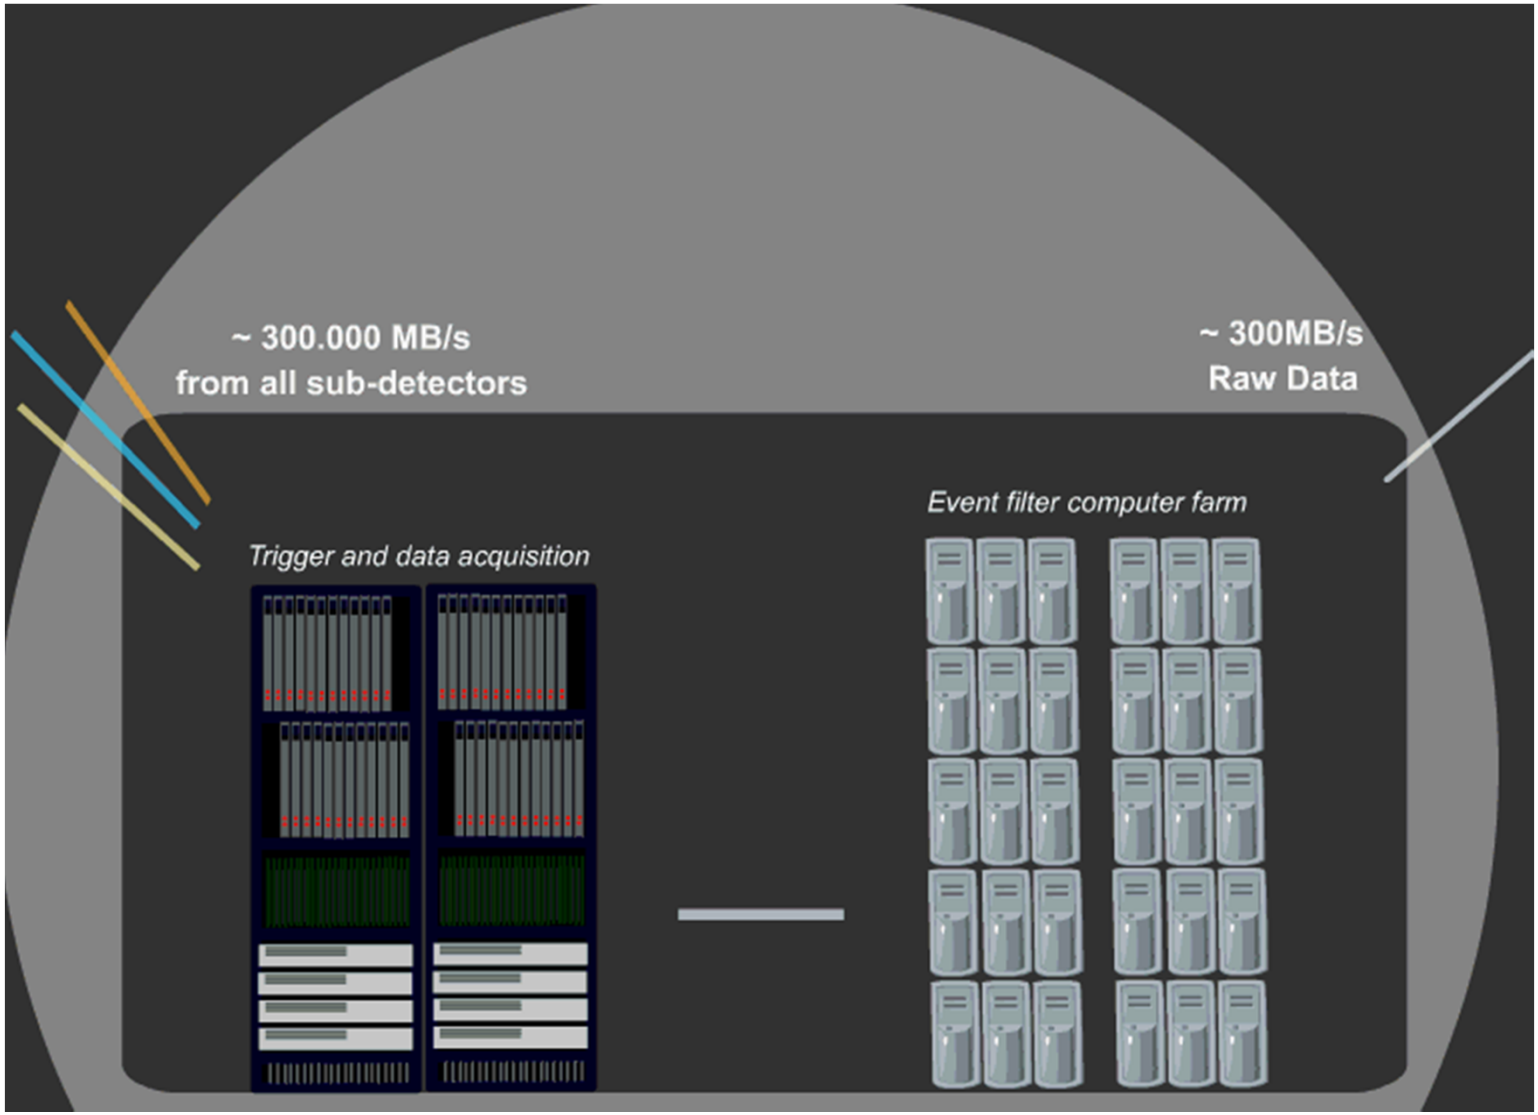

Data acquisition and storage for LHC @ CERN

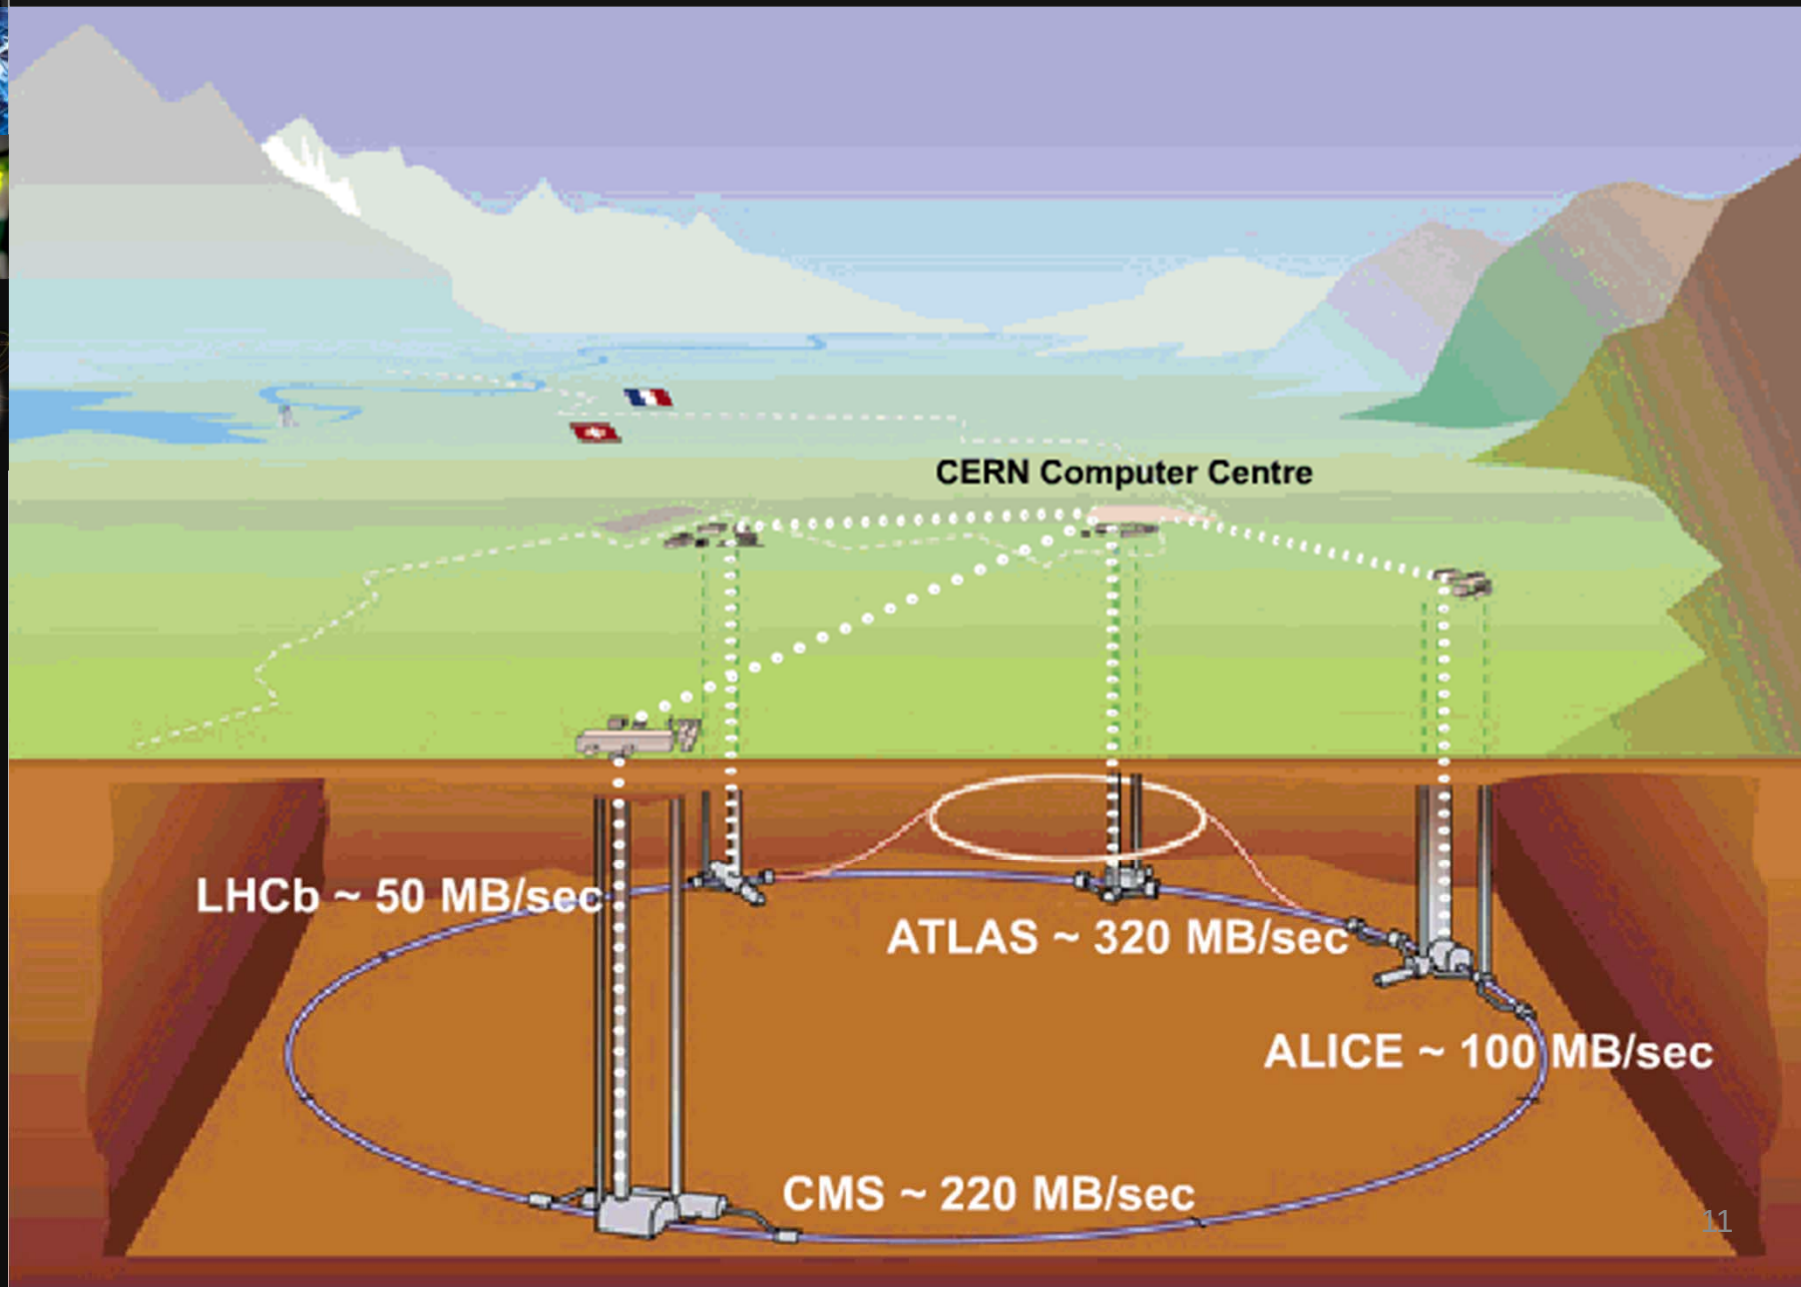

WLCG Worldwide LHC Computing Grid

The CERN Data Centre in Numbers

CERN IT Department

- Data Centre Operations (Tier 0)
  - 24x7 operator support and System Administration services to support 24x7 operation of all IT services.
  - Hardware installation & retirement
    - ~7,000 hardware movements/year; ~1800 disk failures/year
  - Management and Automation framework for large scale Linux clusters

Racks	**828**
Servers	**8938**
