Big Data and Storage Management at the Large Hadron Collider

Dirk Duellmann
CERN IT, Data & Storage Services

Accelerating Science and Innovation

CERN was founded 1954: 12 European States

"Science for Peace"

Today: 21 Member States

~ 2,300 staff
~ 1,000 other paid personnel
> 11,000 users
Budget (2013) ~1,000 MCHF

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

Candidate for Accession: Romania

Associate Members in Pre-Stage to Membership: Serbia

Applicant States for Membership or Associate Membership: Brazil, Cyprus (awaiting ratification), Pakistan, Russia, Slovenia, Turkey, Ukraine

Observers to Council: India, Japan, Russia, Turkey, United States of America; European Commission and UNESCO

Global Science: 11000 scientists

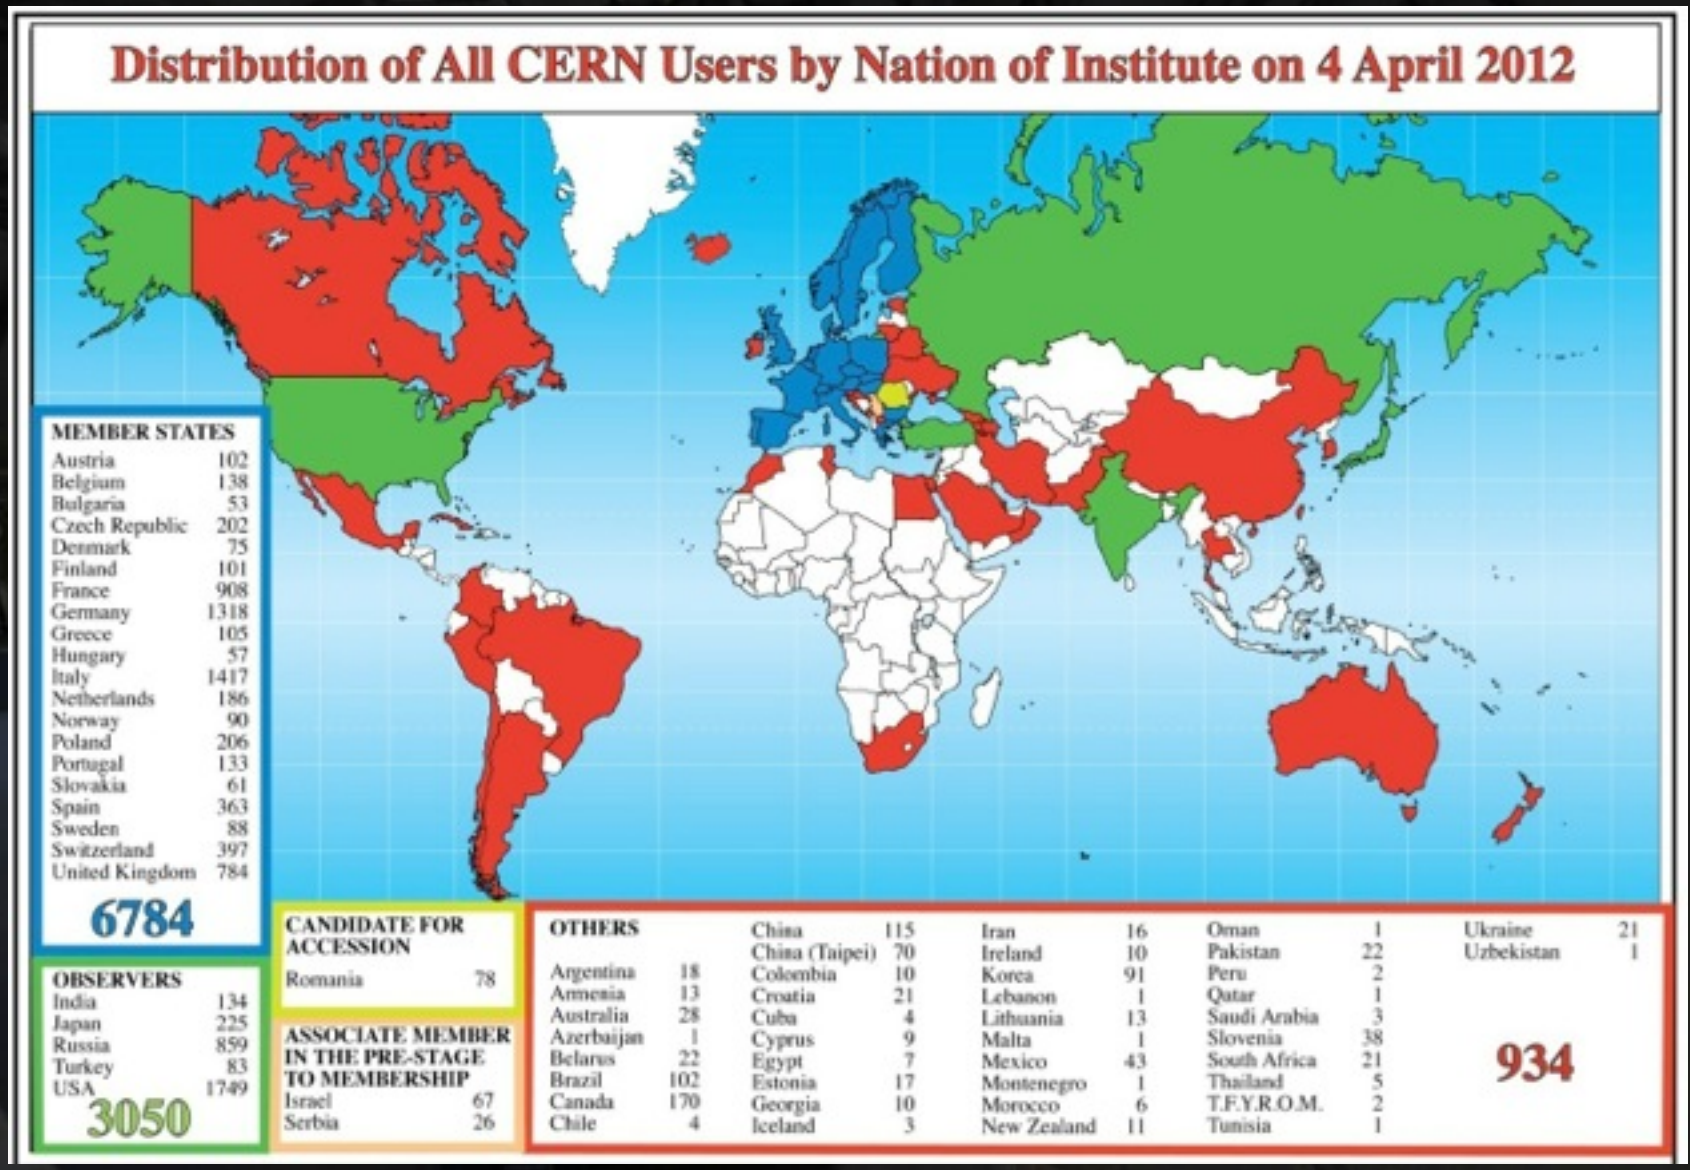

Stars and Planets only account for a small percentage of the universe!

