CERN

European Organization for Nuclear Research
Organisation Européenne pour la Recherche Nucléaire

CERN: Big Science Meets Big Data

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 ~1000 MCHF annual

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

Candidates for Accession: Romania, Israel

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

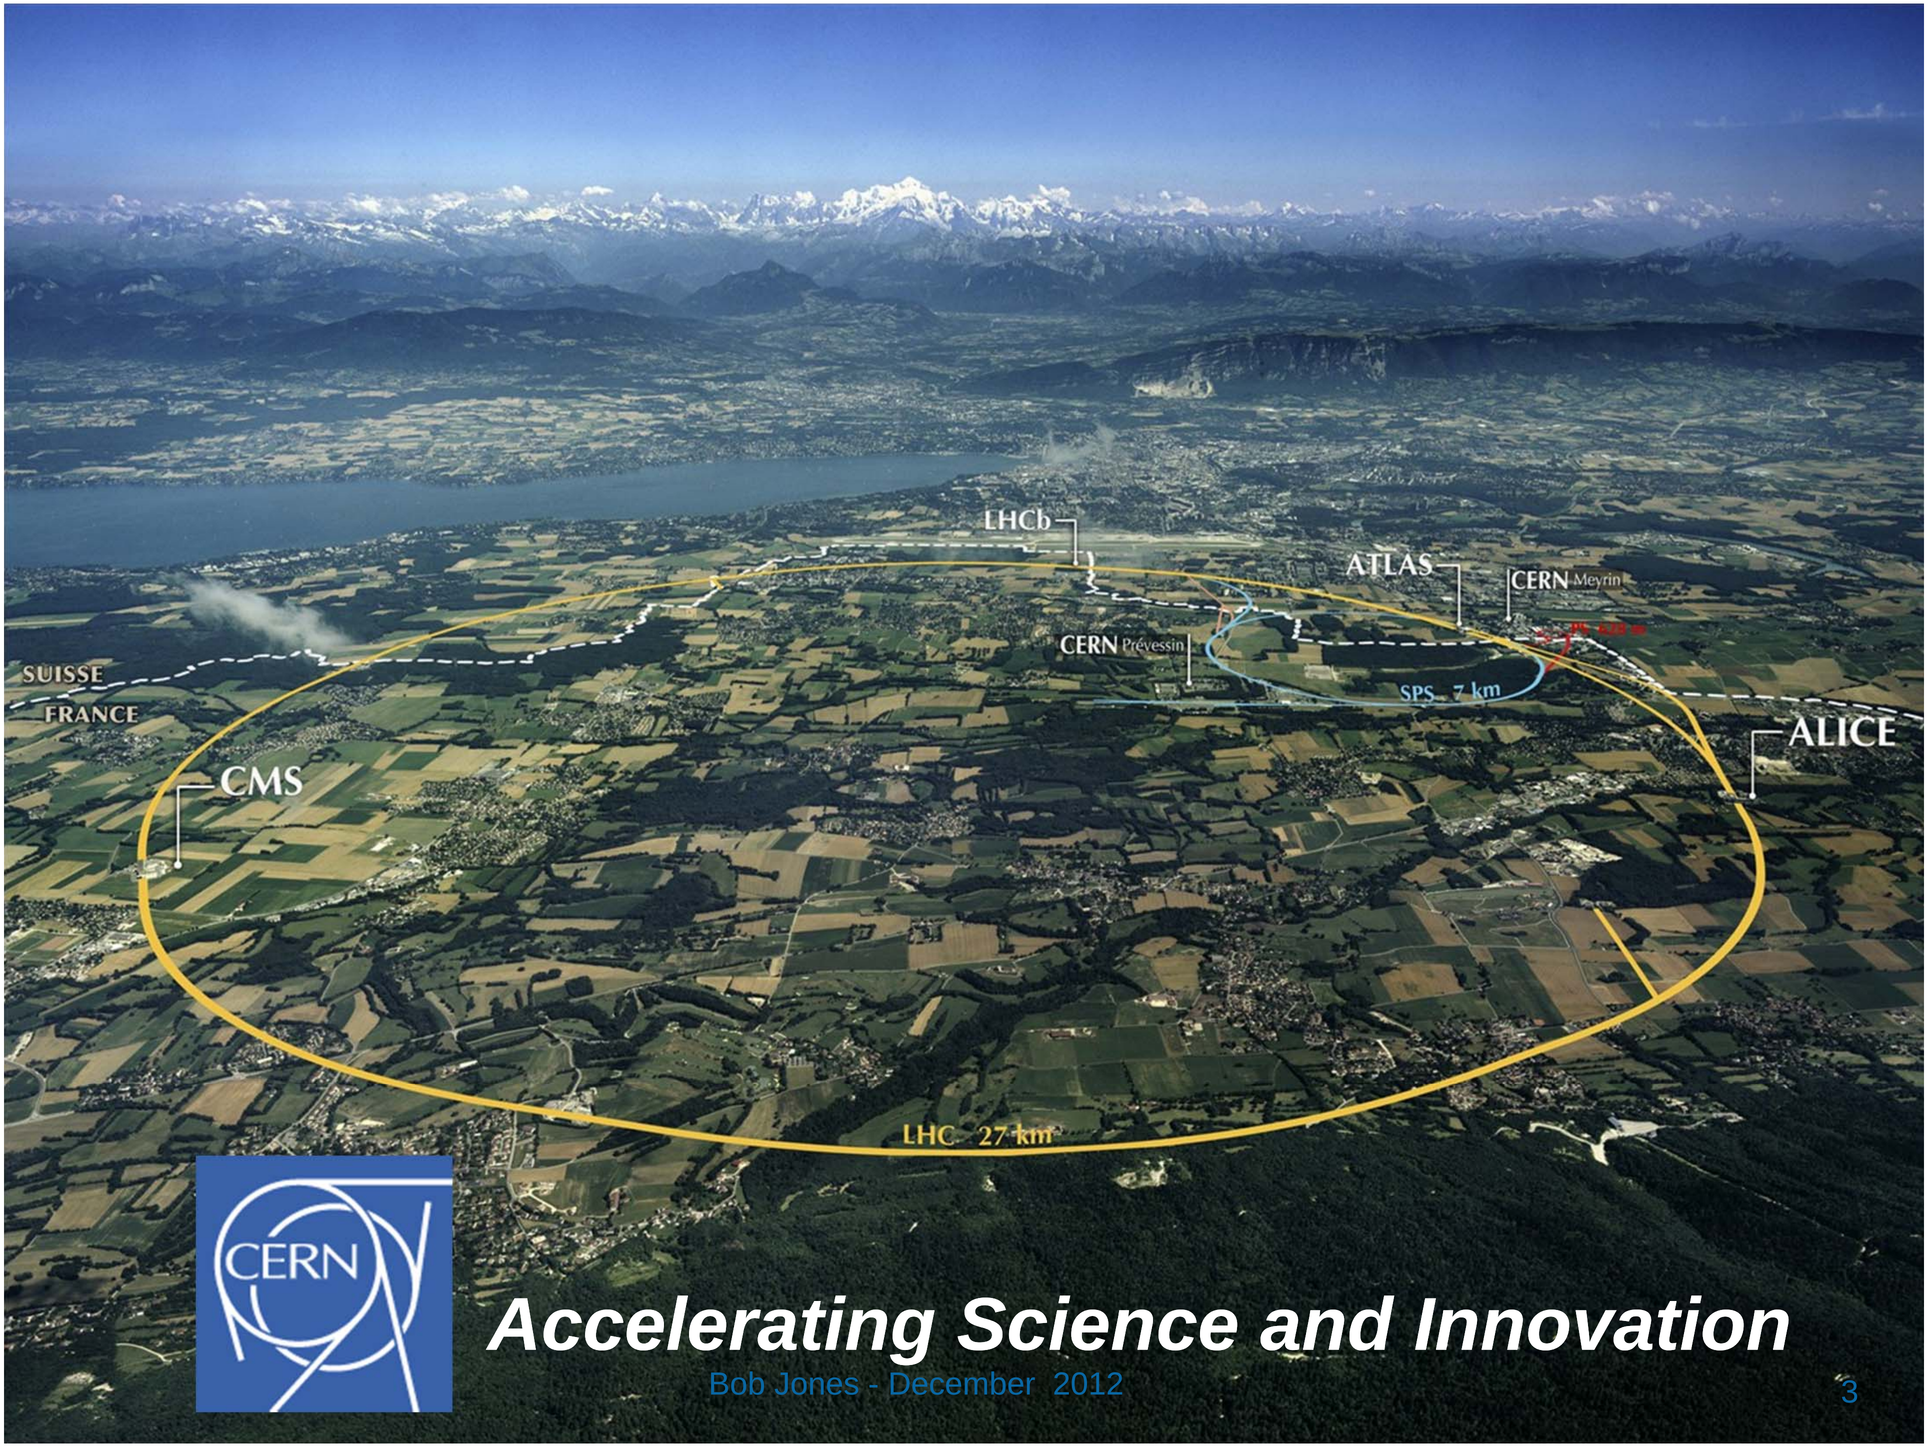