Processors	**15,694**
Cores	**64,238**
HEPSpec06	**482,507**

Disks	**64,109**
Raw disk capacity (TiB)	**63,289**
Memory modules	**56,014**
Memory capacity (TiB)	**158**
RAID controllers	**3,749**

Tape Drives	**160**
Tape Cartridges	**45000**
Tape slots	**56000**
Tape Capacity (TiB)	**34000**

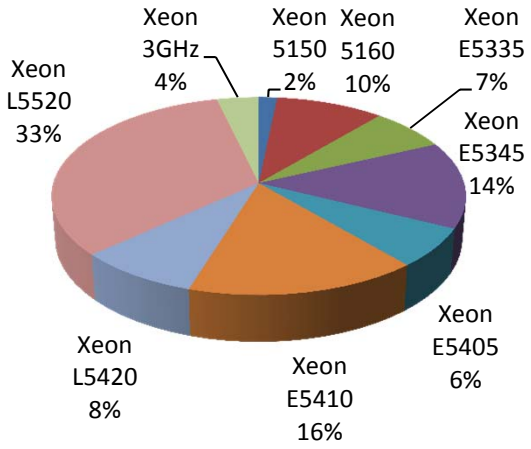

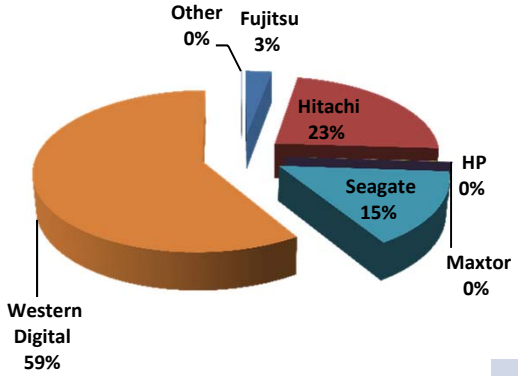

High Speed Routers (640 Mbps → 2.4 Tbps)	**24**
Ethernet Switches	**350**
10 Gbps ports	**2000**
Switching Capacity	**4.8 Tbps**
1 Gbps ports	**16,939**
10 Gbps ports	**558**

