Computing at CERN

Tim Smith CERN/IT

Our Universe is Yours
Notre Univers est le vôtre

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

OED | Oxford English Dictionary
The definitive record of the English language

Quick search: Find word in dictionary GO

Lost for Words? | Advanced search | Help

computing, *n.*

Pronunciation:

Brit. /kəmˈpjuːtɪŋ/ , U.S. /kəmˈpjudɪŋ/

1. The action or an instance of calculating or counting; = computation *n.* 1a.

2. The action or practice of using computers, esp. as a professional or expert; the activity or operation of an electronic computer; (also) = computer science *n.*

An Early “Computer”

- Wim Klein

- Calculating the 73rd root of a 500 digit number took less than 3 minutes…

Computing

- Scientific
  - Compute, Storage, Network
- Technical
  - Design, Operation
- Administrative
  - HR, Finance, Projects
- Desktop
- Collaborative

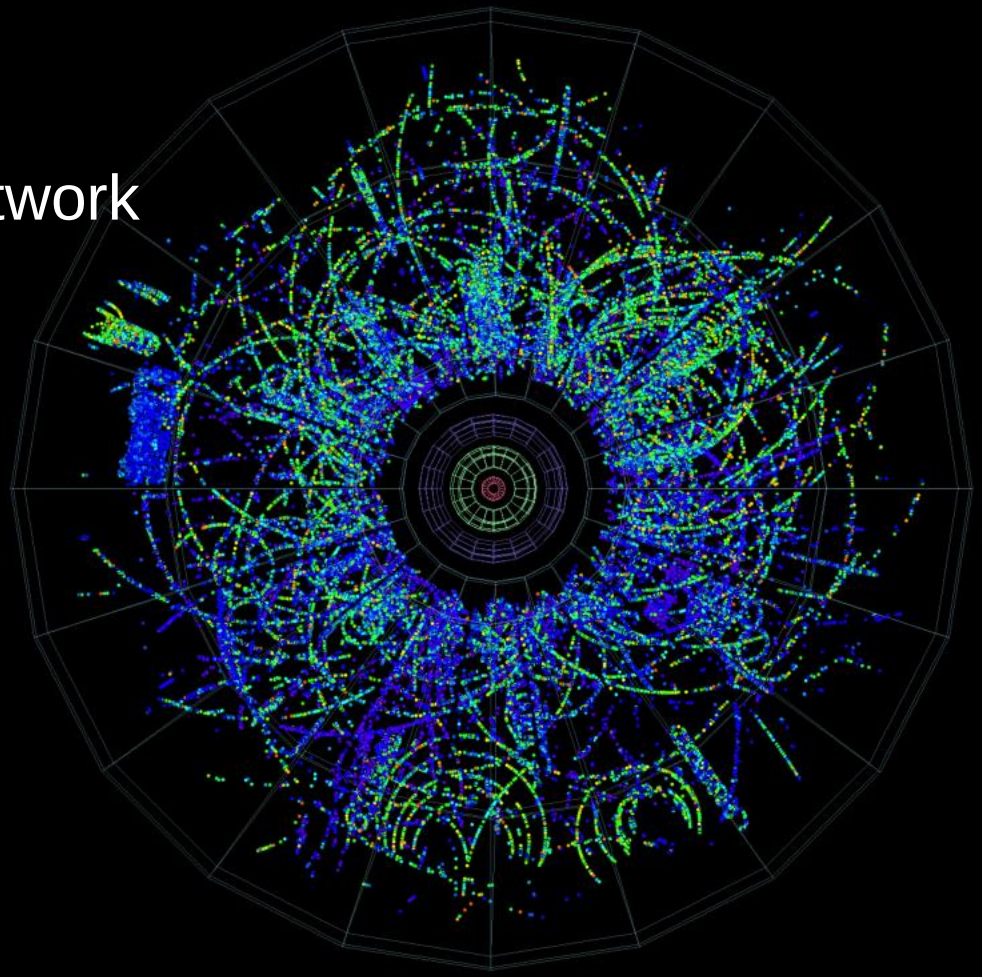

Bubble Chamber

- BEBC 1973-1984
- 6.3 million photos
- 3000 km of film

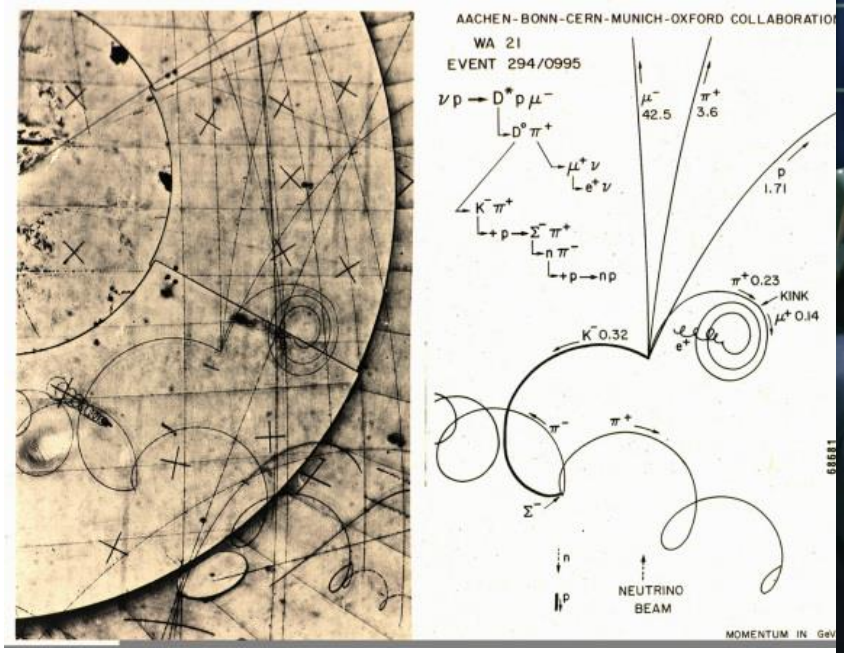

Spark Chambers

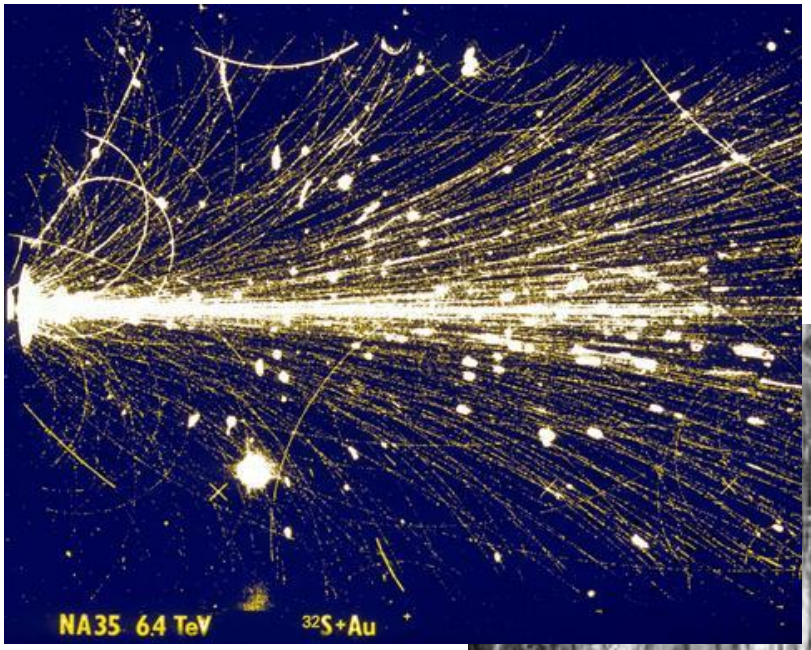

Magnetostrictive readout

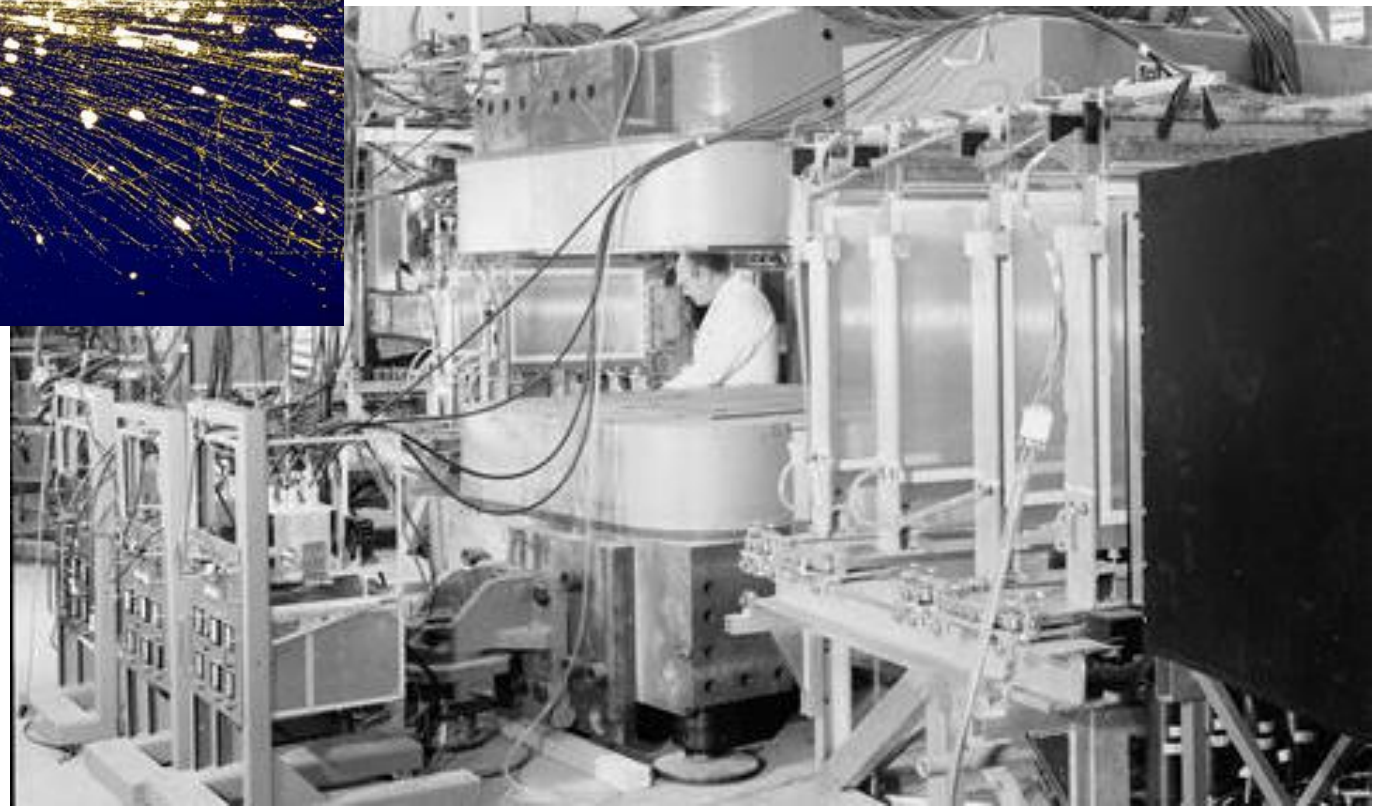

Momentous Events

1960: Vacuum tubes

1965: CDC 6600: #3

Mainframes

New Computer Centre for a New Computer

1983: 2 CDCs and an IBM 3081

Super Computers

1988: Cray XMP

RISC Workstations

Comodity Computing

Farming in a Data Centre

Computing

- Networking

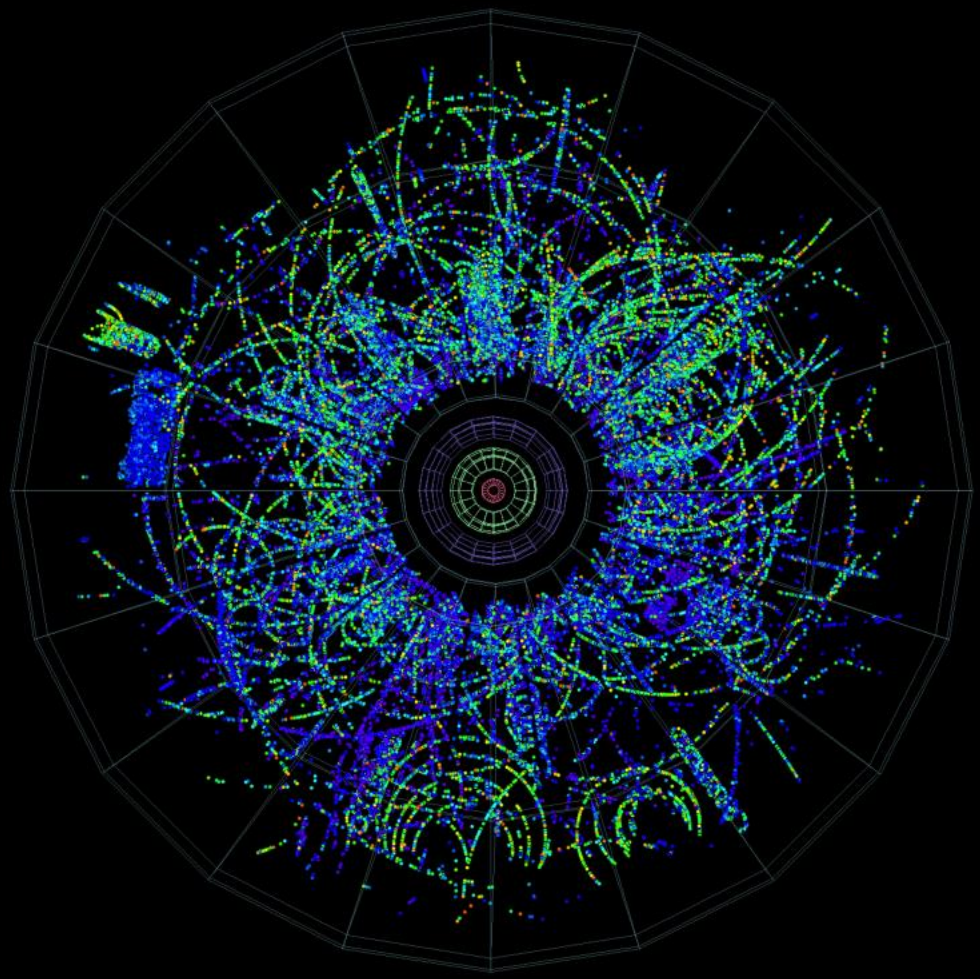

Networking

- Packet-switched network
- 1969 ARPANET
  - US DoD sponsored research at US Universities
  - Aim: communications network to survive a nuclear attack
  - Find next best route if one node obliterated