Accelerating Science and Innovation

Bob Jones - December 2012

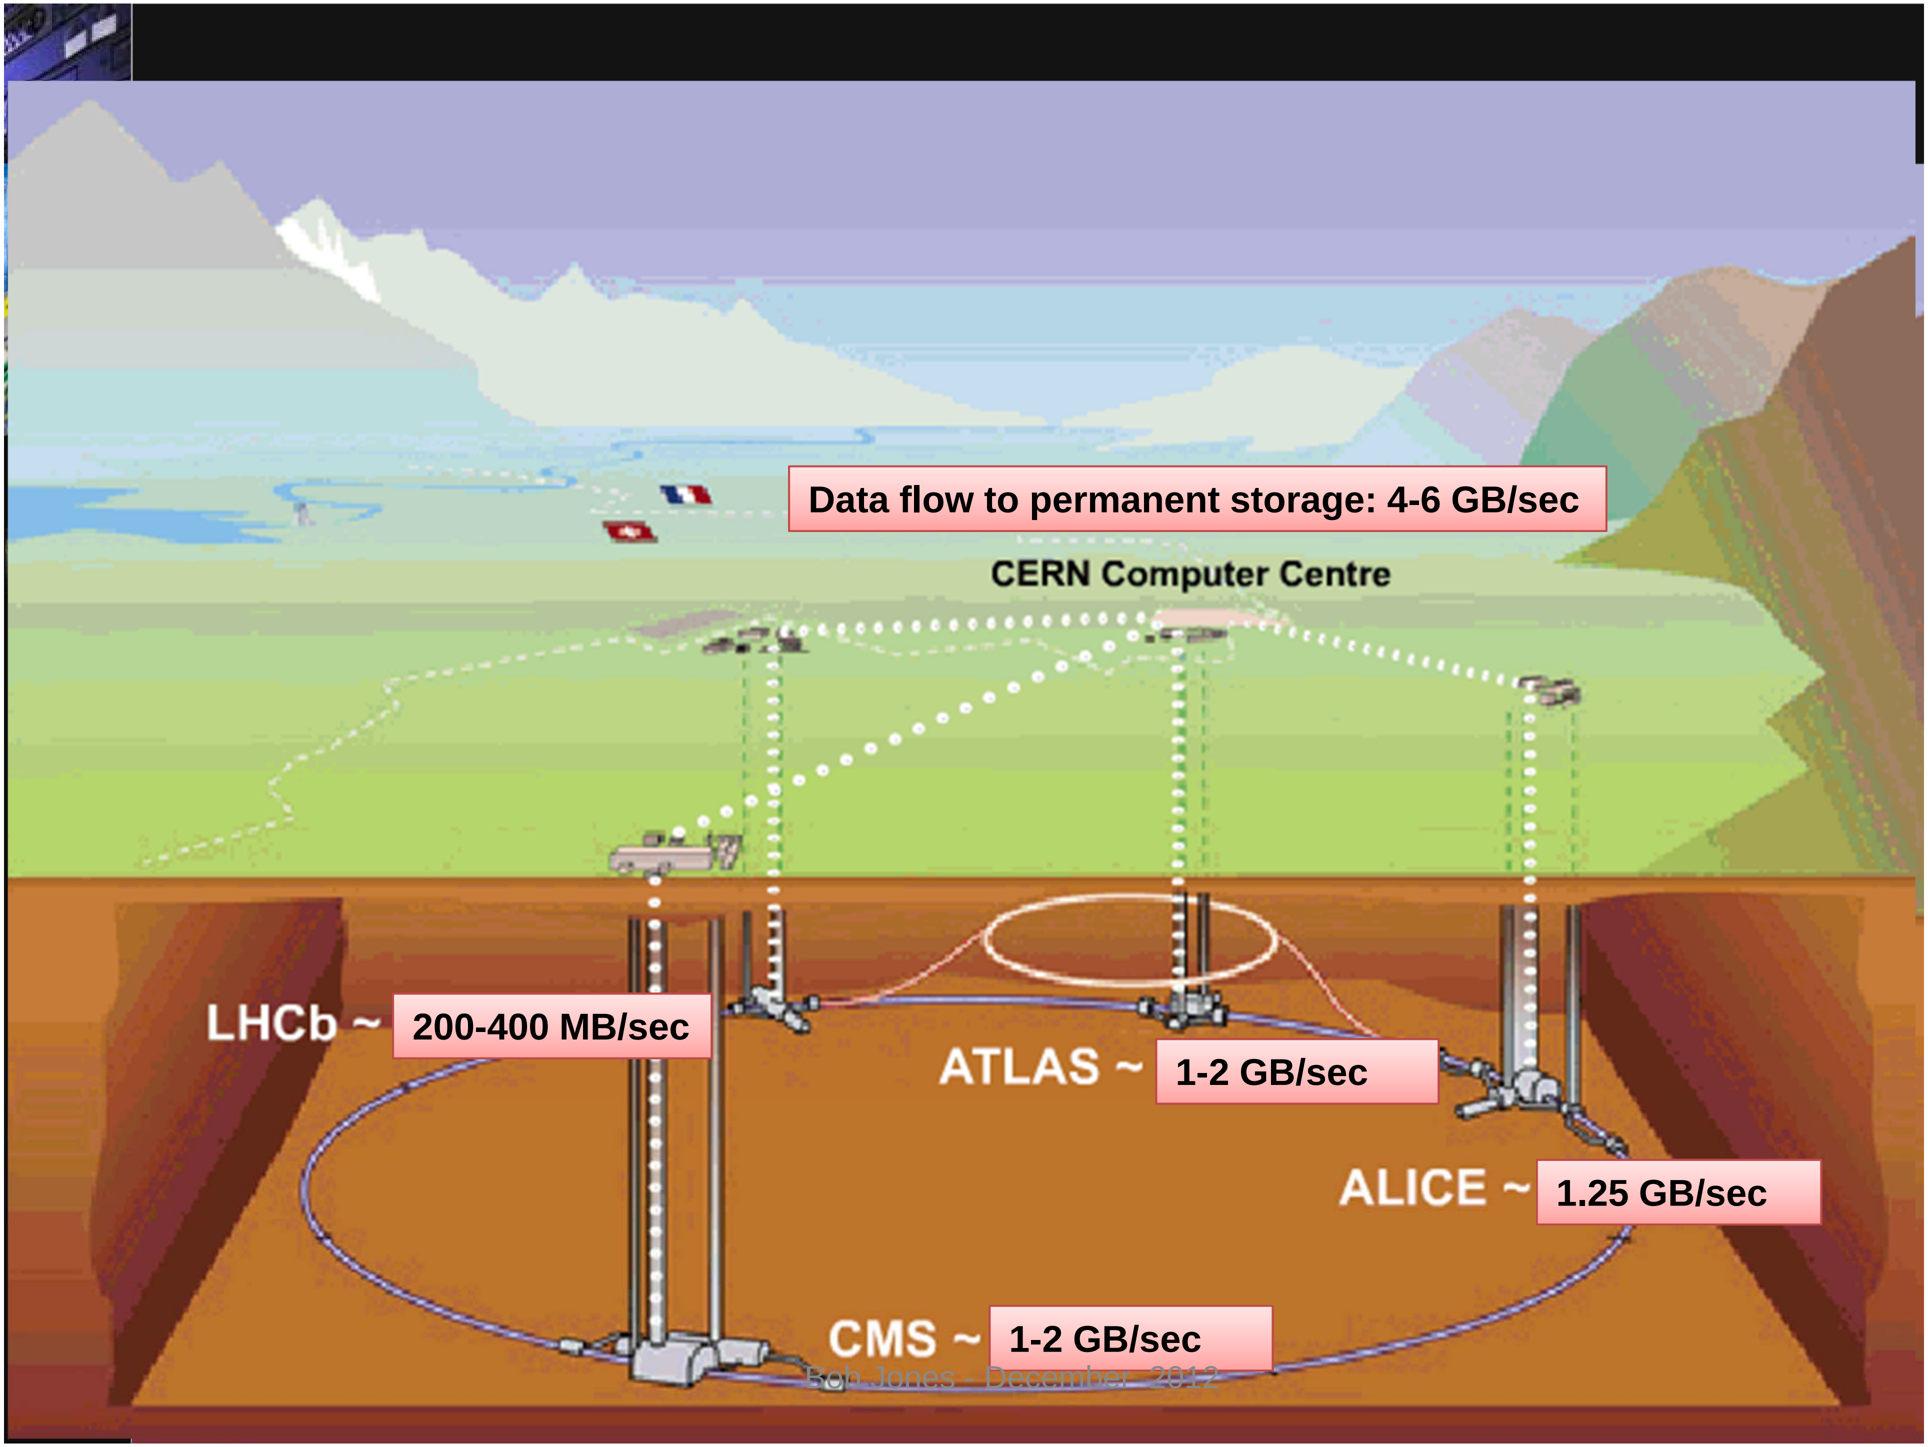

Data flow to permanent storage: 4-6 GB/sec
CERN Computer Centre
LHCb ~ 200-400 MB/sec
ATLAS ~ 1-2 GB/sec
ALICE ~ 1.25 GB/sec
CMS ~ 1-2 GB/sec
Bob Jones - December 2012

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

- Secure access via X509 certificates issued by network of national authorities - International Grid Trust Federation (IGTF)