IT Power Consumption	**2456 KW**
Total Power Consumption	**3890 KW**

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

LHC-scale data processing

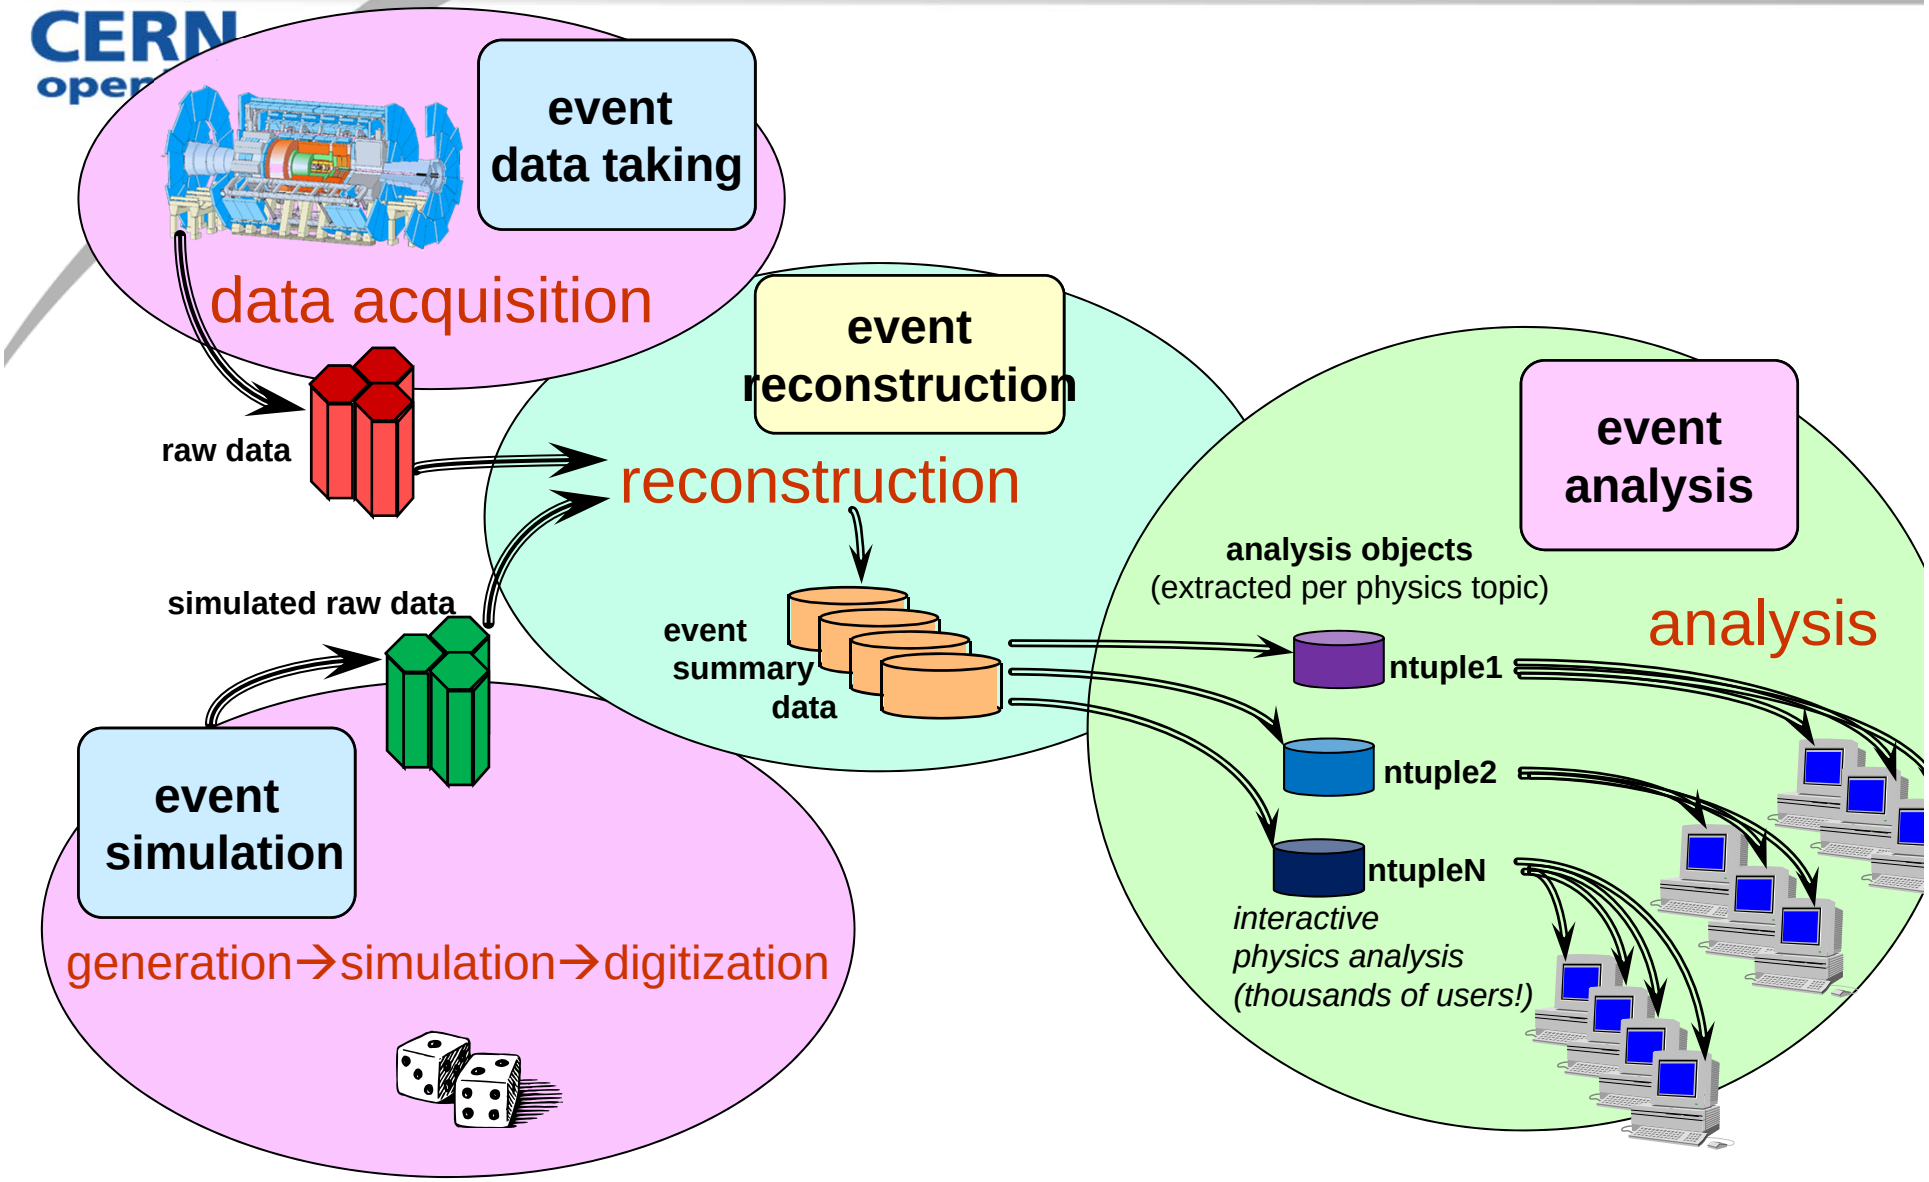

Event data taking

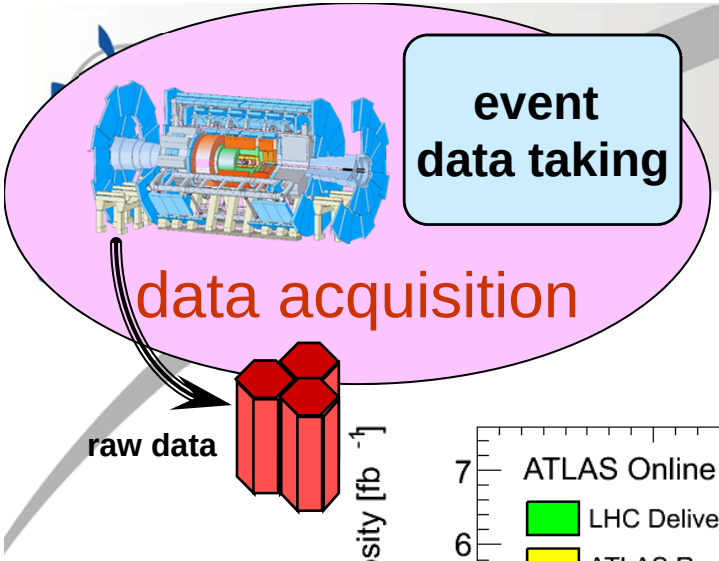

In 2011 LHC delivered 5.61 fb^{-1} of p-p collision data

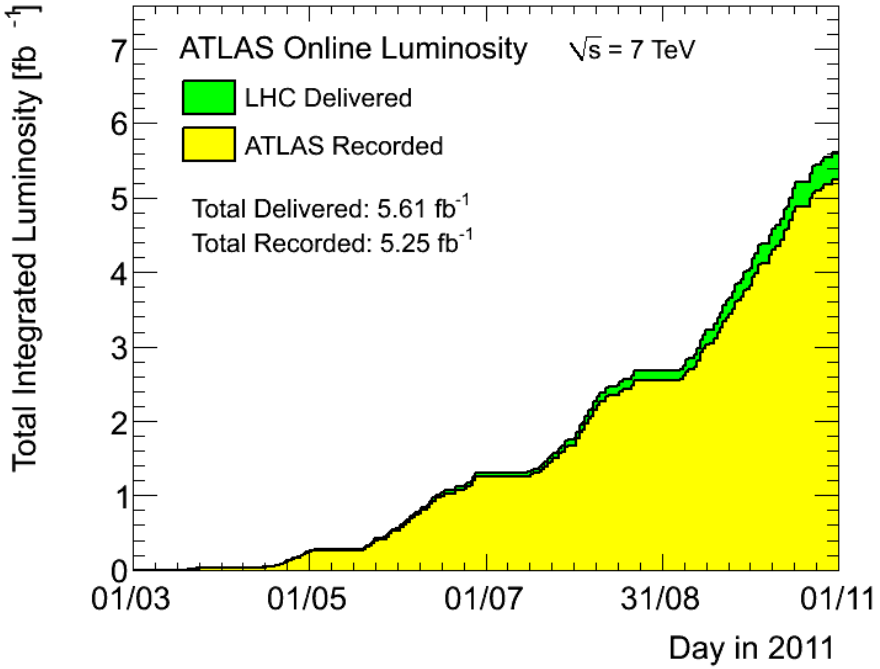

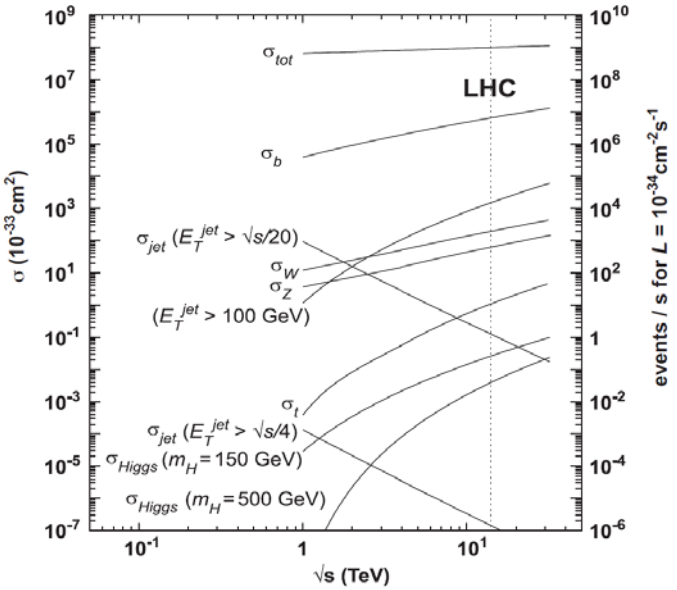

Cross sections and event rates for various processes as a function of the proton-proton center-of-mass energy.

~300 billion inelastic proton-proton interactions

ATLAS uses a flexible trigger menu to determine which events are interesting enough to record…

ATLAS recorded 1.6 billion events in 2011

The LHC Data Challenge

CERN IT Department

- The accelerator will run for 20 years
- Experiments ***are*** producing more than **15 Million Gigabytes** of data each year (about 3 million DVDs – 550 years of movies!)
- LHC data analysis requires a computing power equivalent to **~100,000 of today's fastest PC processors**
- Requires many cooperating computer centres, as CERN can only provide <20% of the capacity

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

WLCG – what and why?

- A distributed computing infrastructure to provide the production and analysis environments for the LHC experiments
- Managed and operated by a worldwide collaboration between the experiments and the participating computer centres
- The resources are distributed – for funding and sociological reasons
- Our task was to make use of the resources available to us – no matter where they are located