- 70s and 80s proliferation
  - US: NASA Science Net, CSnet, Energy Sciences Net, NSFnet
  - FR: CYCLADES
  - UK: Mark I, SERCnet
  - Commercial: Tymnet, CompuServ, BITnet, DECnet
  - Protocols: NCP, X.25 (1976), TCP/IP (1982)
  - CERnet

Networking

- Science without borders
  - Data exchange across the iron curtain
  - 1988 first data connection between China and scientific world – IHEP to CERN

- Truly international Internet
  - 1989 first external TCP/IP connection
  - 1990 principle link US-EU from CERN
    - (1.5Mb/s)
  - 1991 80% of the internet capacity installed in Europe for international traffic was terminated at CERN

CERN Internet Exchange Point

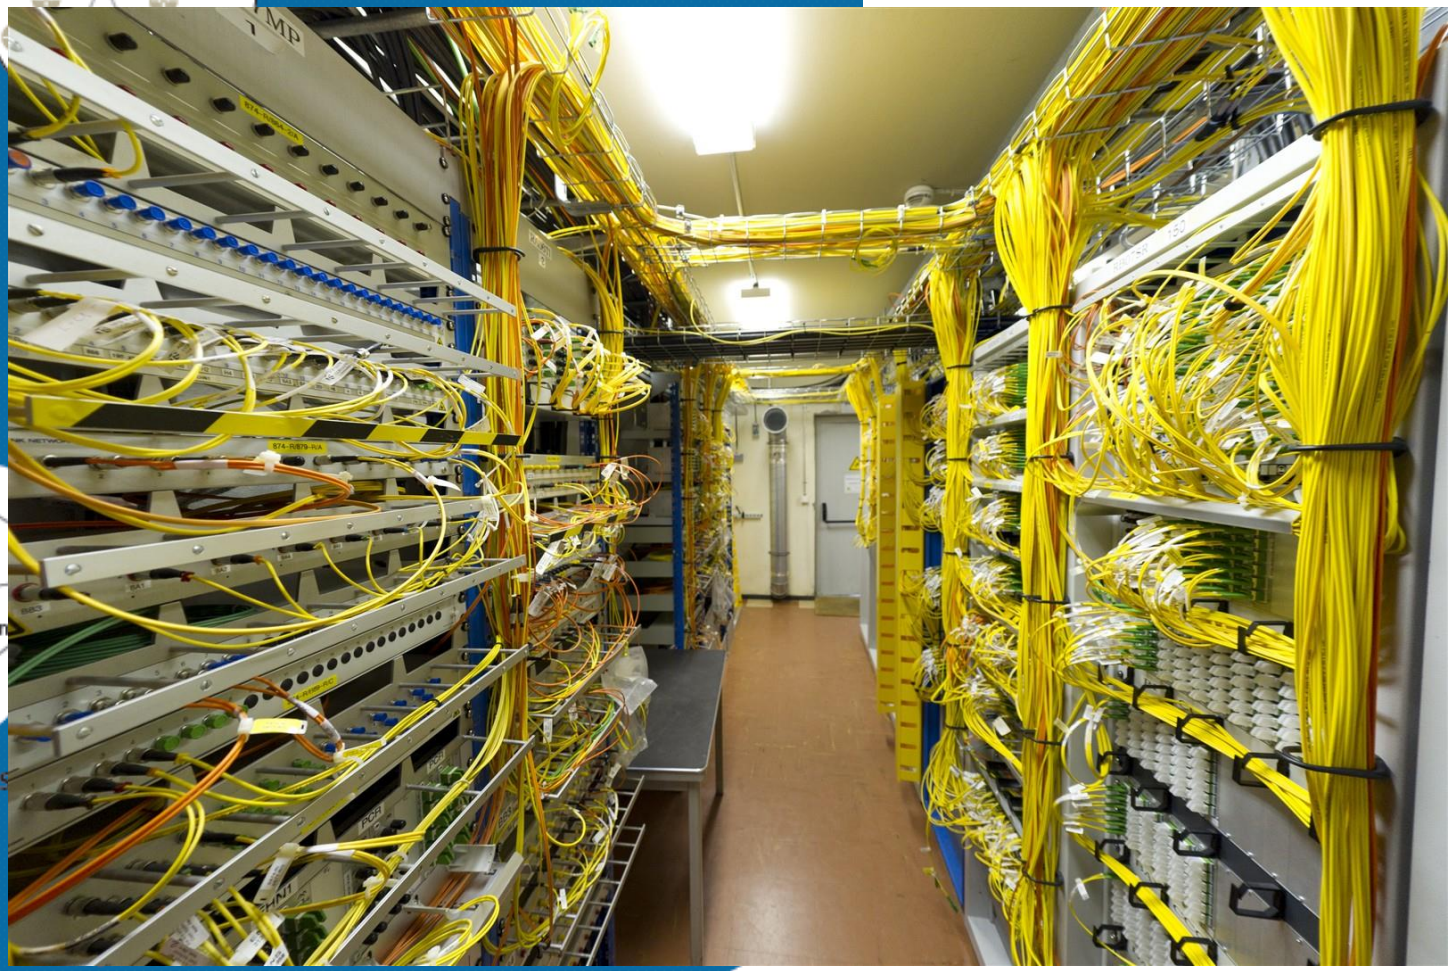

Computing

- Information Management

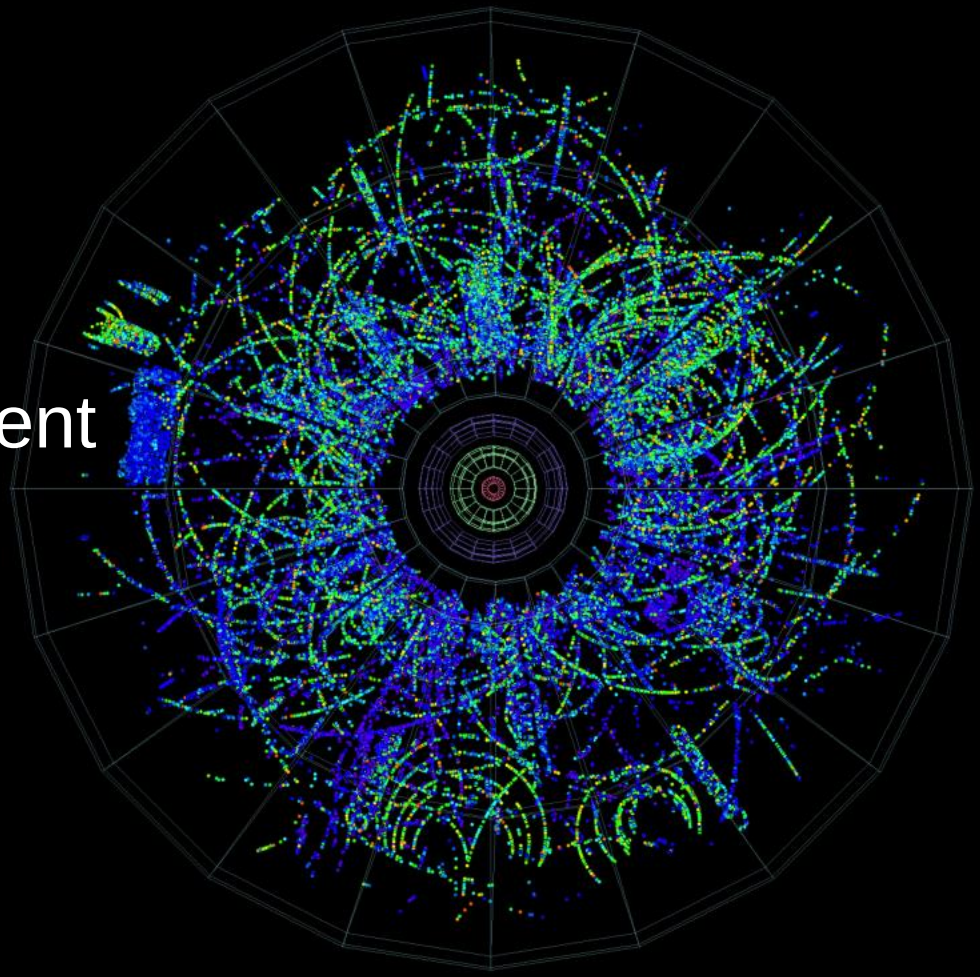

Information Management - *circa 1989*

- Keep track of LHC project and CERN?
  - Researchers turnover ~2 years
- Information about CERN and its experiments
  - Not hierarchical, or centrally controlled
    - A multiply connected web
  - Experts store locally, update independently
  - Community is distributed: remote access
- System to link it all together
- CERN is a model in miniature of the rest of world in a few years time

Vague but Exciting …

Vague but exciting ...

CERN DD/OC

Tim Berners-Lee, CERN/DD

Information Management: A Proposal

March 1989

Information Management: A Proposal

Abstract

This proposal concerns the management of general information about accelerators and experiments at CERN. It discusses the problems of loss of information about complex evolving systems and derives a solution based on a distributed hypertext system.