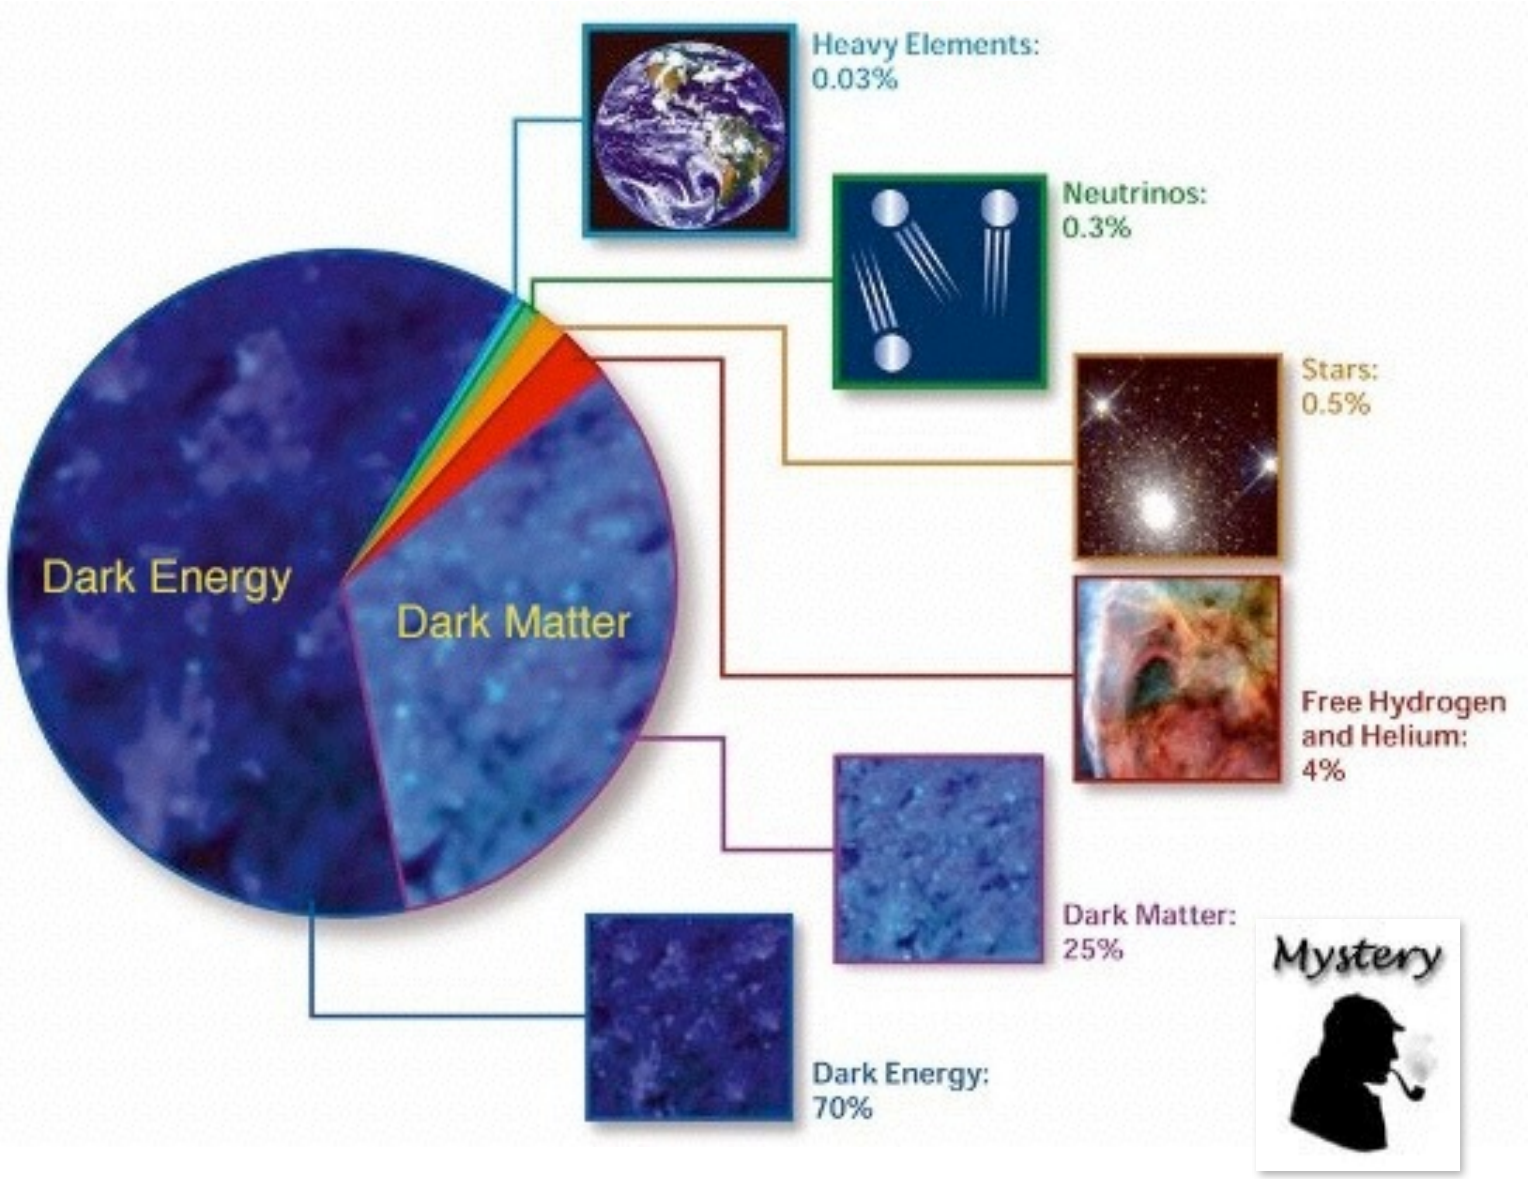

Collisions at the LHC: summary

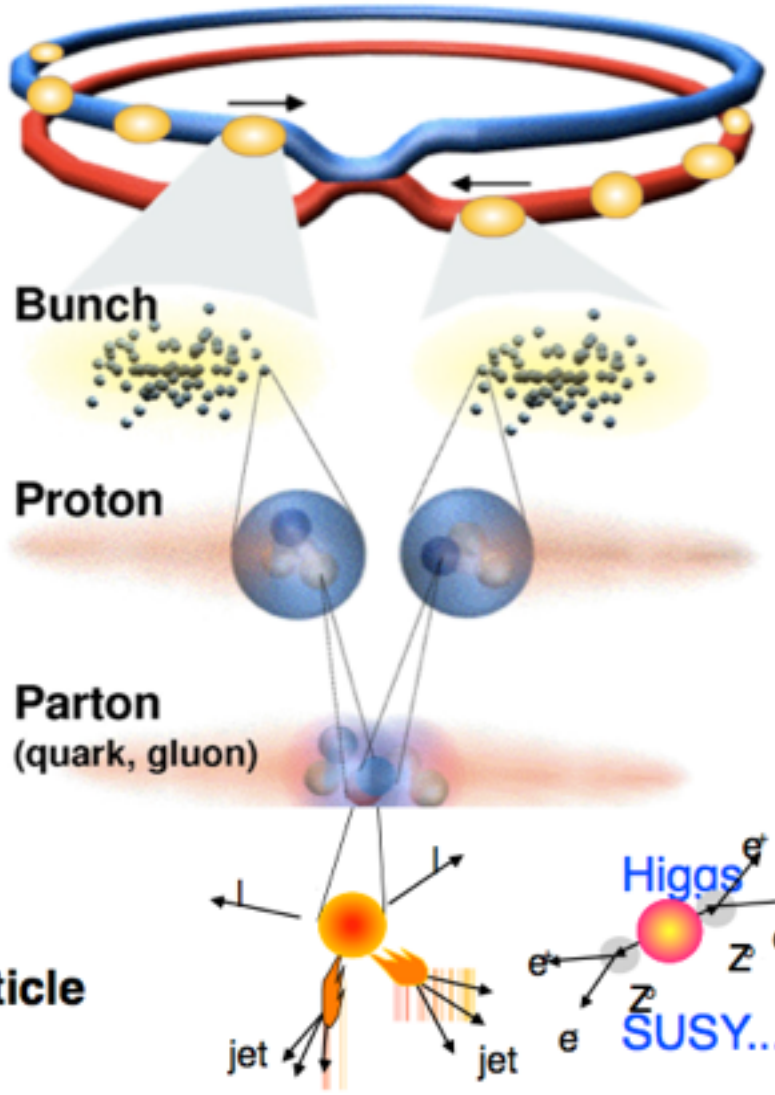

Proton - Proton	2808 bunch/beam
Protons/bunch	10^{11}
Beam energy	7 TeV (7×10^{12} eV)
Luminosity	$10^{34} cm^{-2} s^{-1}$
Crossing rate	40 MHz
Collision rate ≈	10^{7}-10^{9}

New physics rate ≈ .00001 Hz

Event selection:
1 in 10,000,000,000,000

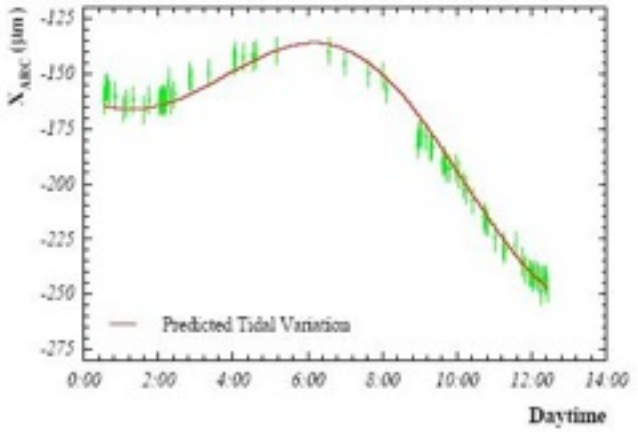

Tides

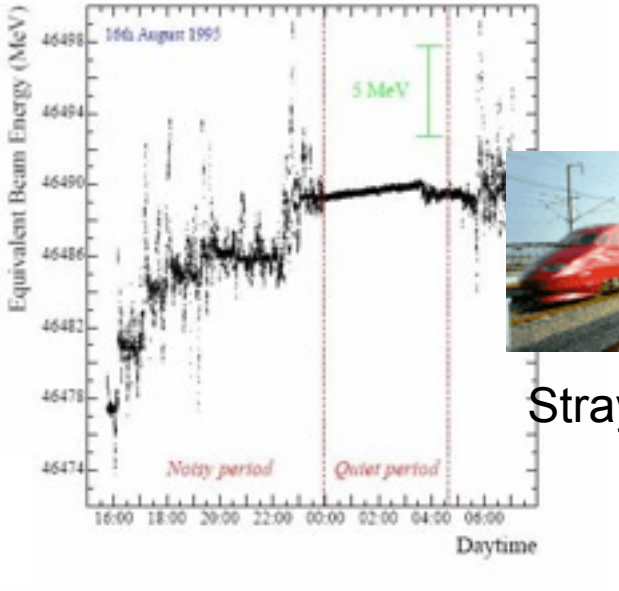

Stray currents

Precision !

