The Worldwide LHC Computing Grid (WLCG)

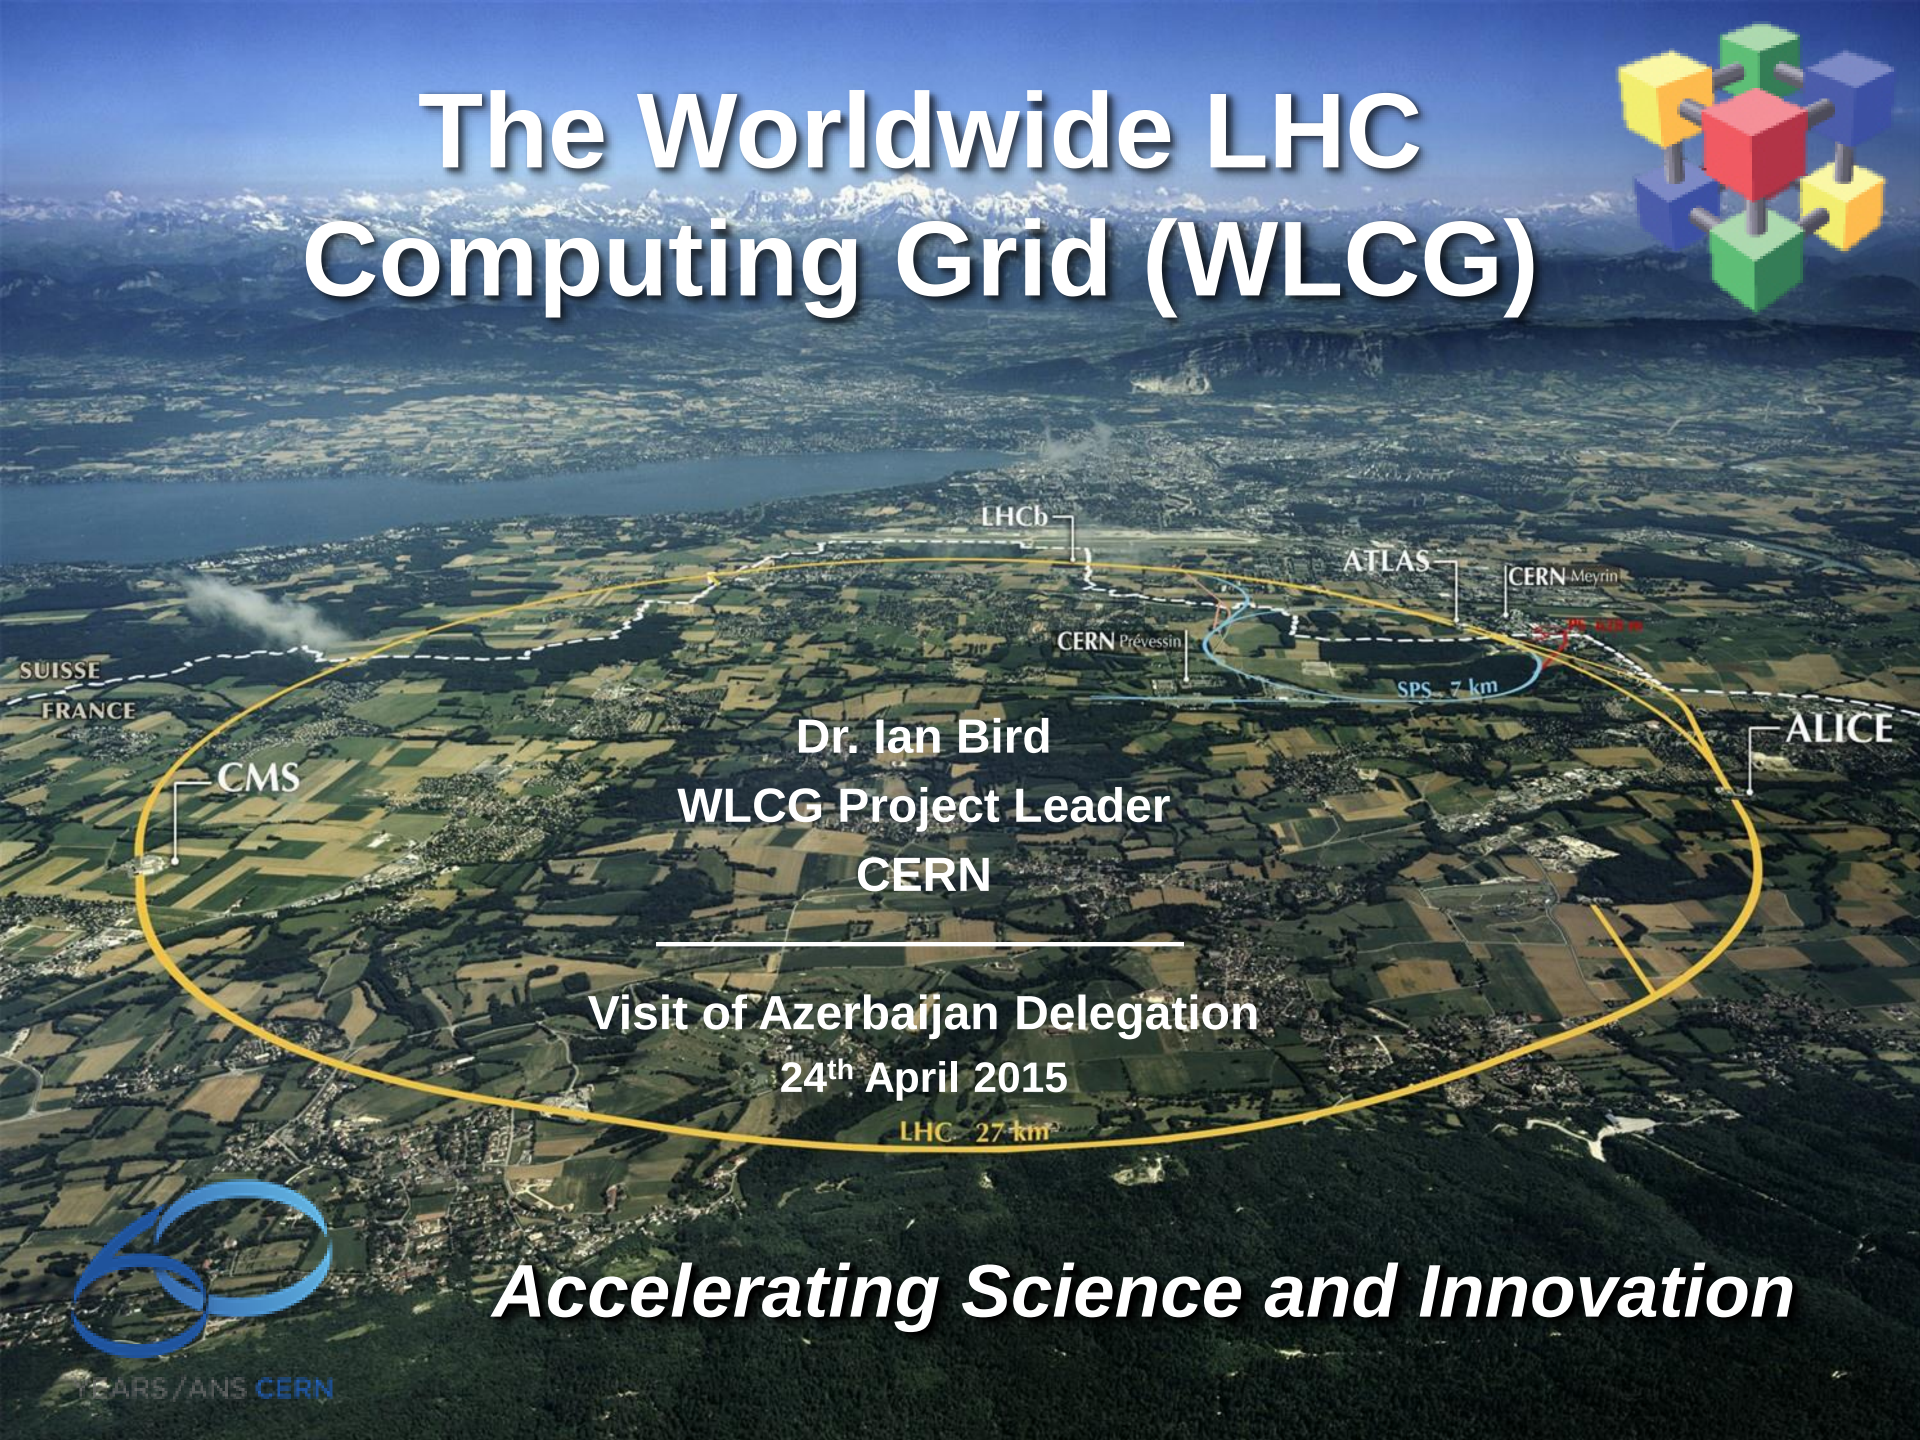

Dr. Ian Bird
WLCG Project Leader
CERN

Visit of Azerbaijan Delegation
24^{th} April 2015

Accelerating Science and Innovation

YEARS / ANS CERN

Tools: LHC and Detectors

pp, B-Physics, CP Violation
(matter-antimatter symmetry)

LHCb
ATLAS
CERN Meyrin
SPS 7 km

General Purpose,
proton-proton, heavy ions
Discovery of new physics:
Higgs, SuperSymmetry

ALICE
CMS

Exploration of a new energy frontier
in p-p and Pb-Pb collisions

Heavy ions, pp
(state of matter of early universe)

How to make a discovery ?

Accelerator ➡ Experiment ➡ Analysis (Computing) ➡ Discovery

Grid / Cloud

Big Data


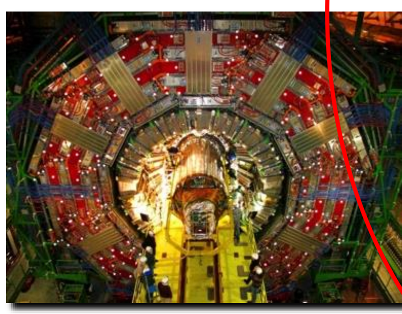

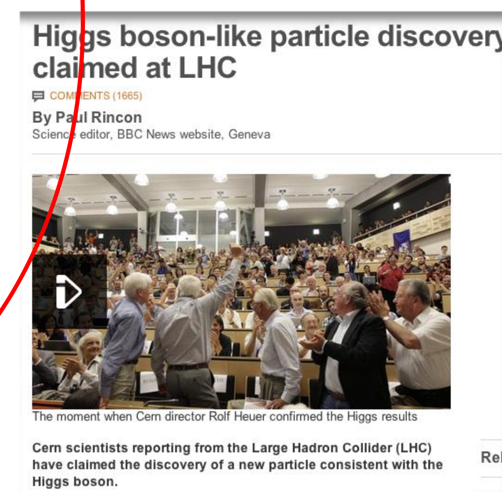

Higgs boson-like particle discovery claimed at LHC

COMMENTS (1665)

By Paul Rincon
Science editor, BBC News website, Geneva

The moment when Cern director Rolf Heuer confirmed the Higgs results

Cern scientists reporting from the Large Hadron Collider (LHC) have claimed the discovery of a new particle consistent with the Higgs boson.

What is the data?