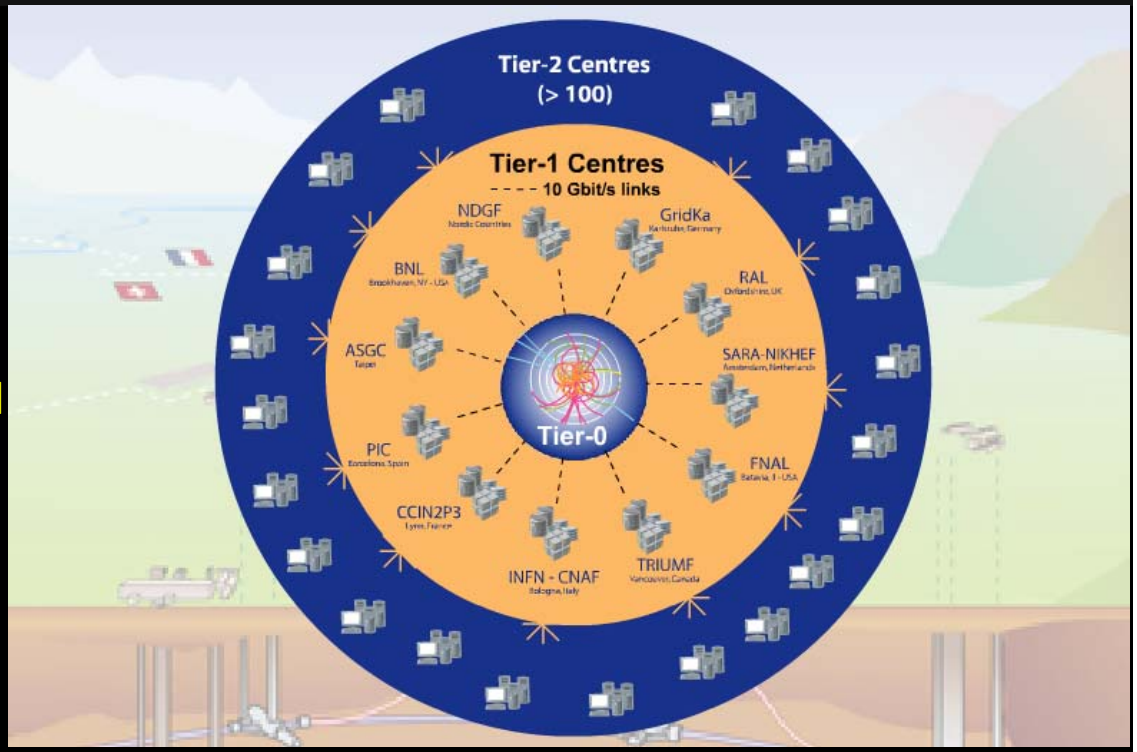

Tier-0 (CERN):
- Data recording
- Initial data reconstruction
- Data distribution