  - http://www.igtf.net/

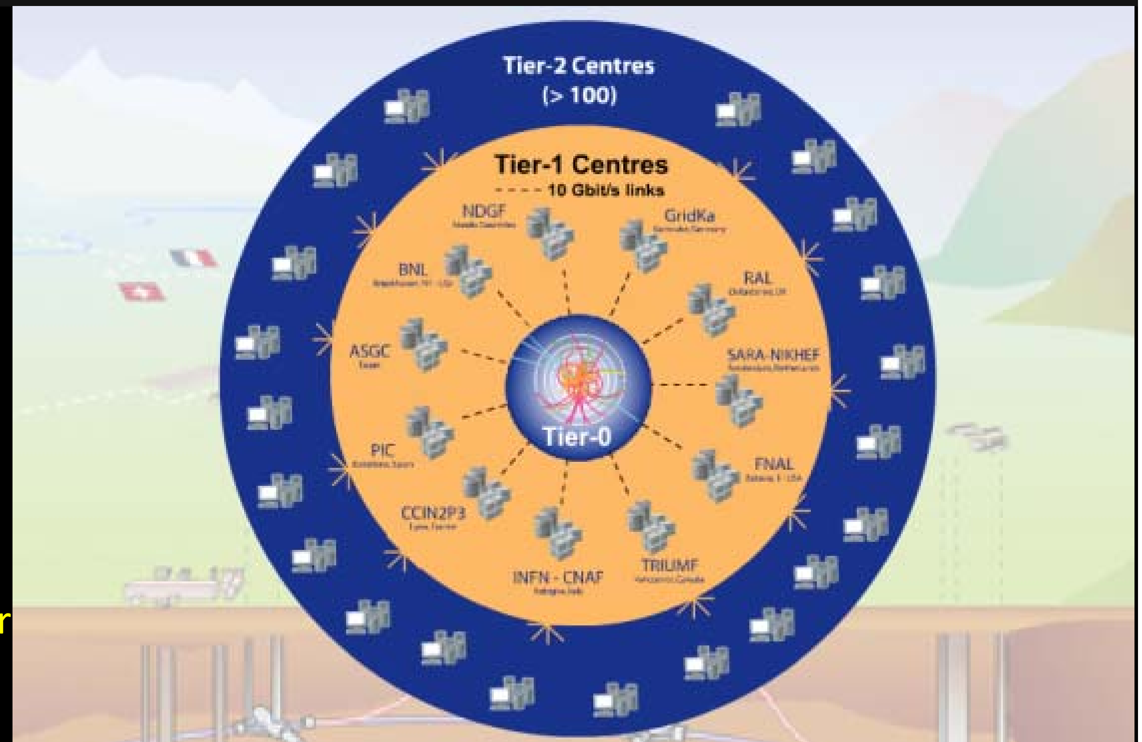

Tier-0 (CERN):
- Data recording
- Initial data reconstruction
- Data distribution

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

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

WLCG: Data Taking

- Castor service at Tier 0 well adapted to the load:
  - Heavy Ions: more than 6 GB/s to tape (tests show that Castor can easily support >12 GB/s); Actual limit now is network from experiment to CC
  - Major improvements in tape efficiencies – tape writing at ~native drive speeds. Fewer drives needed
  - ALICE had x3 compression for raw data in HI runs

HI: ALICE data into Castor > 4 GB/s (red)

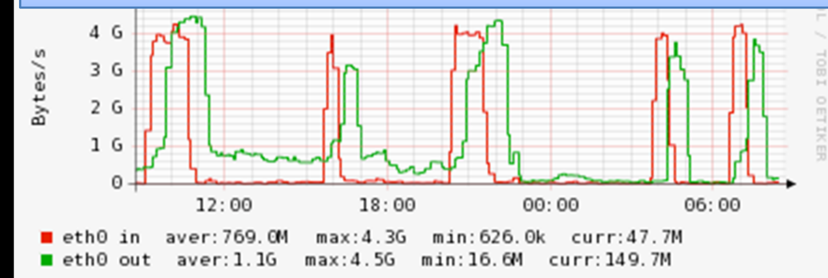

HI: Overall rates to tape > 6 GB/s (r+b)