- 150 million sensors deliver data …40 million times per second

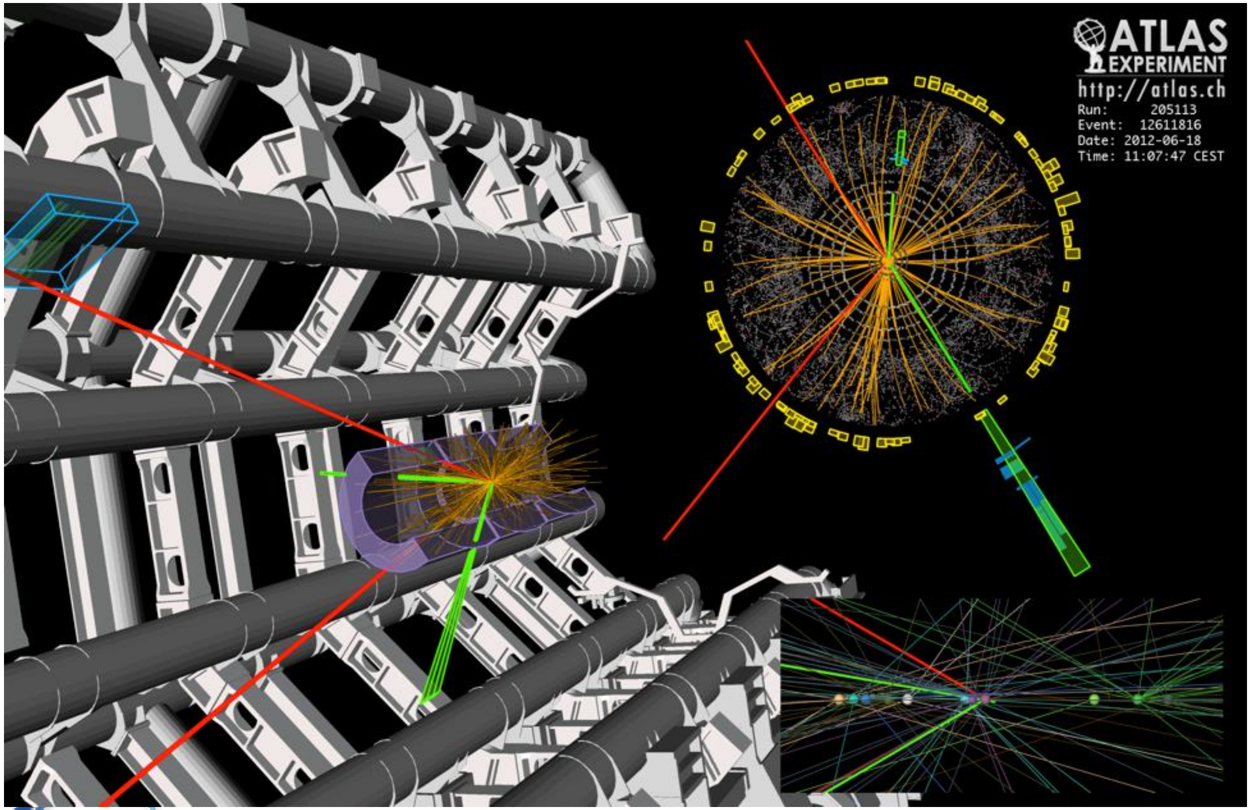

Up to 6 GB/s to be permanently stored after filtering

Data Collection and Archiving at CERN

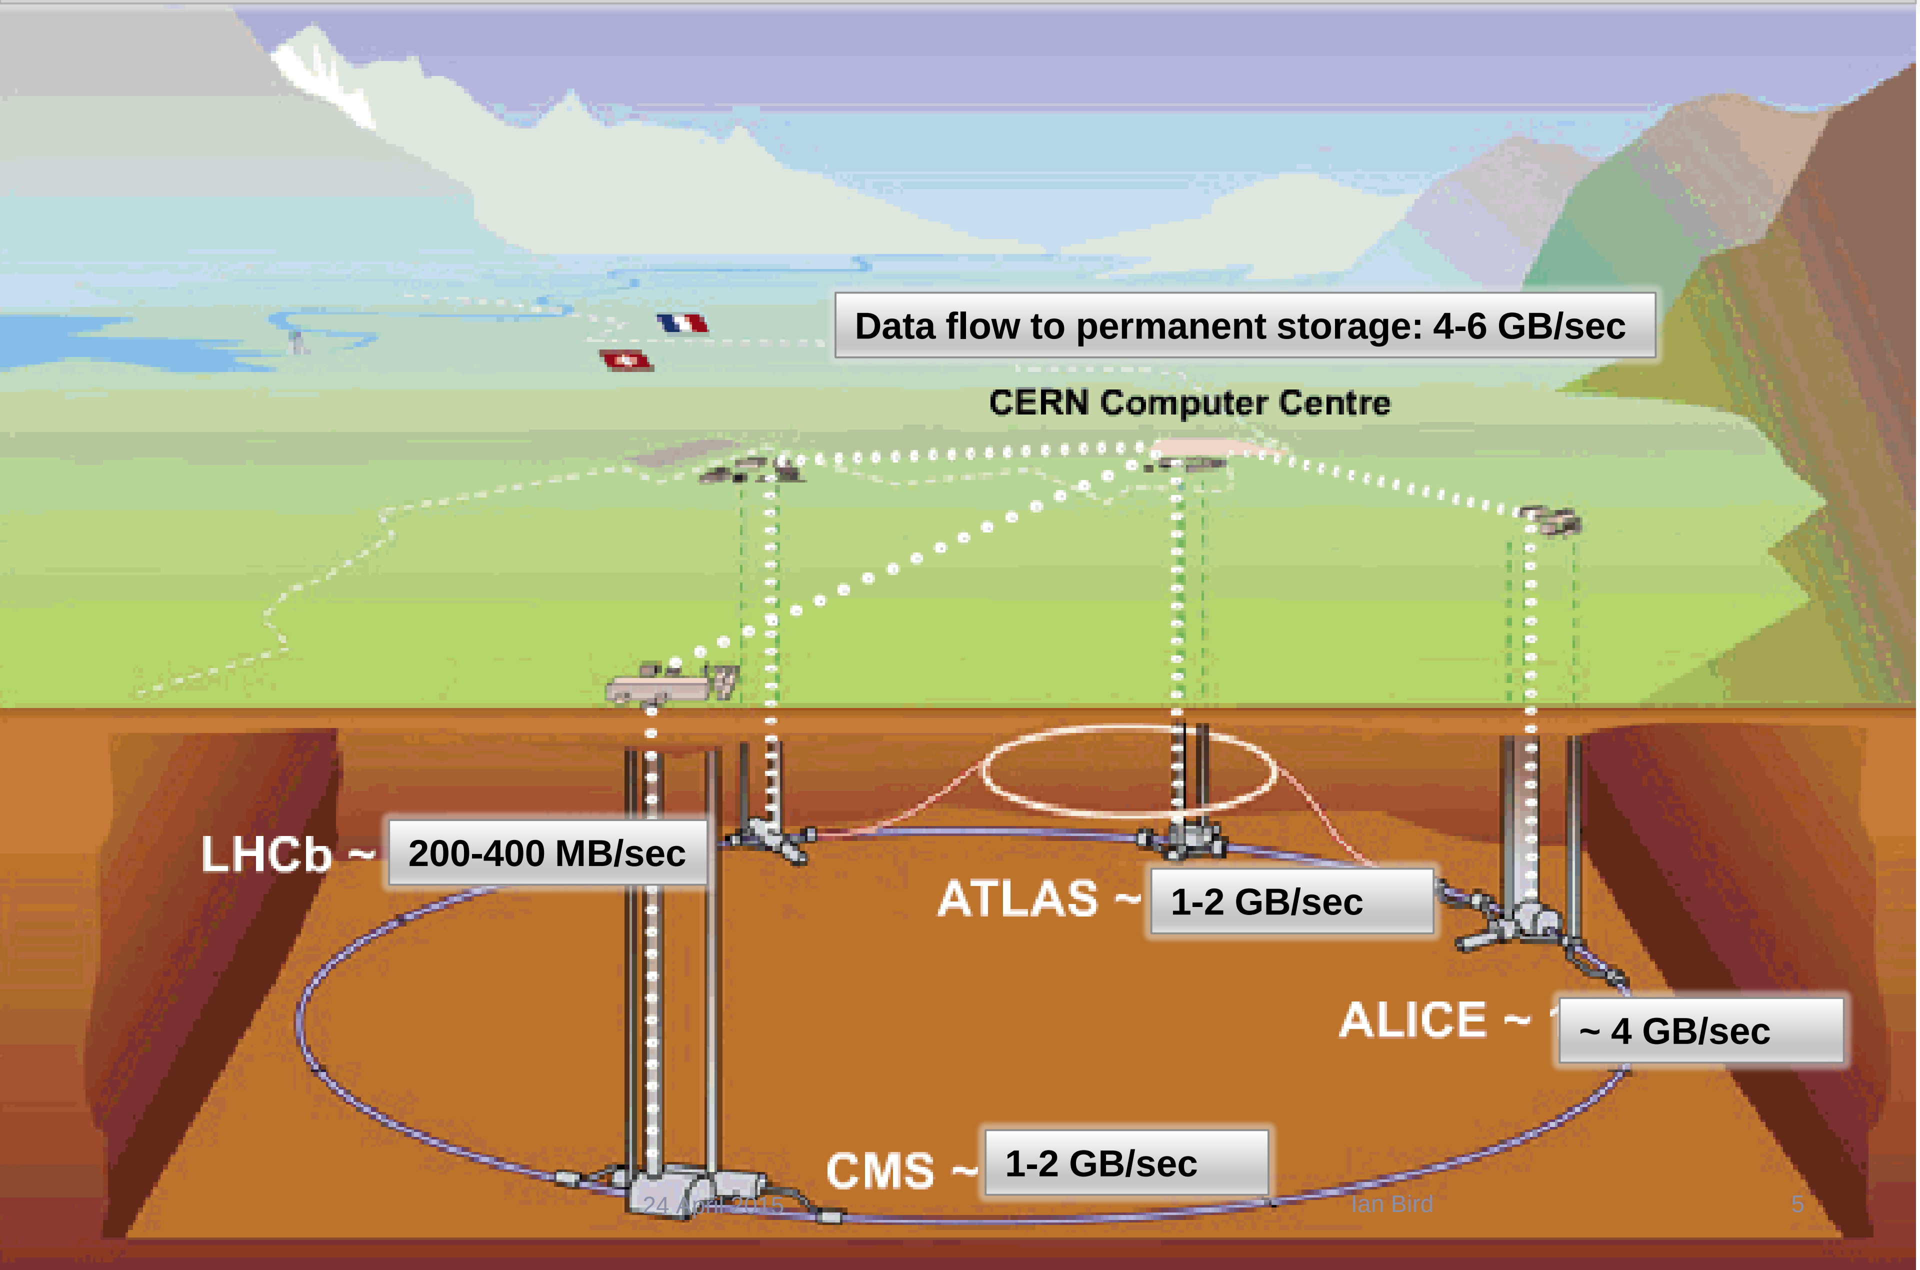

The Worldwide LHC Computing Grid

Tier-0 (CERN): data recording, reconstruction and distribution

Tier-1: permanent storage, re-processing, analysis

Tier-2: Simulation, end-user analysis

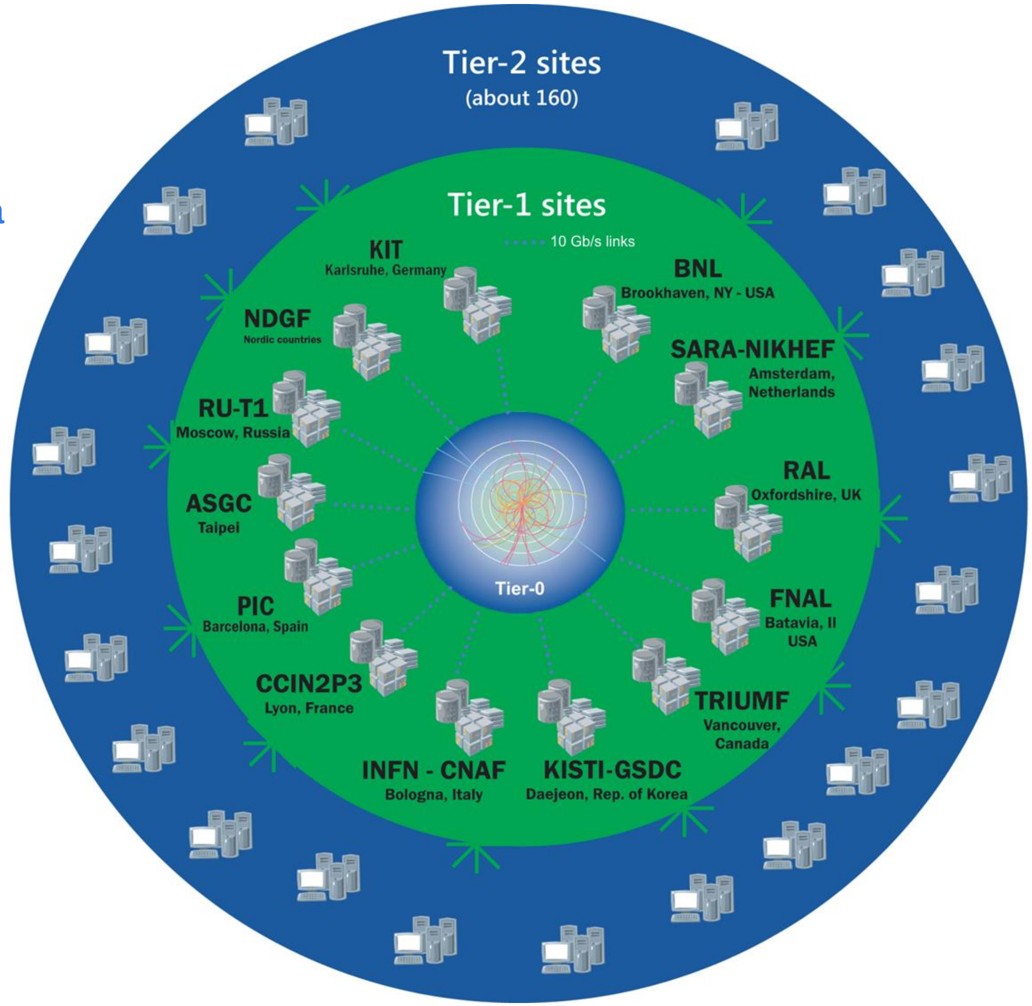

nearly 170 sites, 40 countries

~350’000 cores

500 PB of storage

> 2 million jobs/day