Tier-1 (11 centres):
- Permanent storage
- Re-processing
- Analysis

Tier-2 (~130 centres):
- Simulation
- End-user analysis

LHC Networking

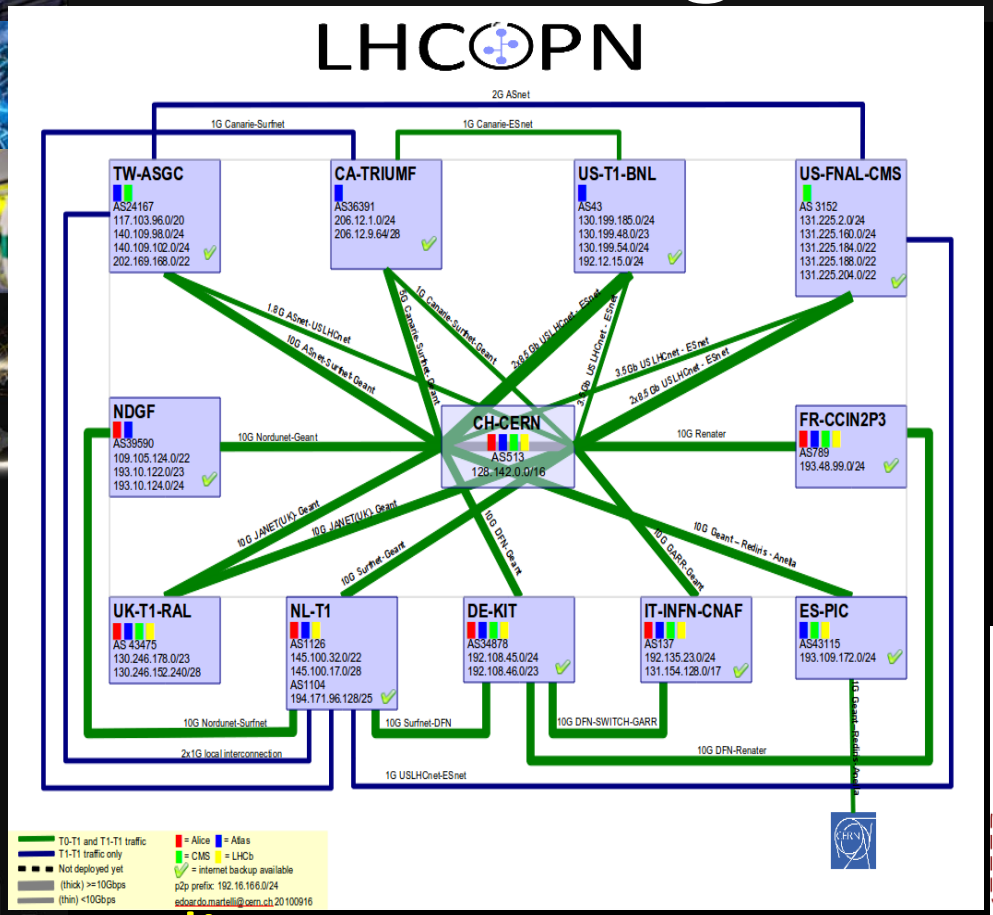

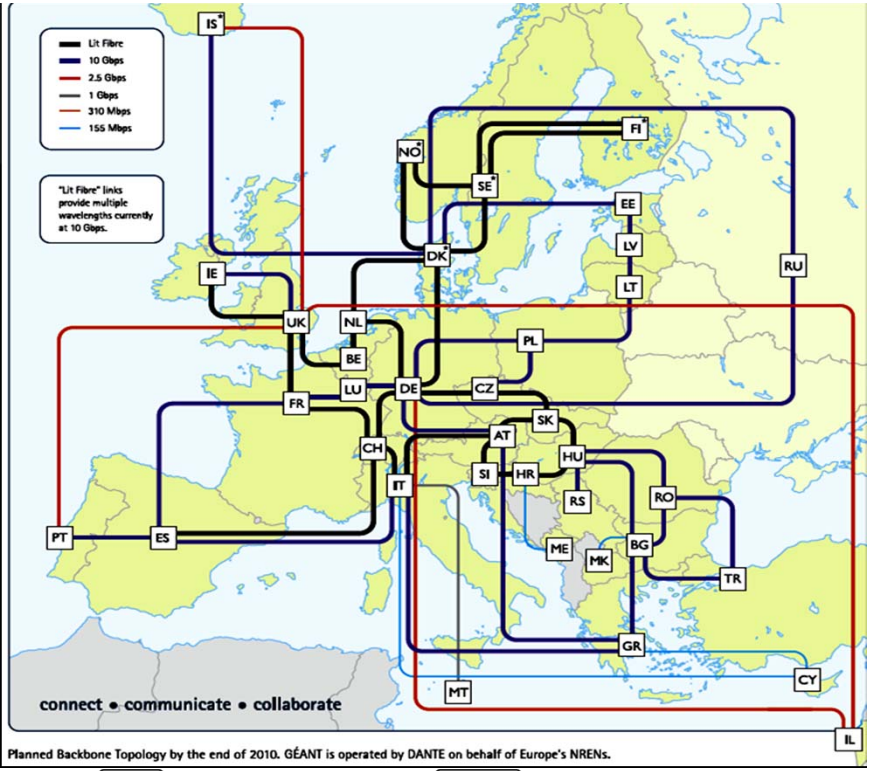

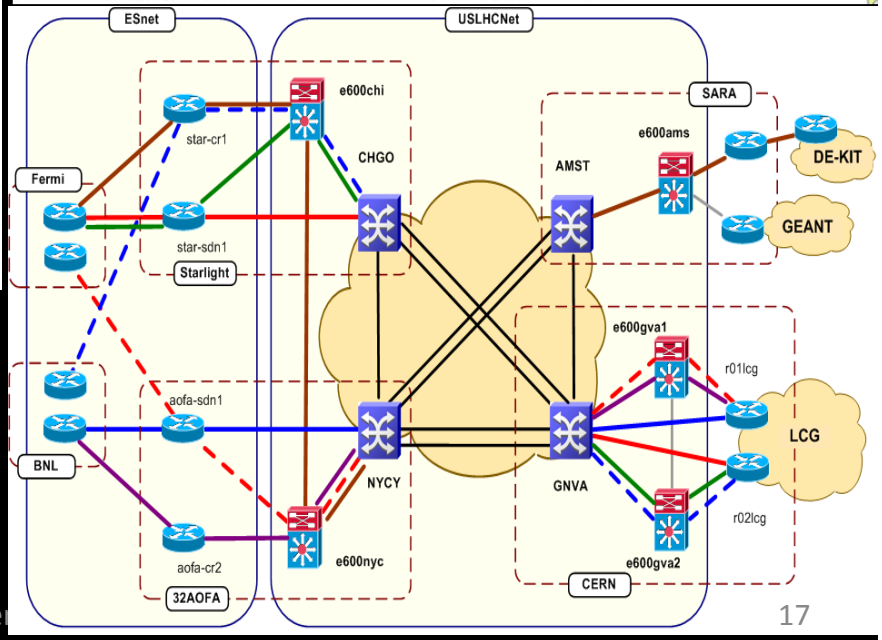

- Relies on
  - OPN, GEANT, US-LHCNet
  - NRENs & other national & international providers