The 27 km long ring is sensitive to <1mm changes

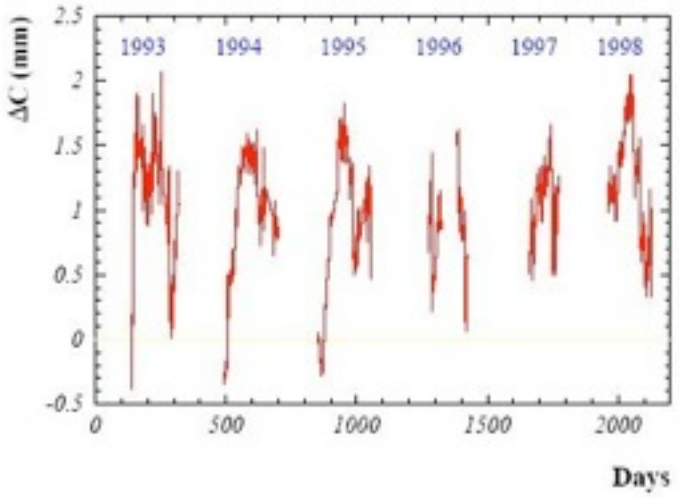

Rainfall

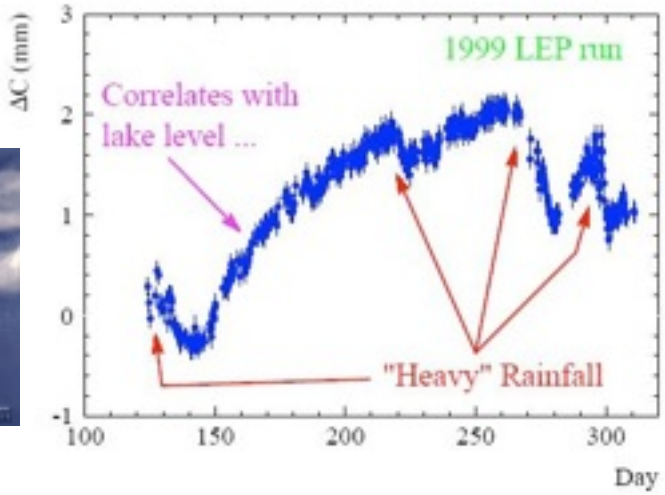

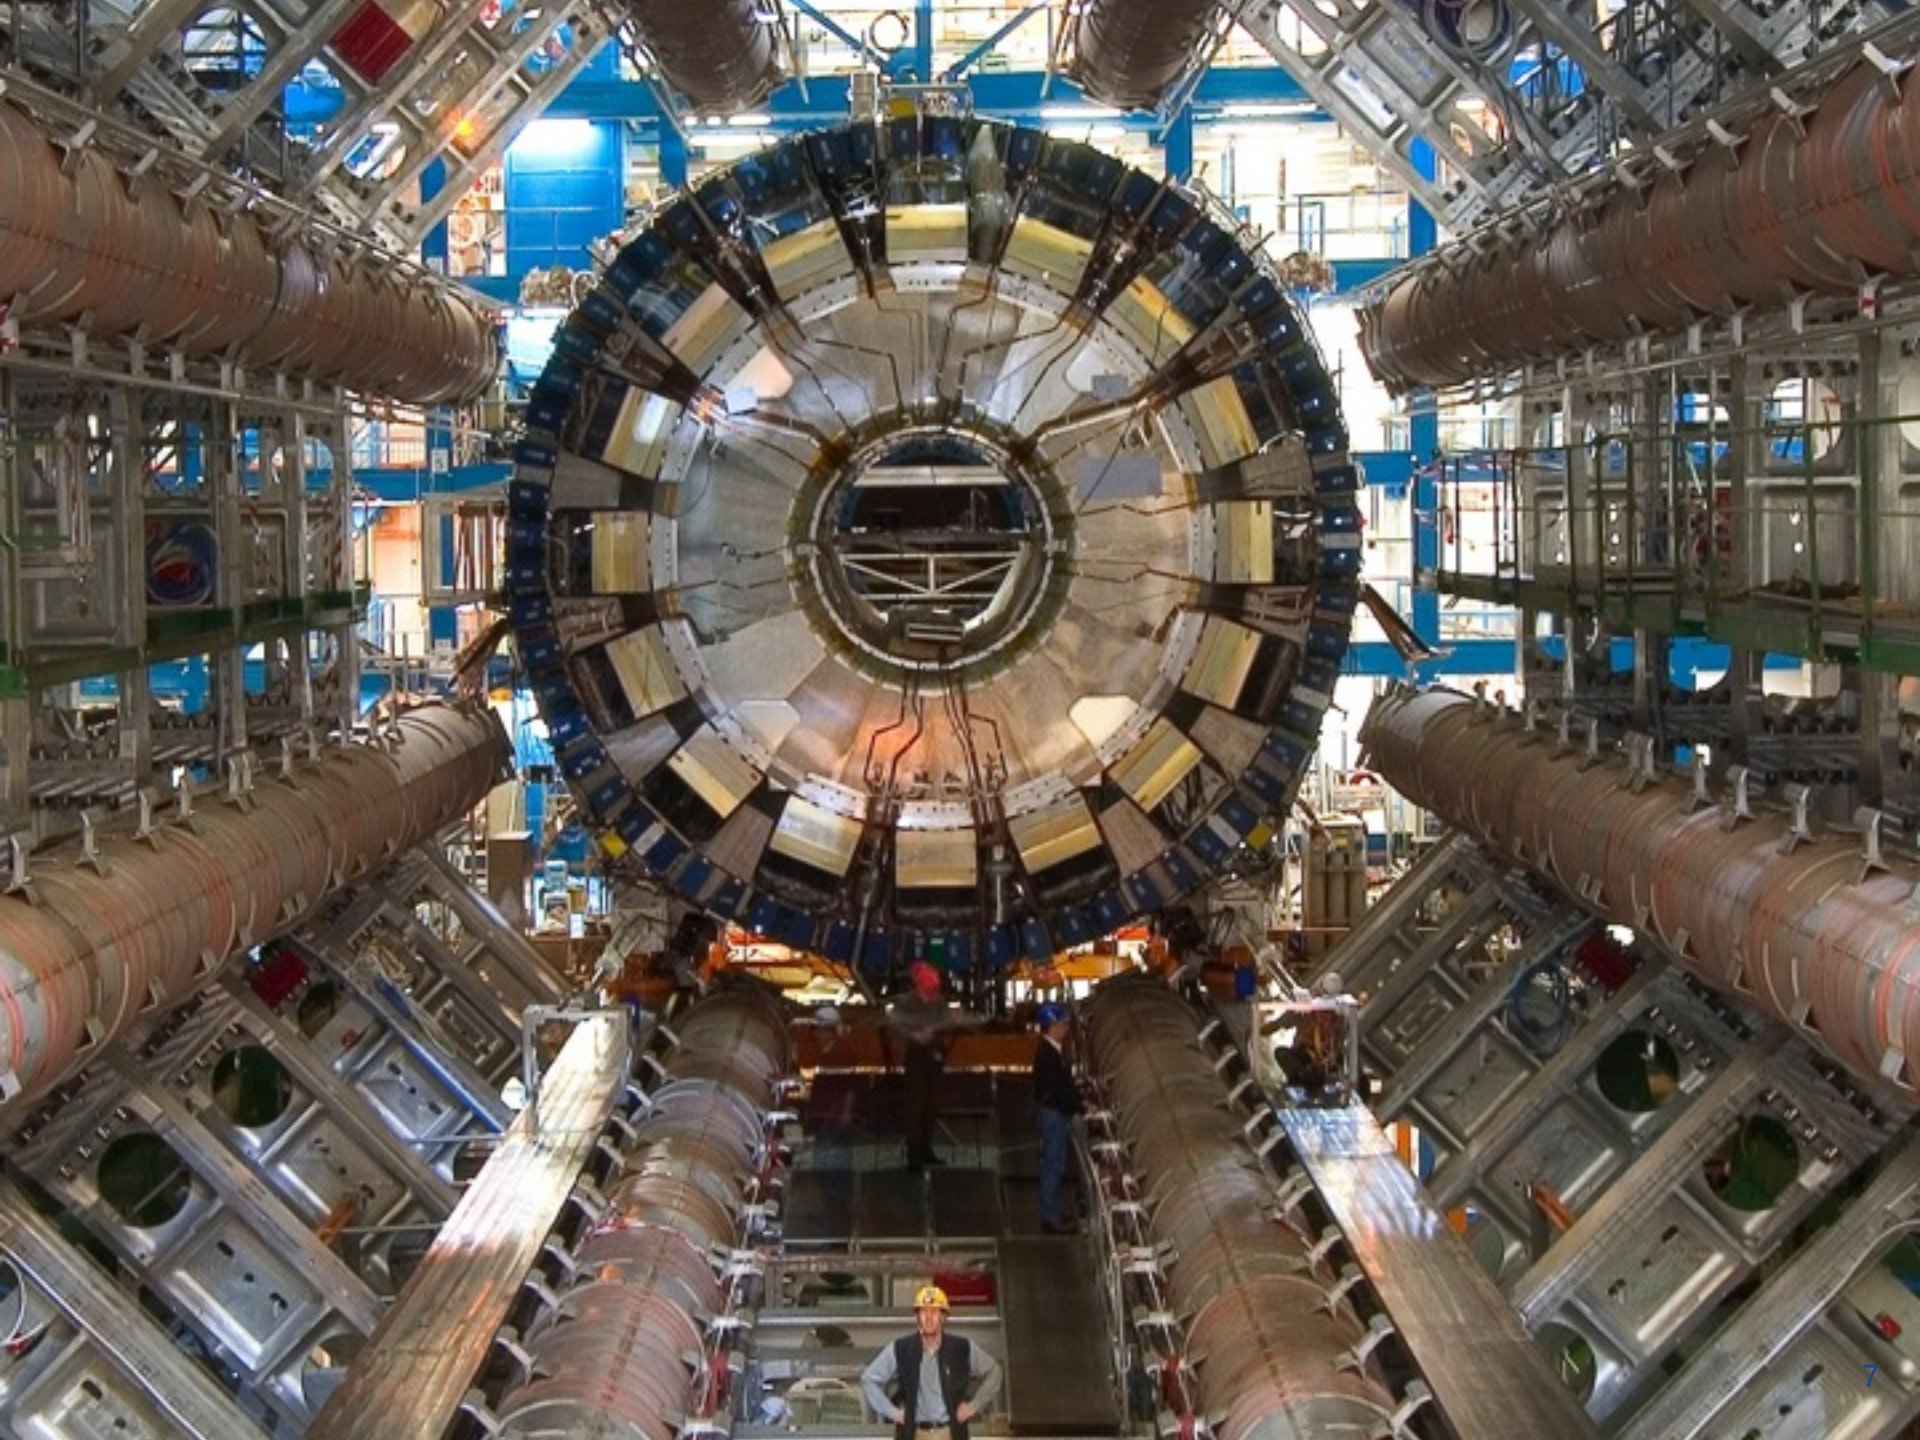

A collision at LHC

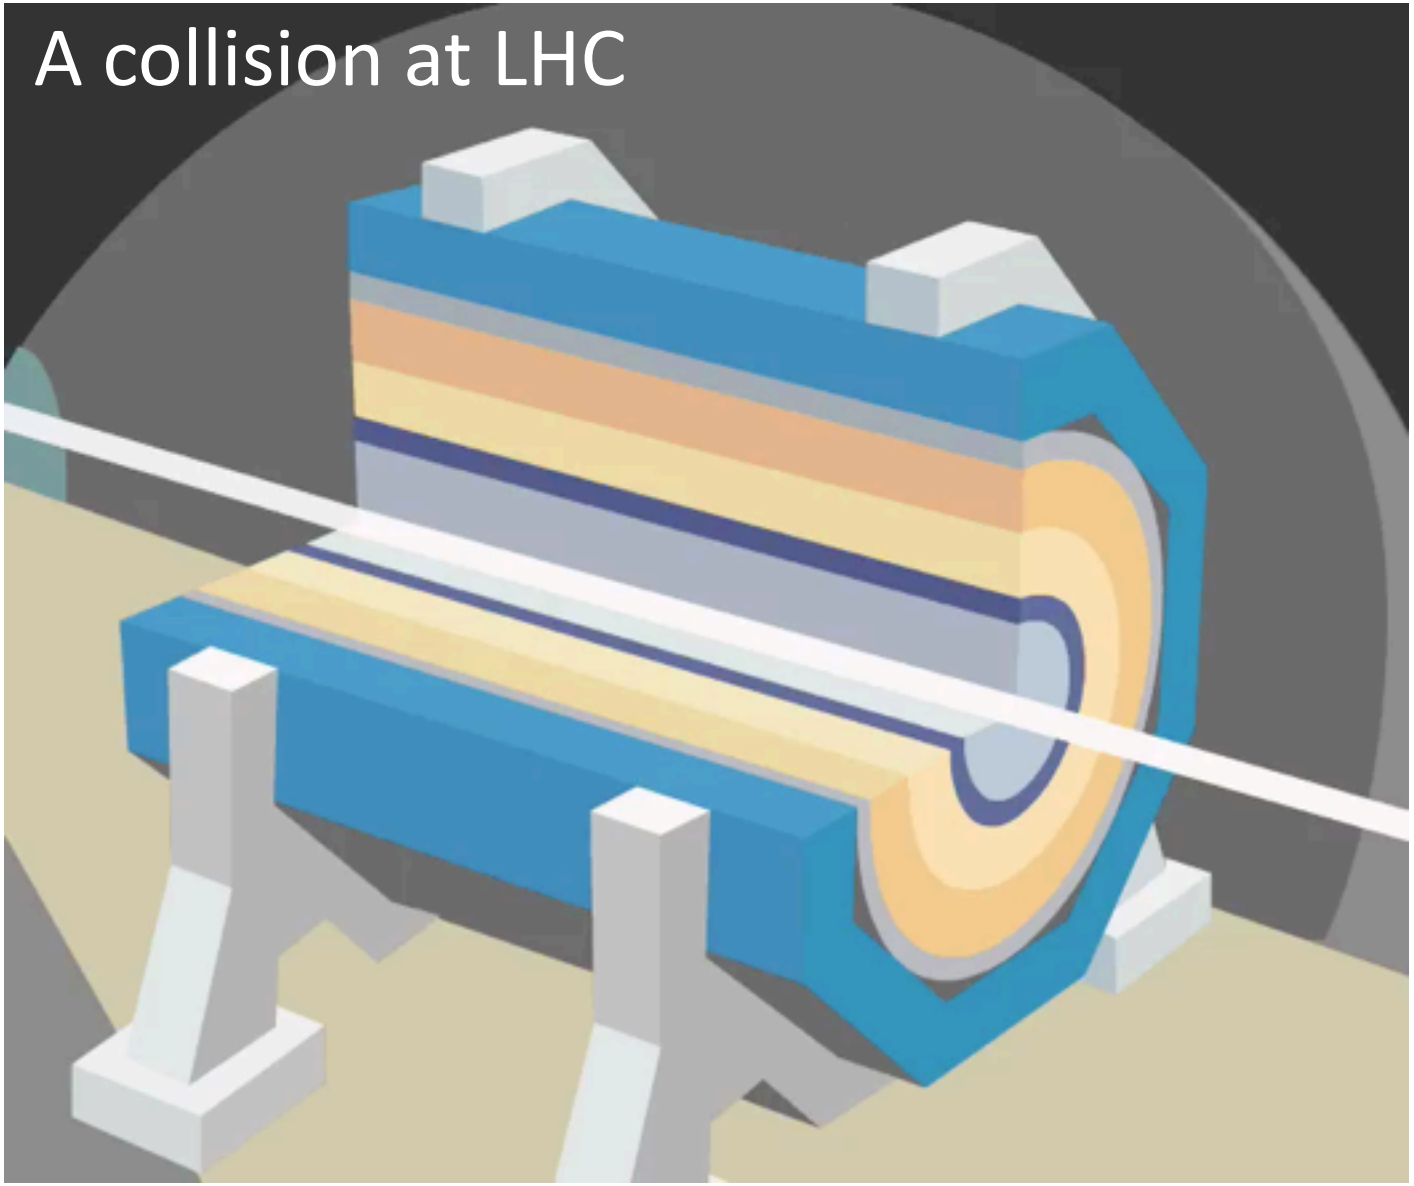

The Data Acquisition for one Detector

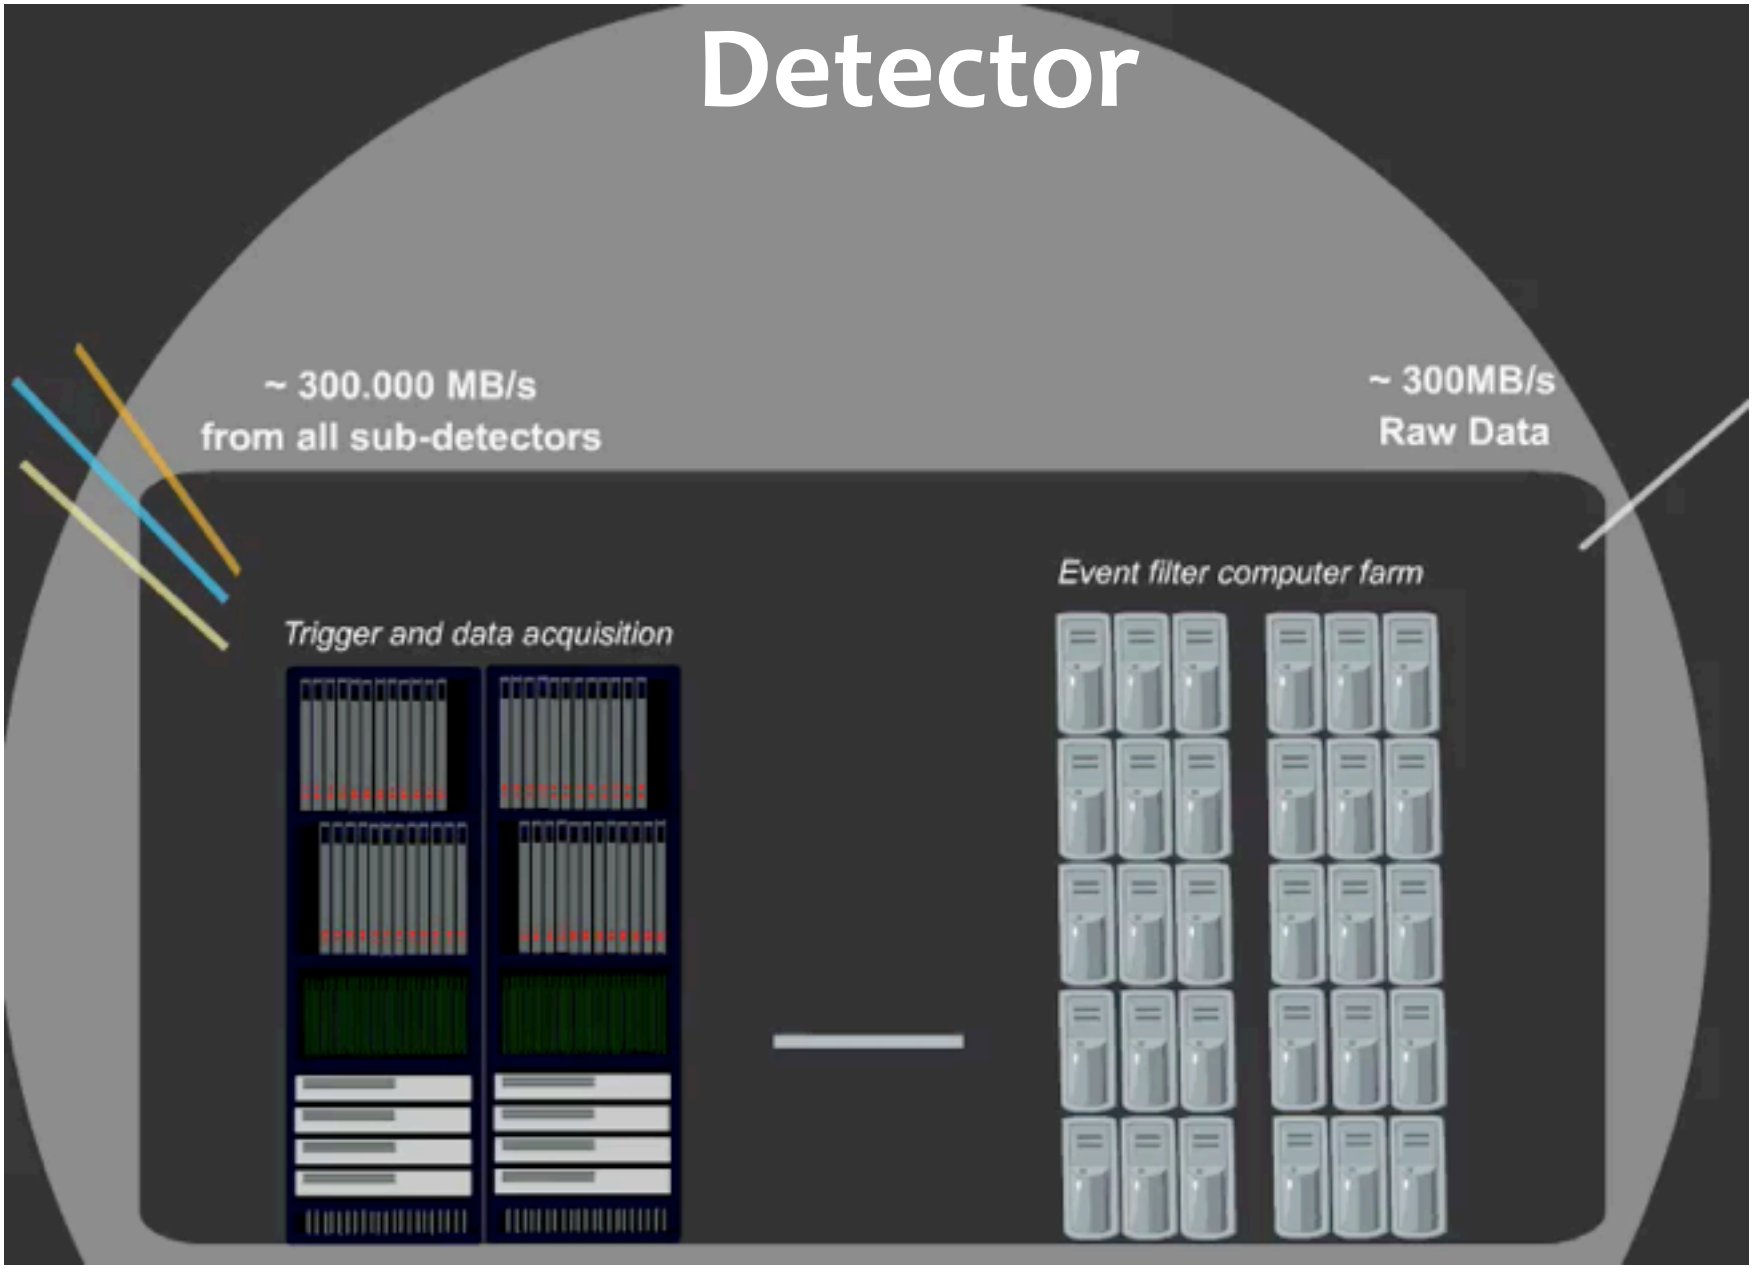

WLCG
Worldwide LHC Computing Grid