10-100 Gb links

WLCG:
An International collaboration to distribute and analyse LHC data

Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists

Truly world-wide

Tier 0 Tier 1 Tier 2

Scale of data today …

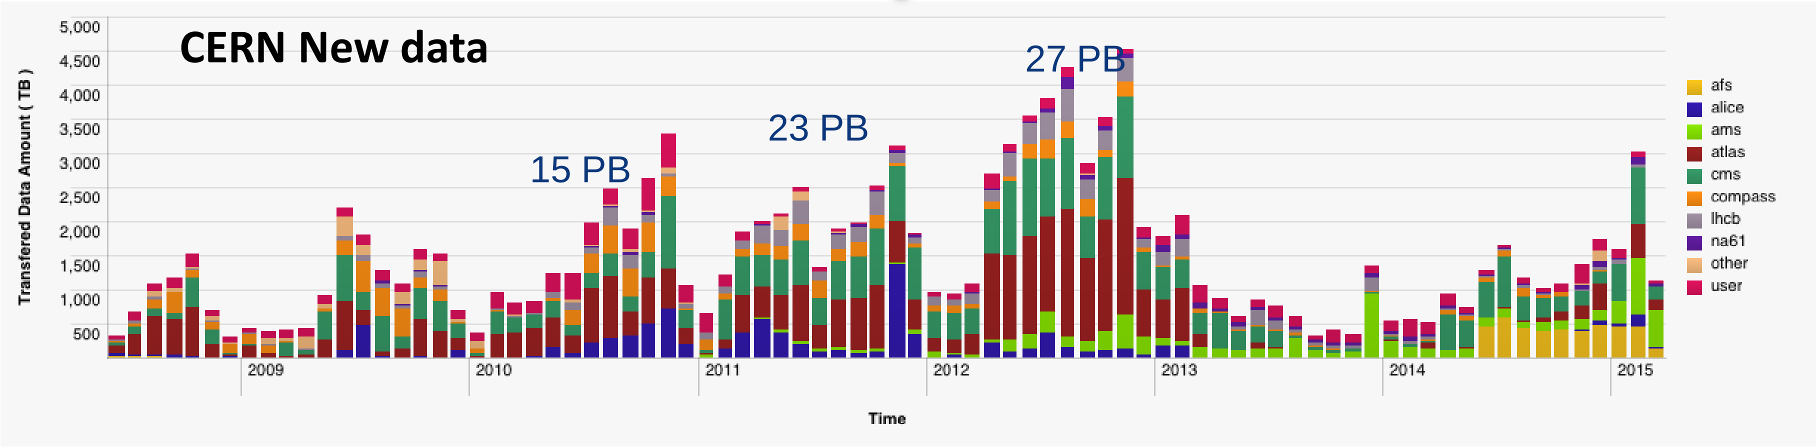

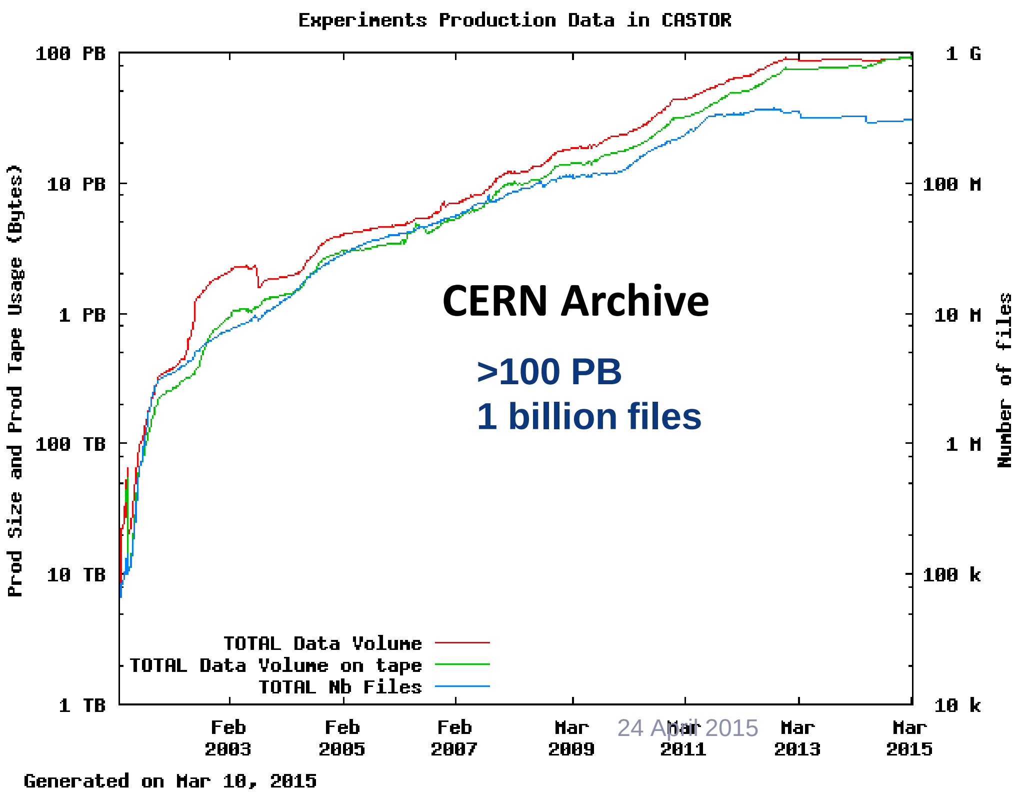

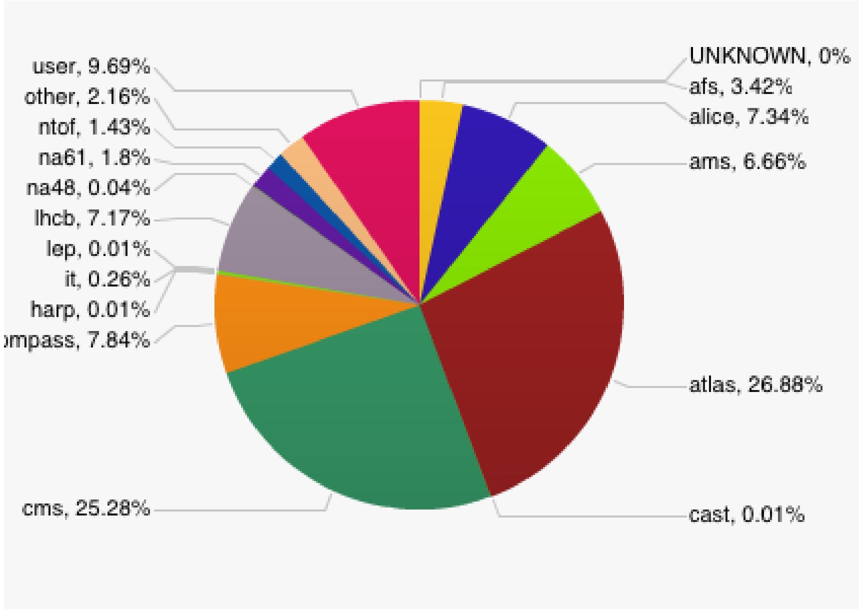

LHC – Big Data…

Few PB of raw data becomes ~100 PB! ➔

- Duplicate raw data
- Simulated data
- Many derived data products
- Recreate as software gets improved
- Replicate to allow physicists to access it