Keywords: Hypertext, Computer conferencing, Document retrieval, Information management, Project control

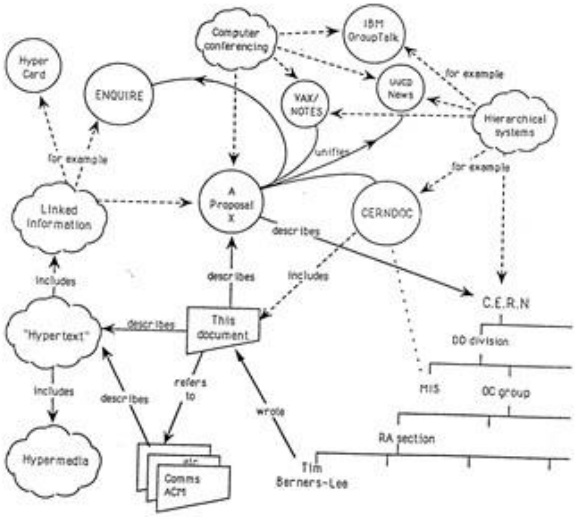

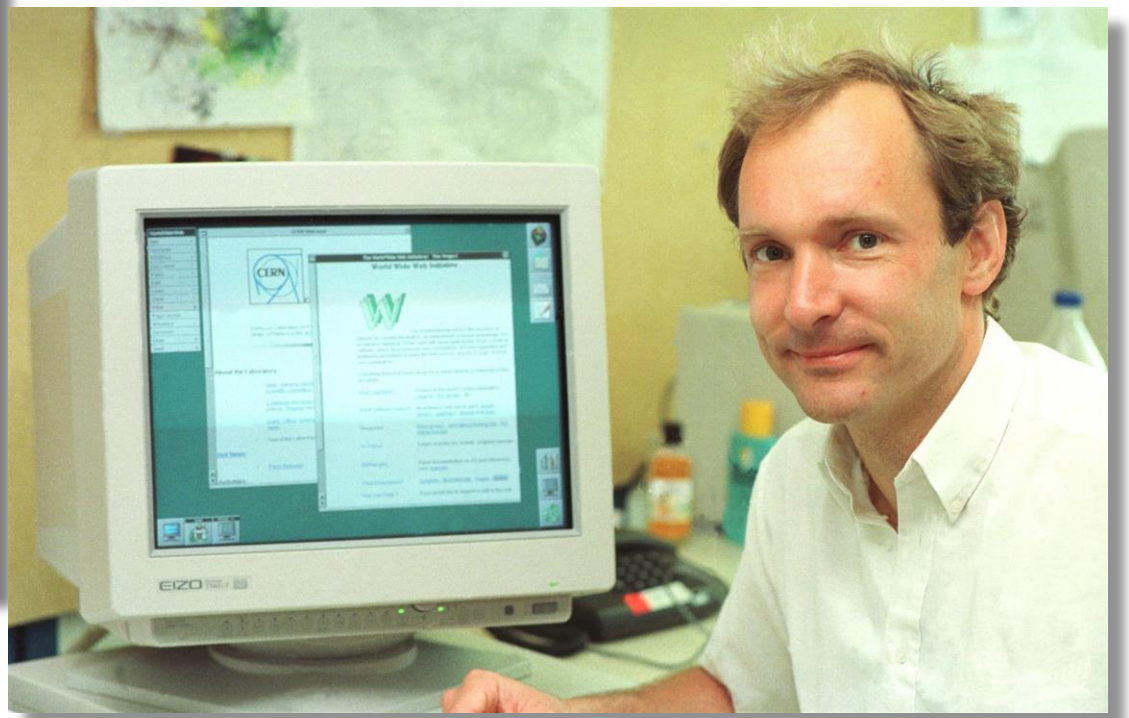

Growth of the Web

- Aug 1991 went public
  - Tim posted project to alt.hypertext and other internet groups

- Dec 1991 First web s
- 1992 rapid expansio
  - Universities and rese

Growth of the Web

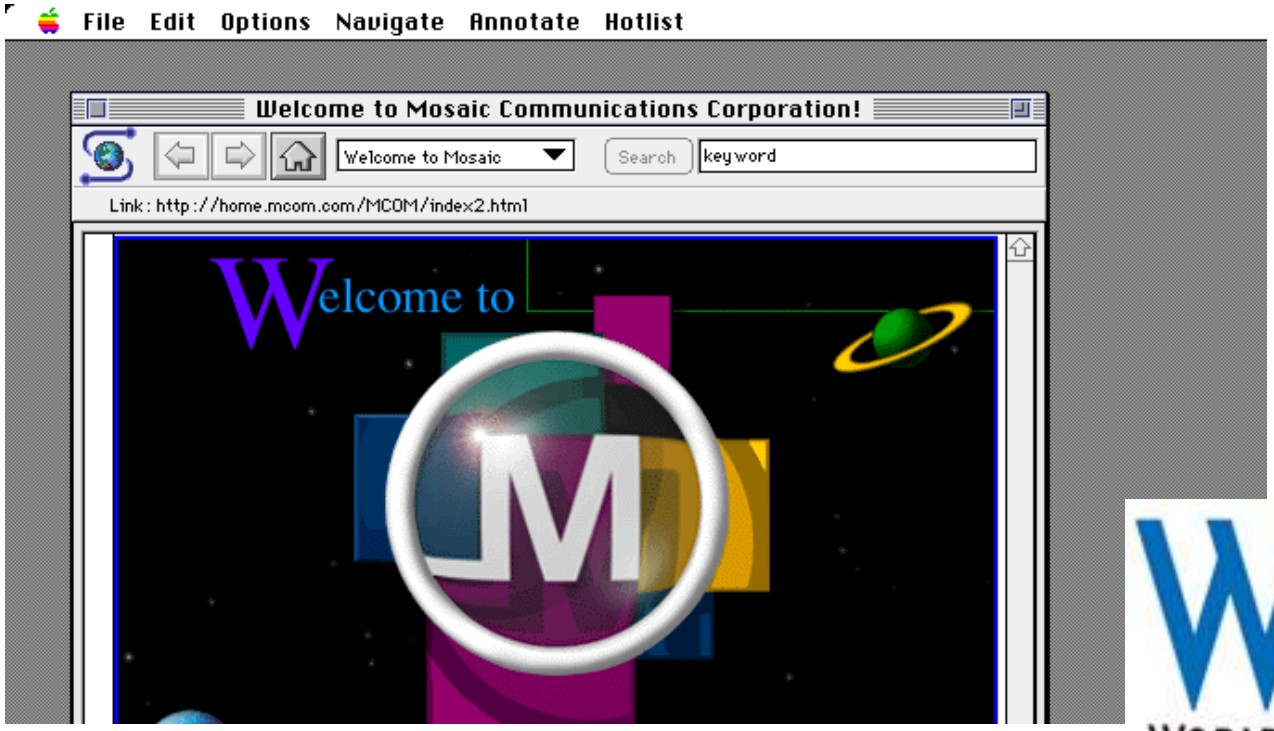

- 1993 rapid expansion across the world
  - National Center for Supercomputing Applications (NCSA) at the University of Illinois released its Mosaic browser

Computing

- The LHC Era

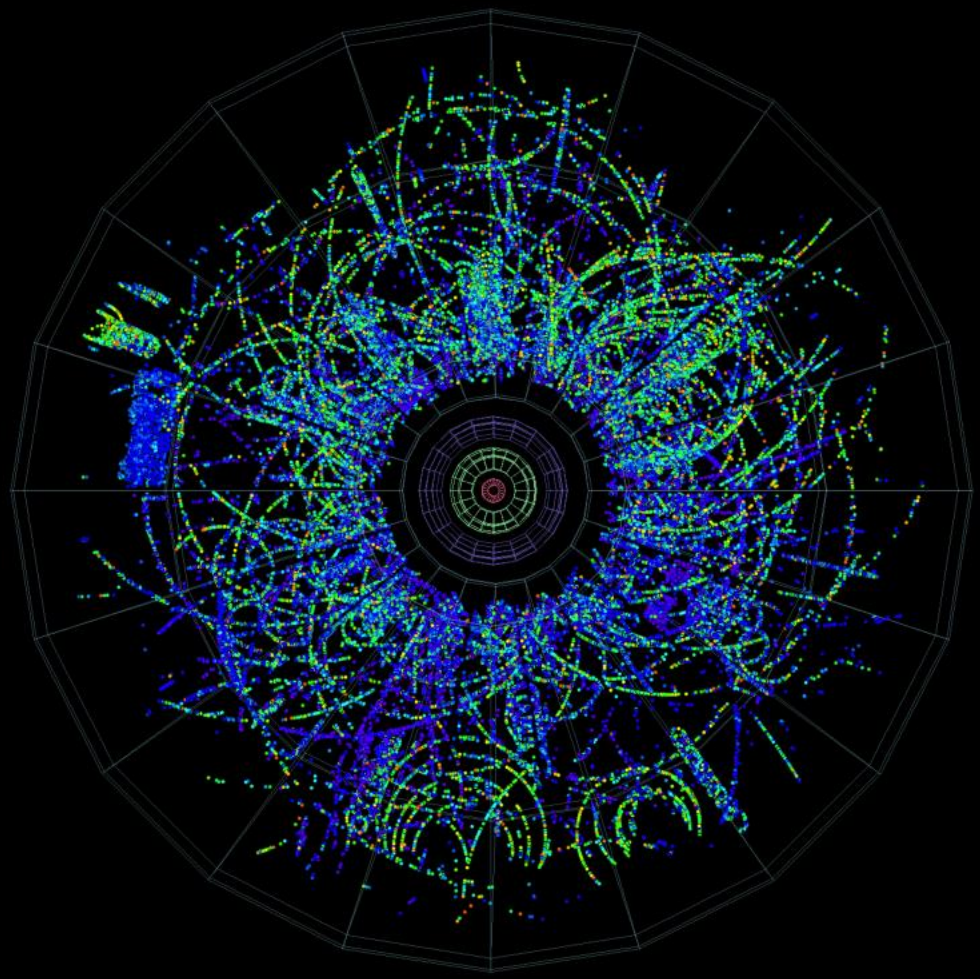

The LHC Data Challenge

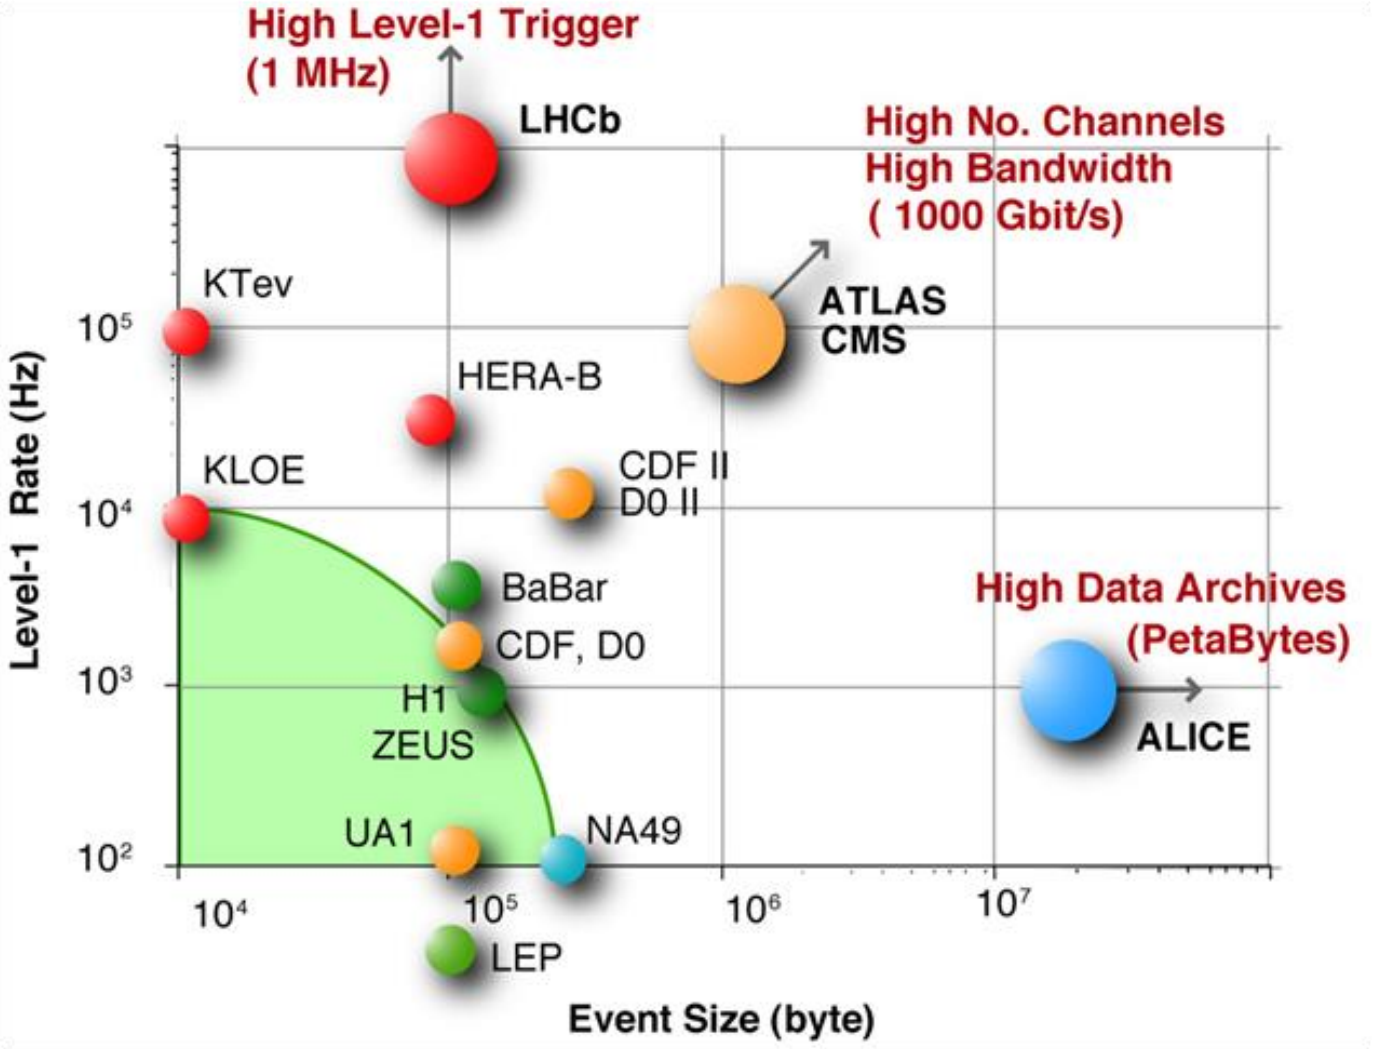

Big Data !

