The Worldwide LHC Computing Grid

Nils Høimyr
IT Department

The ATLAS experiment

CERN IT Department

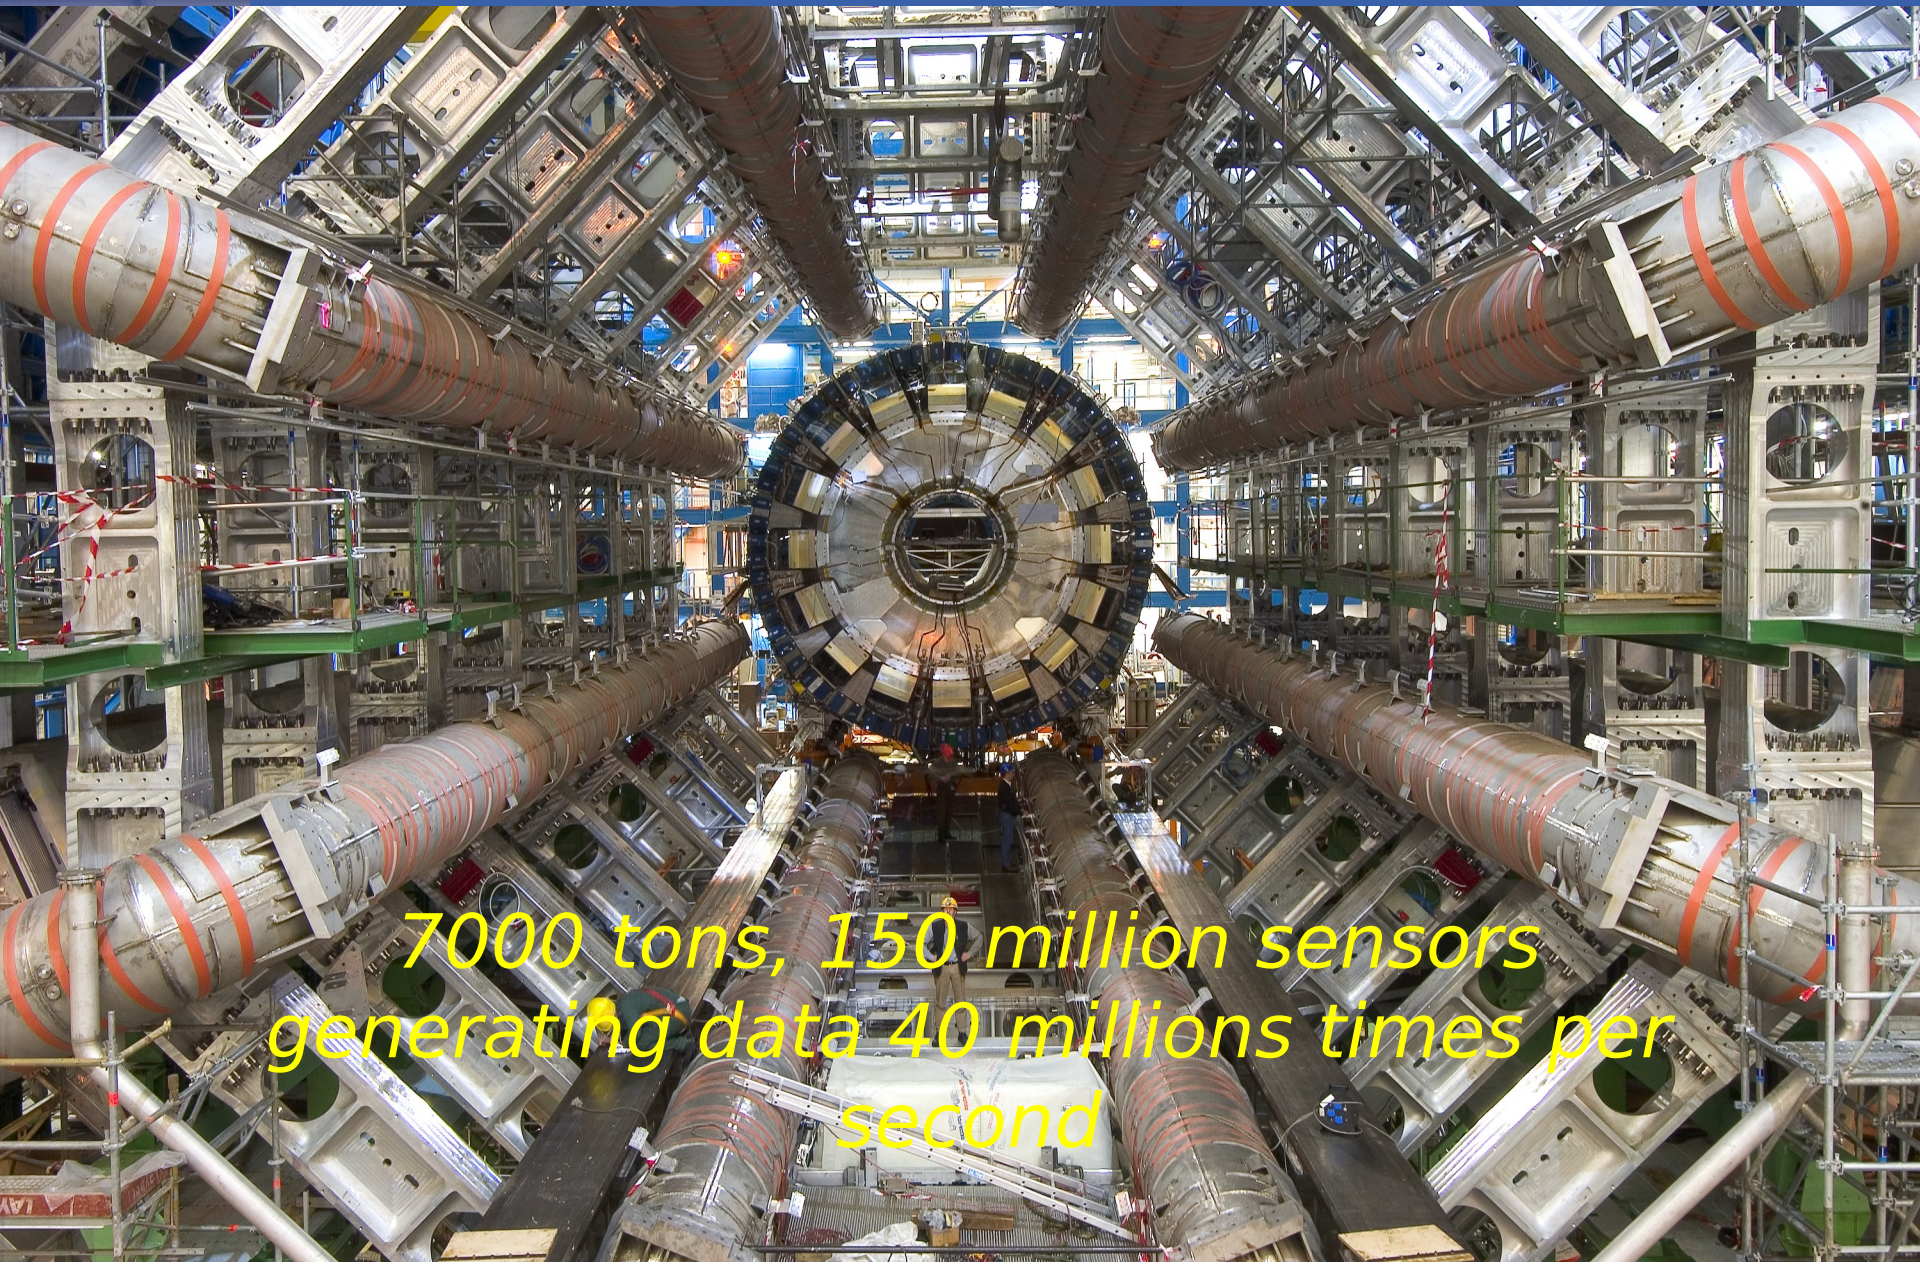

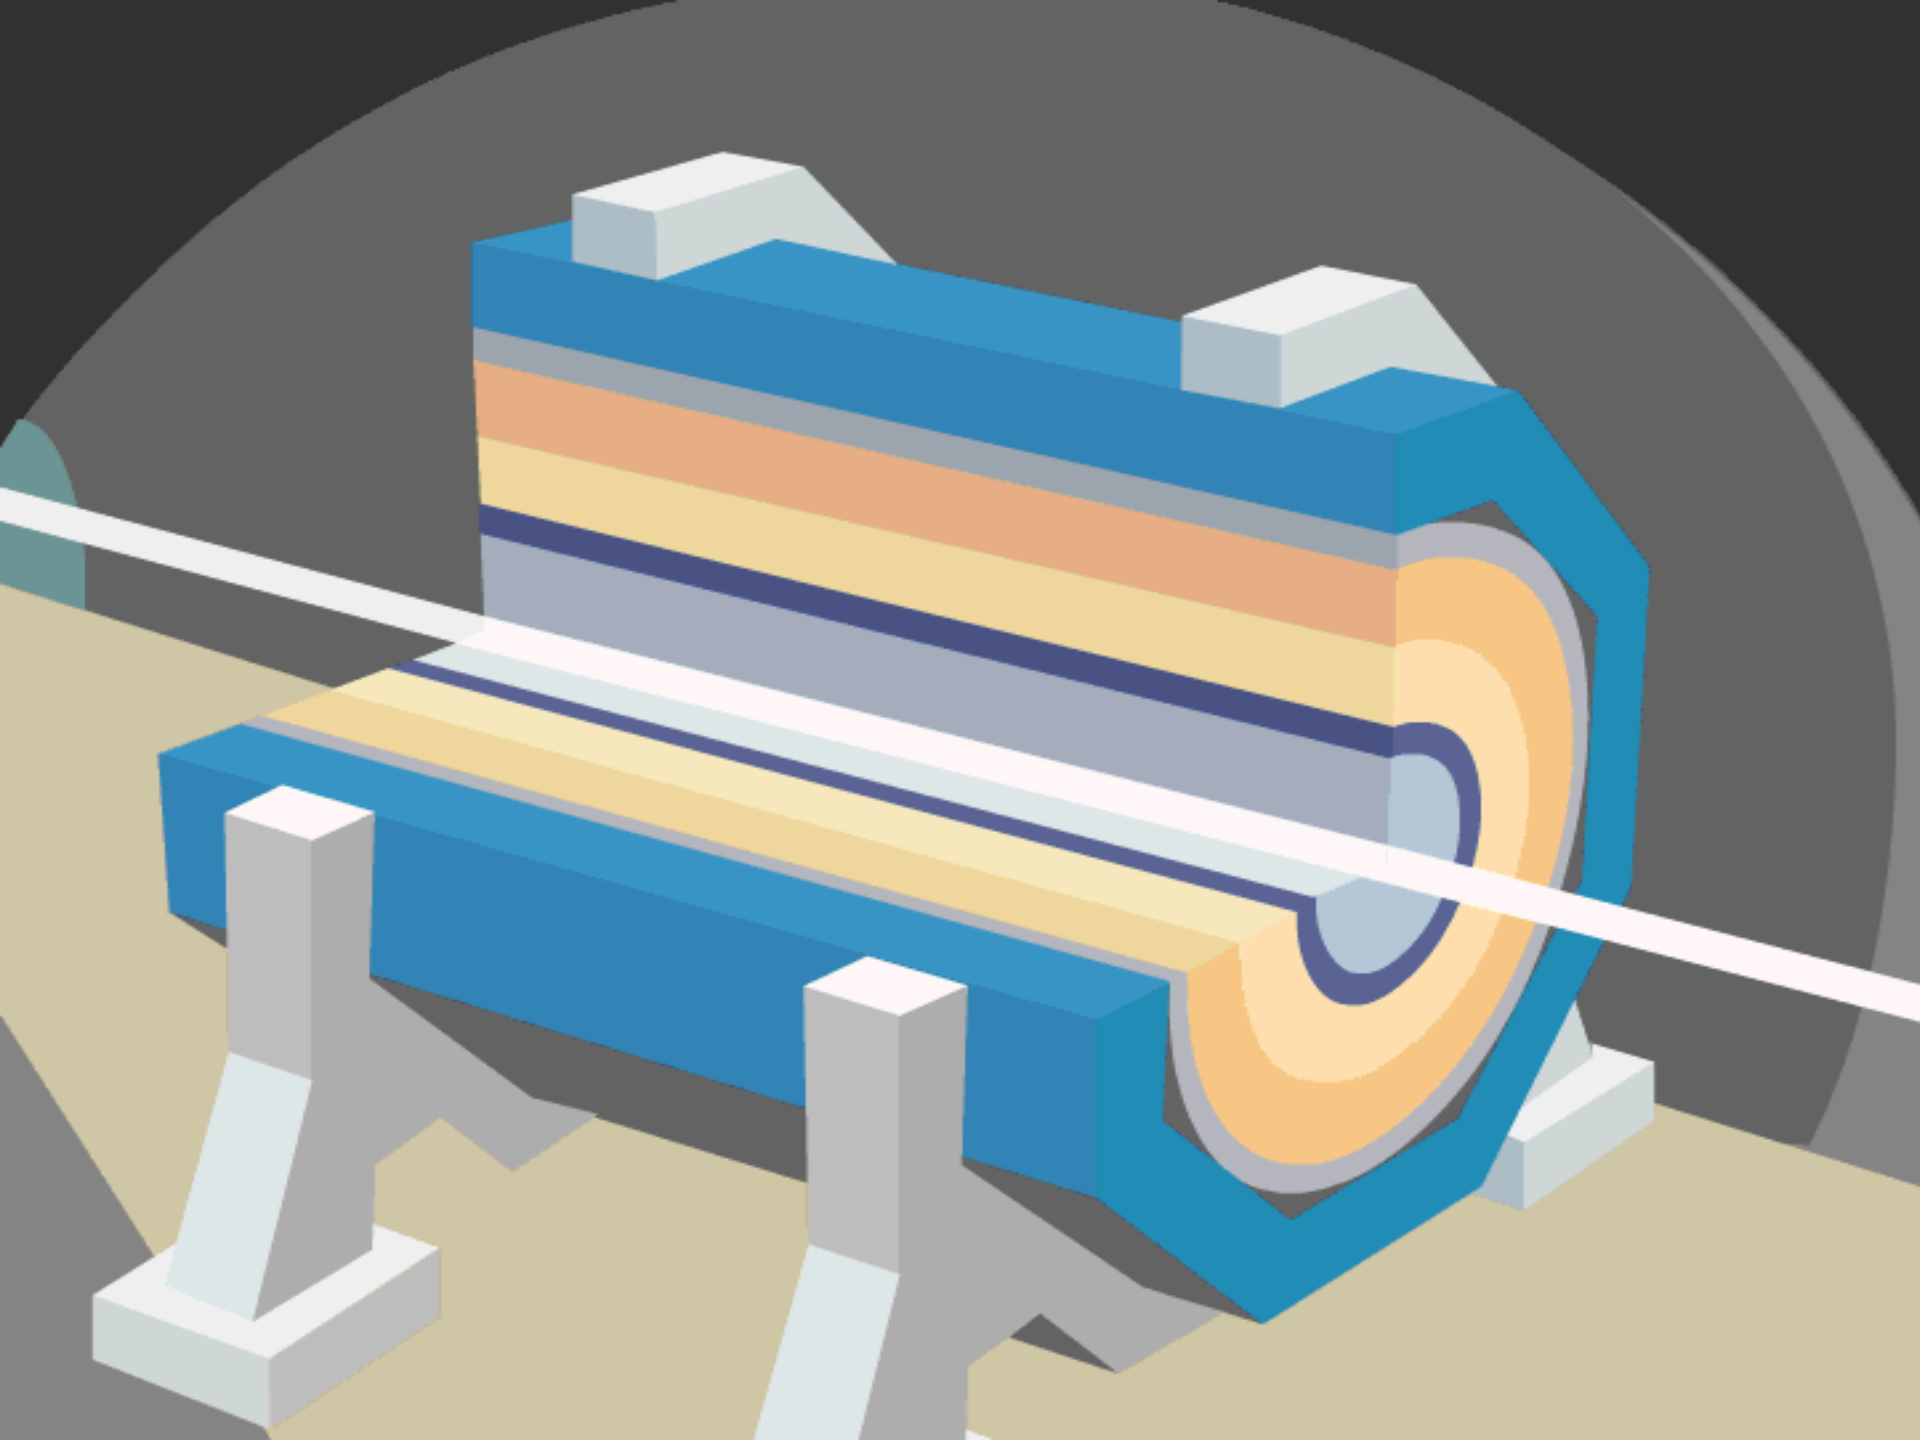

~ 300.000 MB/s
from all sub-detectors

~ 300MB/s
Raw Data

Trigger and data acquisition

Event filter computer farm

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