Tier 0 at CERN: Acquisition, First reconstruction, Storage & Distribution

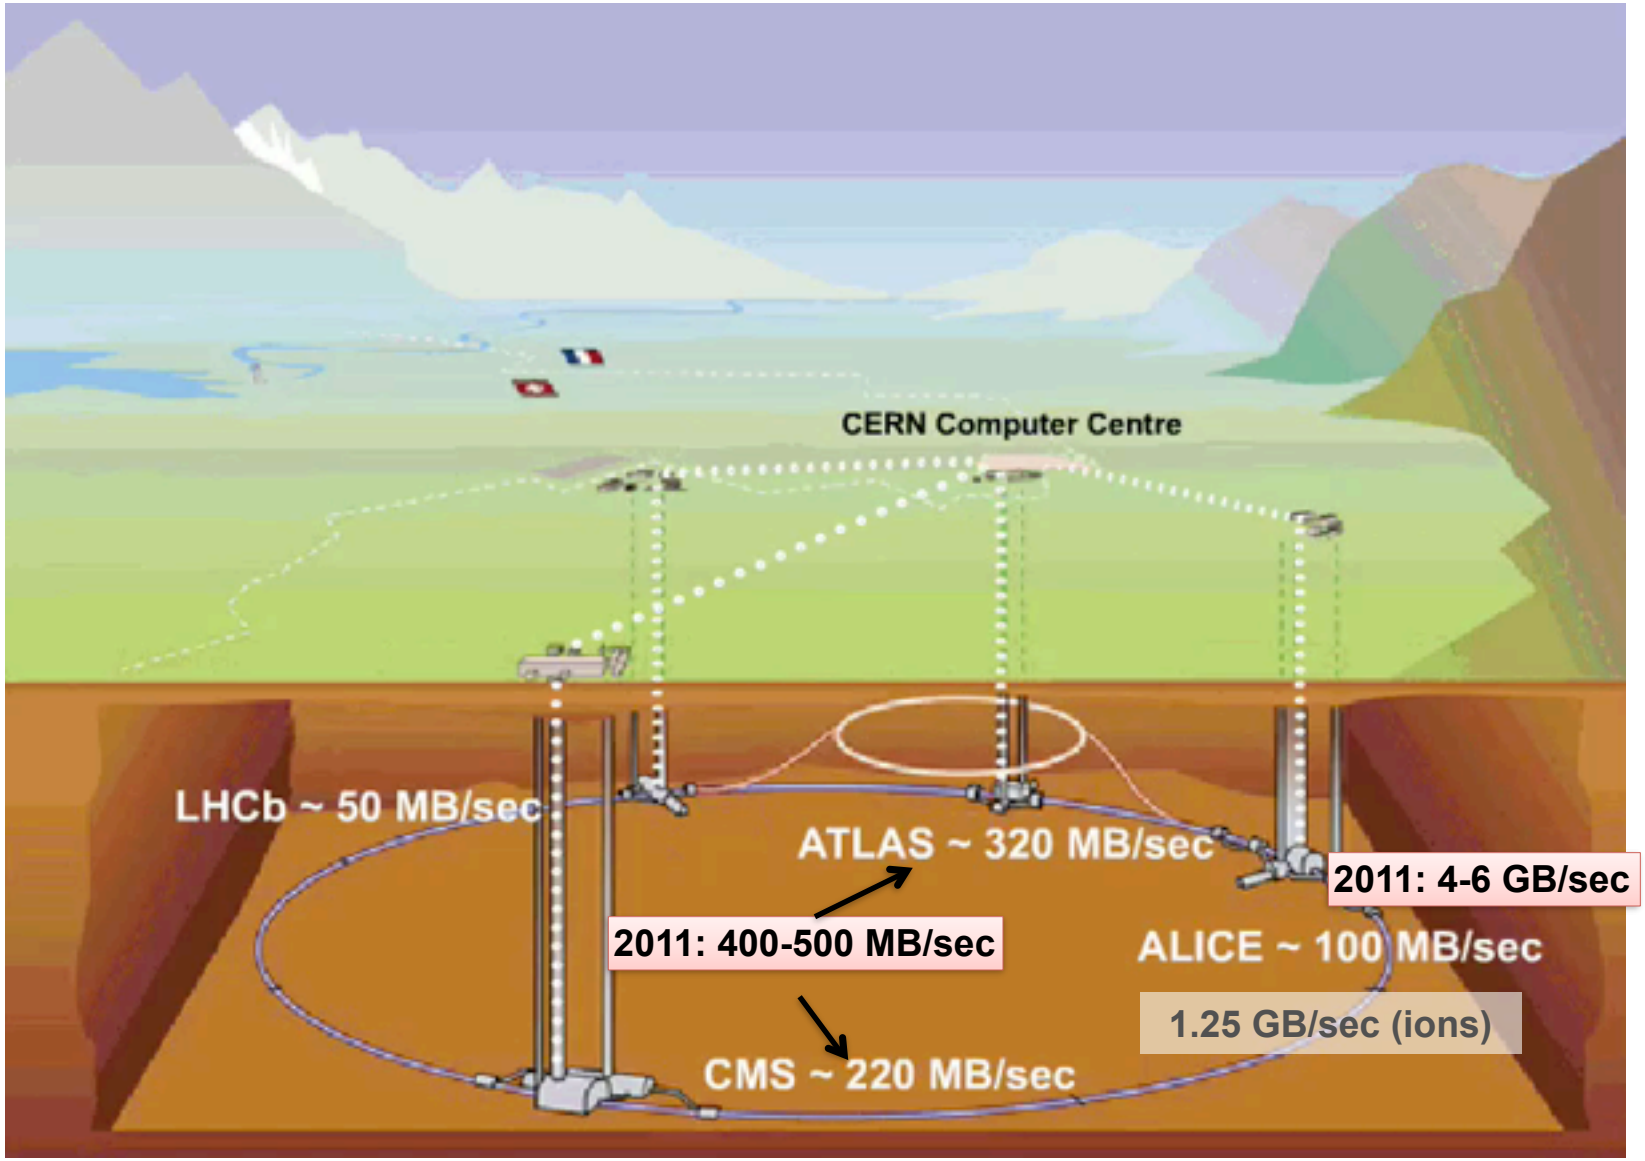

WLCG Worldwide LHC Computing Grid

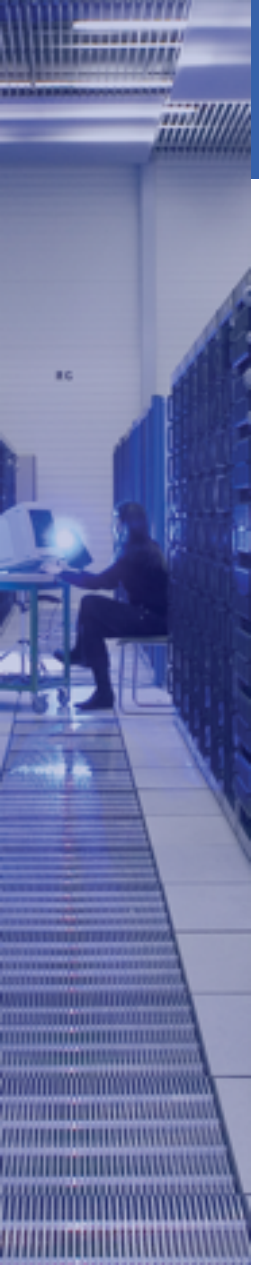

The LHC Data Challenge

CERN IT Department

- The accelerator will run for 20 years
- Experiments are producing about **25 Million Gigabytes** of data each year (about 3 million DVDs – 850 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