WLCG in 2011

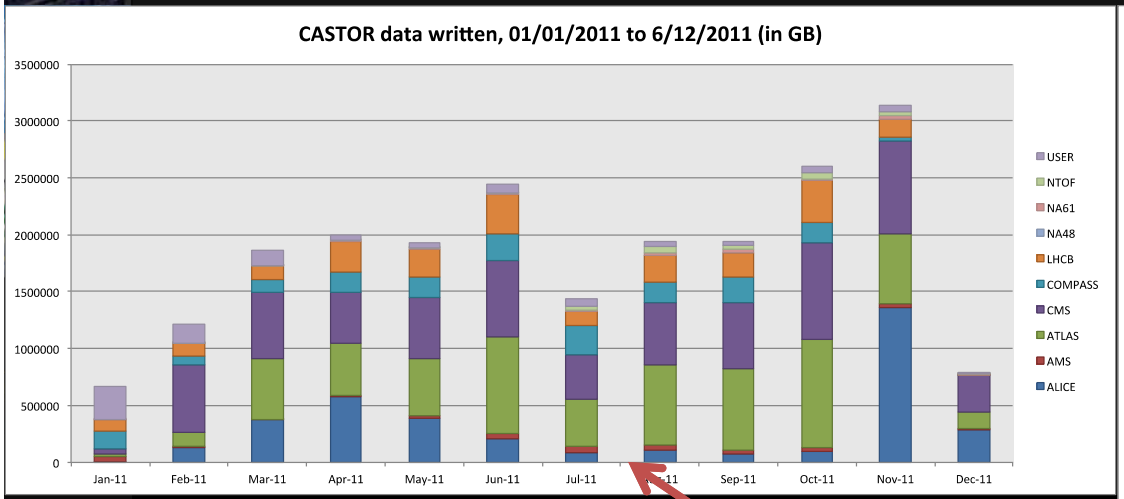

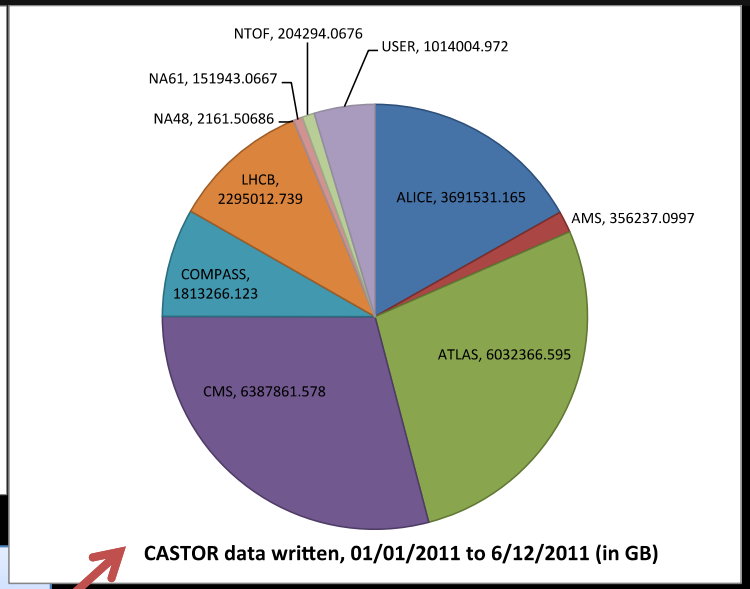

22 PB data written in 2011
More than 6 GB/s to tape during HI run

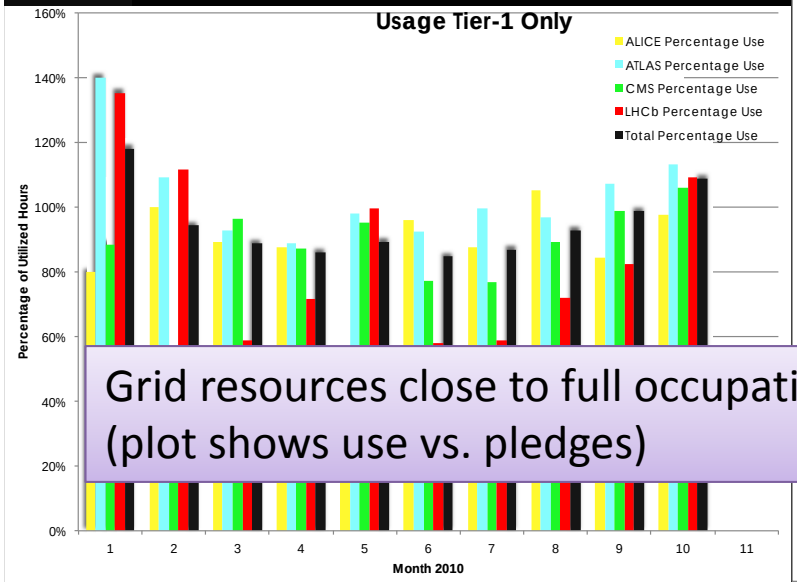

Grid resources close to full occupation
(plot shows use vs. pledges)

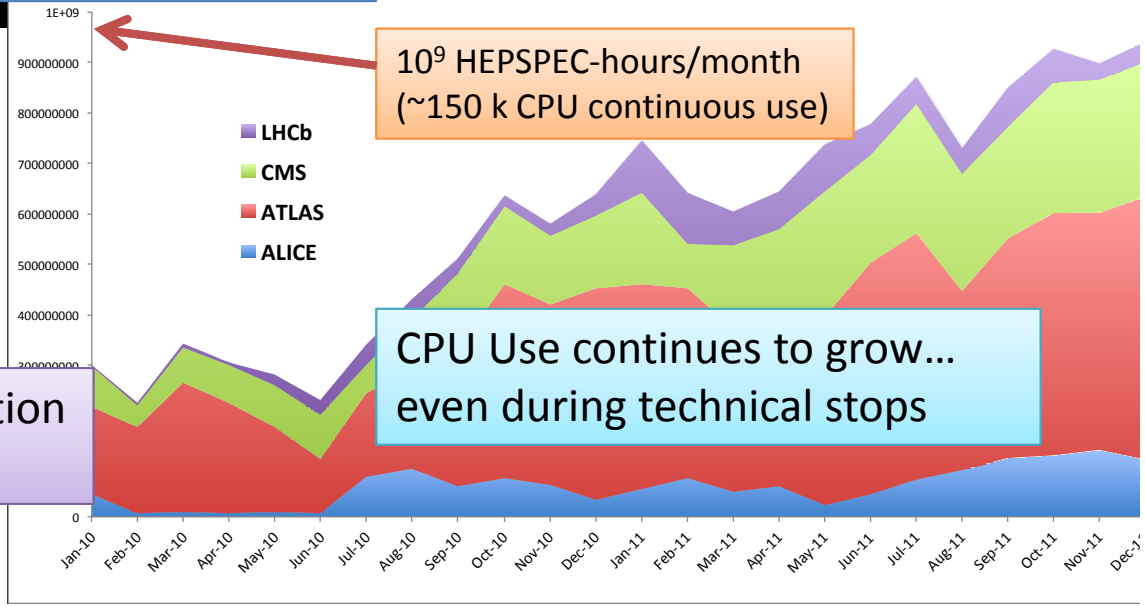

10^9 HEPSPEC-hours/month
(~150 k CPU continuous use)