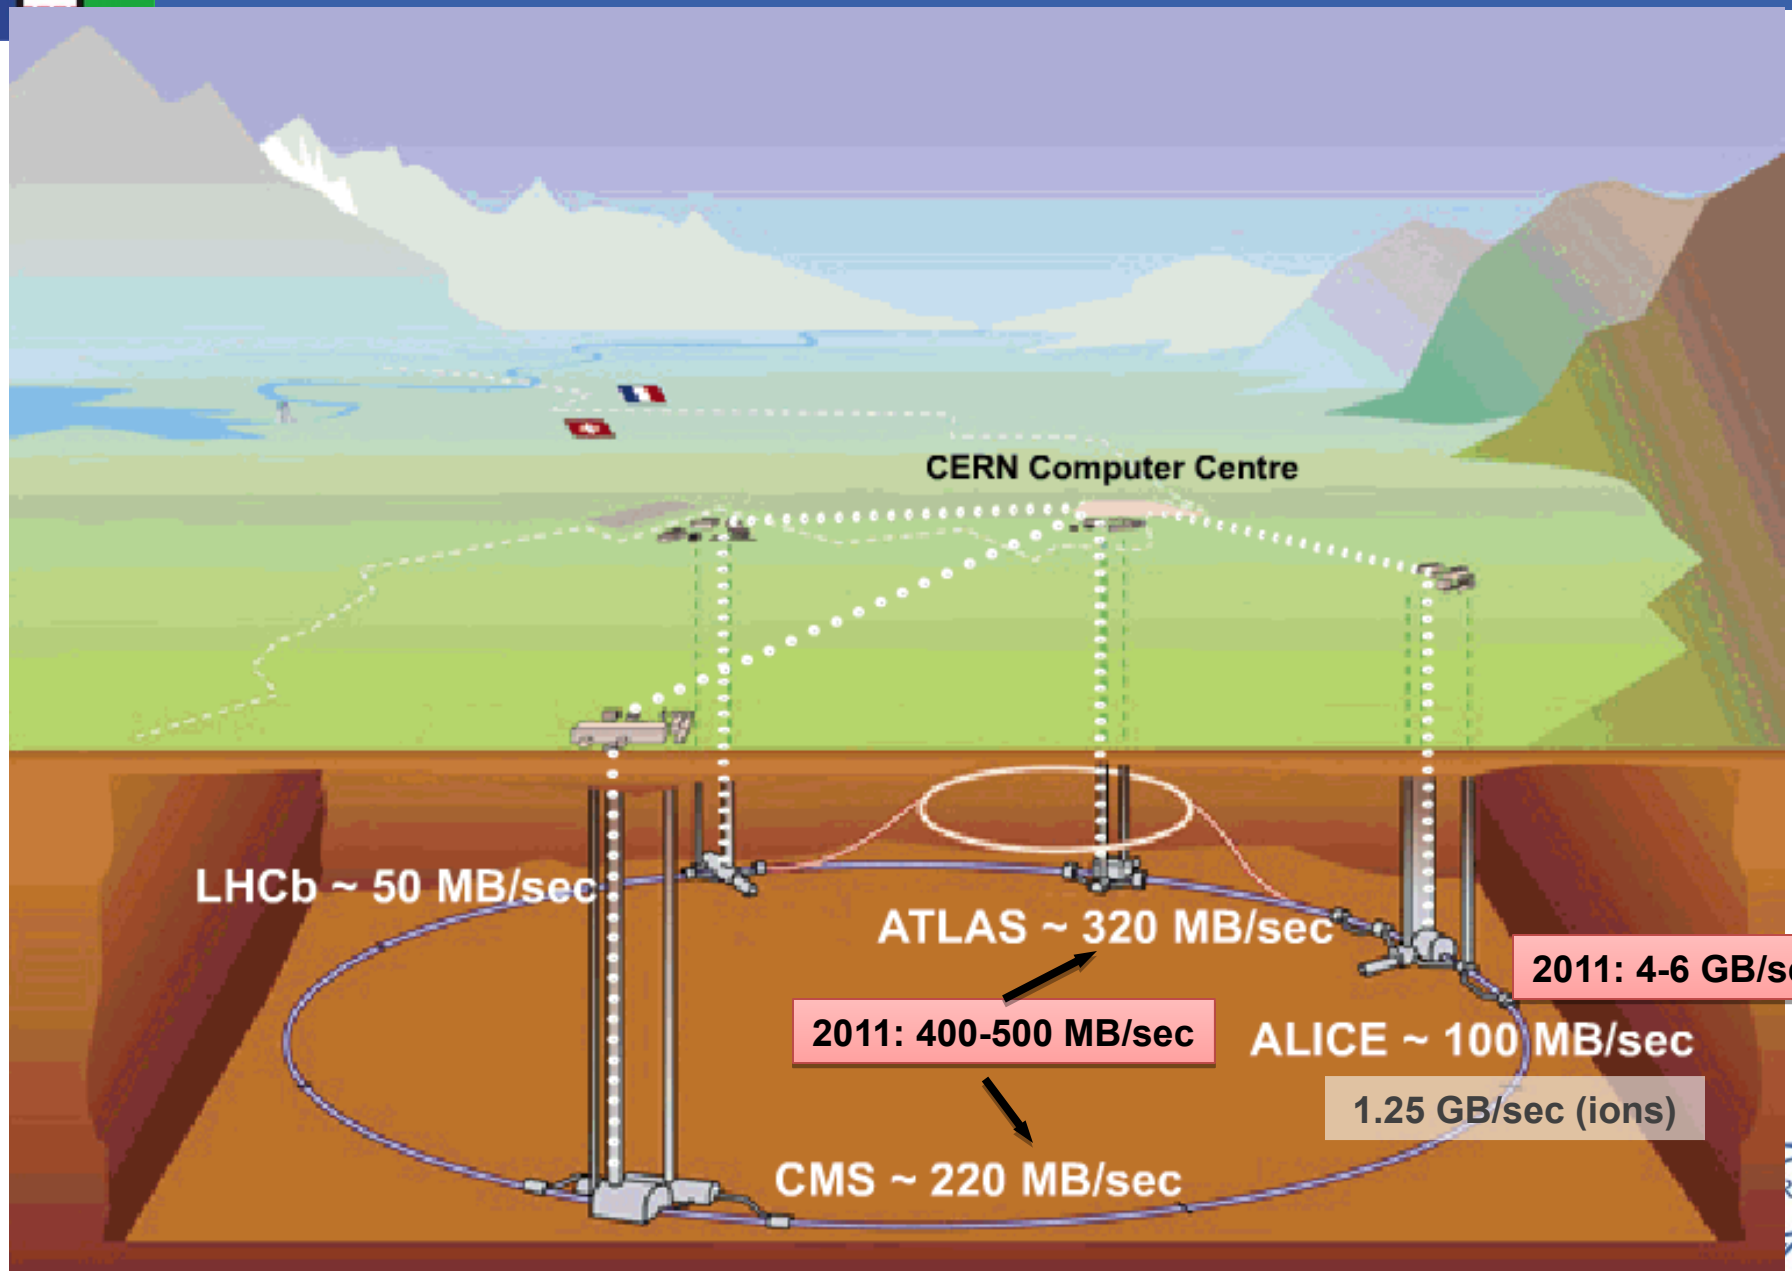

The CERN Data Centre in Numbers

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
- CERN Computer Centre (B513)
  - ~10 000 servers
  - ~90 000 processor cores
  - About 100 Pb of storage
  - 3500 kW power
- Details on hardware:
- http://hwcollect.cern.ch/