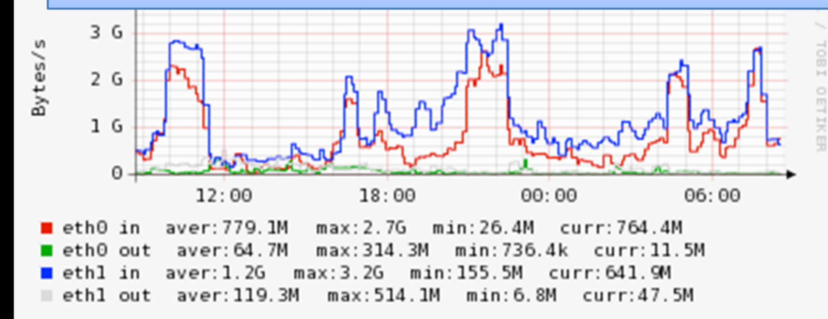

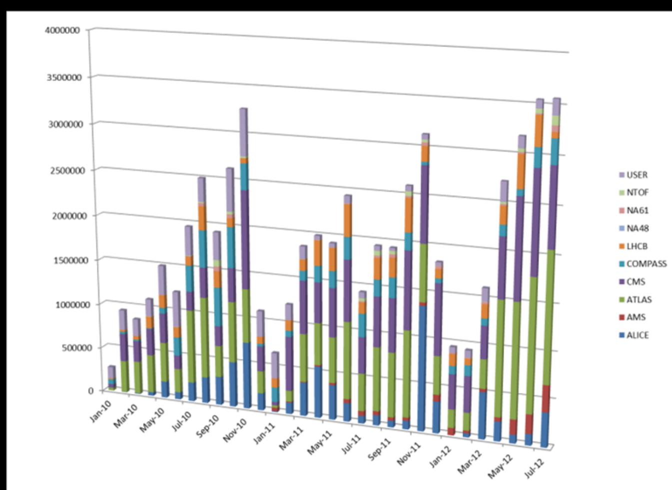

23 PB data written in 2011

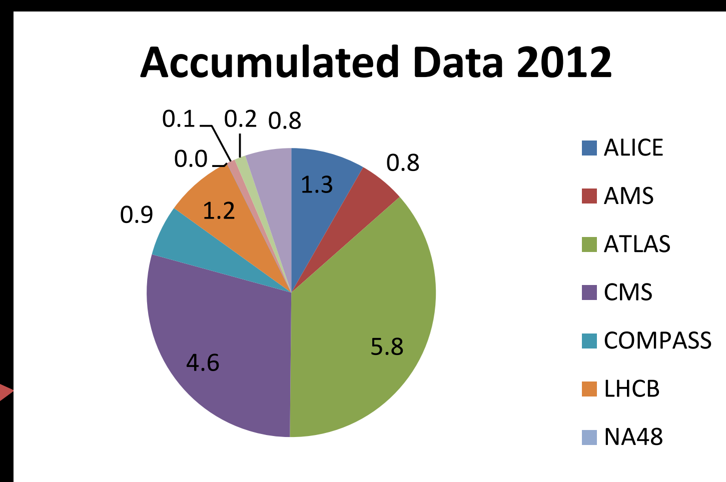

Overall use of WLCG

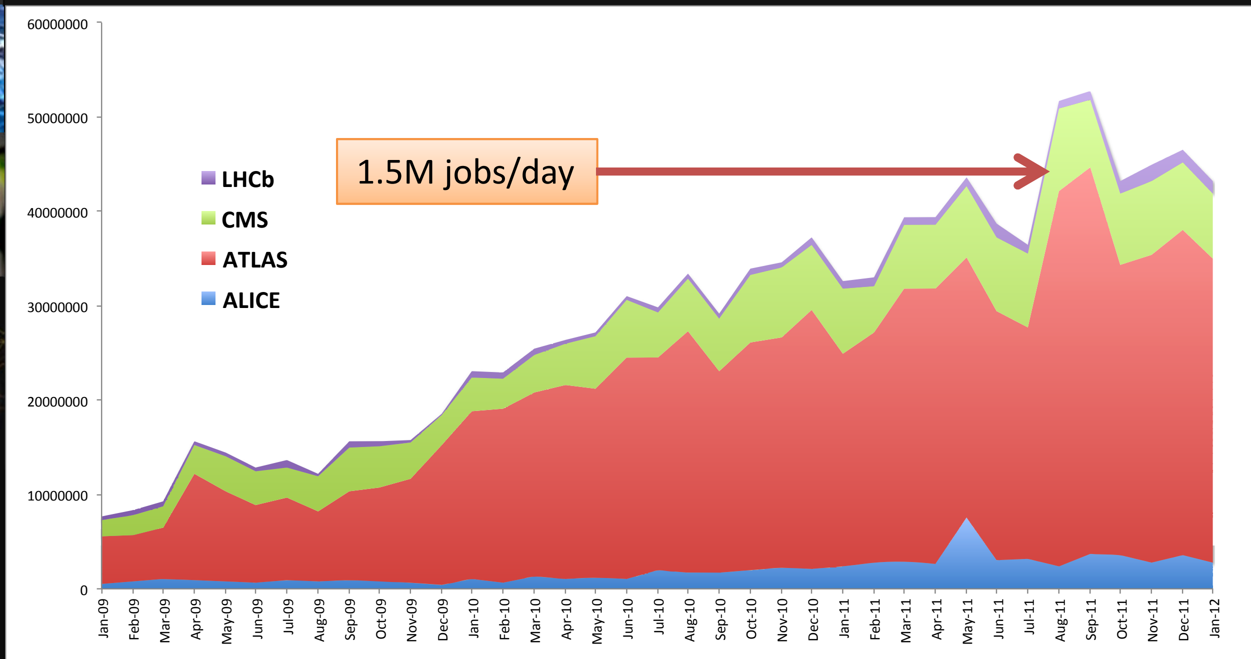

Usage continues to grow even over end of year technical stop

- # jobs/day
- CPU usage

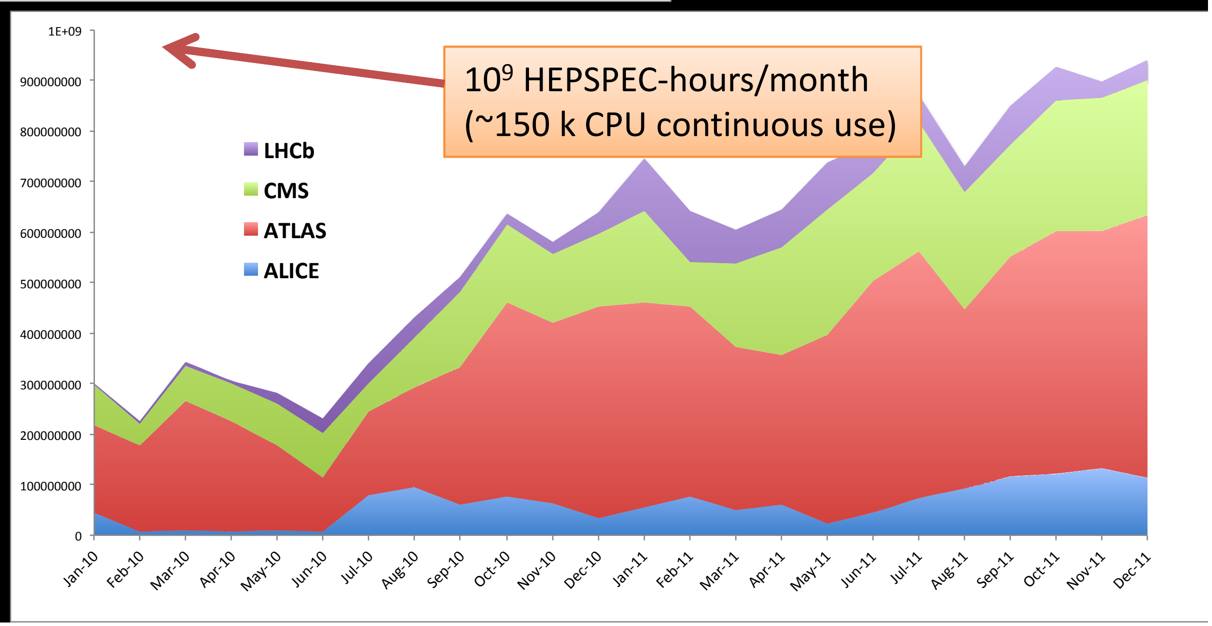

WLCG
Worldwide LHC Computing Grid