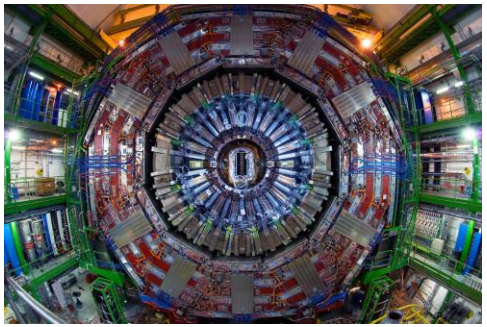

150 million sensors
Generating data 40 million times per second

➔ **Peta Bytes / sec !**

Trigger
Select 100,000 per second

➔ **Tera Bytes / sec !**

HLT / Filter
Select 100 per second

➔ **Giga Bytes / sec !**

Primary Storage

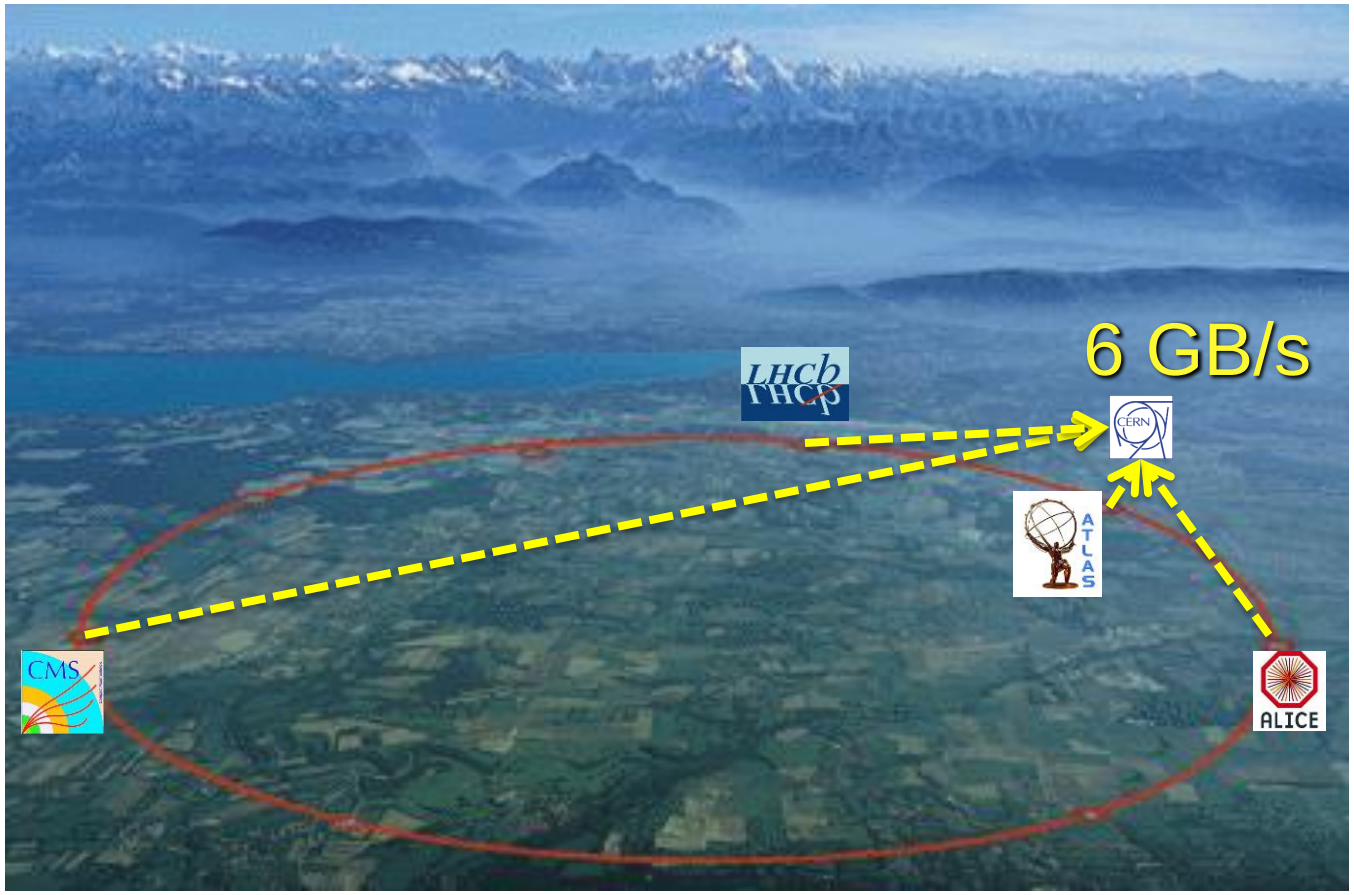

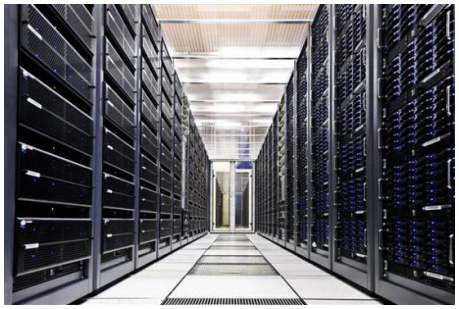

80,000 Disks
88,000 CPU Cores

18,000 1GB NICs
2,700 10GB NICs

The LHC Data Challenge

- Few places can store it
- Processing needs 5x CERN

- HEP community distributed
  - Local funding for computing
- Distributed solution…

- Models of Networked Analysis at Regional Centres

x 2 locations @ CERN

Resource Management - *circa* 1999

- Keep track of LHC project and CERN?
  - Researchers turnover ~2 years
  - Computing power and Data Storage capacity
- Resources of CERN and its experiments
  - Not hierarchical, or centrally controlled
    - A multiply connected web
  - Experts install locally, update independently
  - Community is distributed: remote access
- System to link it all together
- CERN is a model in miniature of the rest of world in a few years time

Solution: the Grid

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

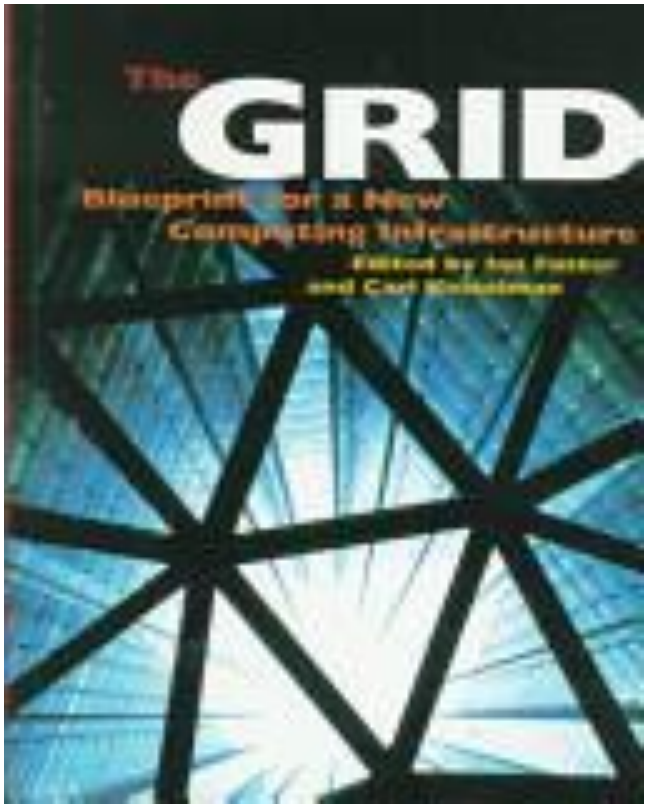

Worldwide LHC Computing Grid

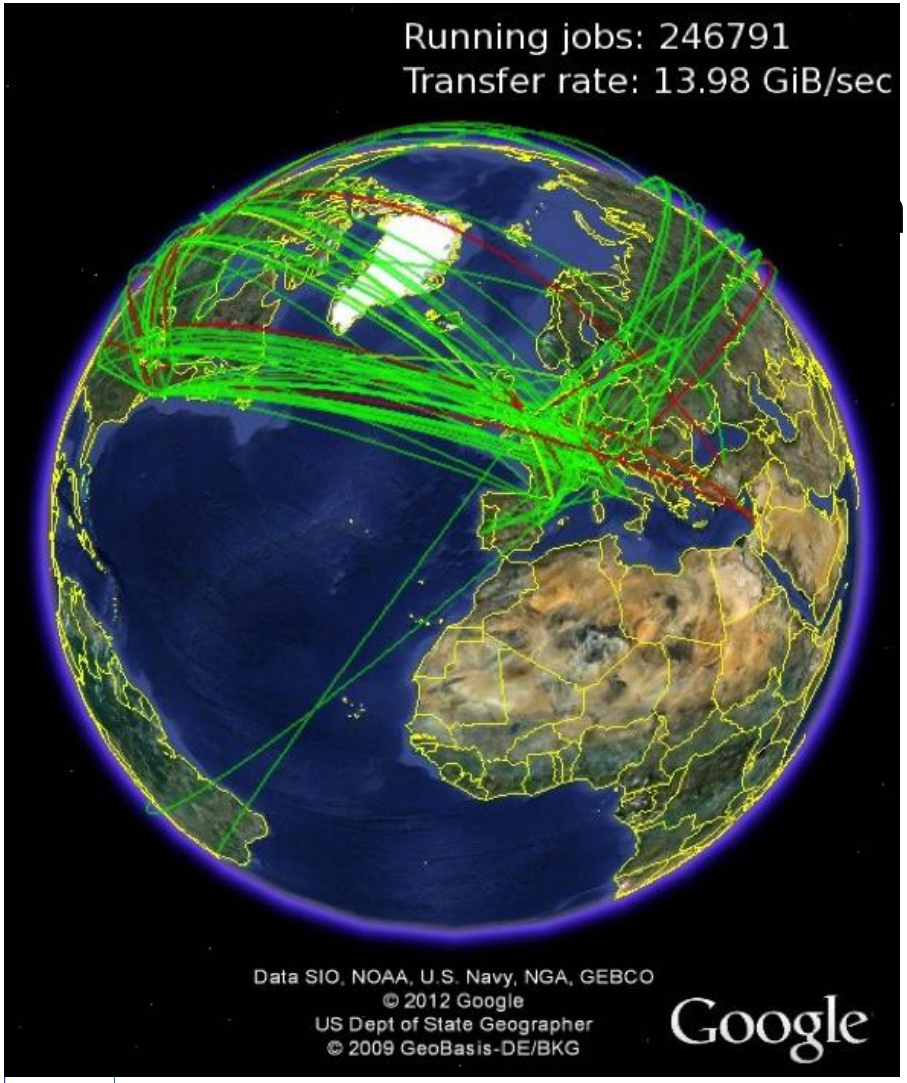

Tier-2s
140 Sites

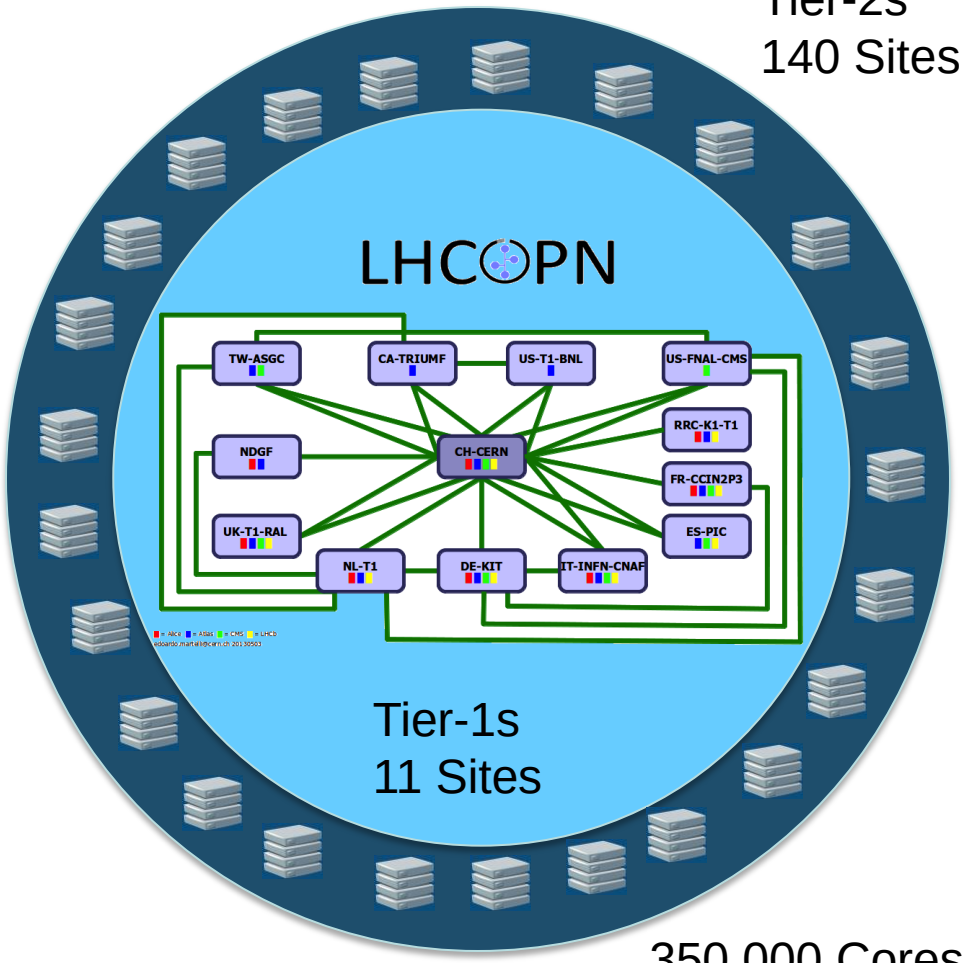

350,000 Cores

What is Grid Middleware?