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

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 ~15% of the capacity

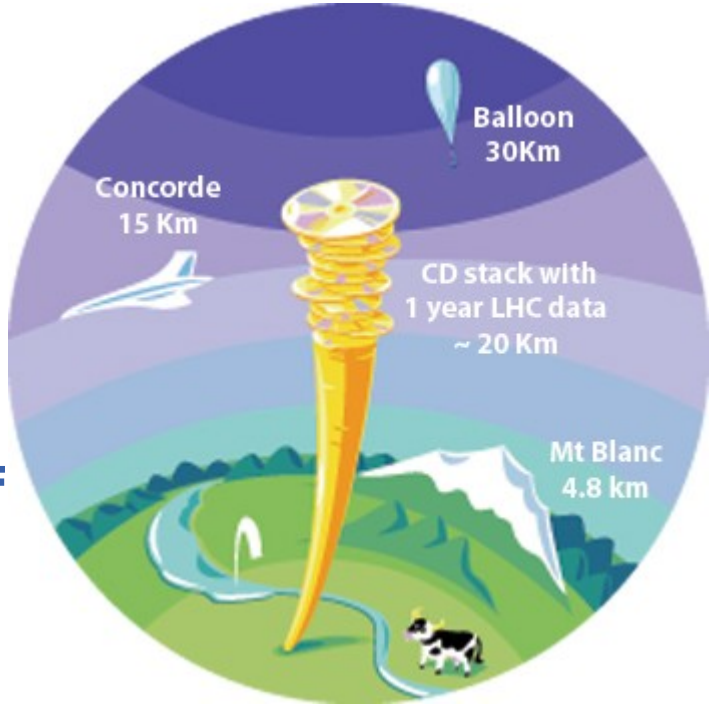

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Solution: the Grid

CERN IT Department

- Use the Grid to unite computing resources of particle physics institutes around the world

The **World Wide Web** provides seamless access to information that is stored in many millions of different geographical locations

The **Grid** is an infrastructure that provides seamless access to computing power and data storage capacity distributed over the globe

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

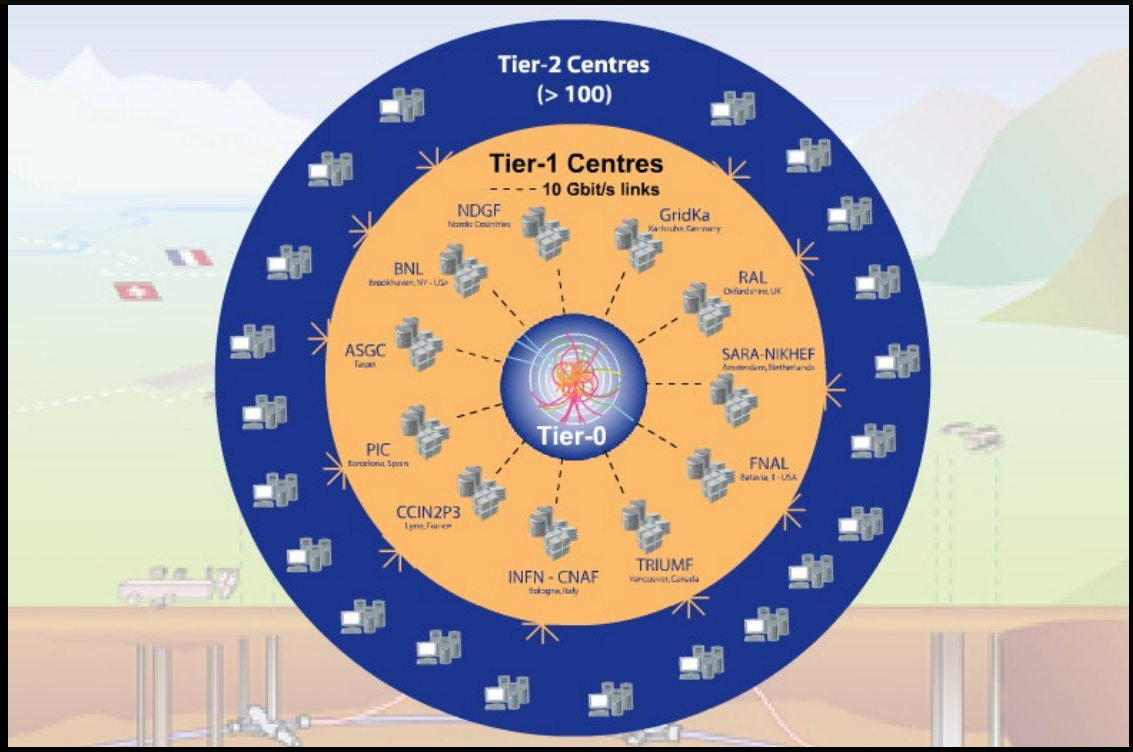

Tier-0 (CERN):
- Data recording
- Initial data reconstruction
- Data distribution

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

Tier-2 (~200 centres):
- Simulation
- End-user analysis

Worldwide LHC Computing Grid

LHC Networking

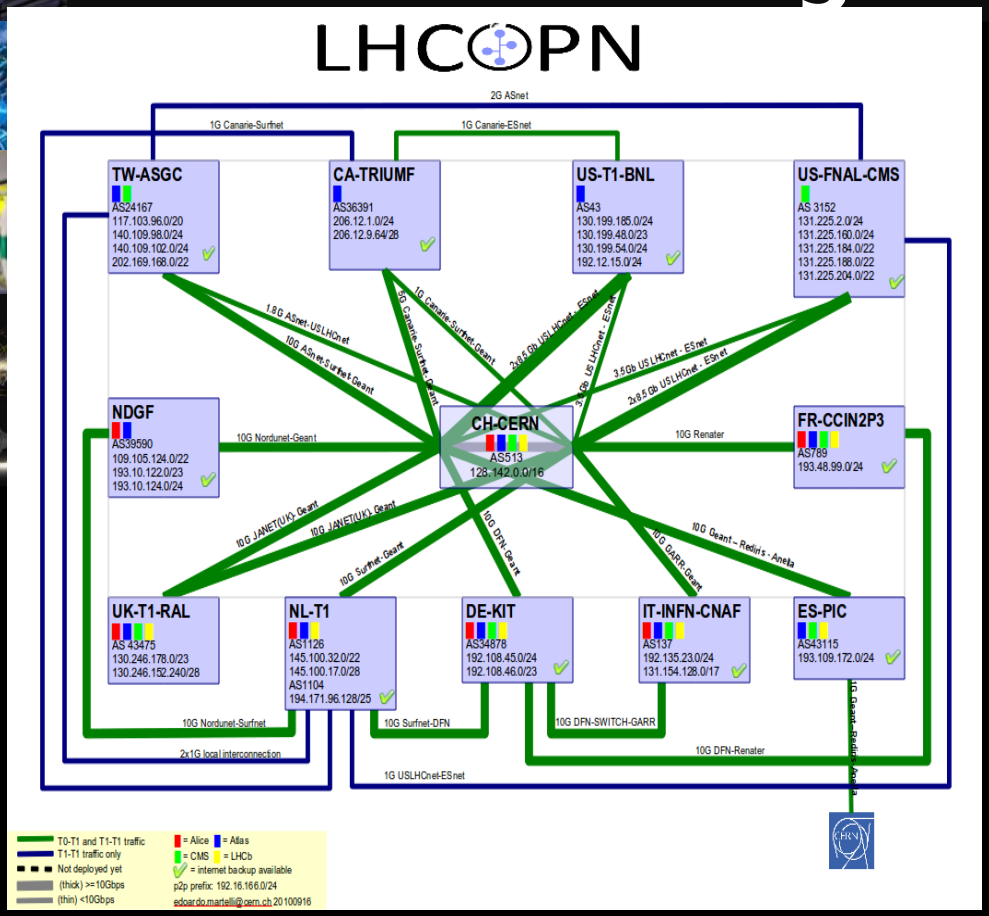

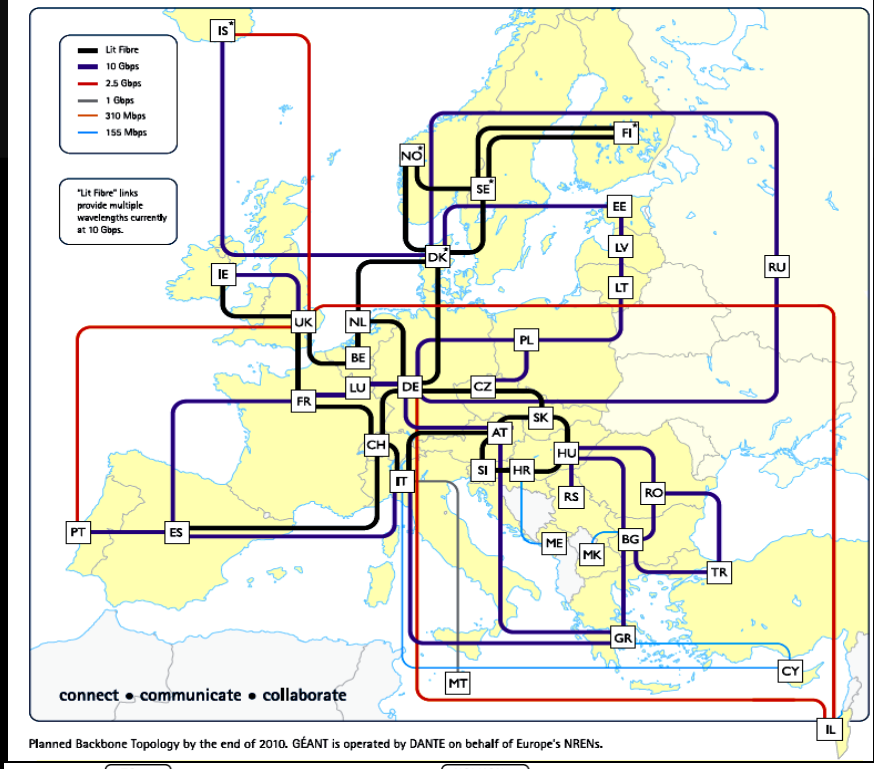

Planned Backbone Topology by the end of 2010. GÉANT is operated by DANTE on behalf of Europe's NRENs.