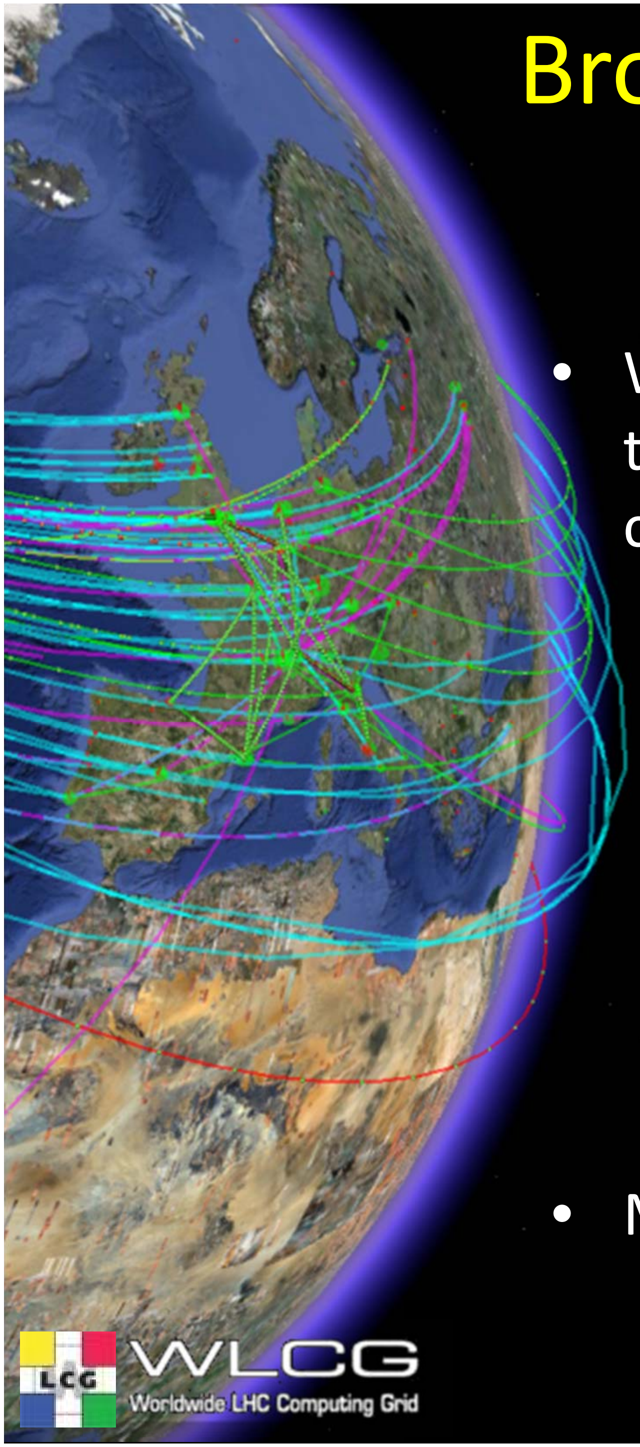

Broader Impact of the LHC Computing Grid

- WLCG has been leveraged on both sides of the Atlantic, to benefit the wider scientific community
  - Europe (EC FP7):
    - Enabling Grids for E-sciencE (EGEE) 2004-2010
    - European Grid Infrastructure (EGI) 2010--
  - USA (NSF):
    - 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
…