- The glue that creates the illusion that a distributed infrastructure is a single resource
  - without enforcing uniformity
  - without central control
- Hide and manage heterogeneity
- Facilitate communication between users and providers

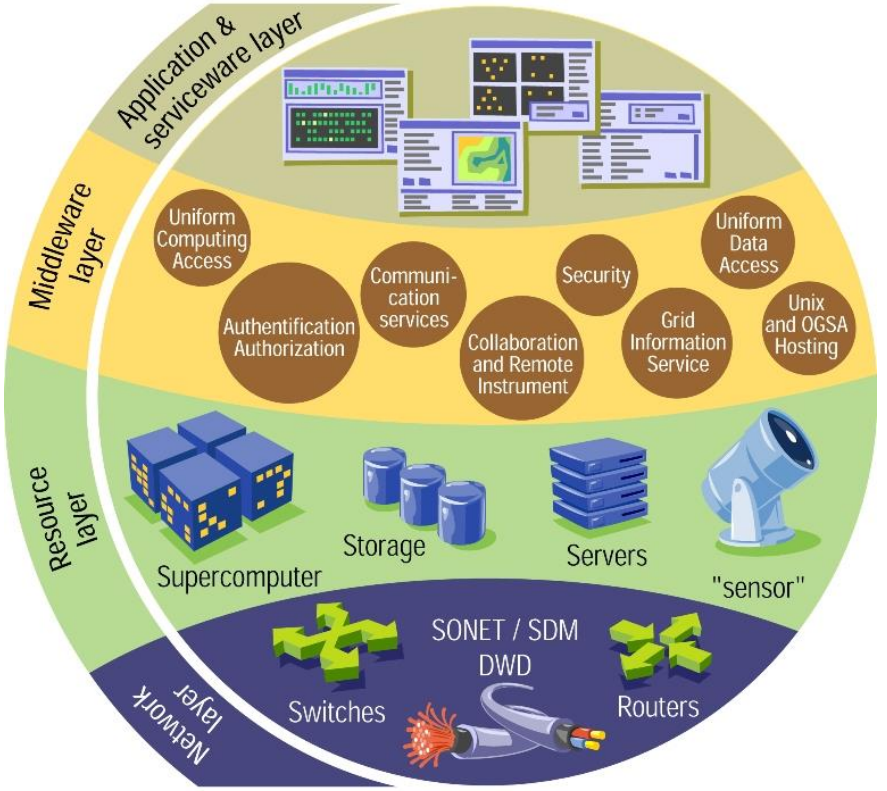

The Grid that Never Sleeps

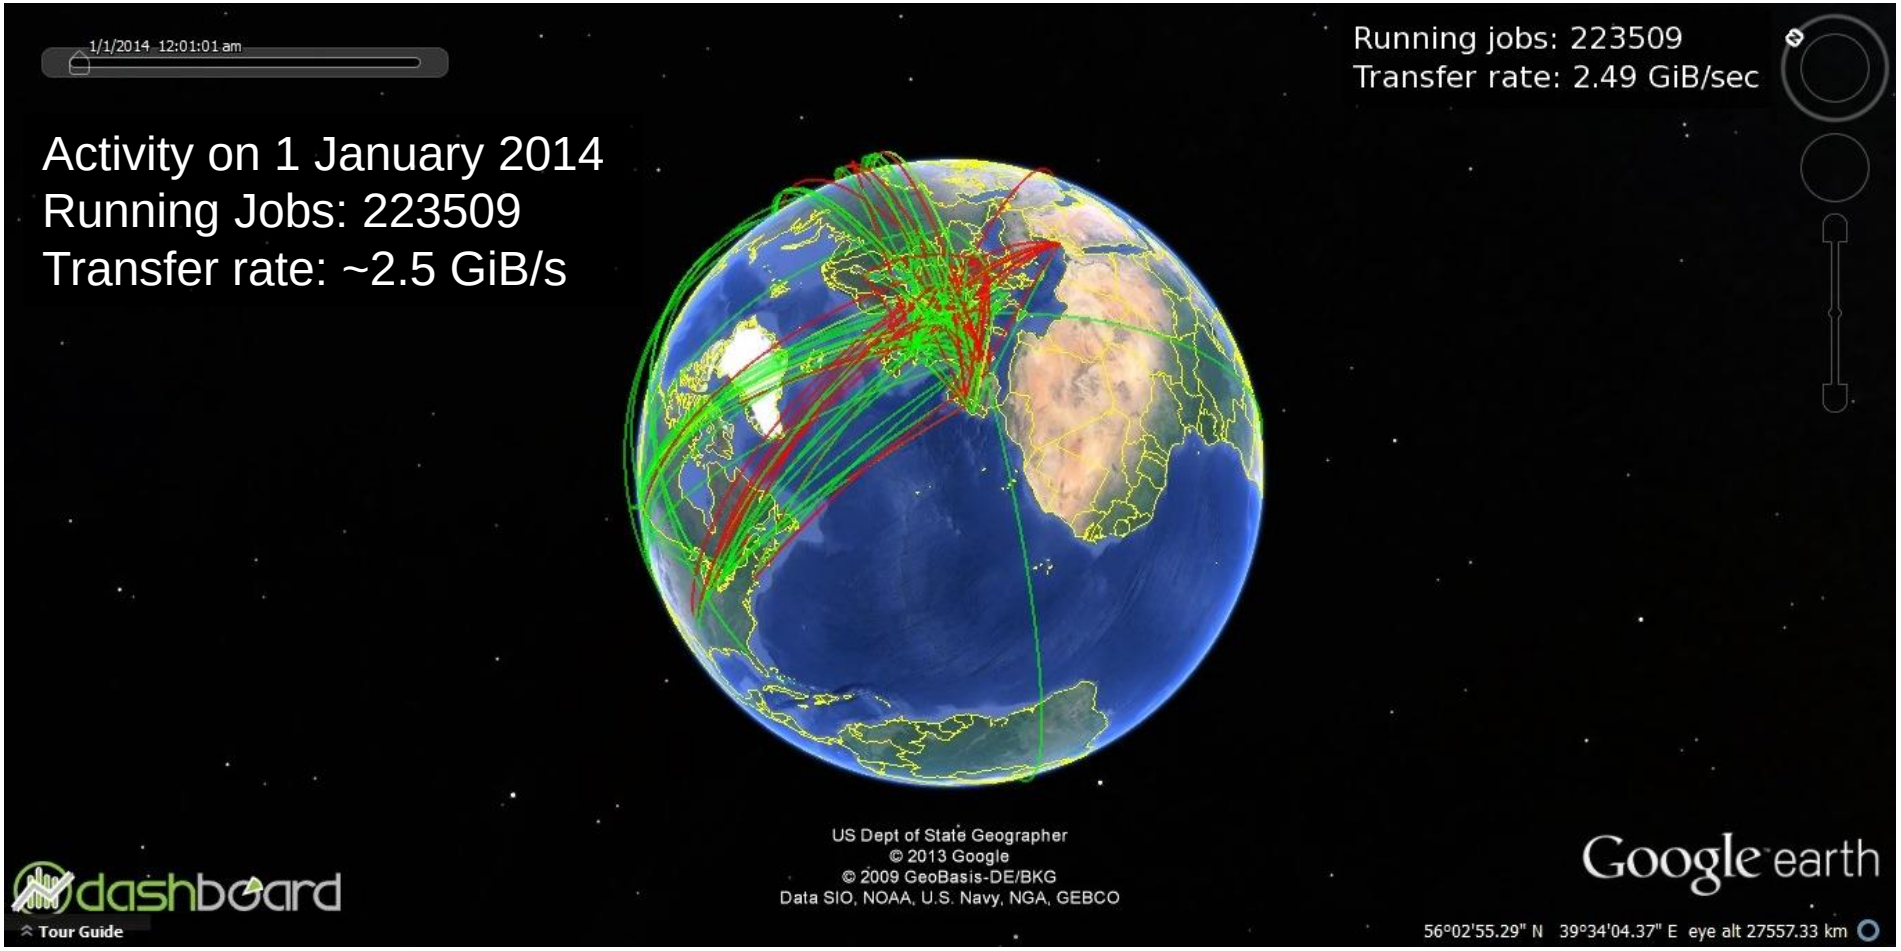

Data Transferred to Tape

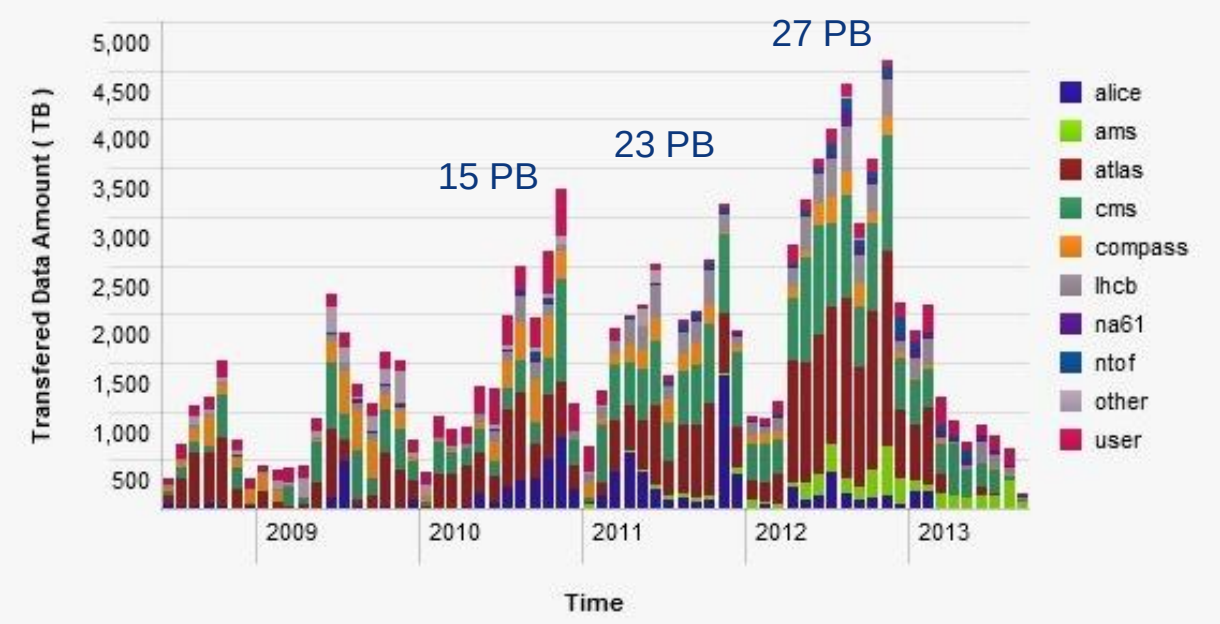

Expected 15TB/y

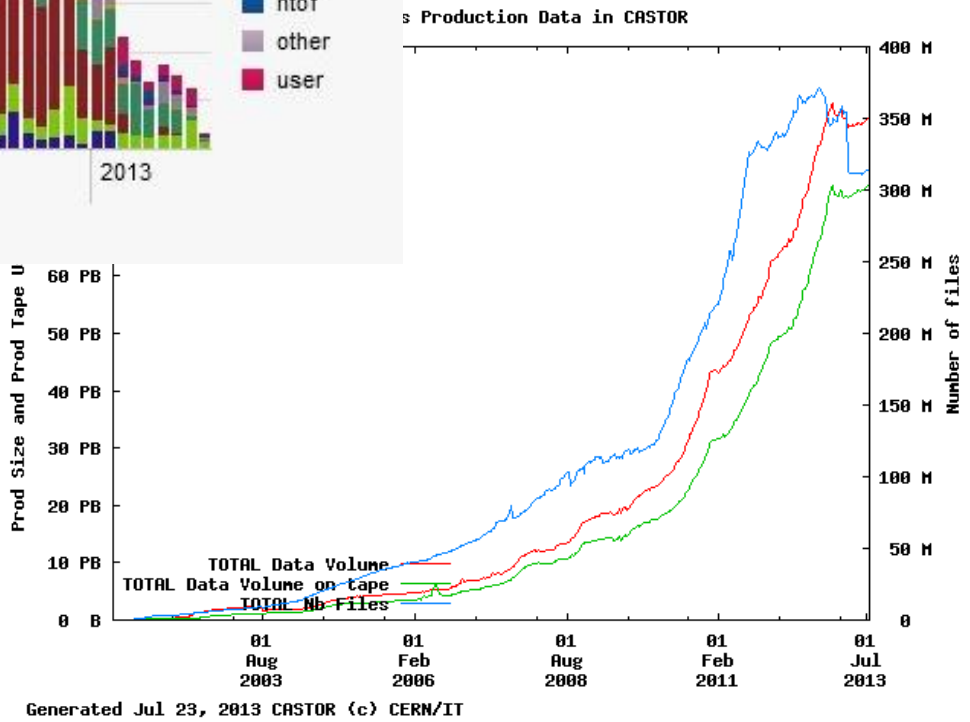

100 PB on tape !

Data Reduction / Analysis

- Publication
- Reduced
- Reconstructed
- Raw

File Size

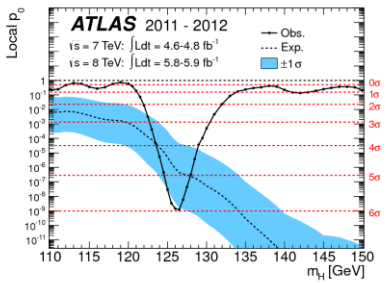

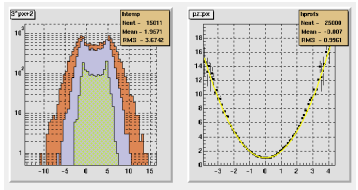

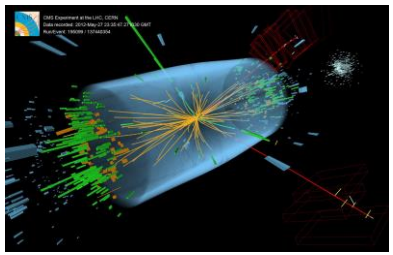

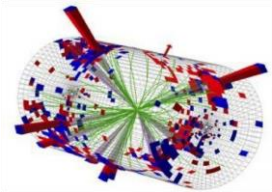

Files

Researchers
T2s, T1s

Analysis Coordinators
T1s

Production Managers
T0, T1s

Data Inflation