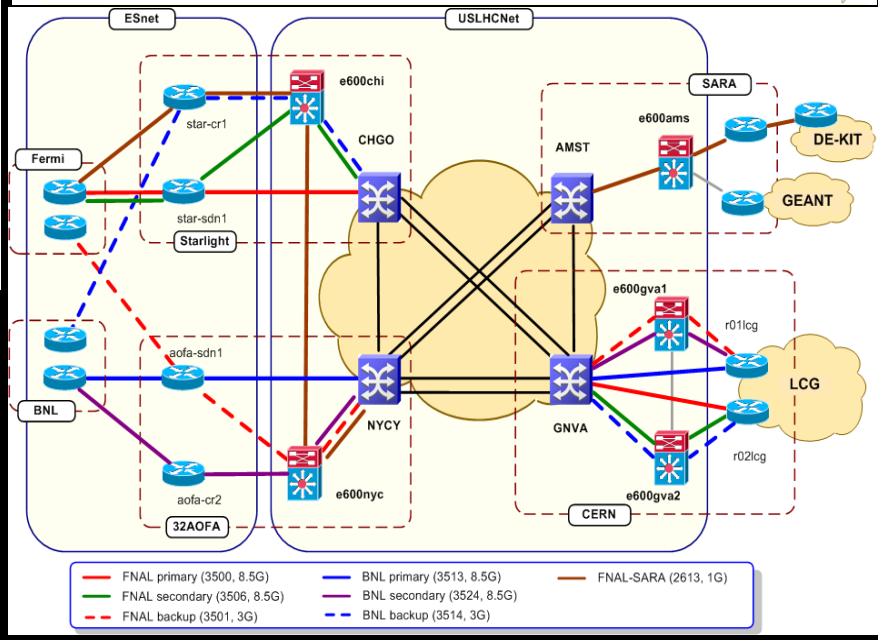

- Relies on
  - OPN, GEANT, US-LHCNet
  - NRENs & other national & international providers