CPU Use continues to grow…
even during technical stops

CPU – 11.2010-10.2011

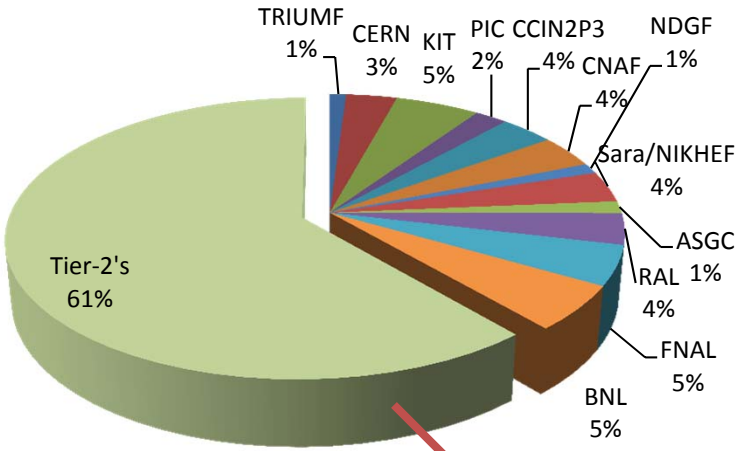

Significant use of Tier 2s for analysis

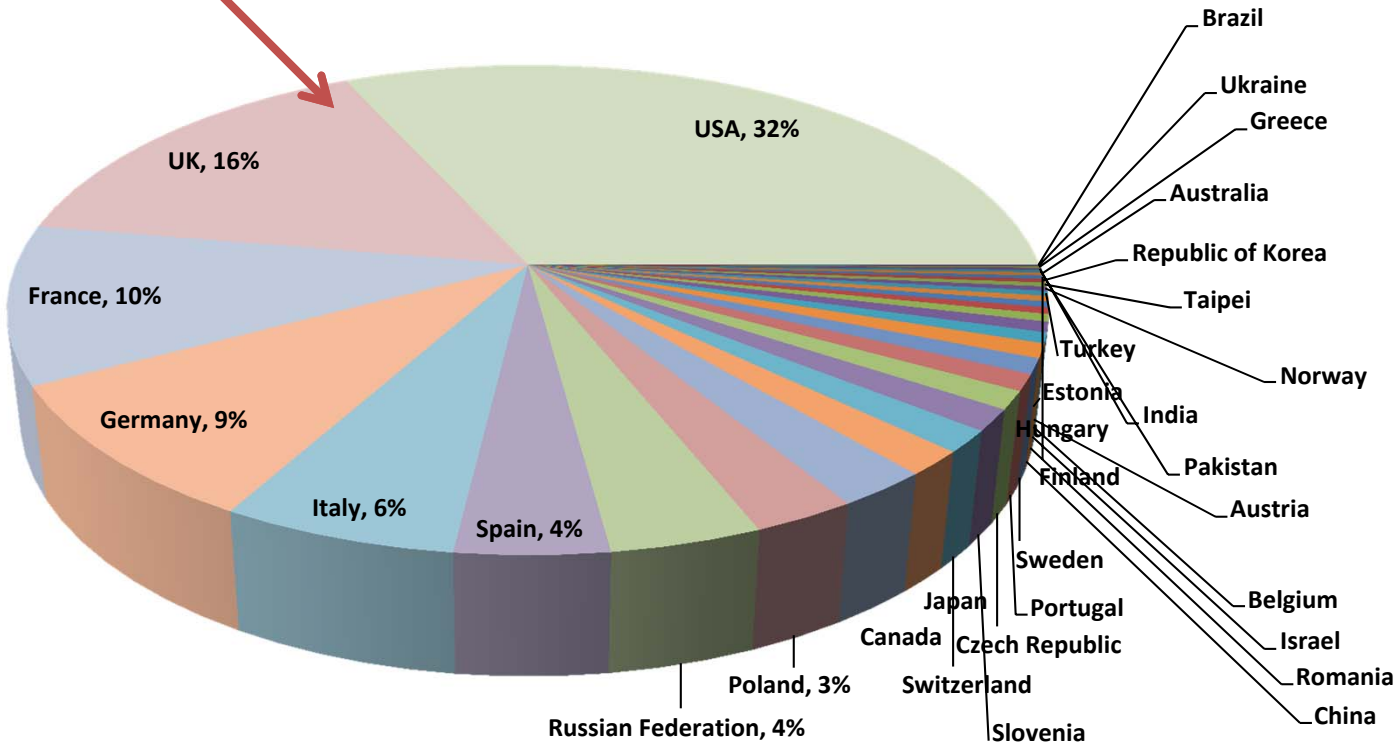

LHC @ 7 TeV

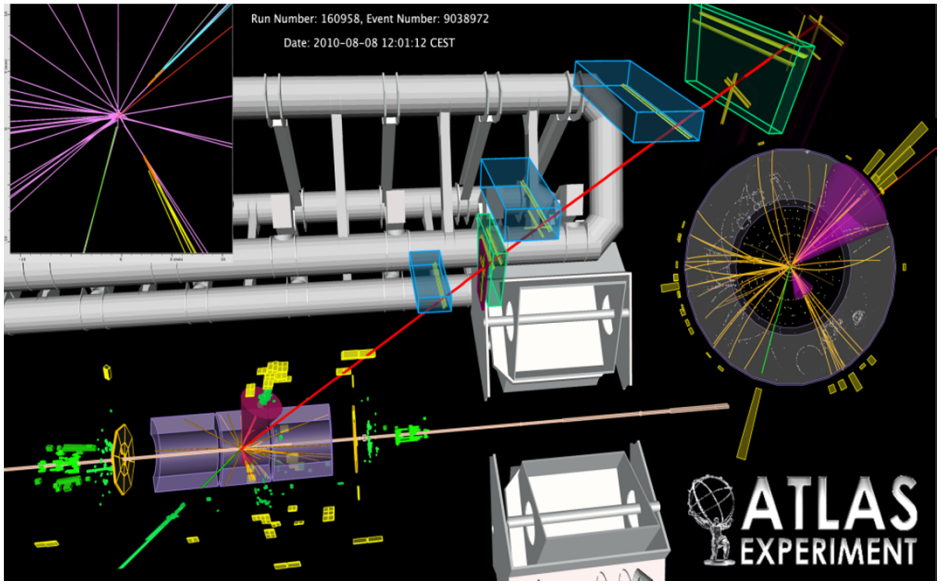
Run Number: 160958, Event Number: 9038972
Date: 2010-08-08 12:01:12 CEST
ATLAS
EXPERIMENT