- Static: Storing 100 PB was a good challenge ☺
- Dynamic: Analysing it means transformation, reduction, transport, replication, regeneration

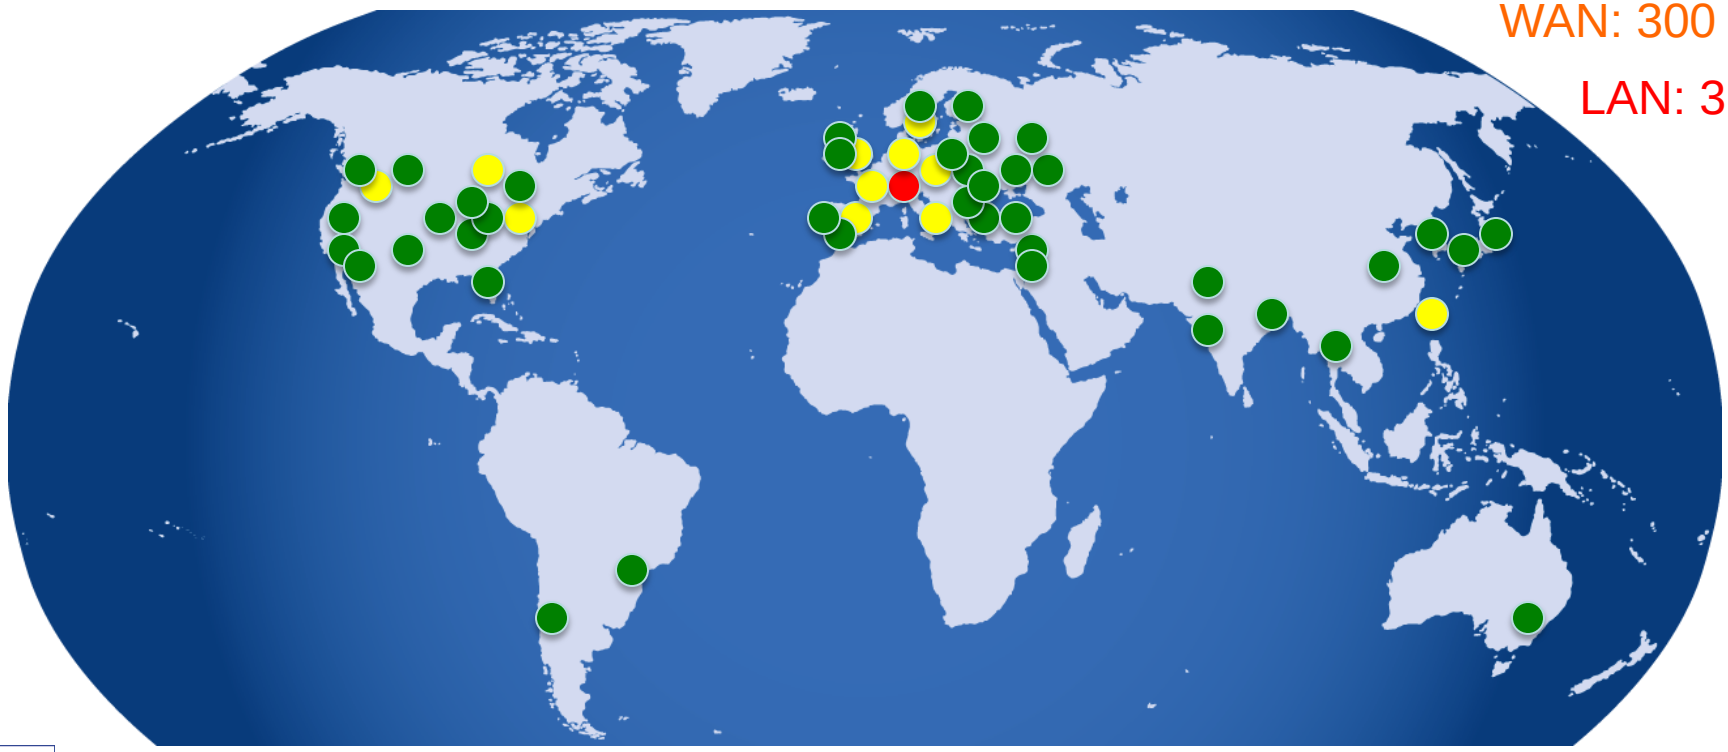

WAN: 300 PB

LAN: 3 EB

Managing a 100 PB Store

- Media Verification
  - Hot / Cold Data
  - Catching and correcting errors while you still can
  - 10% of production drive capacity for 2.6 years

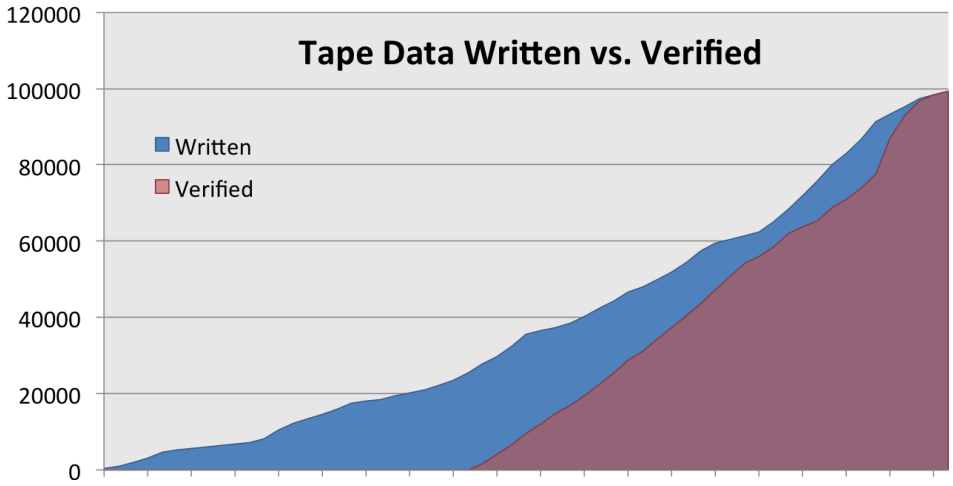

  - (0.000065% data loss)

Managing a 100 PB Store

- Media Migration
  - Drive and Media obsolescence
  - 50% of current drive capacity for 2 years

Shared Infrastructures

- >270 VOs from several scientific domains
  - Astronomy & Astrophysics
  - Civil Protection
  - Computational Chemistry
  - Comp. Fluid Dynamics
  - Computer Science/Tools
  - Condensed Matter Physics
  - Earth Sciences
  - Fusion
  - High Energy Physics
  - Life Sciences
  - .....