Worldwide LHC Computing Grid

Mass storage access

CERN IT Department

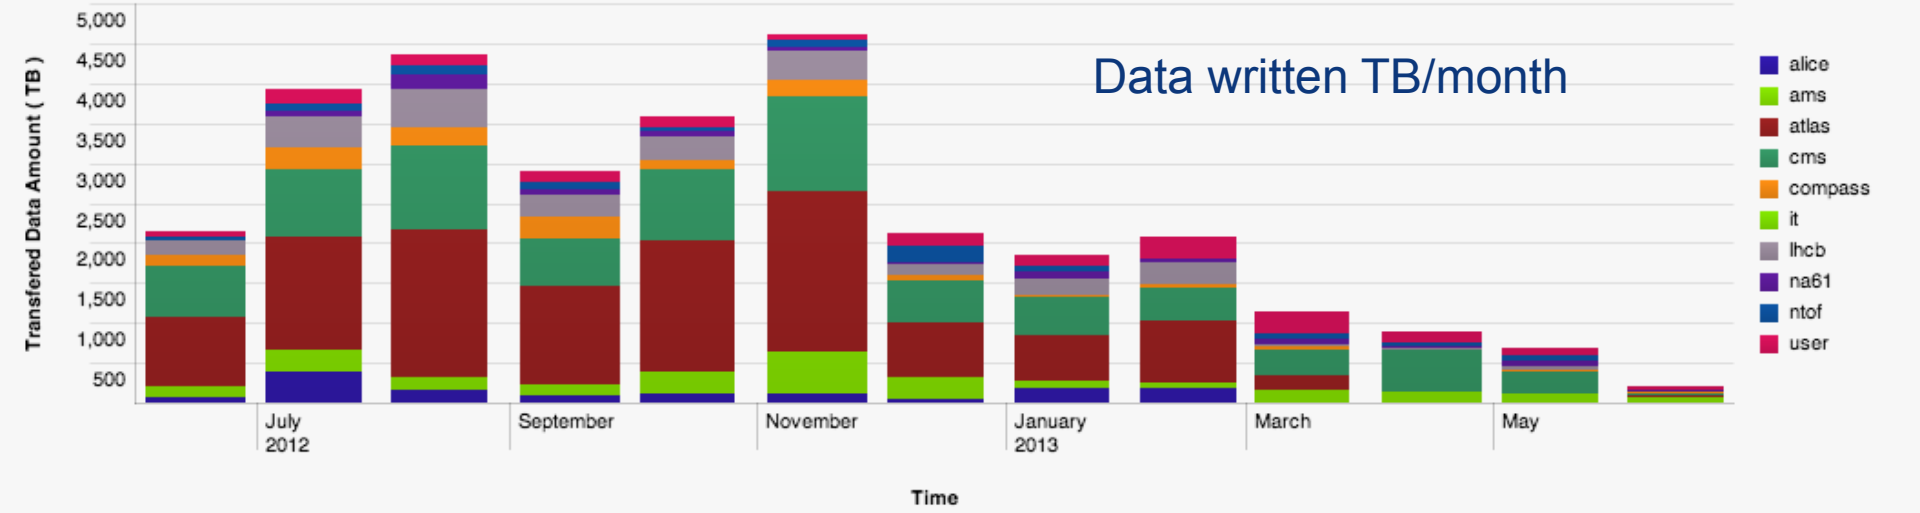

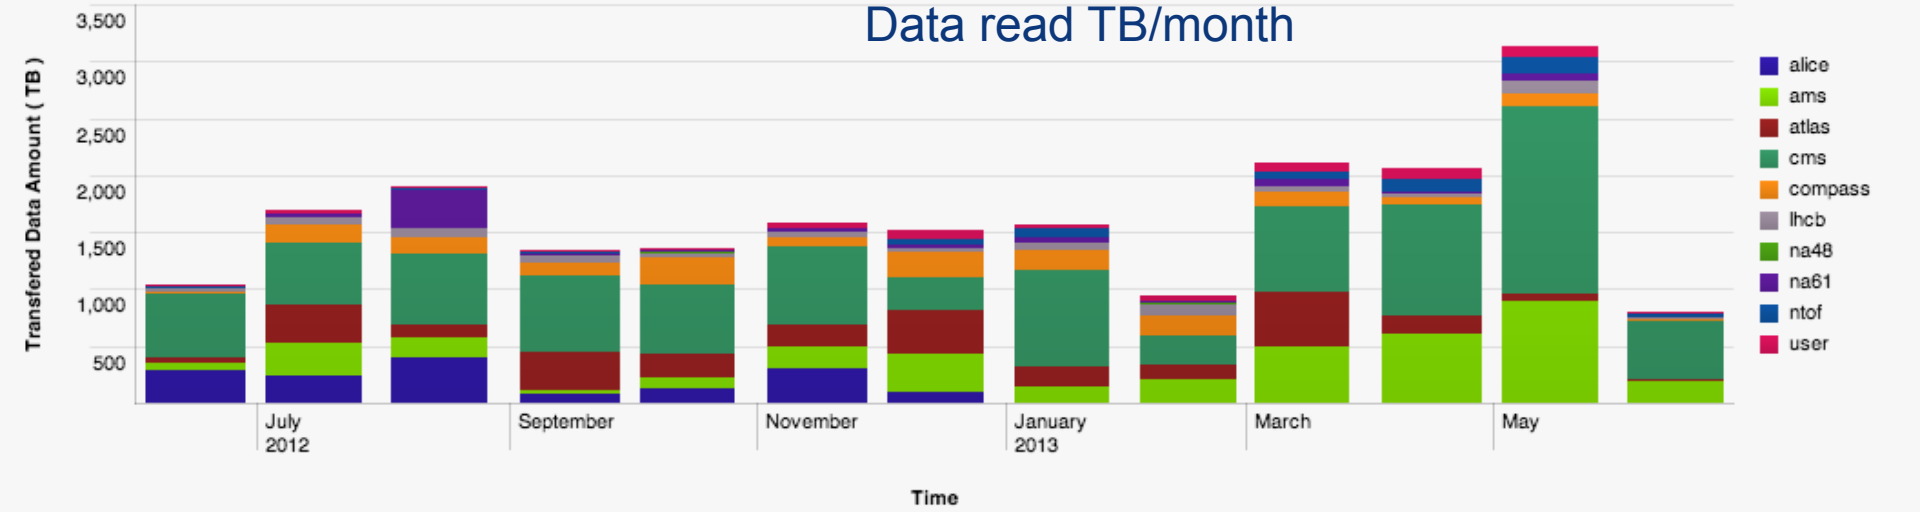

Data transfers

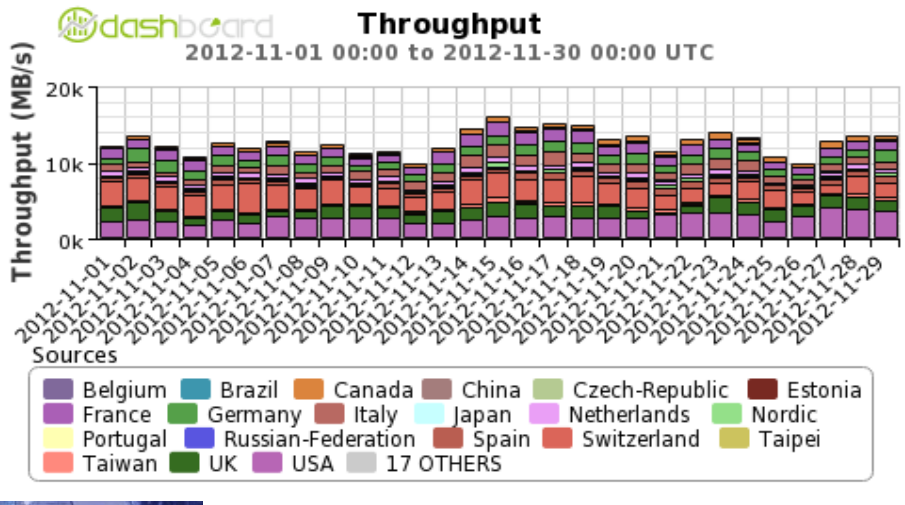

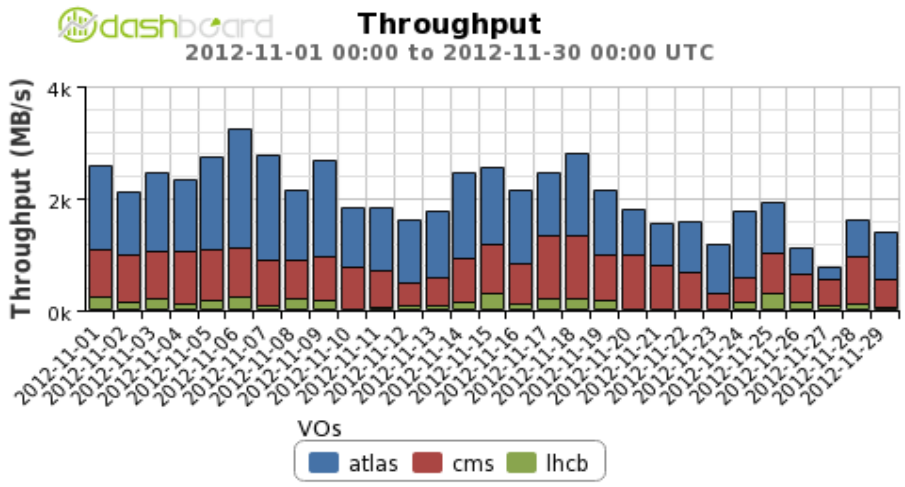

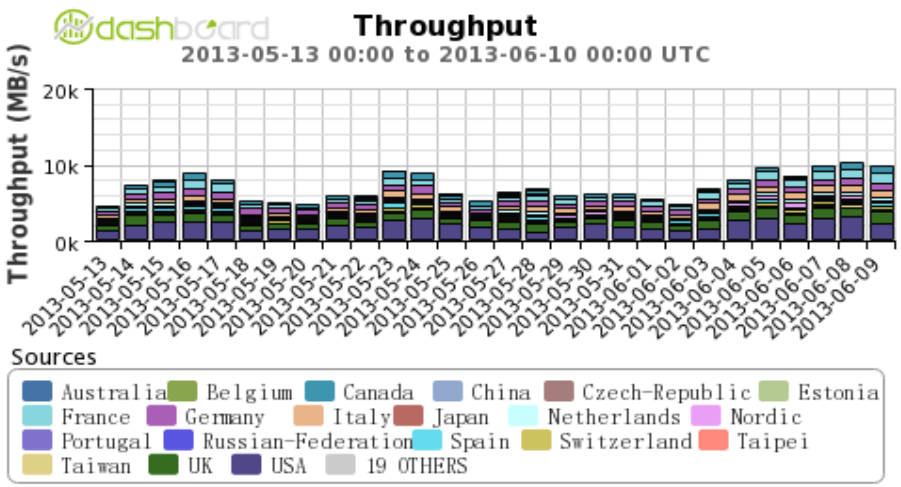

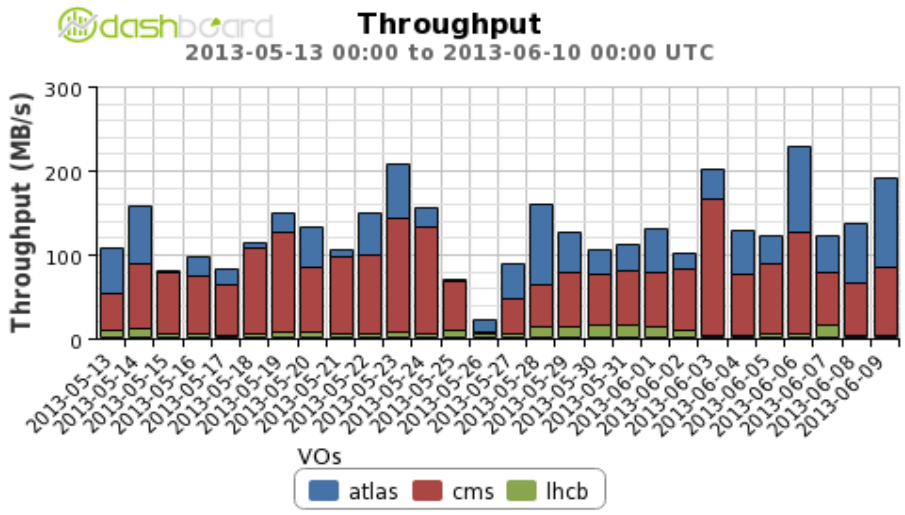

CERN: Data written 2010-12

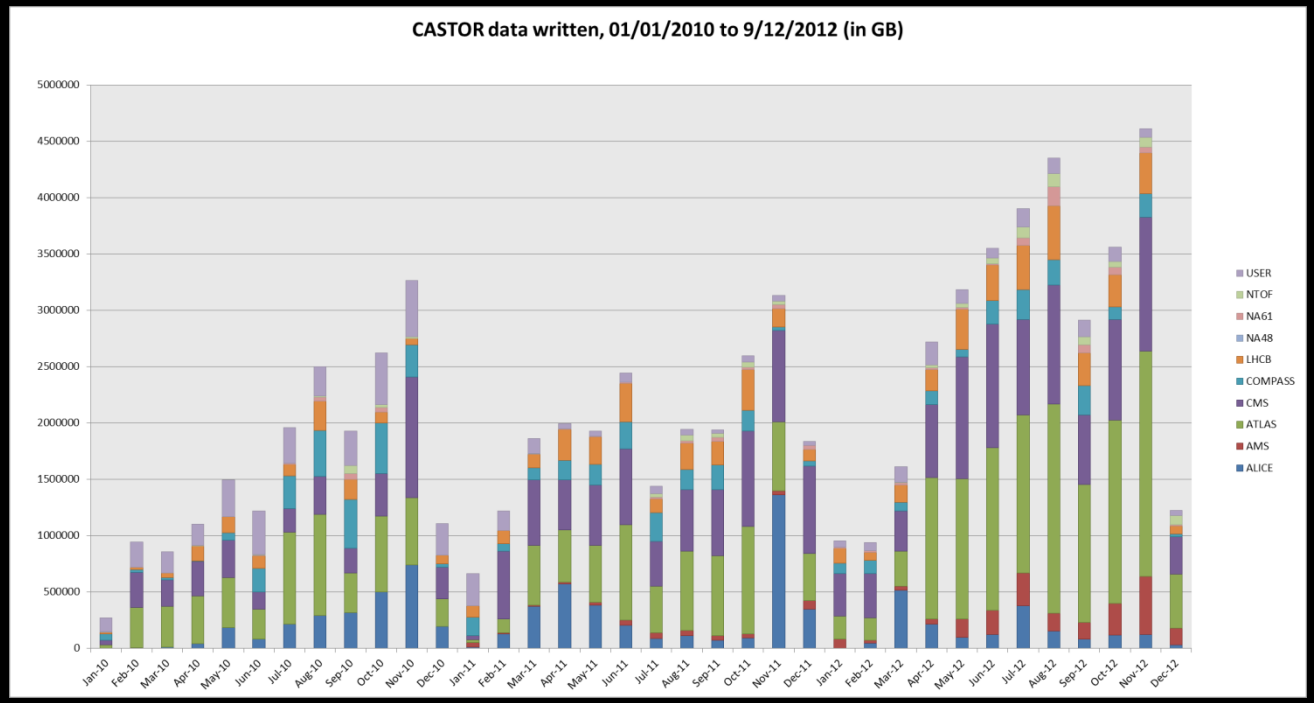

2010-2012 Data written:
Total ~30 PB in 2012
Close to 3-4 PB/month

Data rates increased
3-4 GB/s input
~15 GB/s output

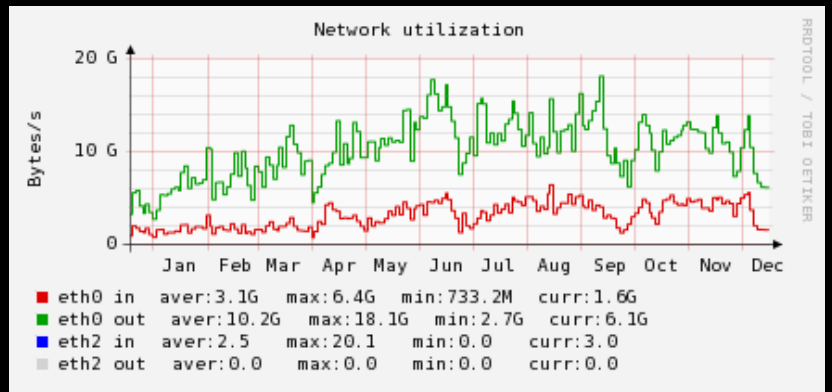

VLCG
Worldwide LHC Computing Grid

ATLAS Computing scale in LHC Run 1

- 150k slots continuously utilized
- ~1.4M jobs/day completed

- More than 5 GB/s transfer rate worldwide

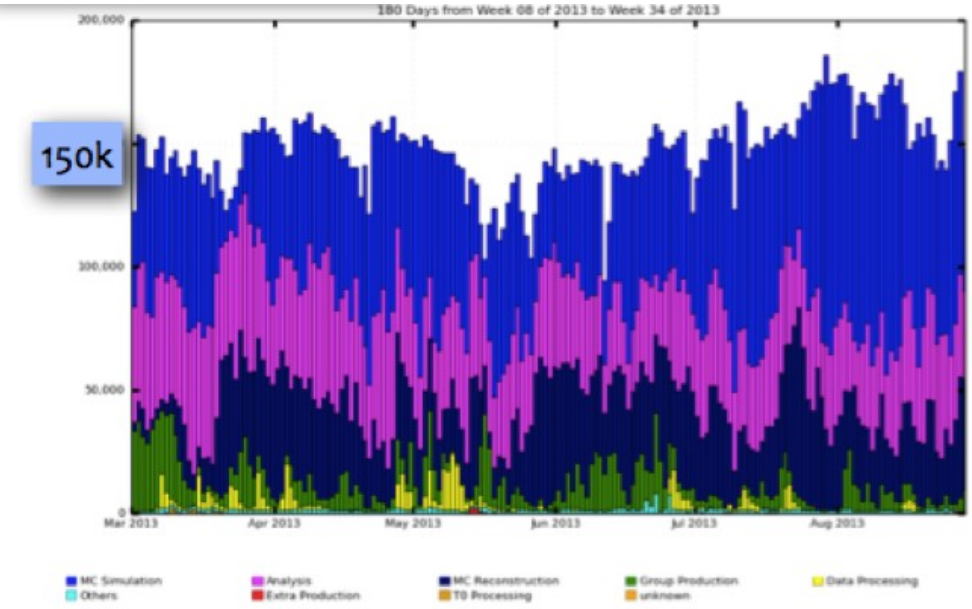

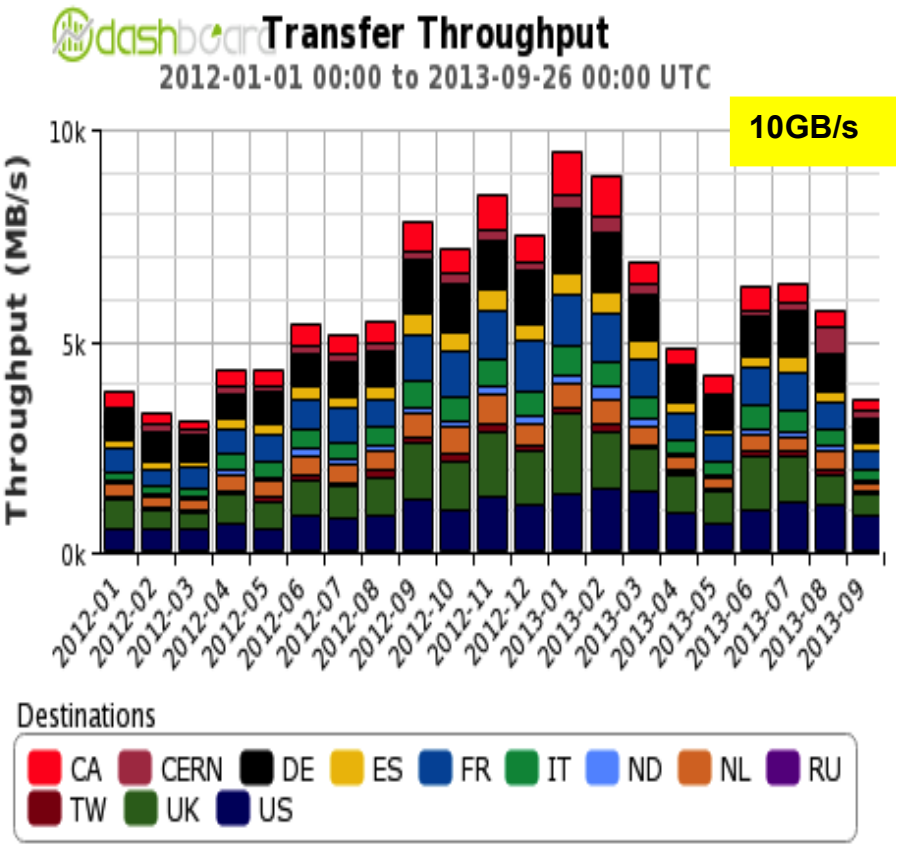

Broader Impact of the LHC Computing Grid

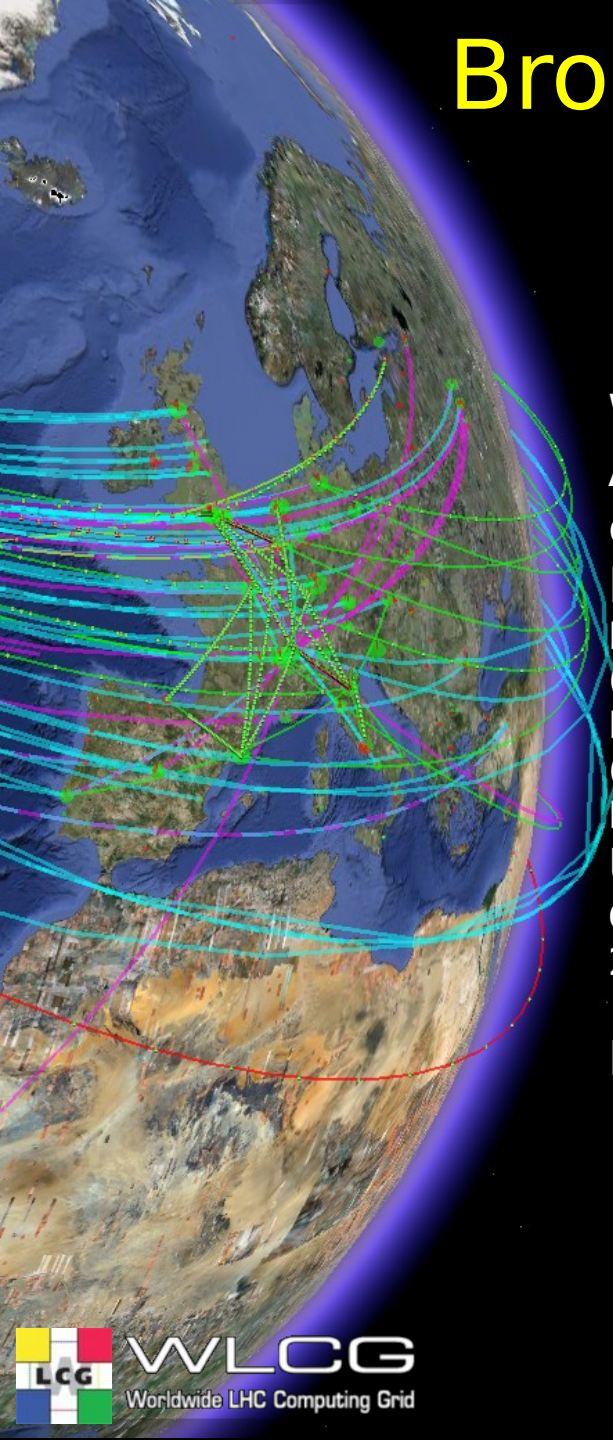

WLCG has been leveraged on both sides of the Atlantic, to benefit the wider scientific community