How to evolve LHC data processing

Making what we have today more sustainable is a challenge

- Big 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 at all levels (including cloud computing services)
- 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

www.cern.ch/openlab

PARTNERS

HP, Intel, Oracle, Siemens

CONTRIBUTOR (2012)

Huawei

Virtuous Cycle

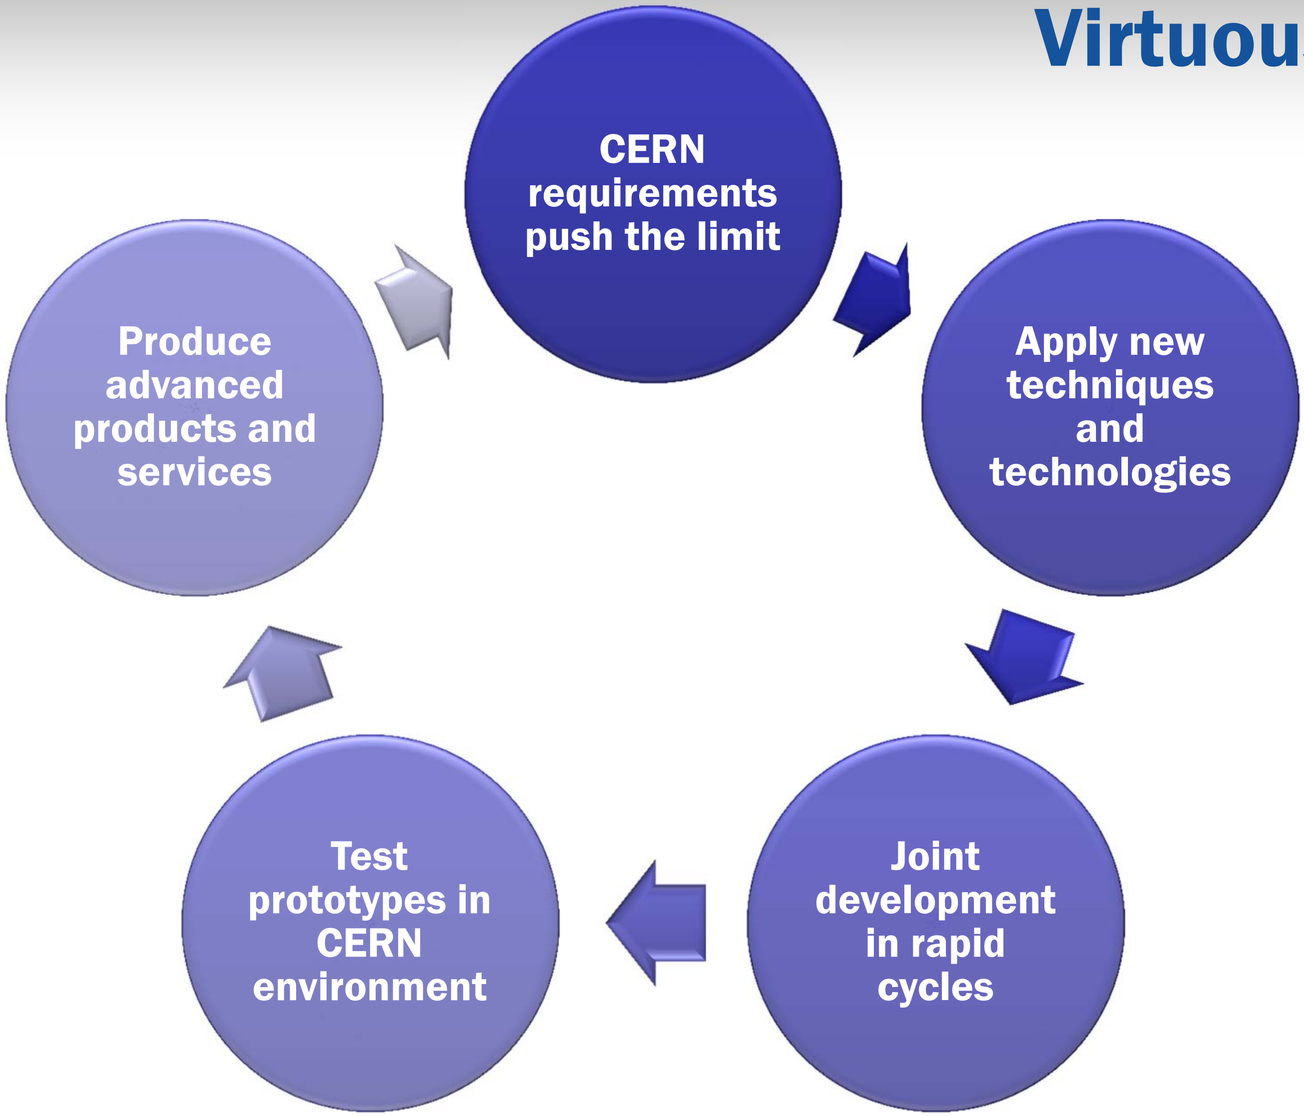

A public-private partnership between the research community and industry

http://openlab.cern.ch

Maaike Limper

Physics Analysis in an ORACLE Database

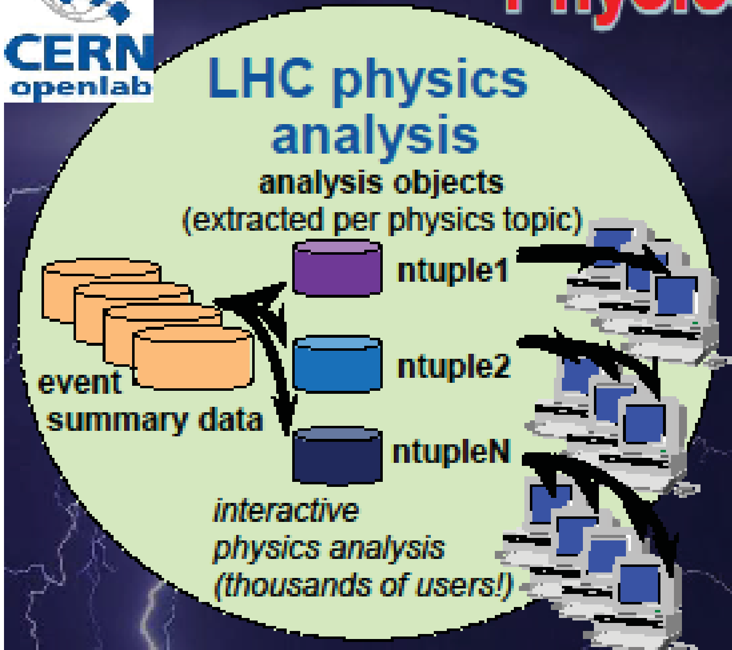

Physics Analysis at LHC is currently done by analyzing information stored in centrally produced ROOT-ntuples
(ROOT: a dedicated C++ analysis framework)

Each physics group determines their ntuple-content based on physics objects and level of detail needed for their analysis (thousands of variables stored per event)

~1000 physics events per ntuple-file, many files need to be processed for each physics analysis, users store files locally or send ROOT-macro to grid

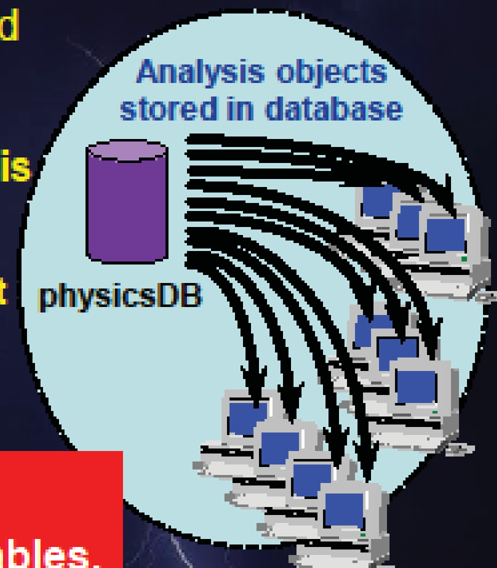

Physics Analysis DB: radically different approach to file-based analysis

Separate tables by physics object type (jets, muons, electrons etc.)

Can decide for each physics topic which tables/columns are relevant

Each user can send their own SQL-query to the physicsDB using PL/SQL functions for specific analysis tasks

A true BIG DATA challenge:

Billions of events, thousands of physics variables, interactive physics analysis by thousands of users

Opportunity to test database tools for optimizing the performance of storing and analyzing complex data