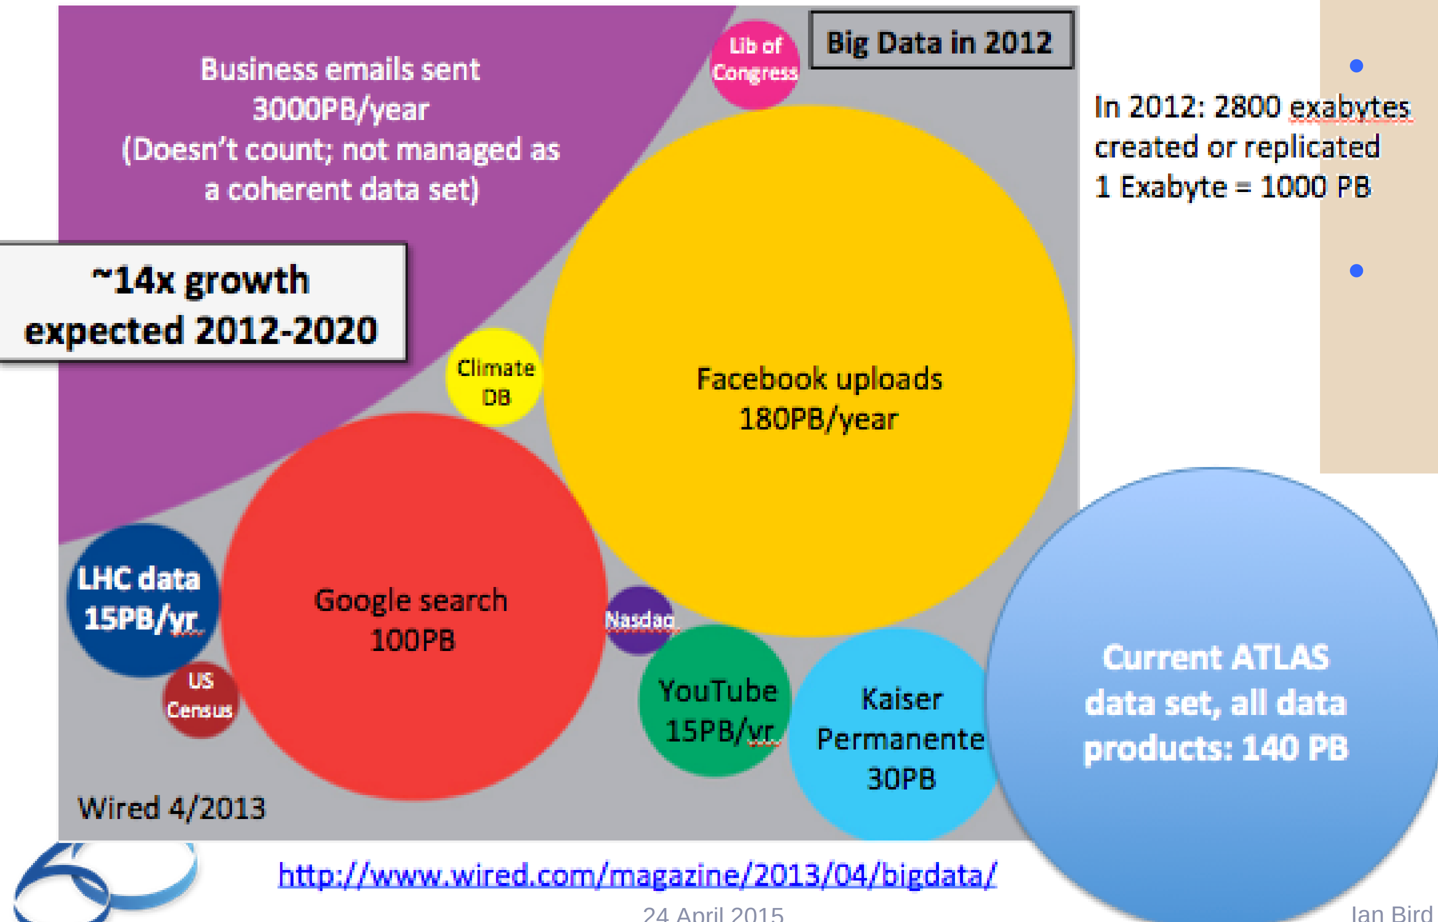

In 2012: 2800 exabytes created or replicated
1 Exabyte = 1000 PB

http://www.wired.com/magazine/2013/04/bigdata/

Tier-1/2 CPU – 01.2012-12.2012

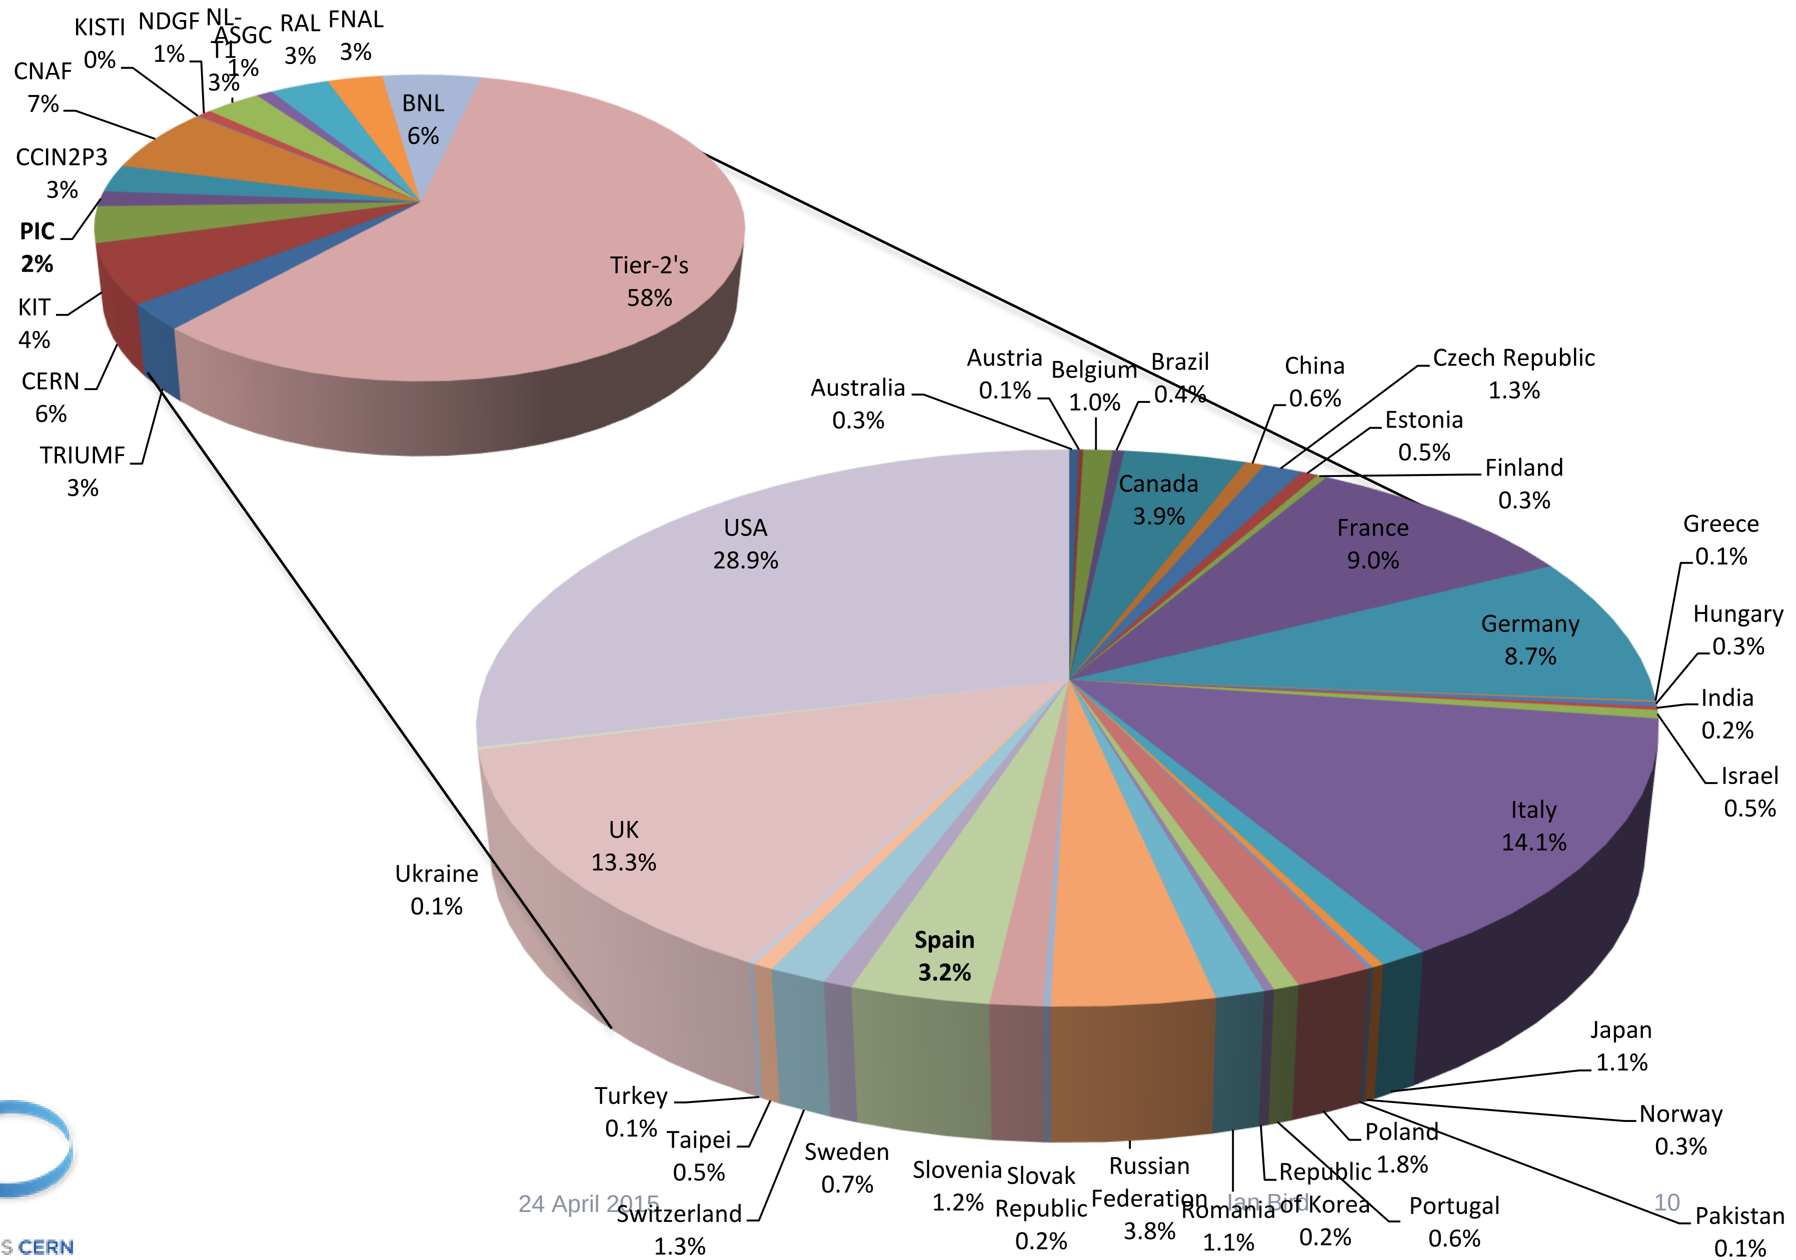

LHC OPN

- Optical Private Network
- Support T0 – T1 transfers
- Some T1 – T1 traffic
- Managed by LHC Tier 0 and Tier 1 sites