Europe:

Enabling Grids for E-sciencE (EGEE) 2004-2010

European Grid Infrastructure (EGI) 2010--

Nordugrid in the Nordic countries

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
- …

WLCG Worldwide LHC Computing Grid

CERN

Scaling CERN Data Center(s) to anticipated Physics needs

CERN Data Center dates back to the 70's

- Upgraded in 2005 to support LHC (2.9 MW)
- Still optimizing the current facility (cooling automation, temperatures, infrastructure)

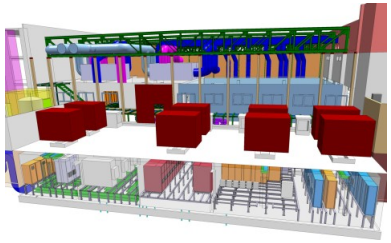

Renovation of the "barn" for accommodating 450 KW of "critical" IT loads (increasing 513 total to 3.5 MW)

Exploitation of 100 KW of remote facility down town

- Understanding costs, remote dynamic management, ensure business continuity

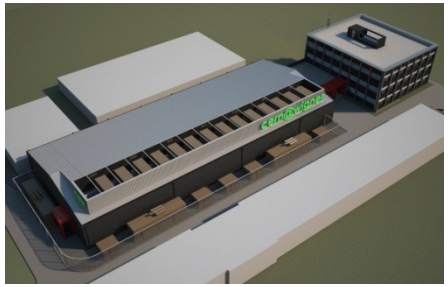

Exploitation of a remote Data center in Hungary

- Max. 2.7 MW (N+1 redundancy)
  - Business continuity
- 100 Gbps connections

Data Management

- Where is LHC in Big Data Terms?

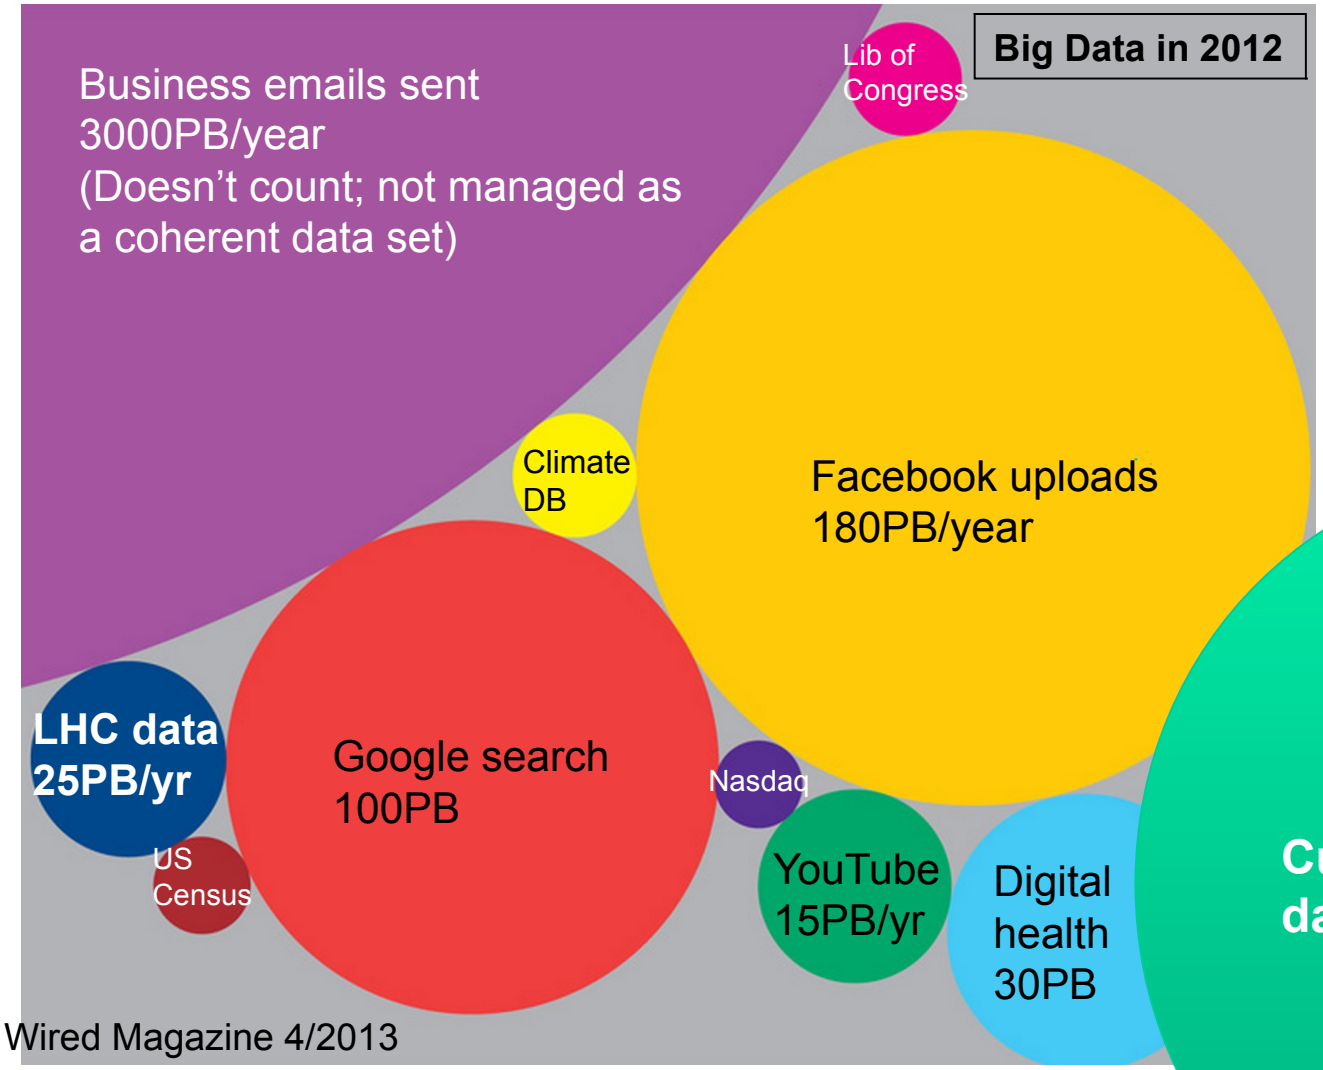

Reputed capacity of NSA's new Utah data center: 5000 PB (50-100 MW, $2 billion)

We are big…

Current LHC data set, all data products: ~300 PB

Computing challenges for Run2

- Computing in LHC Run1 was very successful but Run 2 from 2015 poses new challenges
- Increased energy and luminosity delivered by LHC in Run 2
  - More complex events to process
    - Event reconstruction time (CMS ~2x)
  - Higher output rate to record
    - Maintain similar trigger thresholds and sensitivity to Higgs physics and to potential new physics
    - ATLAS, CMS event rate to storage 2.5x
- Need a substantial increase of computing resources that we probably cannot afford

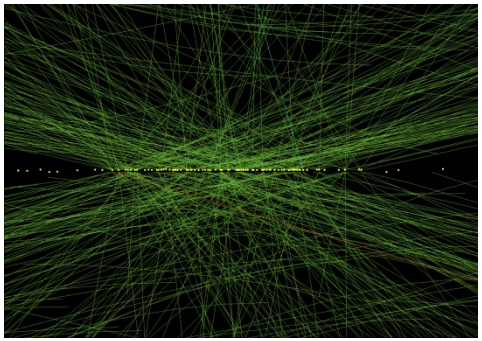

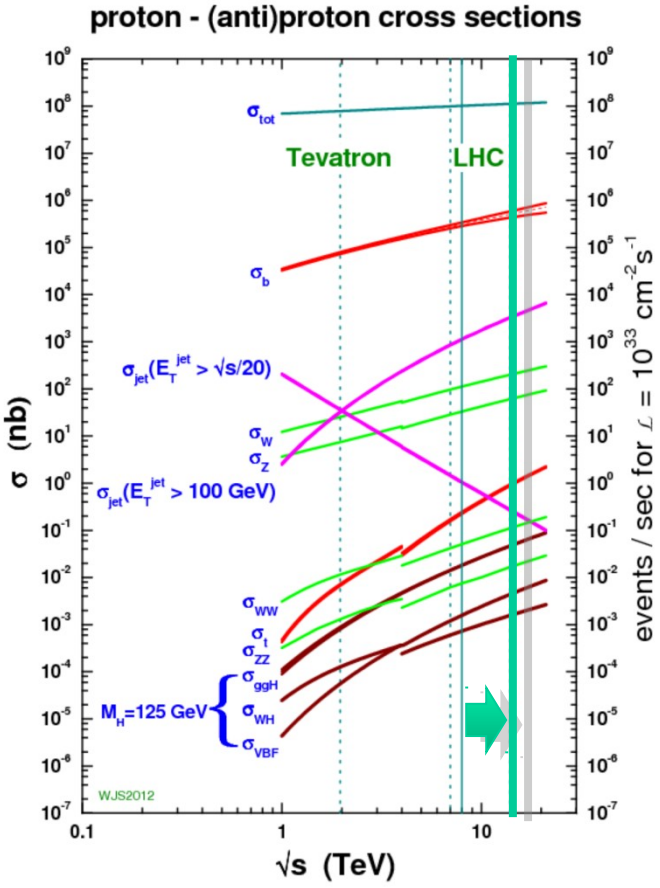

Access to new resources for Run 2

- Access to opportunistic resources
  - HPC clusters, academic or commercial clouds, volunteer computing
- Use HLT farm for offline data processing
  - During extended periods with no data taking and even inter-fill periods
- Adopt advanced architectures
  - Processing in Run1 done under Enterprise Linux on x86 CPUs
  - Many-core processors, low-energy CPUs, GPU environments
  - Challenging heterogeneous environment
  - Parallelization of processing application will be key