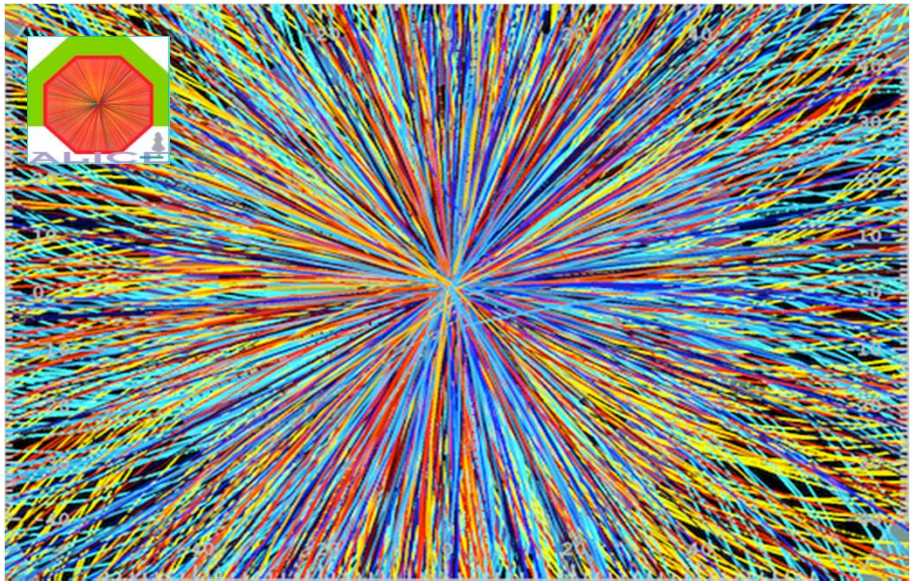

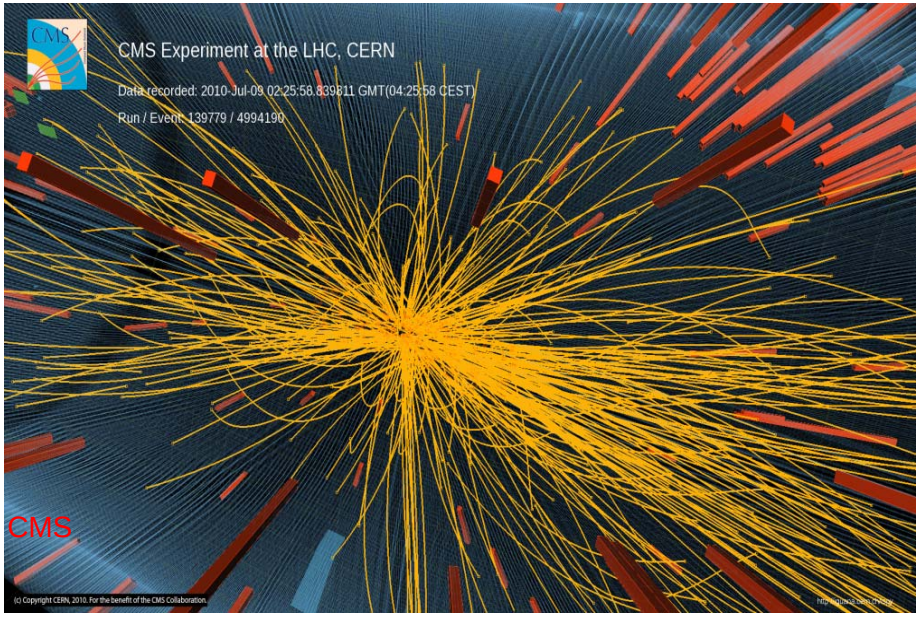
CMS Experiment at the LHC, CERN
Data recorded: 2010-Jul-09 02:25:58.839811 GMT(04:25:58 CEST)
Run / Event: 139779 / 4994190
CMS
(c) Copyright CERN, 2010. For the benefit of the CMS Collaboration.

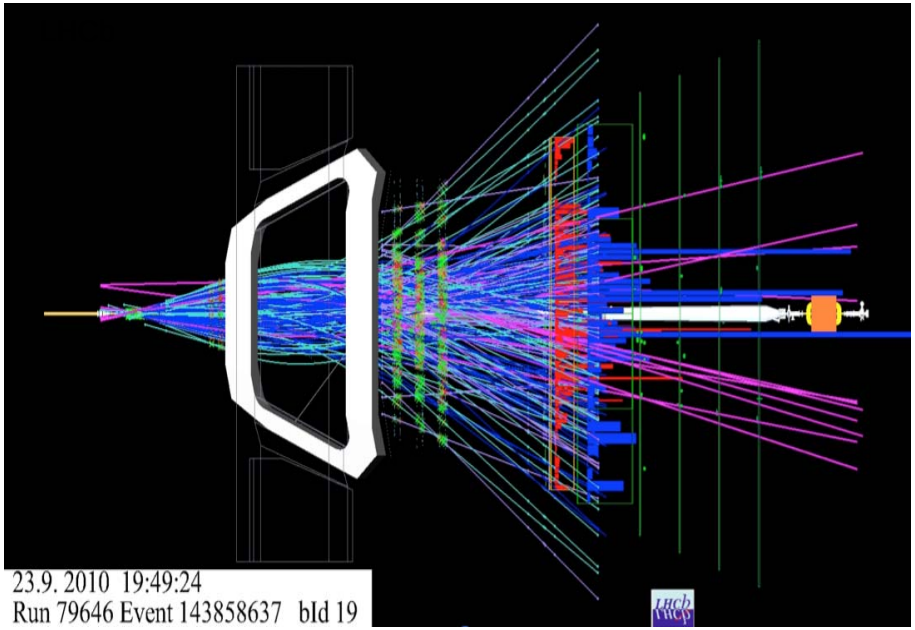
23.9. 2010 19:49:24
Run 79646 Event 143858637 bId 19

How to evolve LHC data processing

Making what we have today more sustainable is a challenge

- Data issues
  - Data management and access
  - How to make reliable and fault tolerant systems
  - Data preservation and open access
- Need to adapt to changing technologies
  - Use of many-core CPUs
  - Global filesystem-like facilities
  - Virtualisation
  - Commercial cloud computing services (IaaS/SaaS/PaaS)
- Network infrastructure
  - Has proved to be very reliable so invest in networks and make full use of the distributed system

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 (2012)

A European cloud computing partnership: big science teams up with big business

HELIX NEBULA THESCIENCECLOUD

Strategic Plan

- Establish multi-tenant, multi-provider cloud infrastructure
- Identify and adopt policies for trust, security and privacy
- Create governance structure
- Define funding schemes

CERN

To support the computing capacity needs for the ATLAS experiment

EMBL

Setting up a new service to simplify analysis of large genomes, for a deeper insight into evolution and biodiversity

esa

To create an Earth Observation platform, focusing on earthquake and volcano research

Atos · Capgemini CONSULTING.TECHNOLOGY.OUTSOURCING · CloudSigma · CSA cloud security alliance · egi · interoute from the ground to the cloud · logica be brilliant together · OpenNebula.org The Open Source Toolkit for Cloud Computing · orange Business Services · SAP · the SERVER LABS the IT architects · sixsq · Telefónica · terradue 2.0 · THALES · Trust IT · T-Systems

Email:contact@helix-nebula.eu Twitter: HelixNebulaSC Facebook: HelixNebula.TheScienceCloud


Thank you for your attention
LHCb
ATLAS
CERN Meyrin
CERN Prévessin
SPS 7 km
SUISSE
FRANCE
CMS
ALICE
LHC 27 km
CERN
Accelerating Science and Innovation