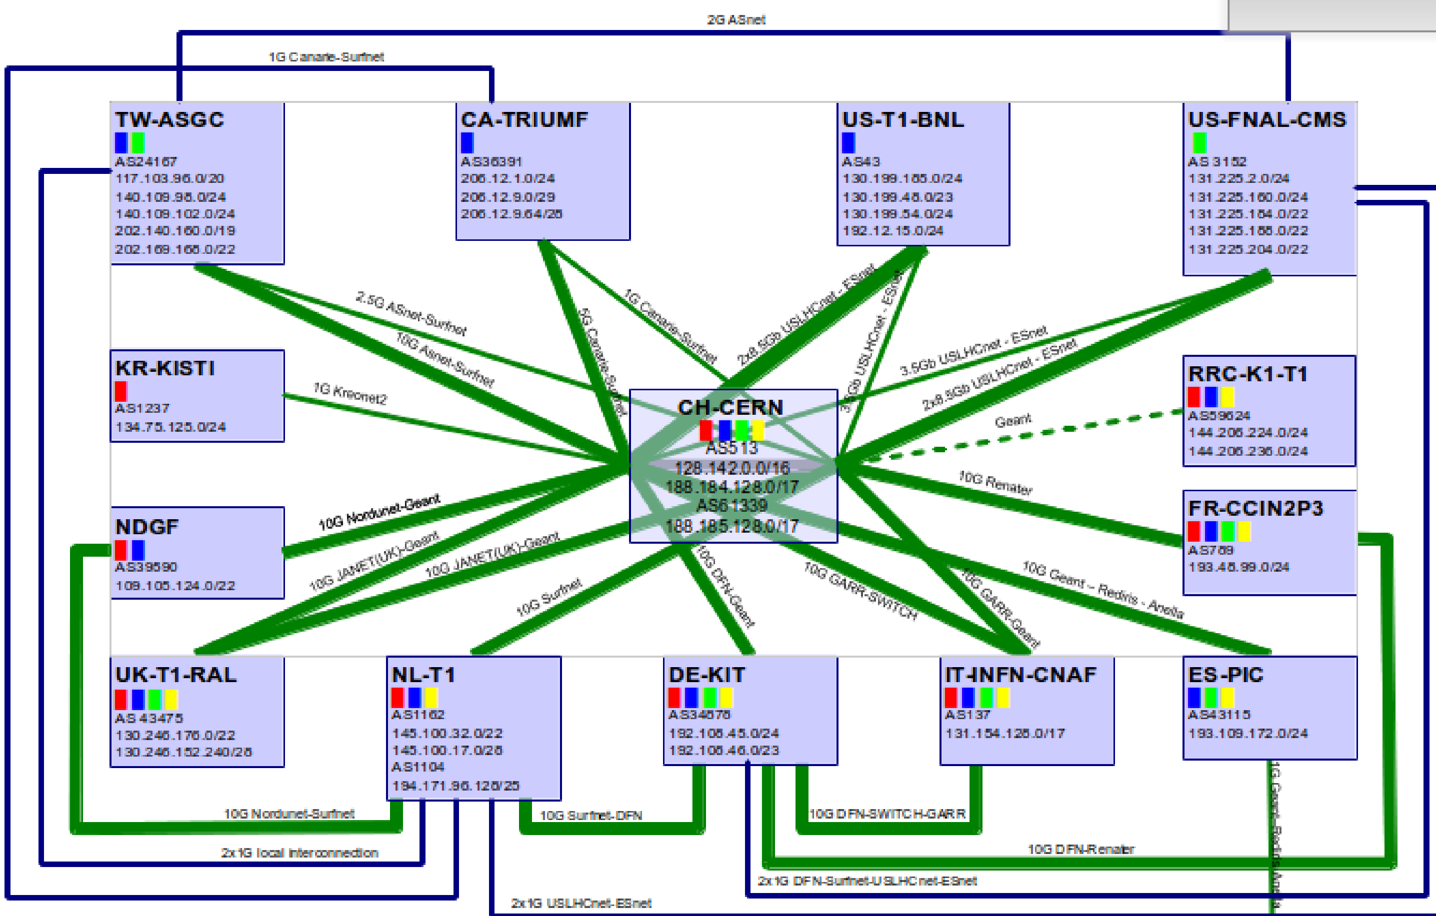

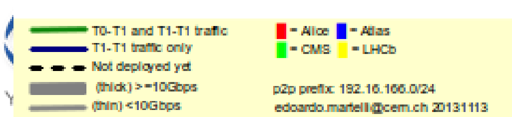

LHCONE: A global infrastructure for the High Energy Physics (LHC and Belle II) data management