Computing Model Evolution

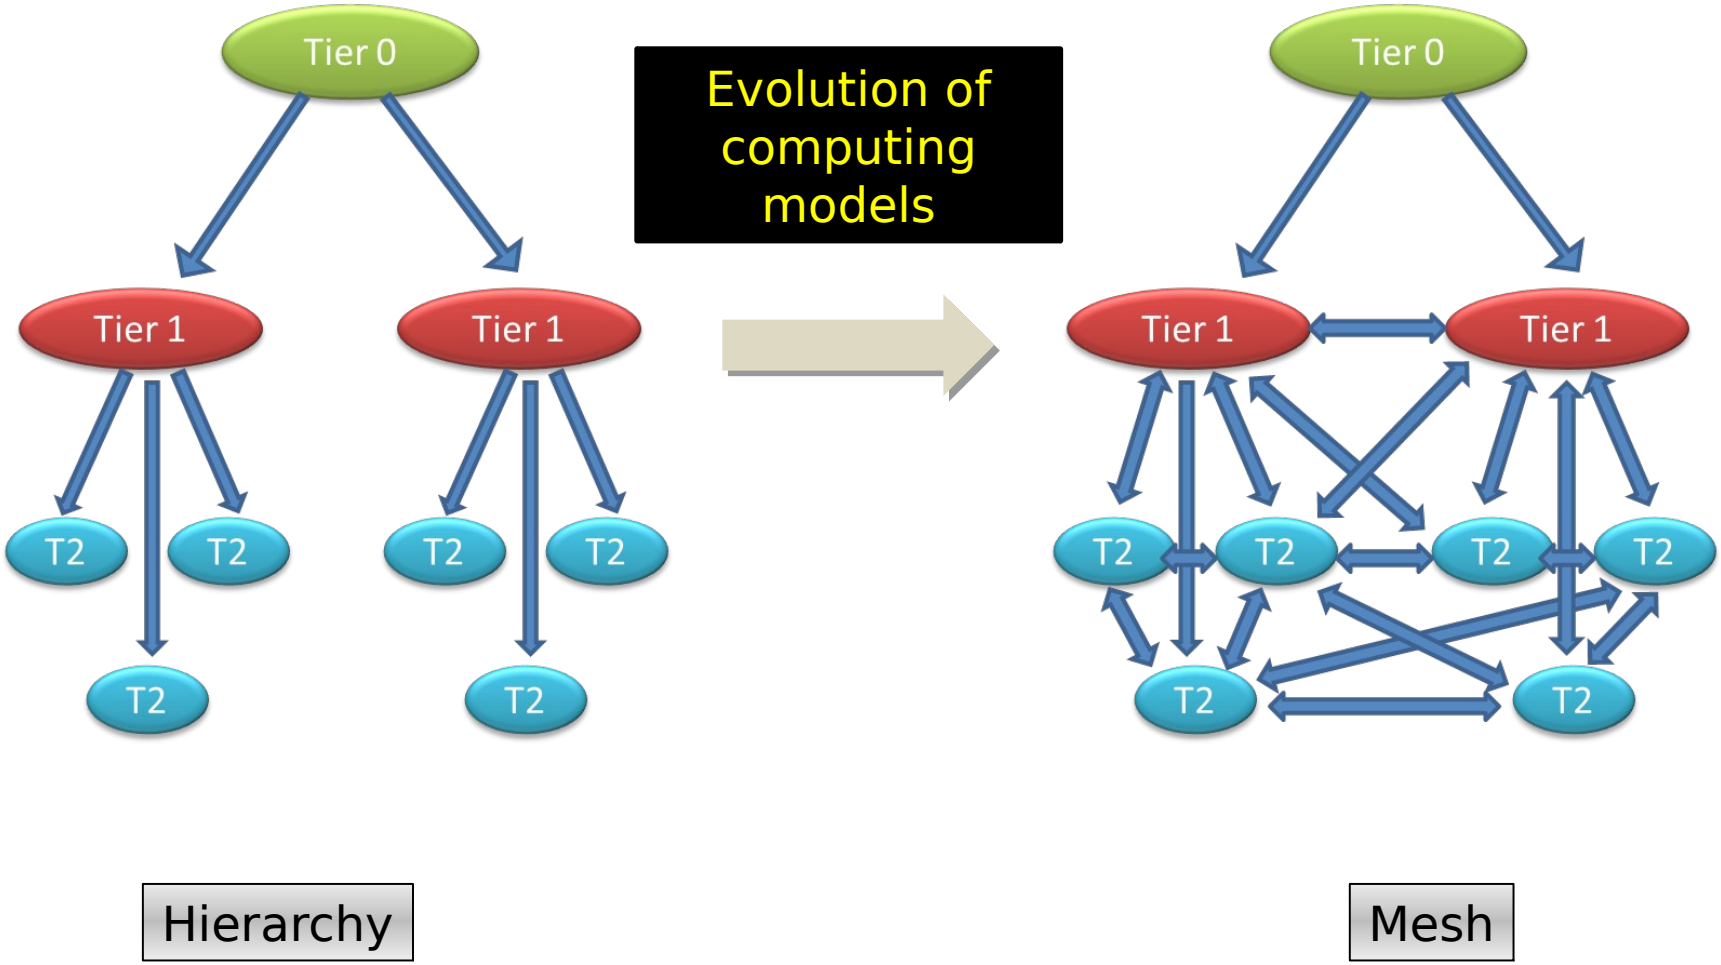

Grid vs Cloud

CERN IT Department

- "Cloud computing" is becoming the standard
  - Web based solutions (http/https and RES)
  - Virtualisation, upload virtual machine images to remote sites
- 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 – lxcloud, now OpenStack
  - 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

From grid to clouds

CERN IT Department

- Data centres providing infrastructure as a service
- Clouds complement and extend the Grid
  - Decrease heterogeneity seen by the user (hardware virtualisation)
  - VMs provide a uniform user interface to resources
    - Isolate software and operating system from physical hardware
    - CernVM and CernVM FS adopted by LHC experiments
- New resources (commercial, research clouds)
- Grid of clouds already tried by the LHC experiments
    - ATLAS ~450k production jobs from Google over a few weeks
    - Tests on Amazon EC2 ~economically viable

Cloud Infrastructure

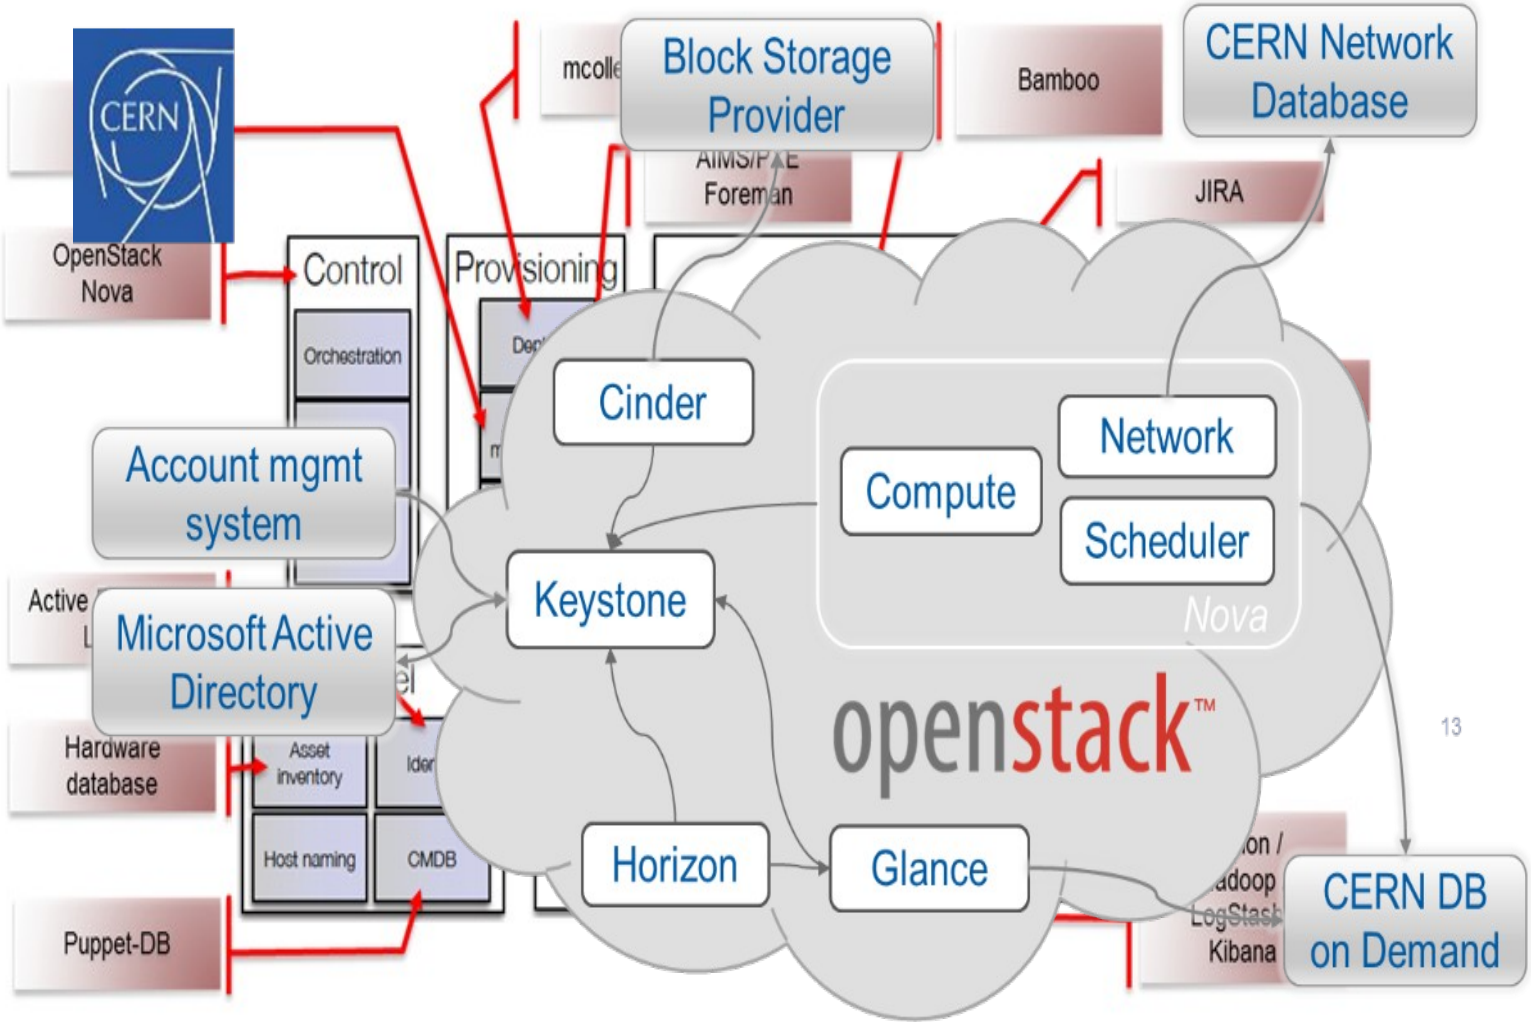

Volunteer grid - LHC@home

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**

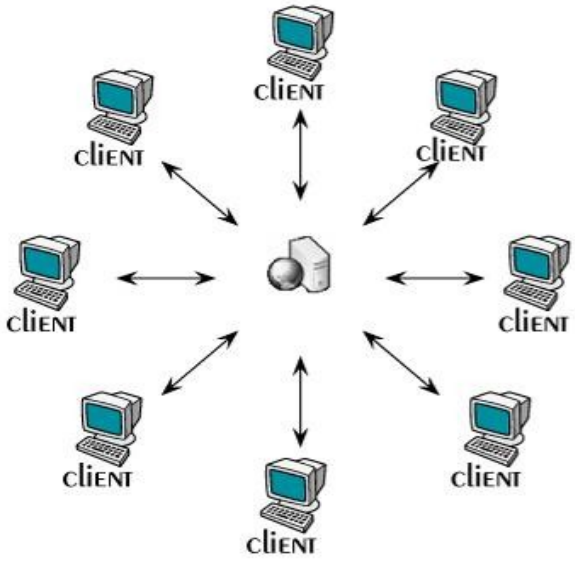

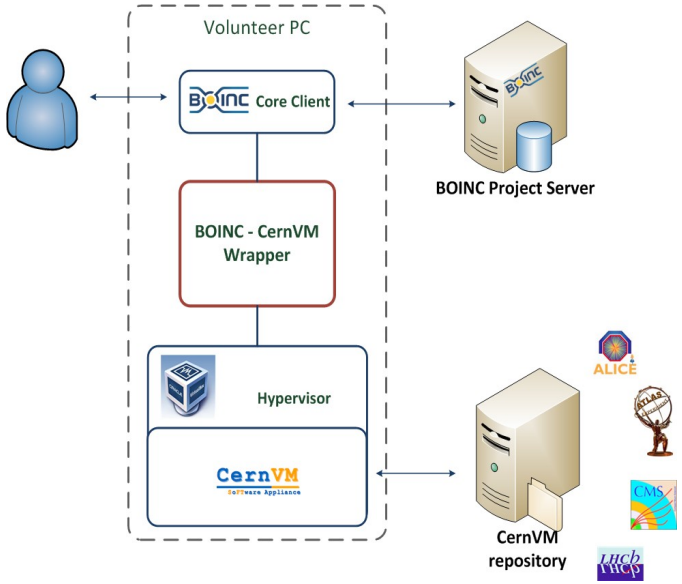

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

International Scientific Collaborations

- Many scientific projects are global collaborations of 100s of partners
- Efficient computing and data infrastructures have become critical as the quantity, variety and rates of data generation keep increasing
- Funding does not scale in the same way
  - Optimization and sharing of resources
- Collaboration with commercial IT companies increasingly important
  - Requirements are not unique anymore