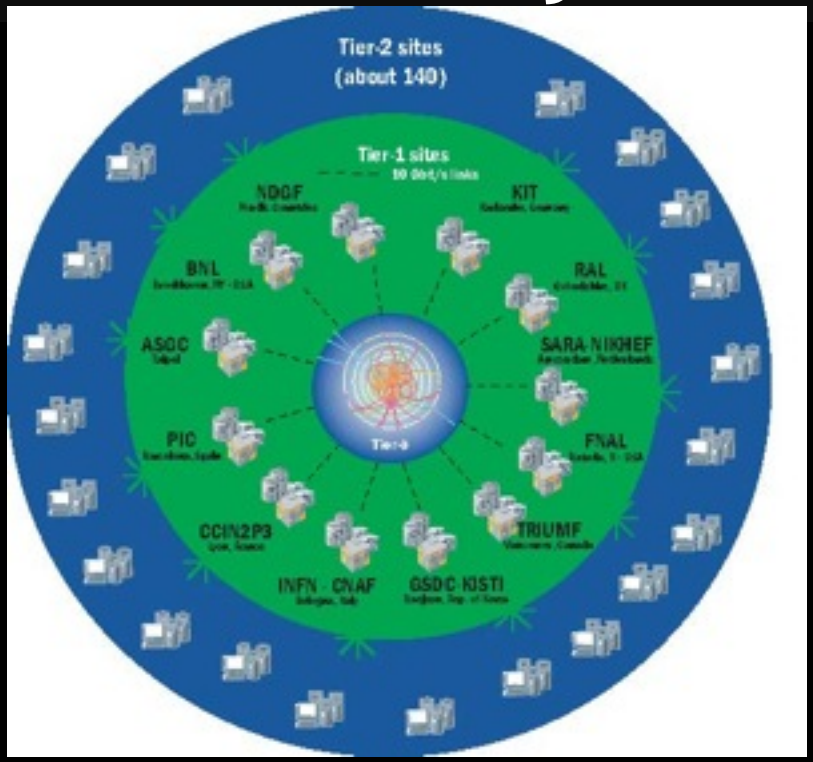

Tier-0 (CERN):
- Data recording
- Initial data reconstruction
- Data distribution

Tier-1 (12 centres + Russia):
- Permanent storage
- Re-processing
- Analysis

Tier-2 (~140 centres):
- Simulation
- End-user analysis

- **~ 160 sites, 35 countries**
- **300000 cores**
- **200 PB of storage**
- **2 Million jobs/day**
- **10 Gbps links**