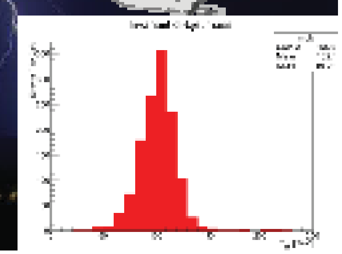

A European Cloud Computing Partnership: big science teams up with big business

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

DLR cnes

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

Atos Capgemini CloudSigma CSA cloud security alliance egi interoute logica OpenNebula.org orange Business Services SAP the SERVER LABS sixsq Telefónica terradue 2.0 THALES Trust IT T-Systems

Email:contact@helix-nebula.eu Twitter: HelixNebulaSC Web: www.helix-nebula.eu

Initial flagship use cases

ATLAS High Energy Physics Cloud Use	Genomic Assembly in the Cloud	SuperSites Exploitation Platform
	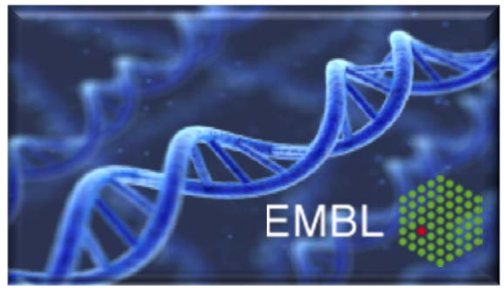	
To support the computing capacity needs for the ATLAS experiment	A new service to simplify large scale genome analysis; for a deeper insight into evolution and biodiversity	To create an Earth Observation platform, focusing on earthquake and volcano research

- **Scientific challenges with societal impact**
- **Sponsored by user organisations**
- ***Stretch* what is possible with the cloud today**

Conclusions

- Big Data Management and Analytics require a solid organisational structure at all levels

- Corporate culture: our community started preparing more than a decade before real physics data arrived

- Now, the LHC data processing is under control but data rates will continue to increase (dramatically) for years to come: **Big Data at the Exabyte scale**

- CERN is working with other science organisations and leading IT companies to address our growing big-data needs


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