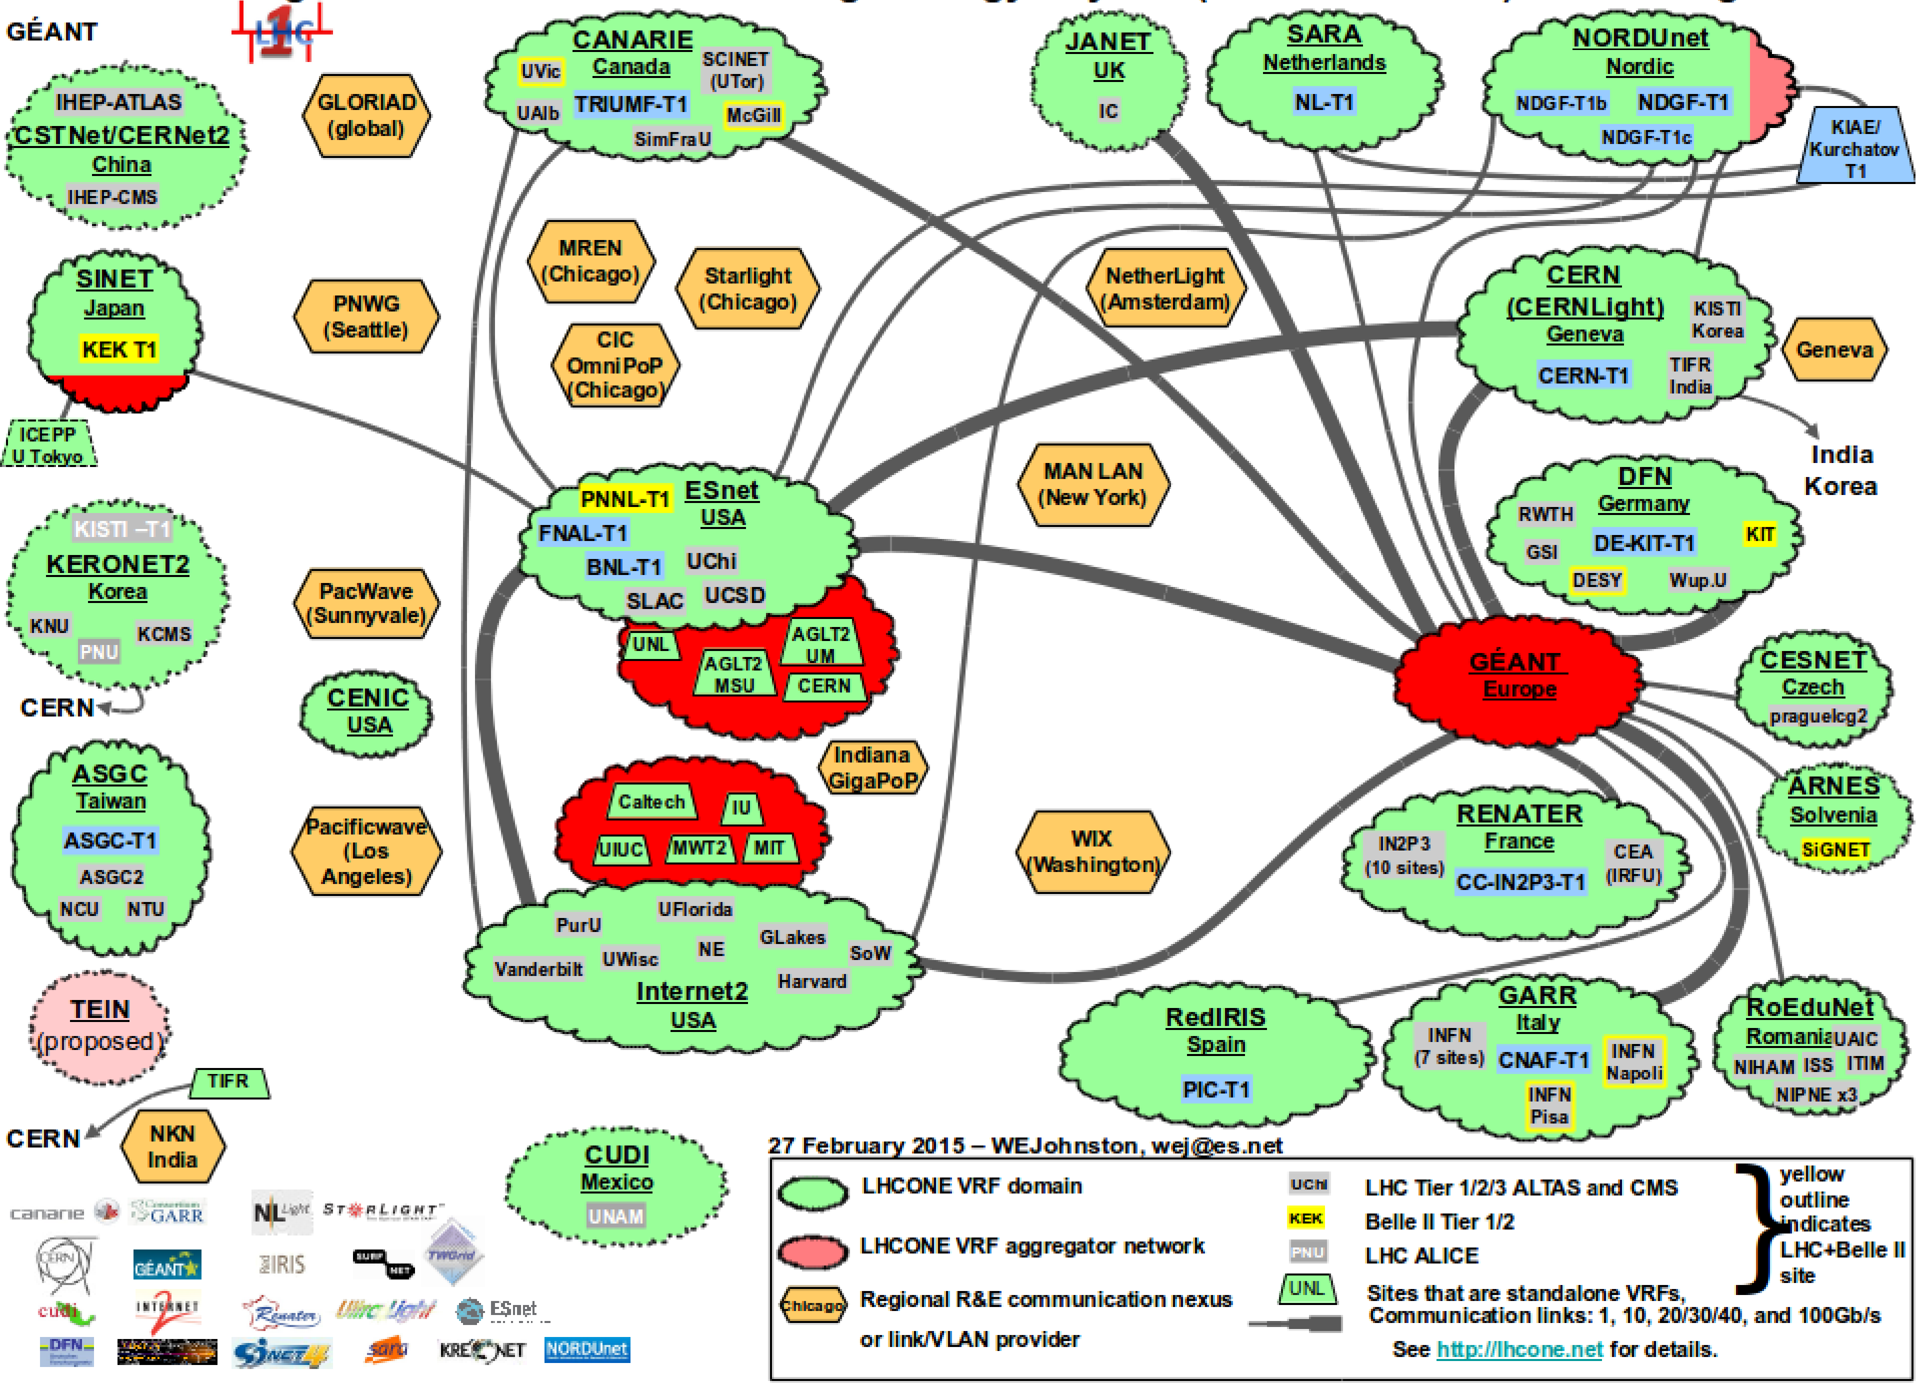

The LHC timeline

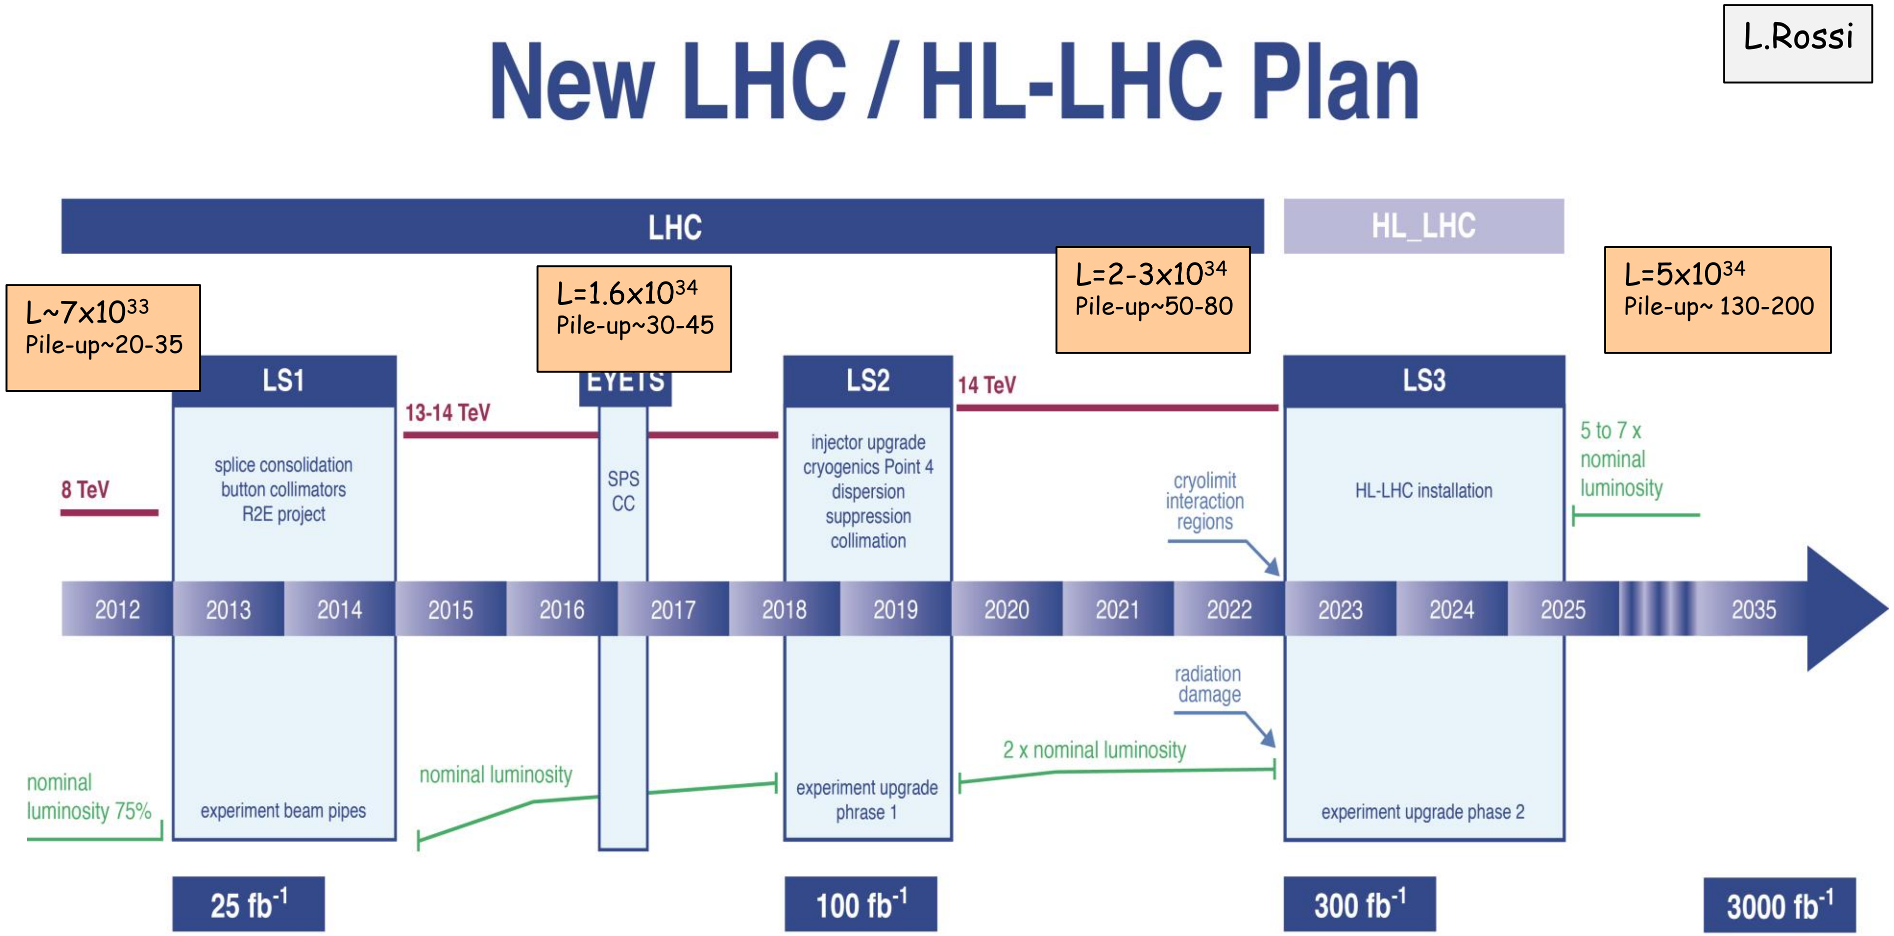

Scale of data tomorrow …

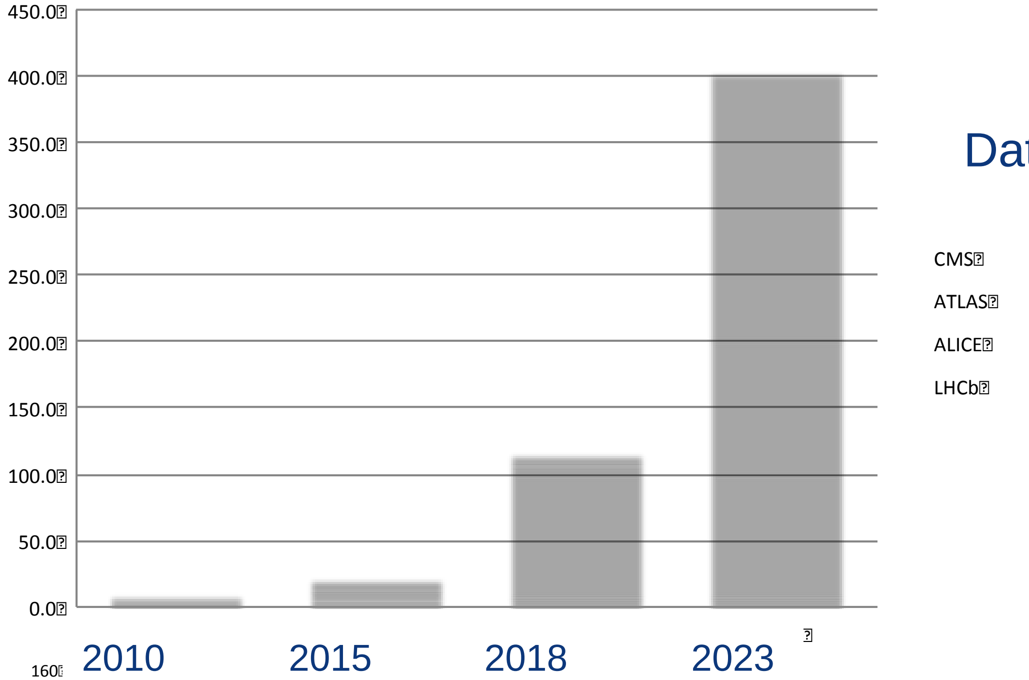

Data: ~25 PB/yr → 400 PB/yr

10 Year Horizon

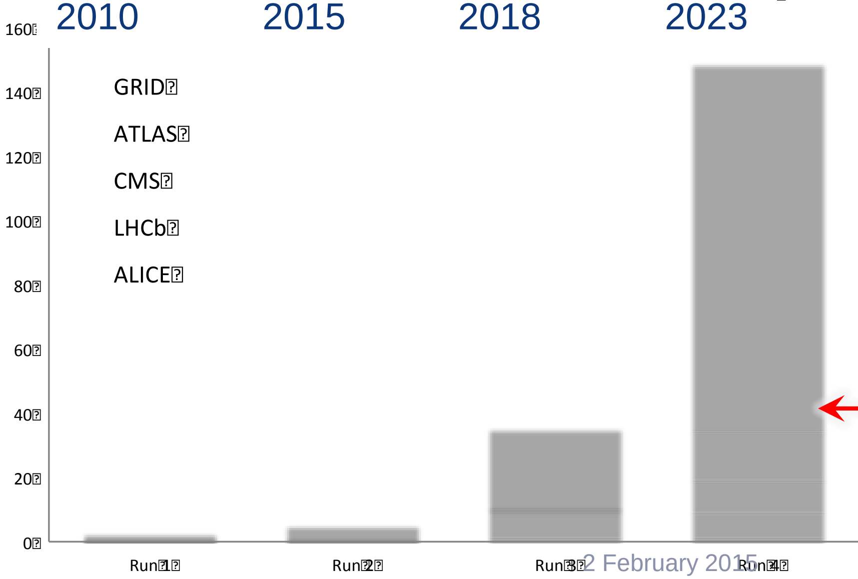

Compute: Growth > x50

What we think is affordable unless we do something differently

Evolution and challenges

- WLCG Grid → *Federated Distributed High Throughput Computing*
  - Easily make use of dedicated and opportunistic resources
    - Grids; clouds, desktop cycles
    - Private, commercial, hybrid,
    - Full time or occasional access
- Challenges:
  - Huge increases in data volumes and processing needs
    - 25 PB/year in 2012 → 400 PB/year in 2024
  - Software complexity and performance
    - Modern CPU architectures require significant software re-engineering
  - Existing computing models will not scale on the 10-year timescale
  - Must live within ~flat budgets