Data 2008-2013

CERN Tape Writes

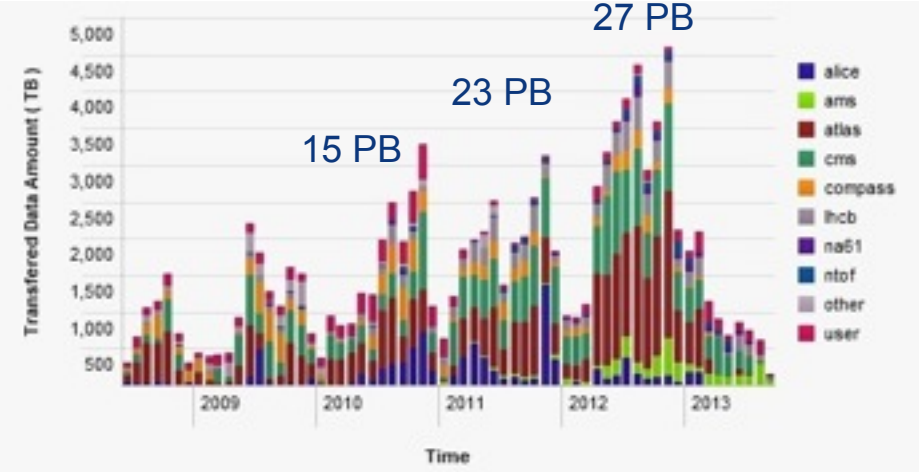

Tape Usage Breakdown

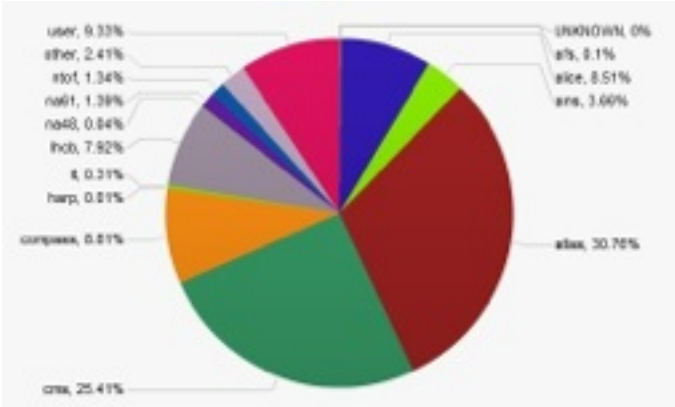

CERN Tape Verification

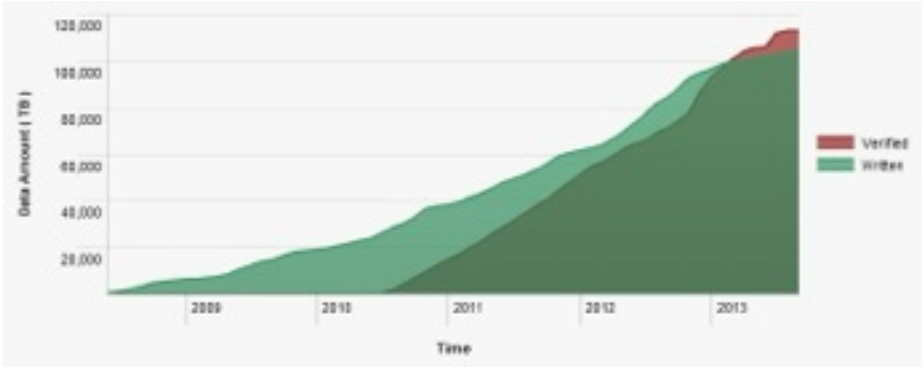

Data Loss: ~65 GB over 69 tapes
Duration: ~2.5 years

CERN Tape Archive

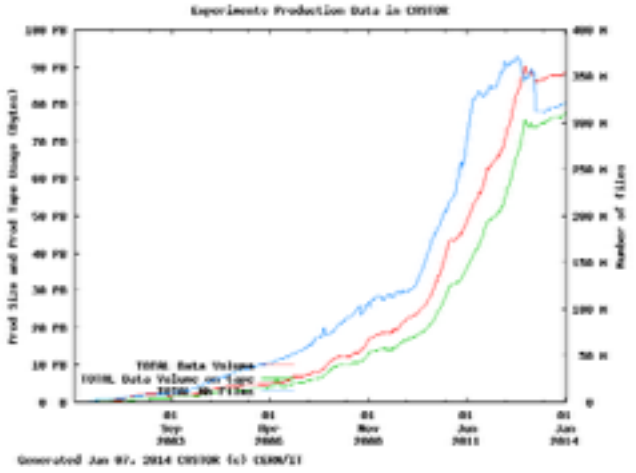

Data transfers

- Global transfer rates are always significant (12-15 GB/s) – permanent on-going workloads
- CERN export rates driven (mostly) by LHC data export

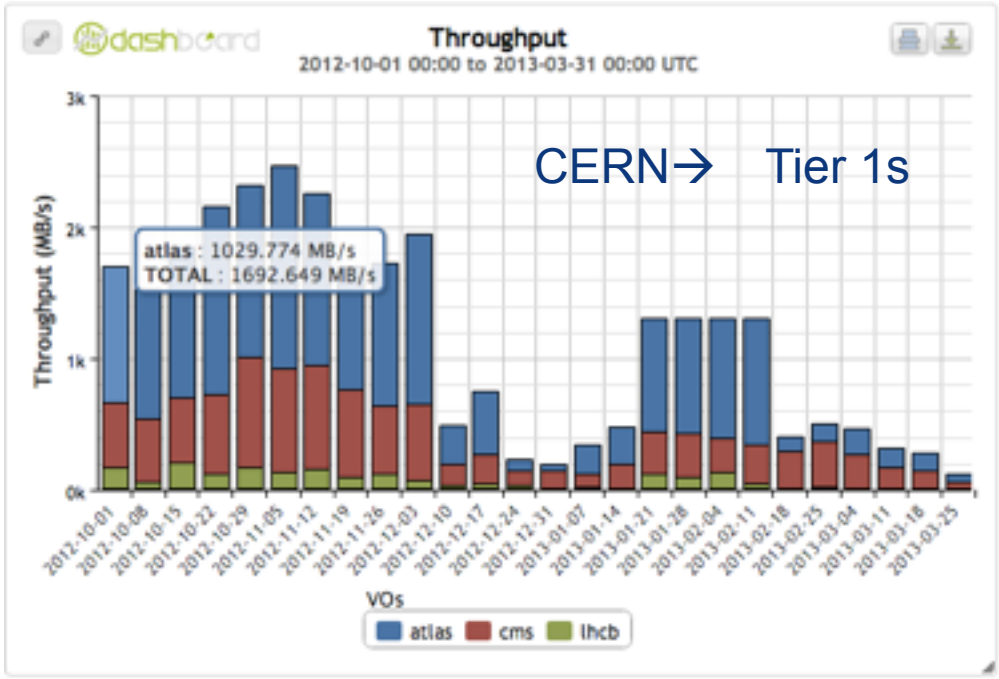

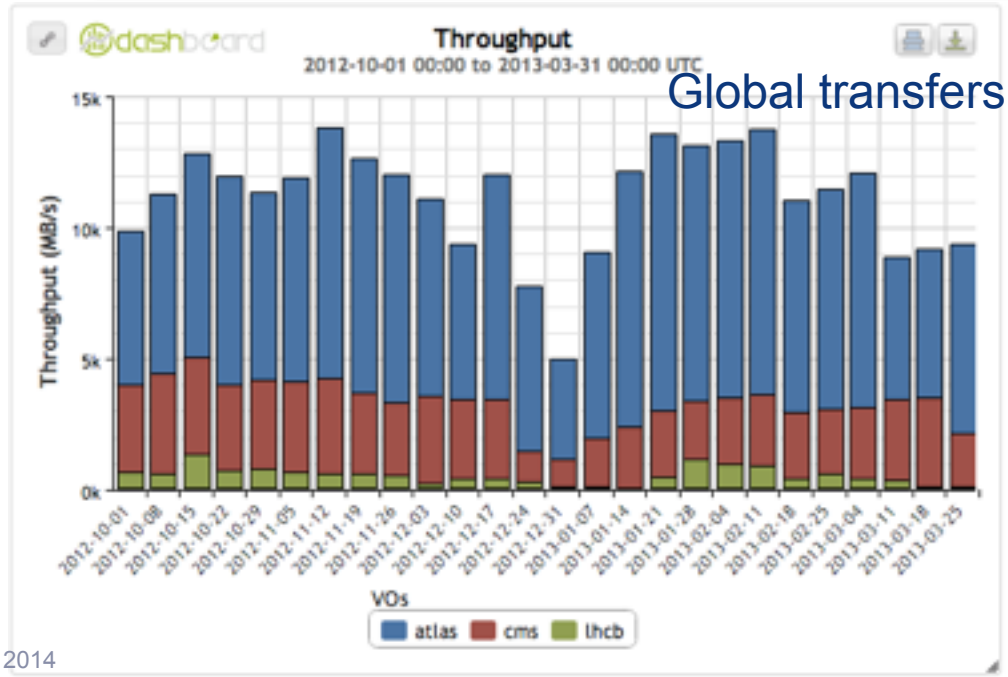

No stop for the computing !