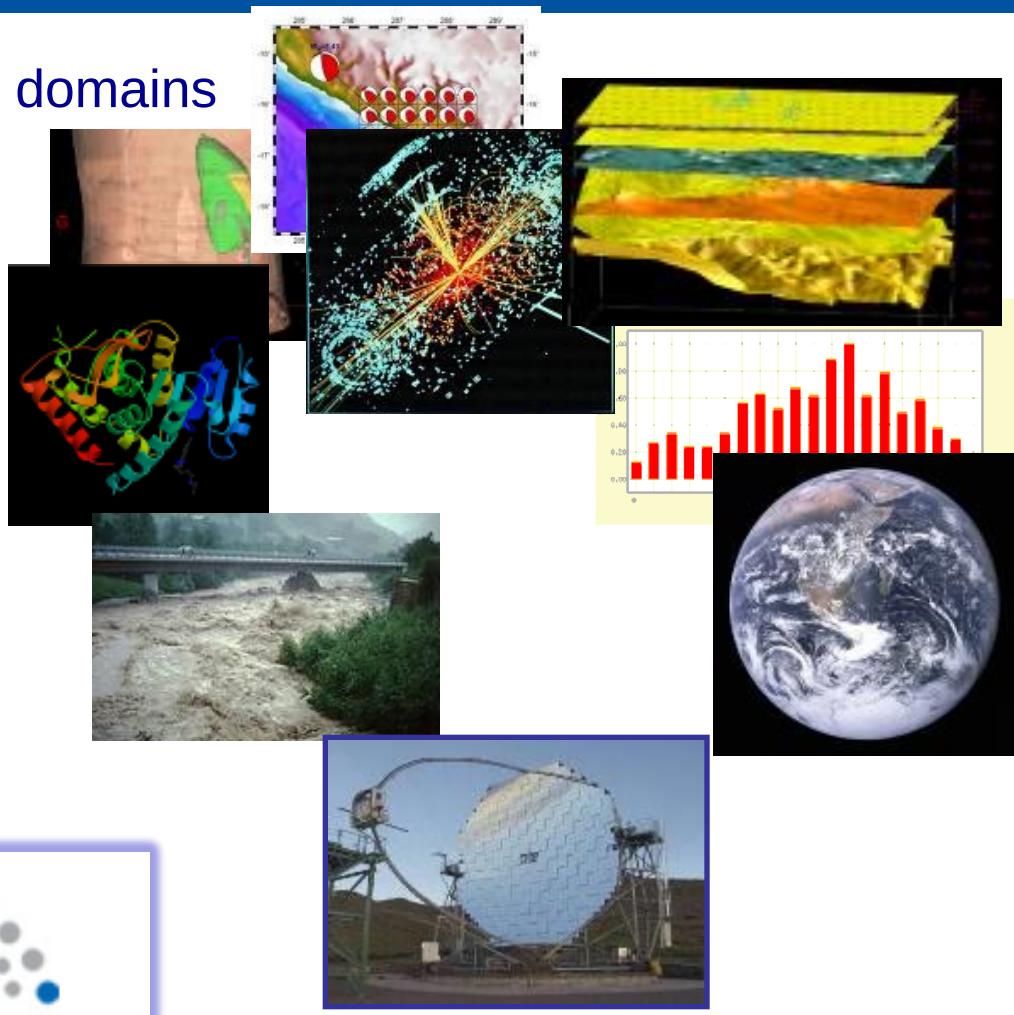

LHC Data Outlook

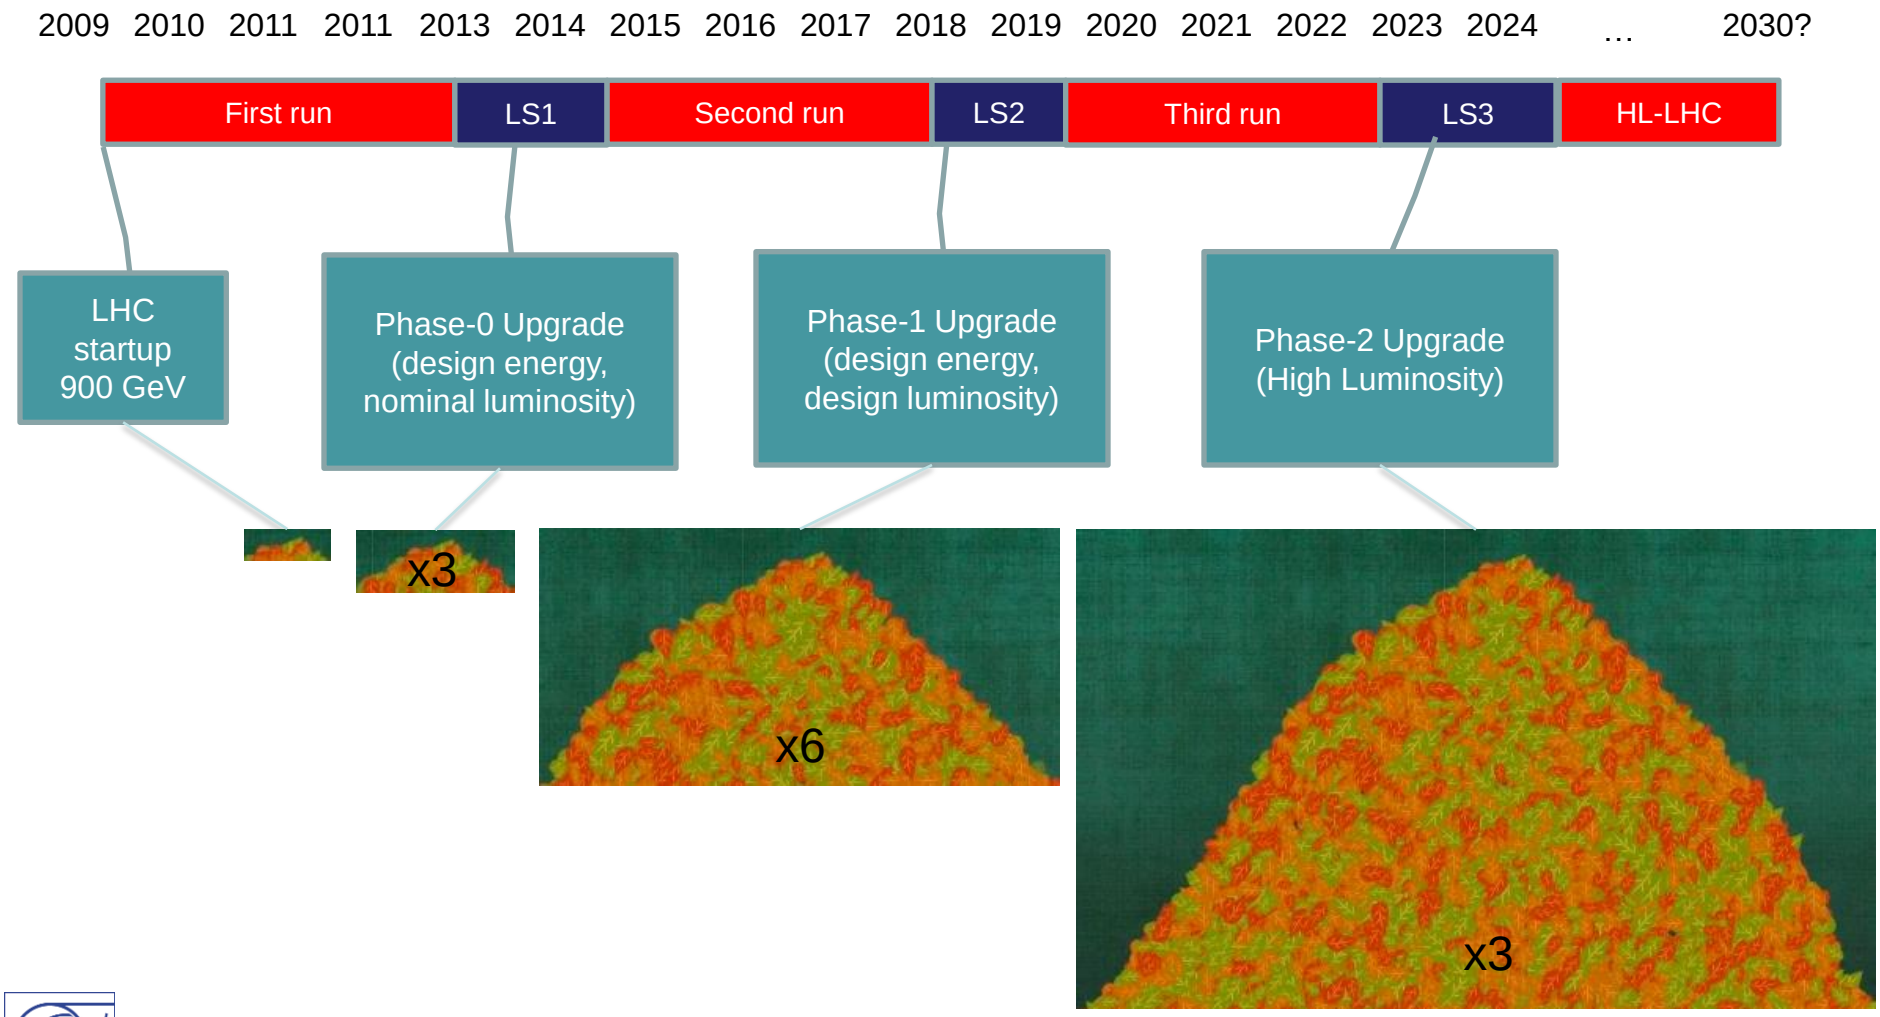

Computing Model Evolution

- Flattening the tiered structure
  - Strict hierarchy of connections becomes more of a mesh
- Network: a resource to schedule
  - Sending data directly to worker node over wide area network
- Pilots jobs
  - Building up an enormous batch queue
  - LHC experiments use the same read-only environment centrally distributed to nearly half a million processor cores
- AAA: any data, any time, any where ☺