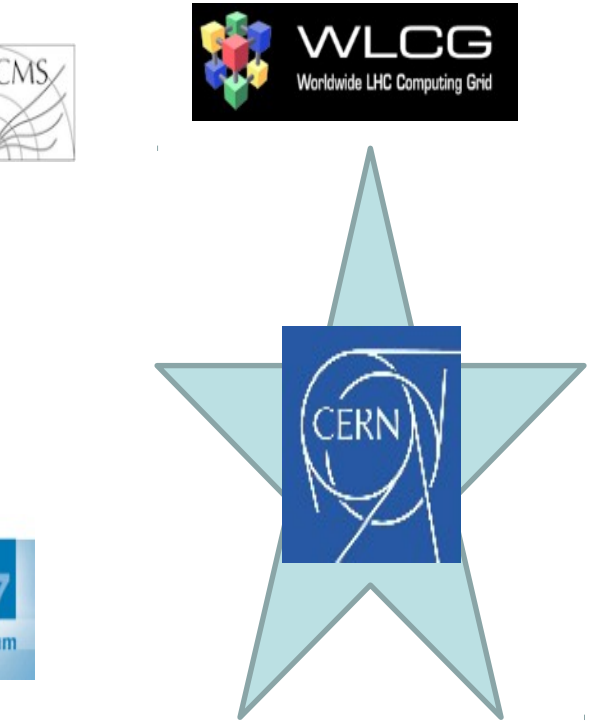

CERN Biomedical Facility

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

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

HUAWEI

ORACLE

SIEMENS

CONTRIBUTOR

rackspace

ASSOCIATE

Yandex

Multi-Core Platforms

CERN IT Department

- Very high-level of parallelism, both online and offline
- Parallelism exploited sending multiple independent computational jobs on separate physical or virtual nodes
- Multiple physical cores partitioned across virtual machines
- Current software not written to exploit multi-cores
- Very important issue with detector simulations, very CPU intensive tasks

- CERN is CPU based, GPUs little used
- Project to exploit multi-core platforms using vectorization techniques
- Build experience with data and task parallelism using Cilk and Haswell (CilkPlus and GCC)
- Requires rewriting the software: Geant 5

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

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 accellerators, 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

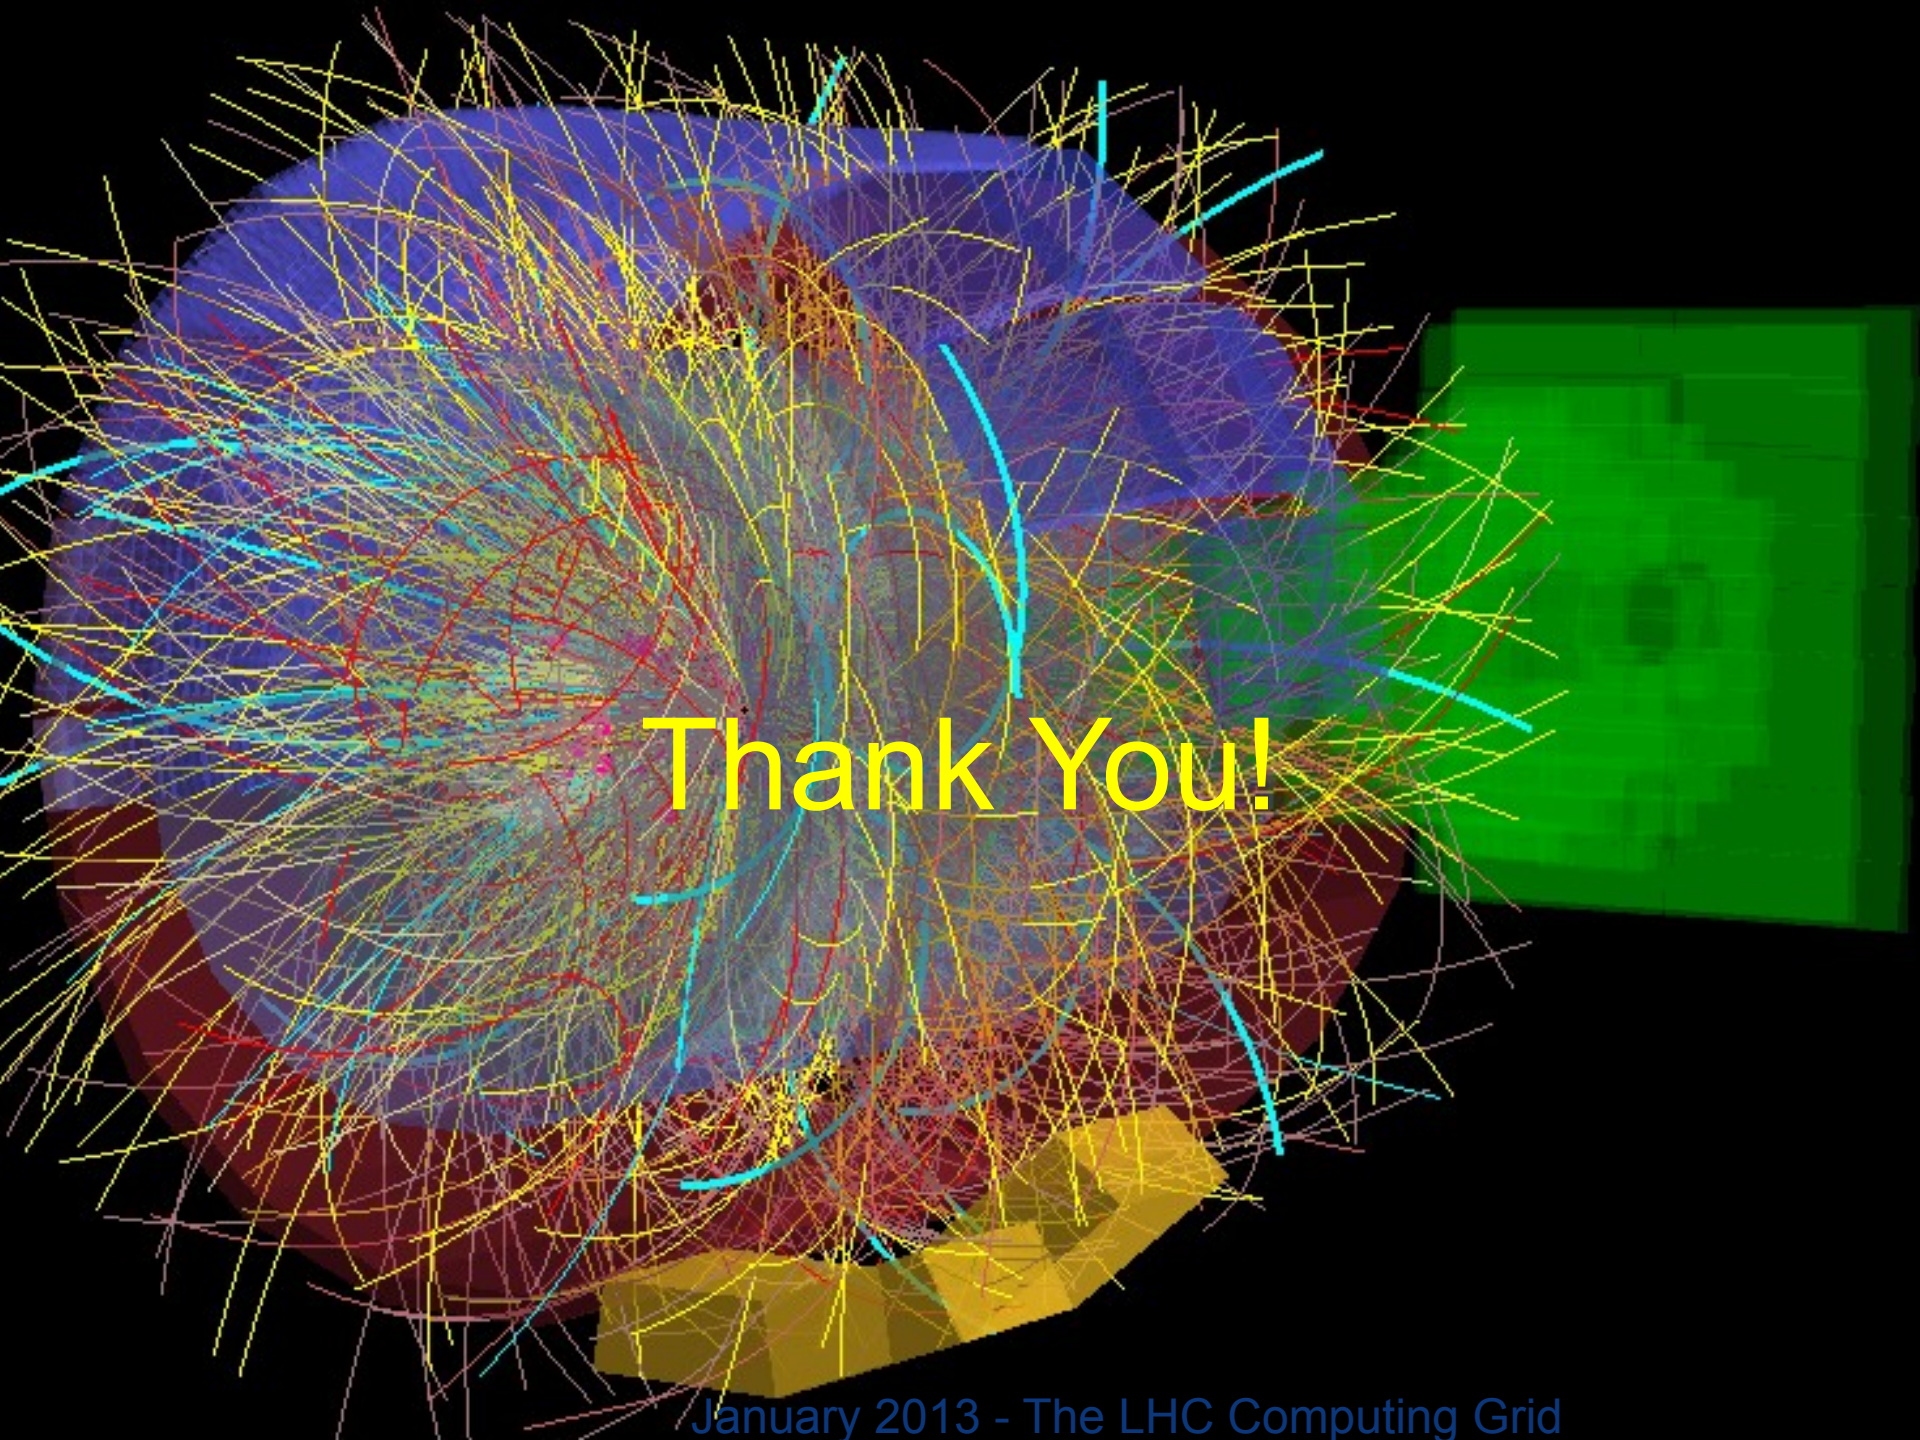
Thank You!

More information

www.cern.ch/openlab

GridCafé
www.gridcafe.org

WLCG
Worldwide LHC Computing Grid
www.cern.ch/lcg

NORDUGRID
Grid Solution for Wide Area Computing and Data Handling
www.nordugrid.org

www.eu-egi.org/

lhcathome.cern.ch

Nils.Hoimyr@cern.ch

cern.ch/it