Activity on 1 January 2014
Running Jobs: 223509
Transfer rate: ~2.5 GiB/s

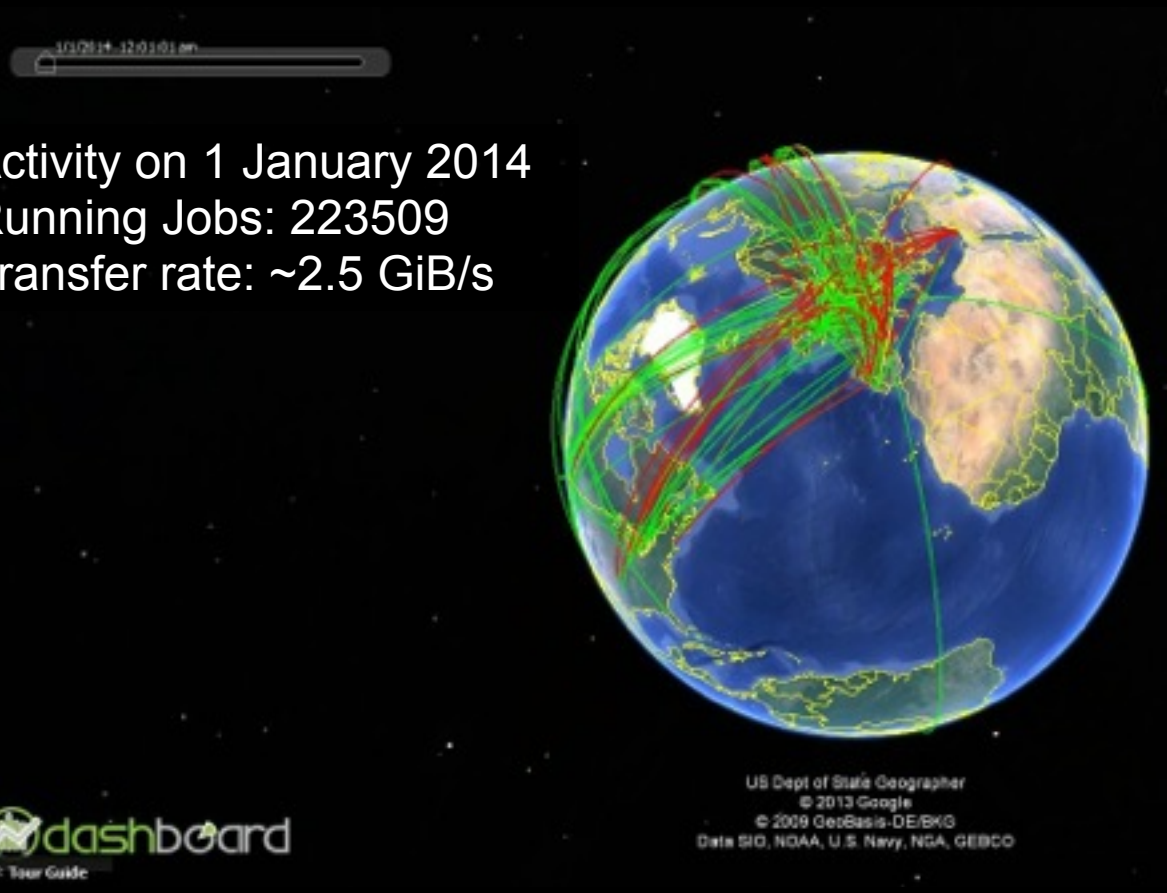

WLCG
Worldwide LHC Computing Grid

Processing on the Grid

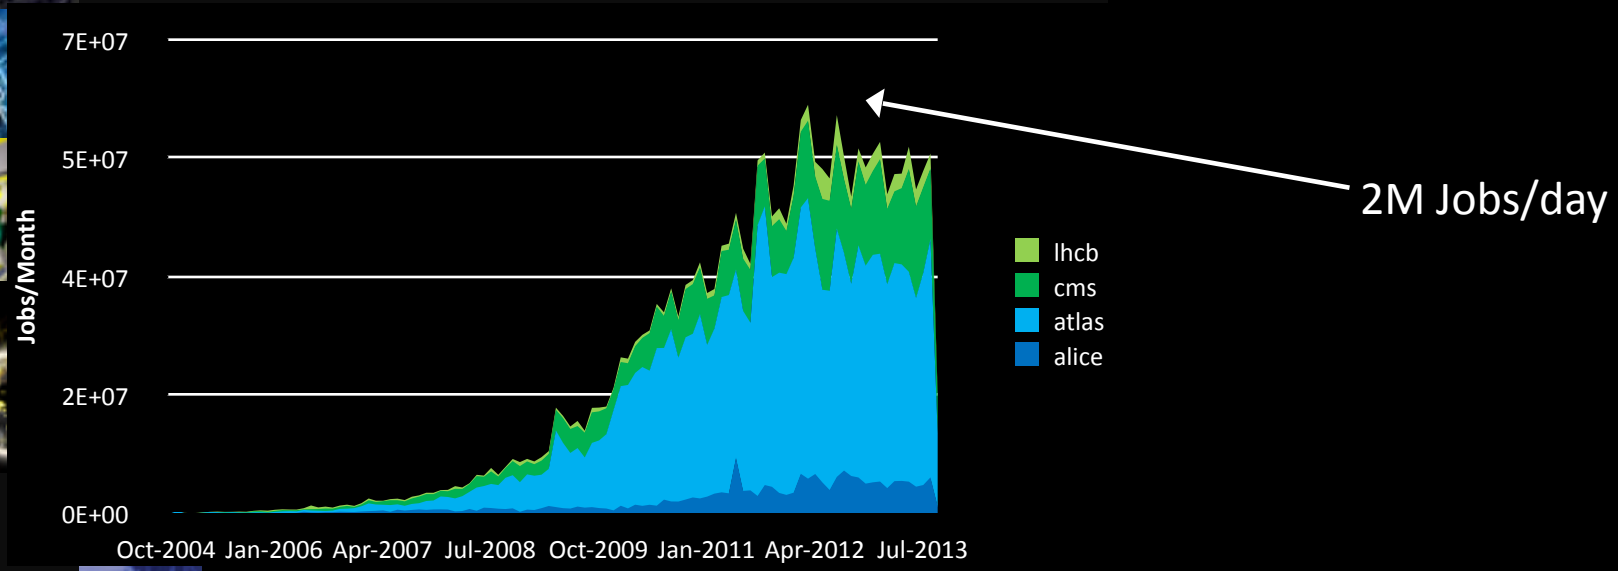

1.4 10^9 HEPSPEC06/Month
(210 K CPU continuous use)

Close to full capacity

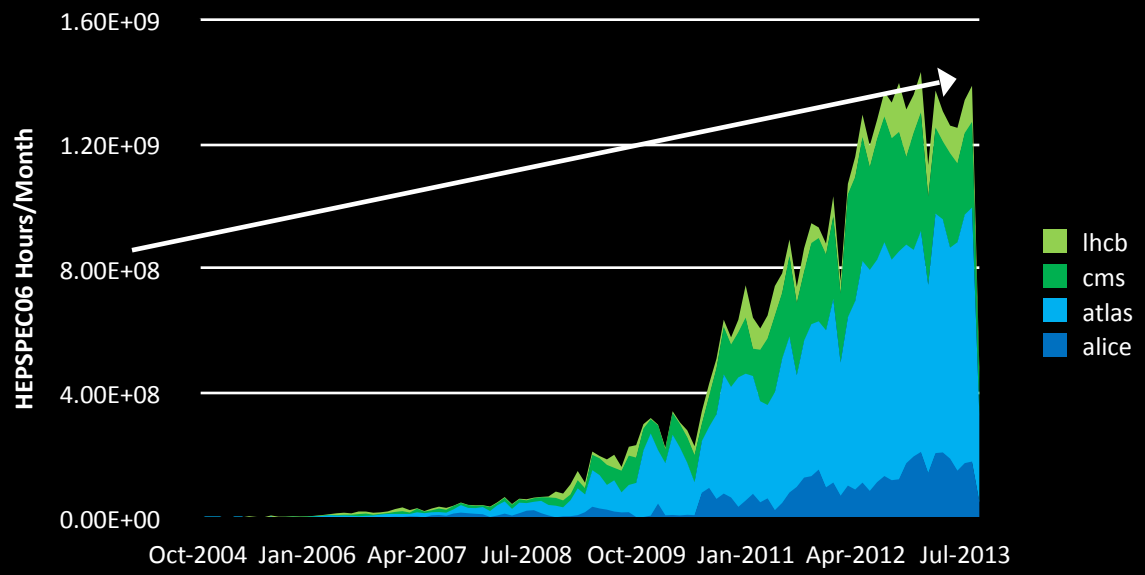

Broader Impact of the LHC Computing Grid

- WLCG has been leveraged on both sides of the Atlantic, to benefit the wider scientific community
  - Europe:
    - Enabling Grids for E-sciencE (EGEE) 2004-2010
    - European Grid Infrastructure (EGI) 2010--
  - USA:
    - Open Science Grid (OSG) 2006-2012 (+ extension?)
- Many scientific applications →

Archeology
Astronomy
Astrophysics
Civil Protection
Comp. Chemistry
Earth Sciences
Finance
Fusion
Geophysics
High Energy Physics
Life Sciences
Multimedia
Material Sciences
…

Virtuous Cycle

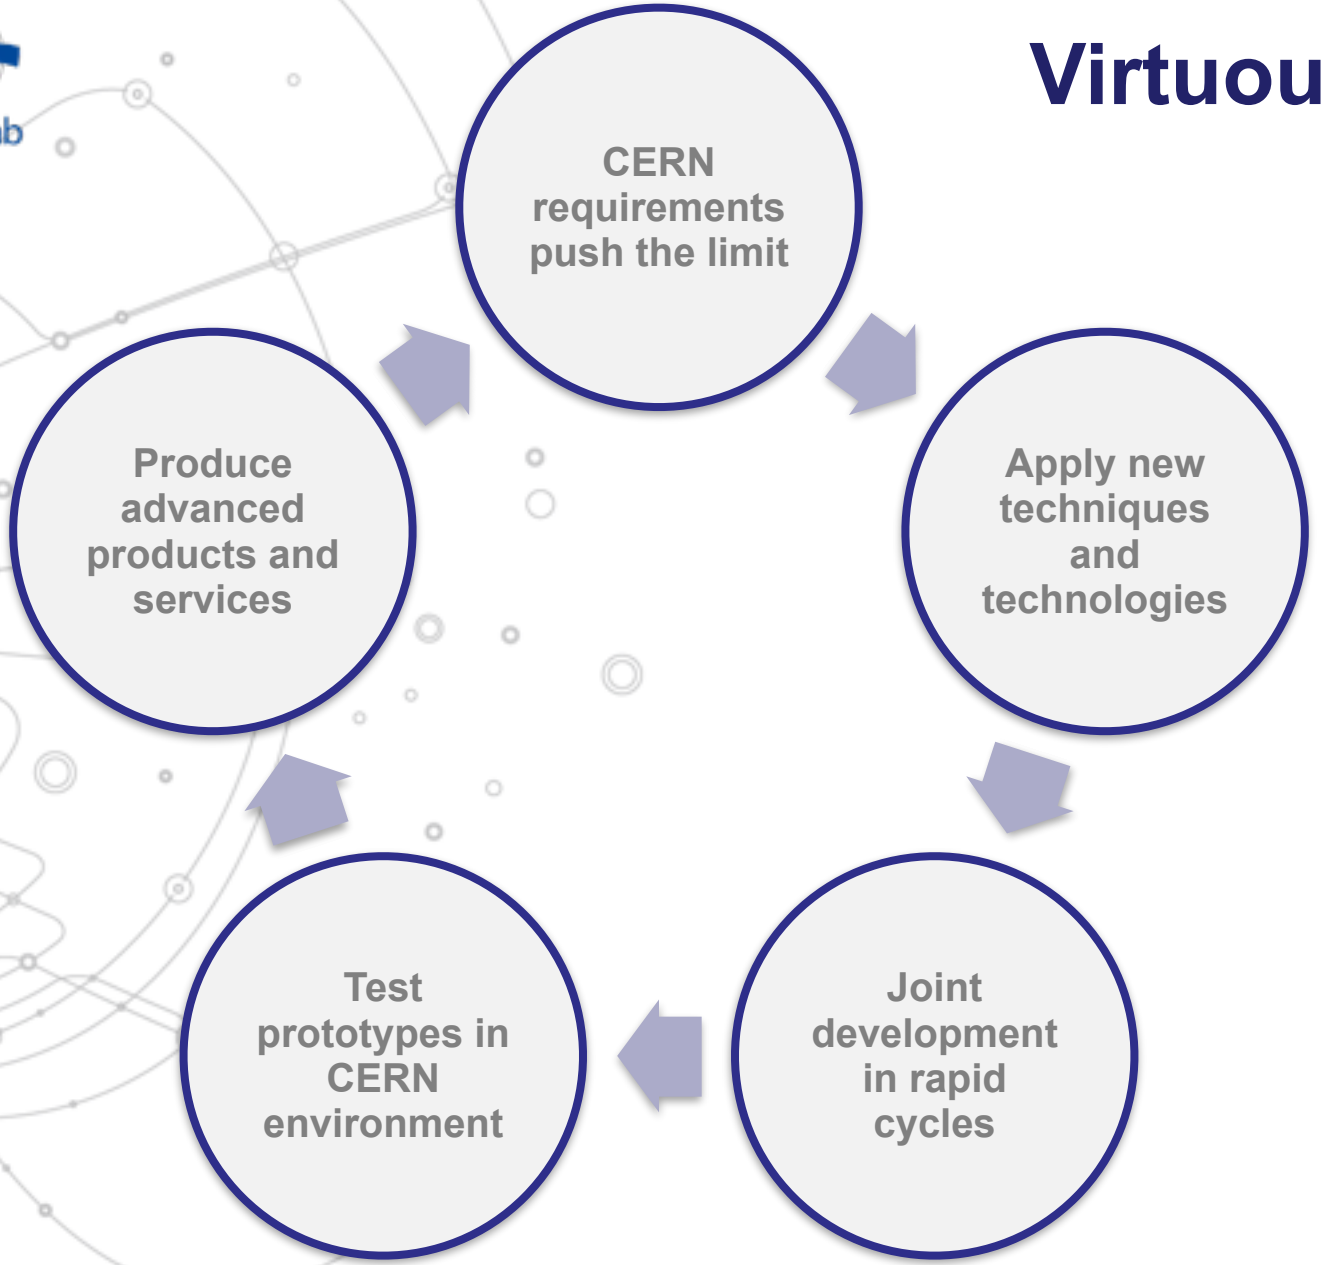

A public-private partnership between the research community and industry

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

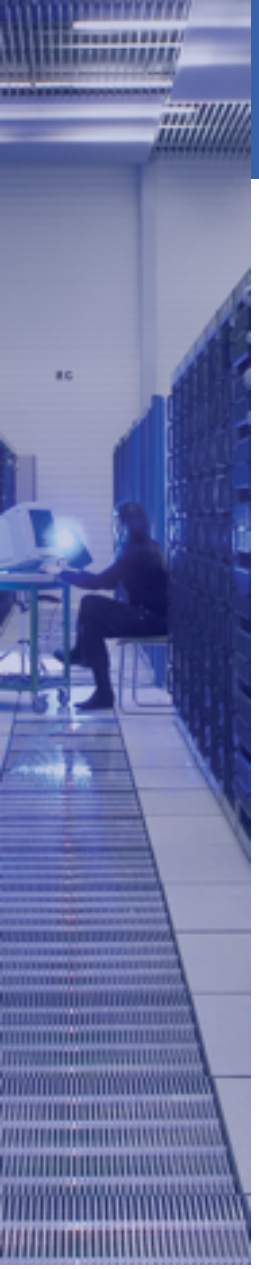

The CERN Data Centre in Numbers

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	*1127*
Servers	*10,070*
Processors	*17,259*
Cores	*90,948*
HEPSpec06	*744,277*