Clouds in LHC

- CERN and many WLCG sites are now using cloud technologies to provision their compute clusters
  - Together with "devops" toolchains to manage the scale we are now at
  - Many are deploying Openstack – global community
- Cloud provisioning
  - Better cluster management and flexibility
  - Can run existing grid services on top – but don't really need to
- LHC experiments also manage HLT farms with Openstack
  - Allows them to switch between DAQ and processing

CERN Cloud

- Remodel IT services on Cloud layered models
  - Manpower, Server utilisation and Processes
  - IaaS, PaaS, SaaS
- Move to commonly used open source tools
  - Focus on strong communities and momentum
  - Stop re-inventing tools elsewhere
- Implement clouds at scale
  - Aim for 90% infrastructure virtualised
  - Exploit ecosystem solutions rather than writing from scratch
  - Request to delivery in a coffee break

Openstack - Federation

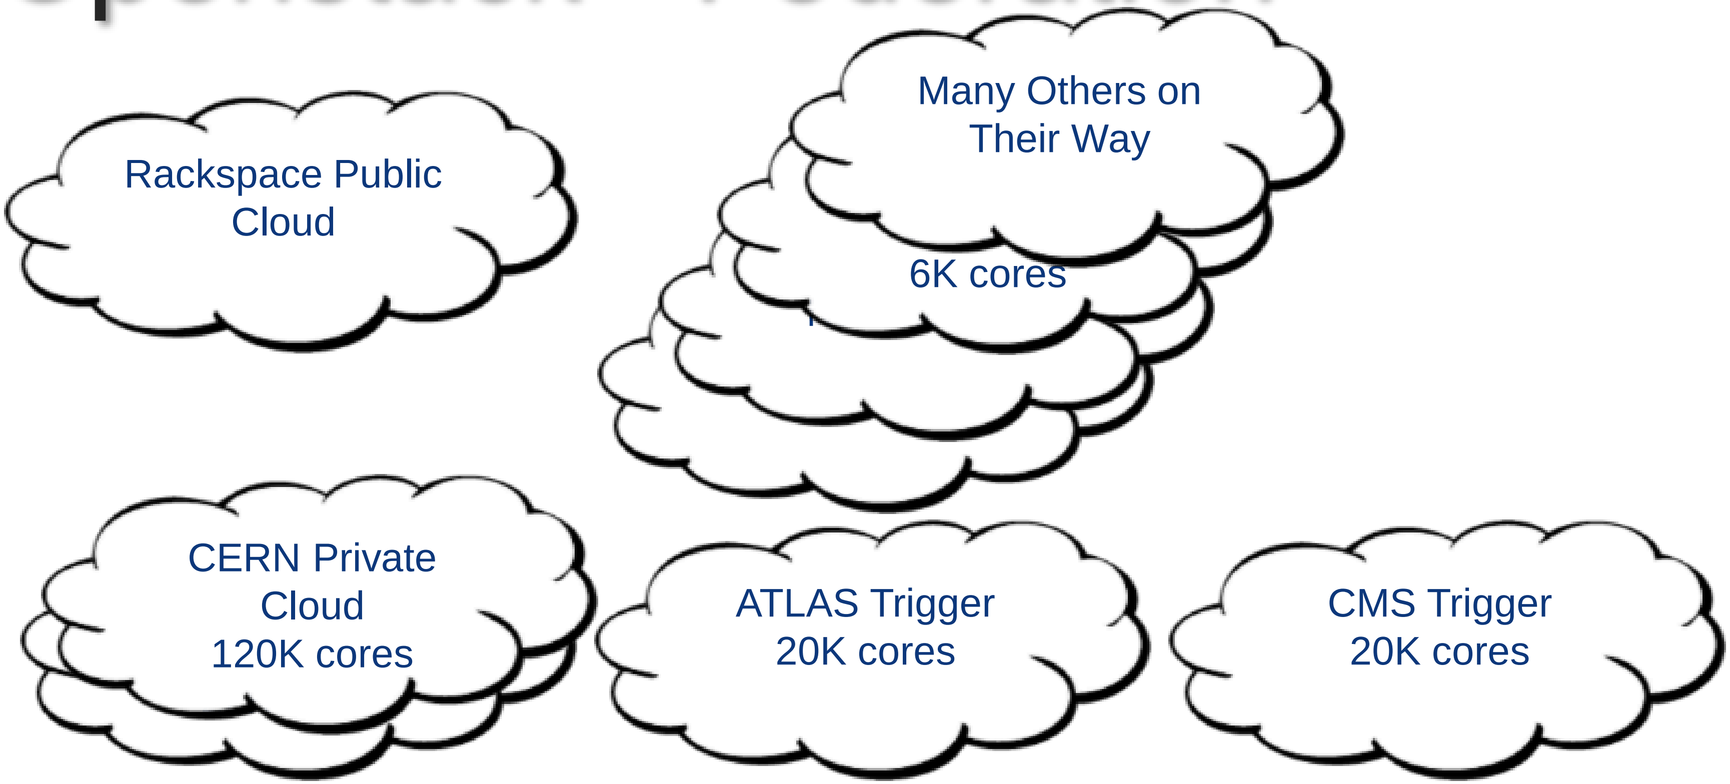

- Share resources, images, accounts between clouds ?
- In collaboration with Rackspace in CERN-openlab
- All contributions are to OpenStack upstream so will appear in all OpenStack clouds at all the sites

CERN openlab in a nutshell

- A science – industry partnership … ion with over a decade of success
- Evaluate state-of-the-art technolo… vironment and improve them
- Test in a research environment to… many business sectors tomorrow
- Train next generation of enginee…
- Disseminate results and outreac…

PARTNERS

HP

HUAWEI

intel

ORACLE

SIEMENS

CONTRIBUTOR

rackspace
the open cloud company

ASSOCIATE

Yandex

YEARS / ANS CERN

Collaboration - Education

- CERN openlab
  - Intel, Huawei, Oracle, Rackspace, Siemens, Yandex

  http://cern.ch/openlab

- CERN School of Computing

  http://cern.ch/csc

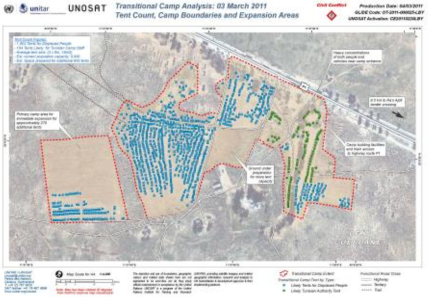

- UNOSAT

  http://cern.ch/unosat

- Citizen Cyber Science Collaboration
  - Involving the General Public

Connectivity (100 Gbps)