Disks	*75,718*
Raw disk capacity (TiB)	*113,852*
Memory modules	*64035*
Memory capacity (TiB)	*312*
RAID controllers	*3,091*

Tape Drives	*120*
Tape Cartridges	*52000*
Tape slots	*66000*
Data on Tape (PiB)	*75*

High Speed Routers	*29*
Ethernet Switches	*874*
10 Gbps/100Gbps ports	*1396/74*
Switching Capacity	*6 Tbps*
1 Gbps ports	*27984*
10 Gbps ports	*5664*

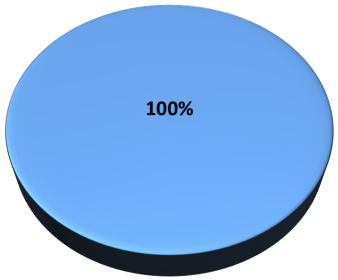

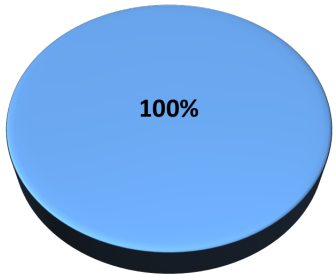

IT Power Consumption	*2392 KW*
Total Power Consumption	*3929 KW*

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

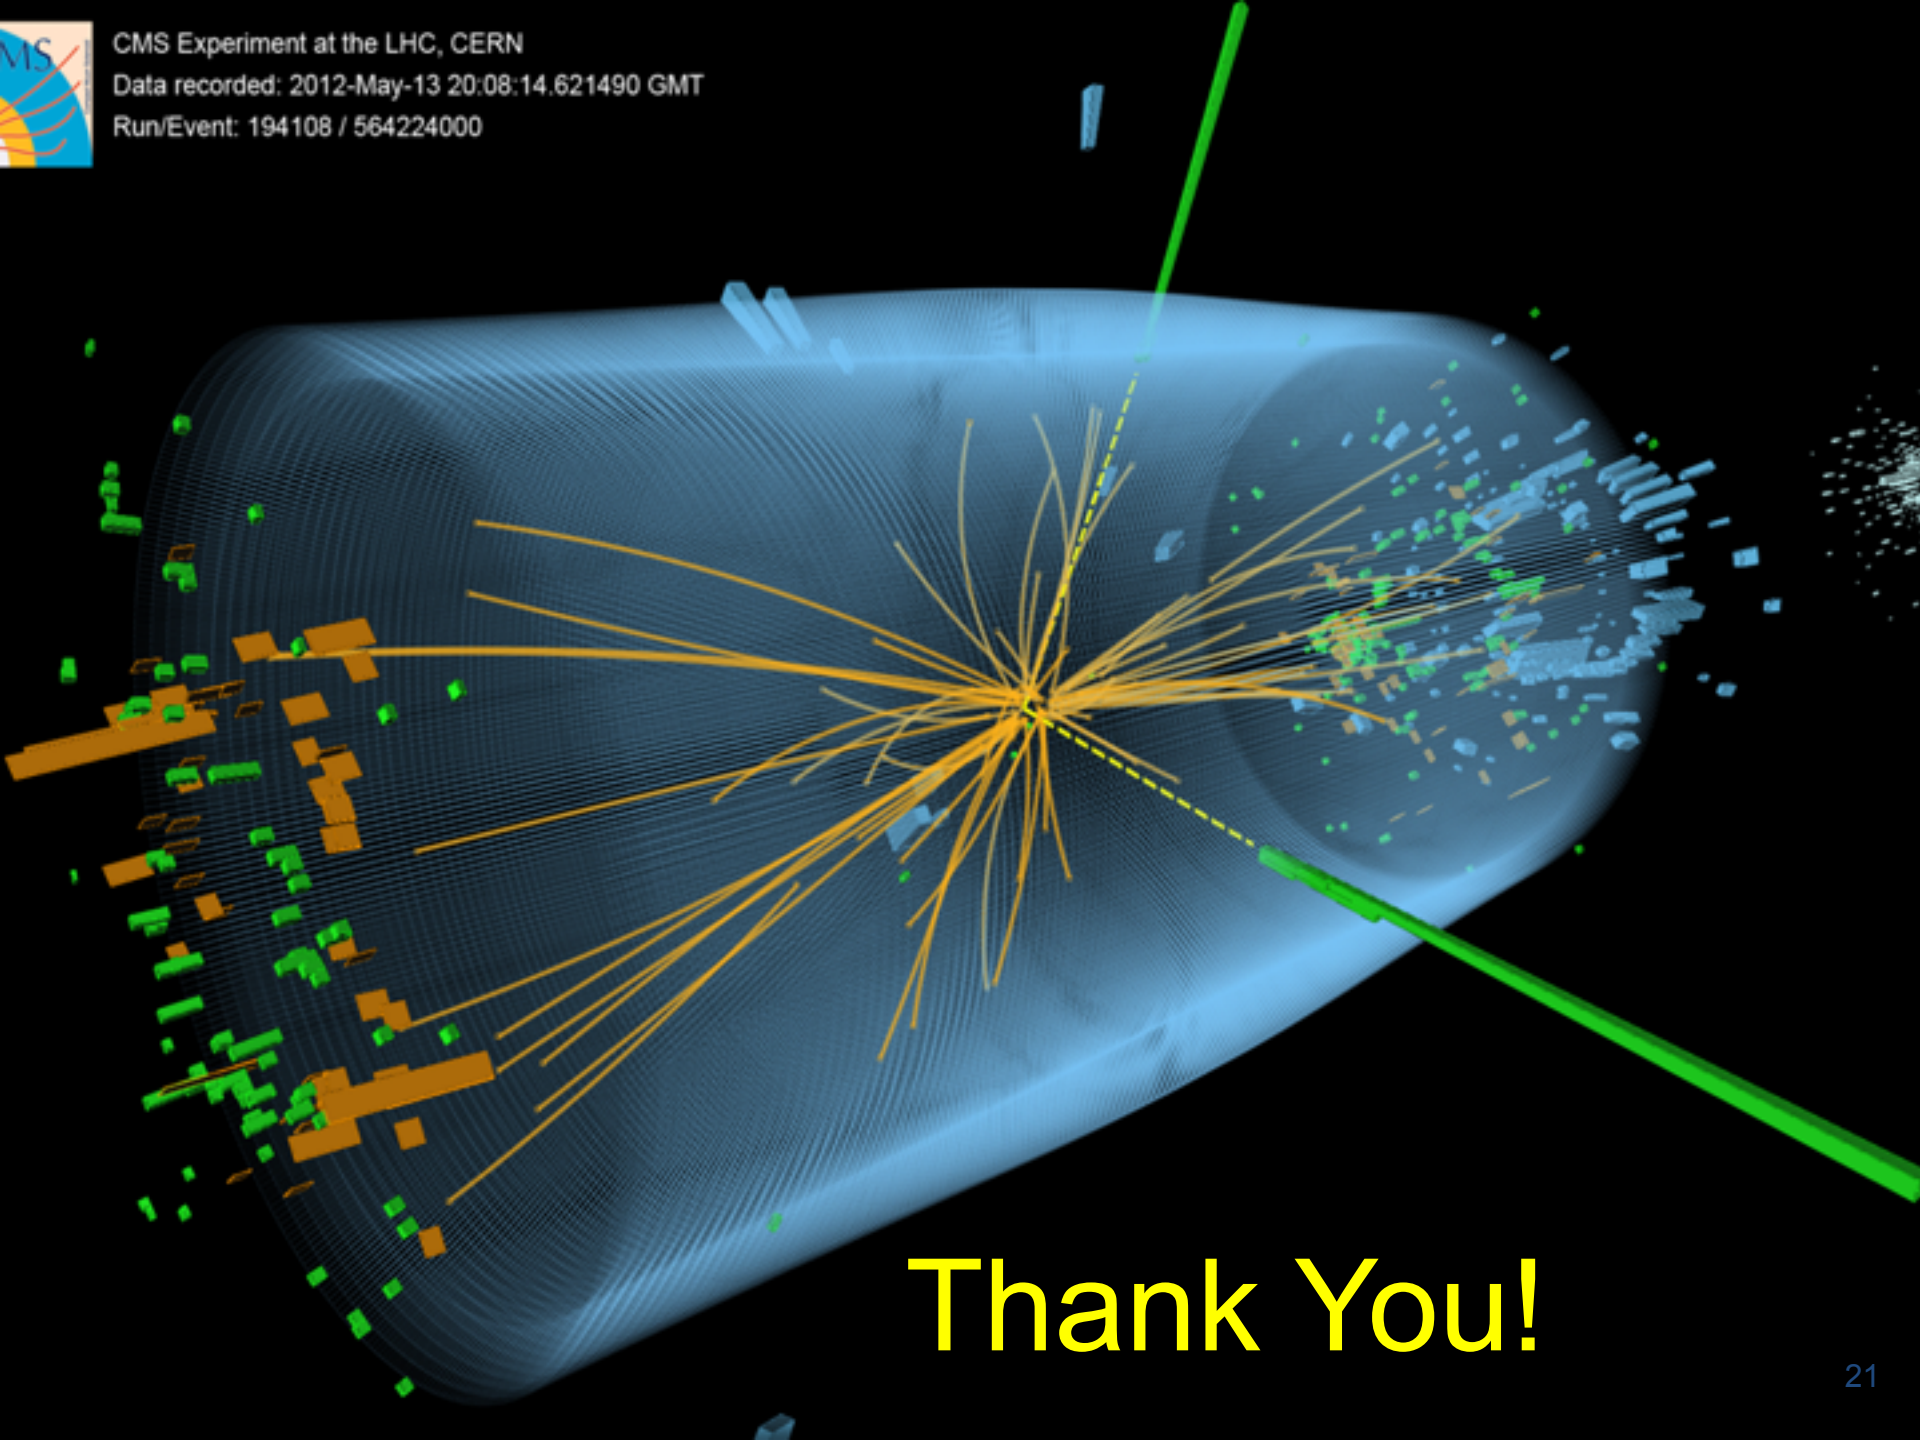
CMS Experiment at the LHC, CERN
Data recorded: 2012-May-13 20:08:14.621490 GMT
Run/Event: 194108 / 564224000

Thank You!


CERN