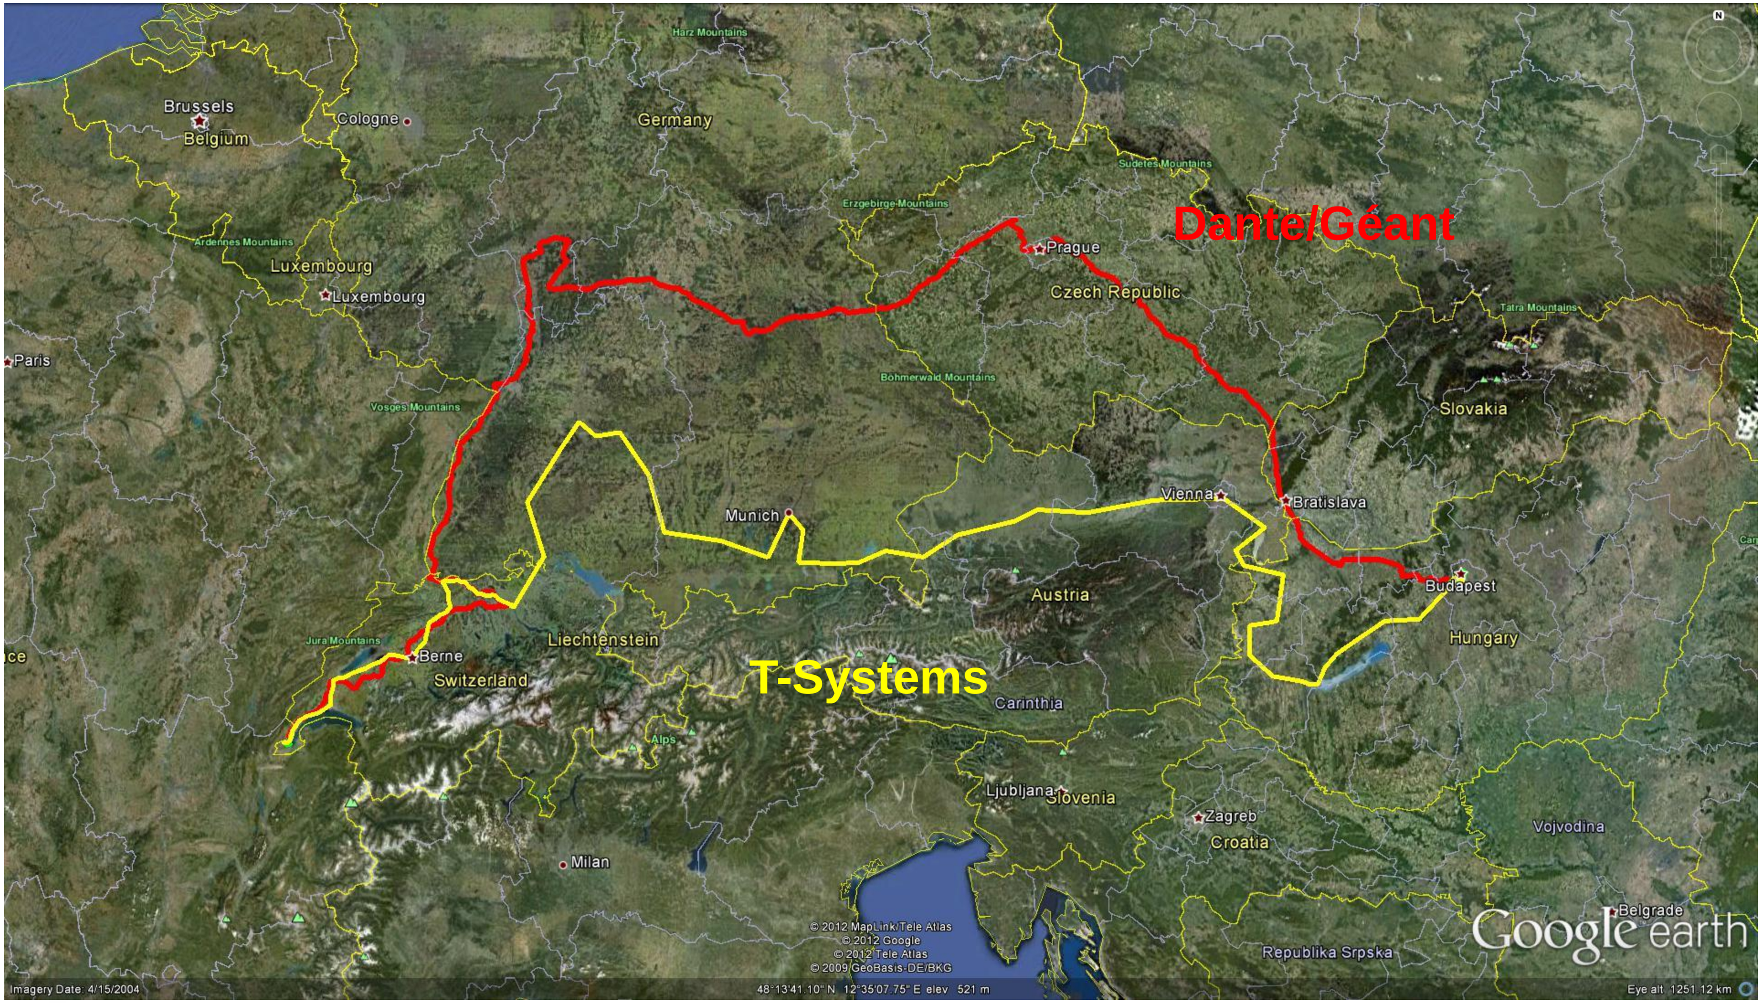

Tier 0

Overview: Data Centre

a day ago to a few seconds ago

MEYRIN DATA CENTRE

	last_value
Number of Cores in Meyrin	109,191
Number of Drives in Meyrin	71,670
Number of Memory Modules in Meyrin	74,148
Number of 10G NIC in Meyrin	4,436
Number of 1G NIC in Meyrin	20,955
Number of Processors in Meyrin	20,068
Number of Servers in Meyrin	10,931
Total Disk Space in Meyrin (TB)	123,082
Total Memory Capacity in Meyrin (TB)	415

WIGNER DATA CENTRE

	last_value
Number of Cores in Wigner	43,280
Number of Drives in Wigner	23,173
Number of Memory Modules in Wigner	21,614
Number of 10G NIC in Wigner	1,399
Numer of 1G NIC in Wigner	5,062
Number of Processors in Wigner	5,412
Number of Servers in Wigner	2,709
Total Disk Space in Wigner (TB)	71,725
Total Memory Capacity in Wigner (TB)	172

NETWORK AND STORAGE

	last_value
Tape Drives	104
Tape Cartridges	24,914
Data Volume on Tape (TB)	101,834
Free Space on Tape (TB)	27,633
Routers (GPN)	134
Routers (TN)	27
Routers (Others)	96
Star Points	638
Switches	3,549

BATCH JOBS

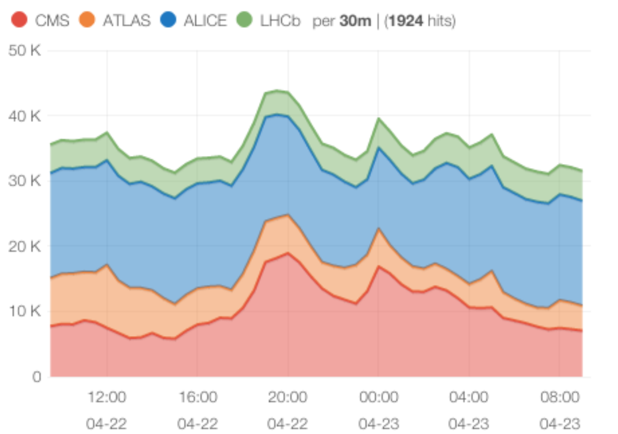

EOS ACTIVE DATA TRANSFERS

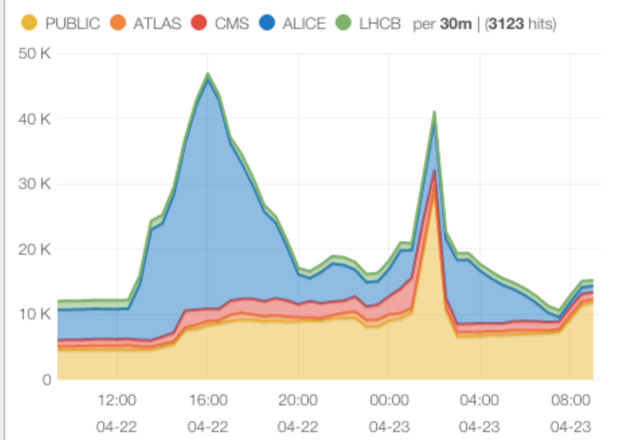

VM CREATED

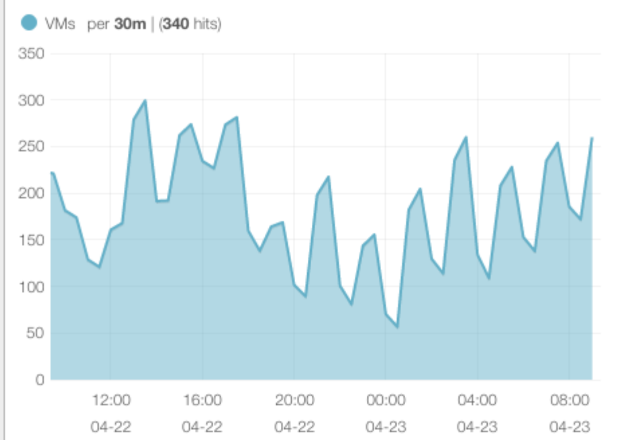

Commodity Internet Consortium

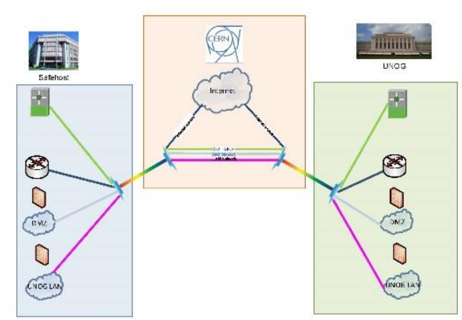

The CERN GVA Data Centre in Numbers

- 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	9808
Processors	18.018
Cores	93,937
HEPSpec06	744,277

Disks	65,617
Raw disk capacity (TiB)	99,329
Memory modules	64035
Memory capacity (TiB)	344
RAID controllers	3,091

Tape Drives	141
Tape Cartridges	50,623
Tape slots	66000
Data on Tape (PiB)	99
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	2.3 MW
Total Power Consumption	2.8 MW

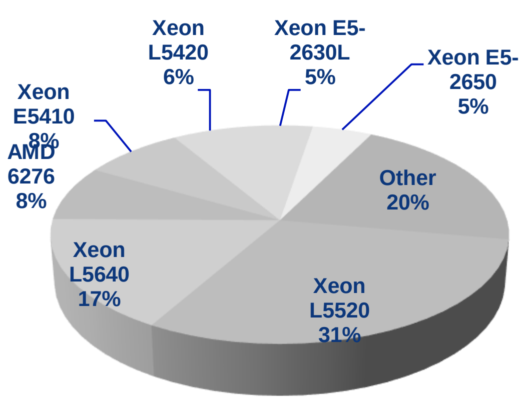

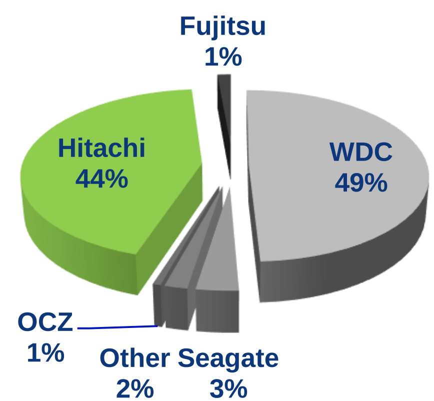

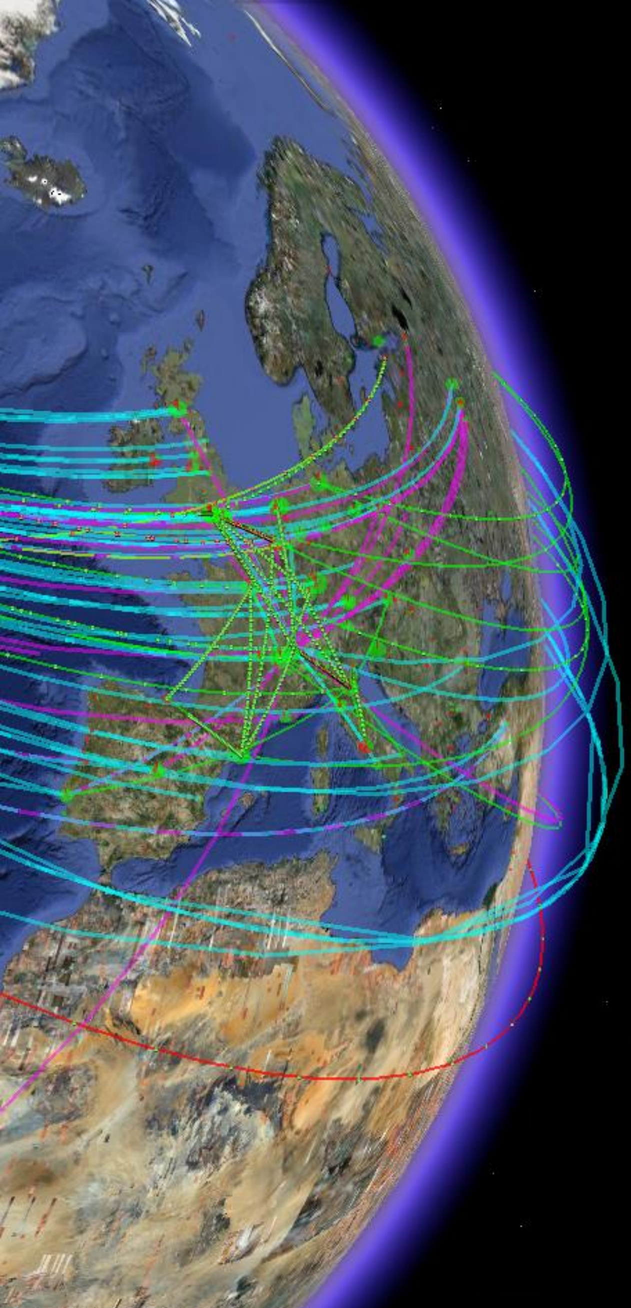

Conclusions

- First years of LHC data – WLCG has helped deliver physics rapidly
- Just the start of a >20 year exploration of new physics
- Entering a phase of consolidation and evolution
- Challenges for computing – scale & complexity – will continue to increase dramatically