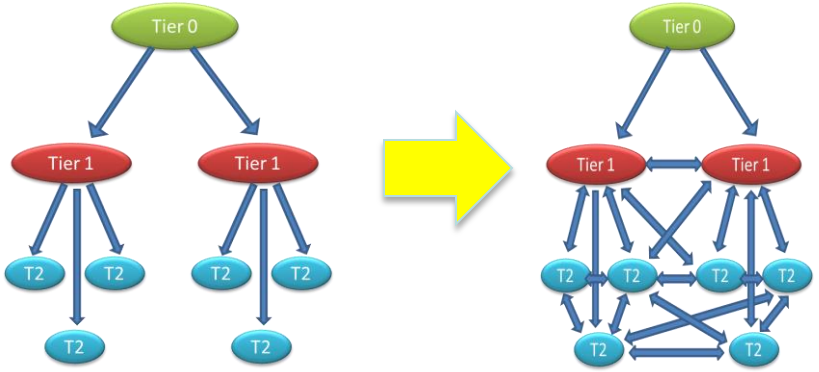

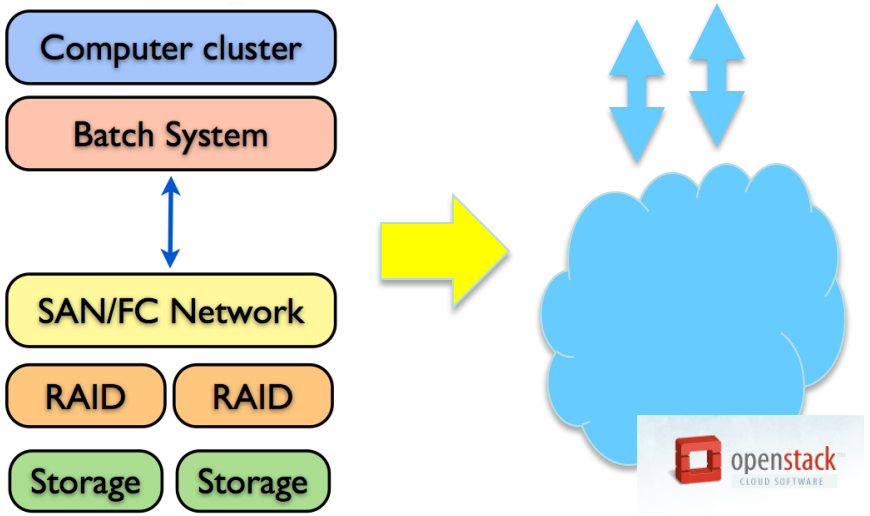

A New Computer Centre

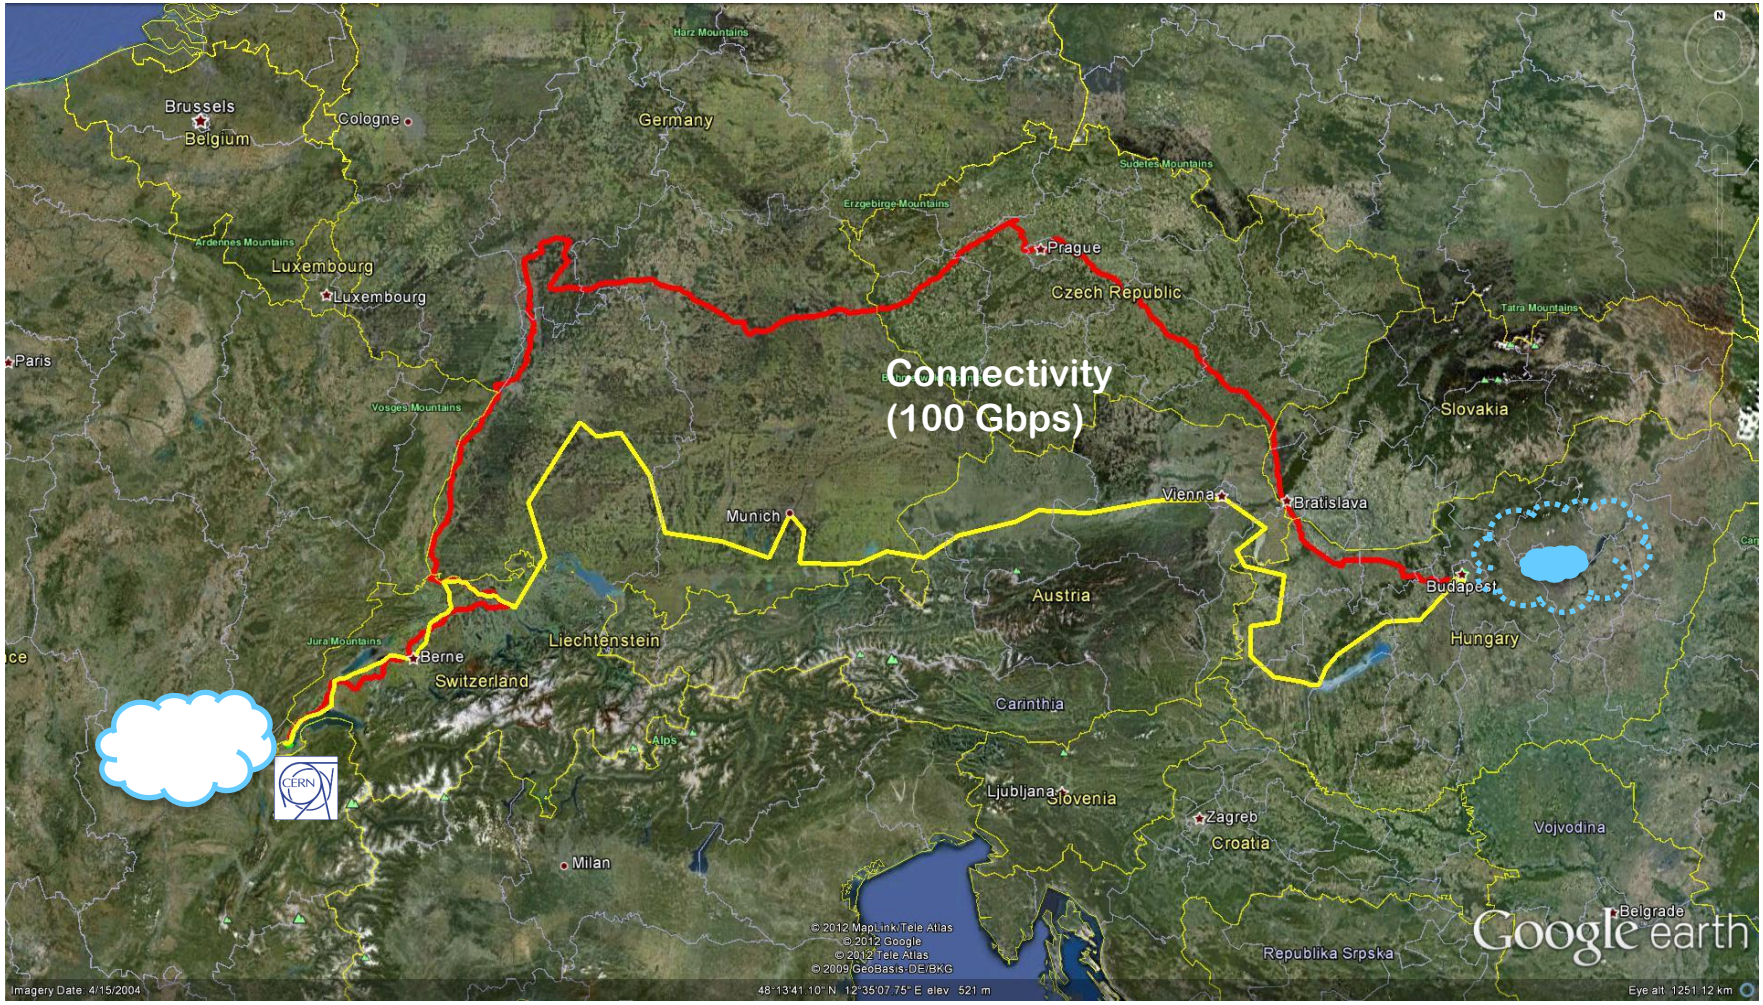

2015: 15k servers, 300k VMs

Computing

- Technical

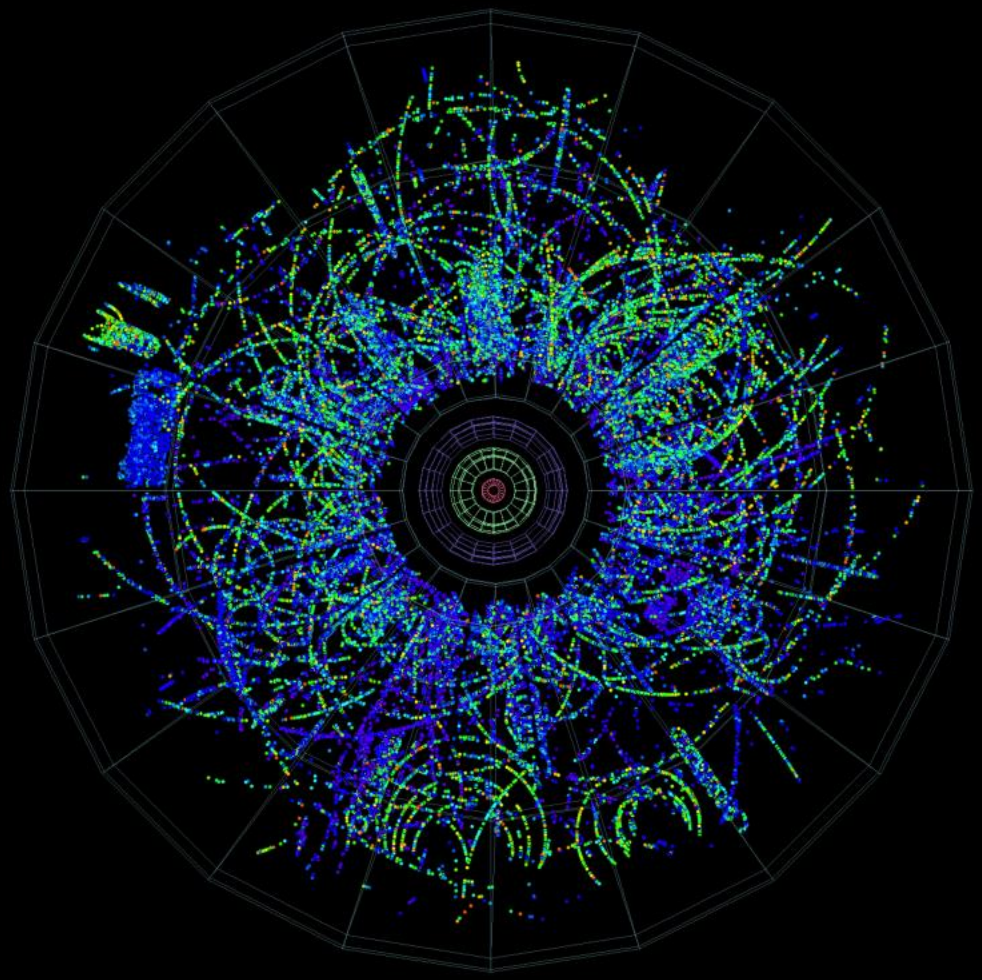

Design

- Magnet Design

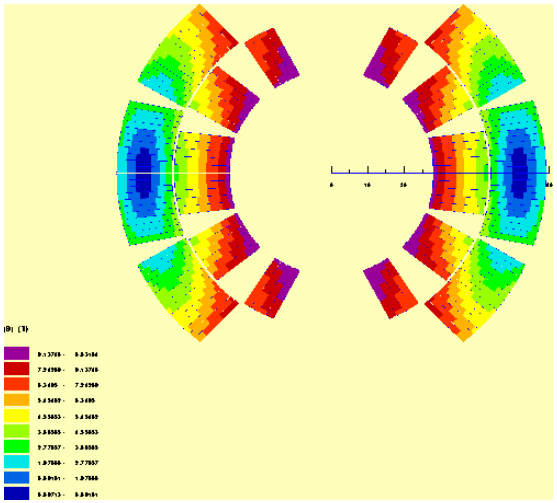

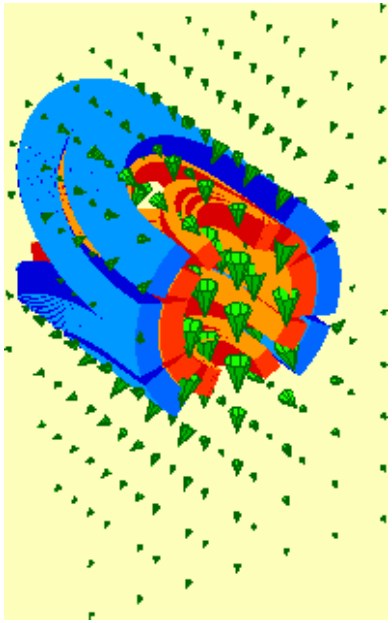

- Volunteer Computing
  - LHC@home
  - SixTrack
    - Simulates particles accelerating around the 27 km LHC to find their orbit stability

Operations

15-Nov-2012 11:50:45 | Fill #: 3288 | Energy: 4000 GeV | I(B1): 1.87e+14 | I(B2): 1.88e+14

	ATLAS	ALICE	CMS	LHCb
Experiment Status	PHYSICS	PHYSICS	PHYSICS	THANKS!!
Instantaneous Lumi [(ub.s)^-1]	3993.7	2.511	4080.1	395.9
BRAN Luminosity [(ub.s)^-1]	4083.7	1.632	4078.3	210.5
Fill Luminosity (nb)^-1	87568.7	61.0	89584.7	6829.4
BKGD 1	0.422	0.589	2.586	0.892
BKGD 2	83.489	177.284	3.361	4.286
BKGD 3	1.517	5.974	15.079	1.312

LHCb VELO Position: IN | Gap: -0.0 mm | STABLE BEAMS | TOTEM: STANDBY

Performance over the last 24 Hrs — Updated: 11:50:41

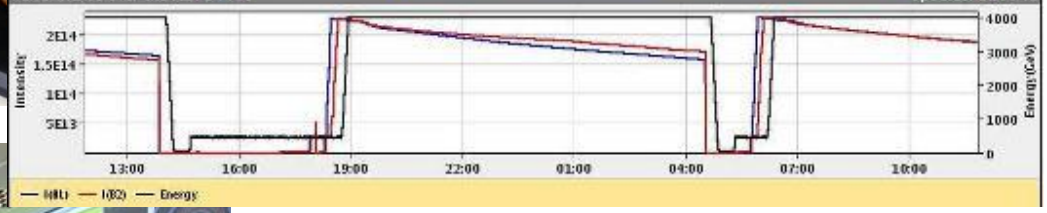

Computing

- Administrative
- Software

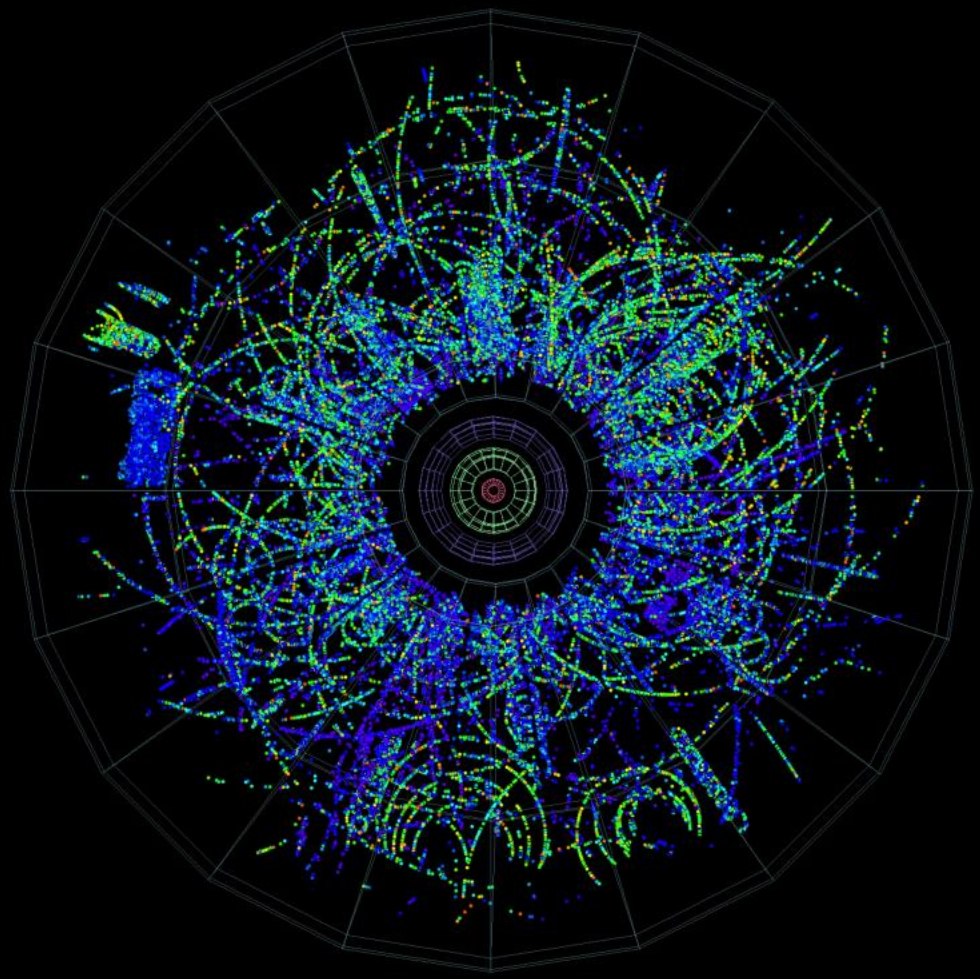

Administrative Information Services

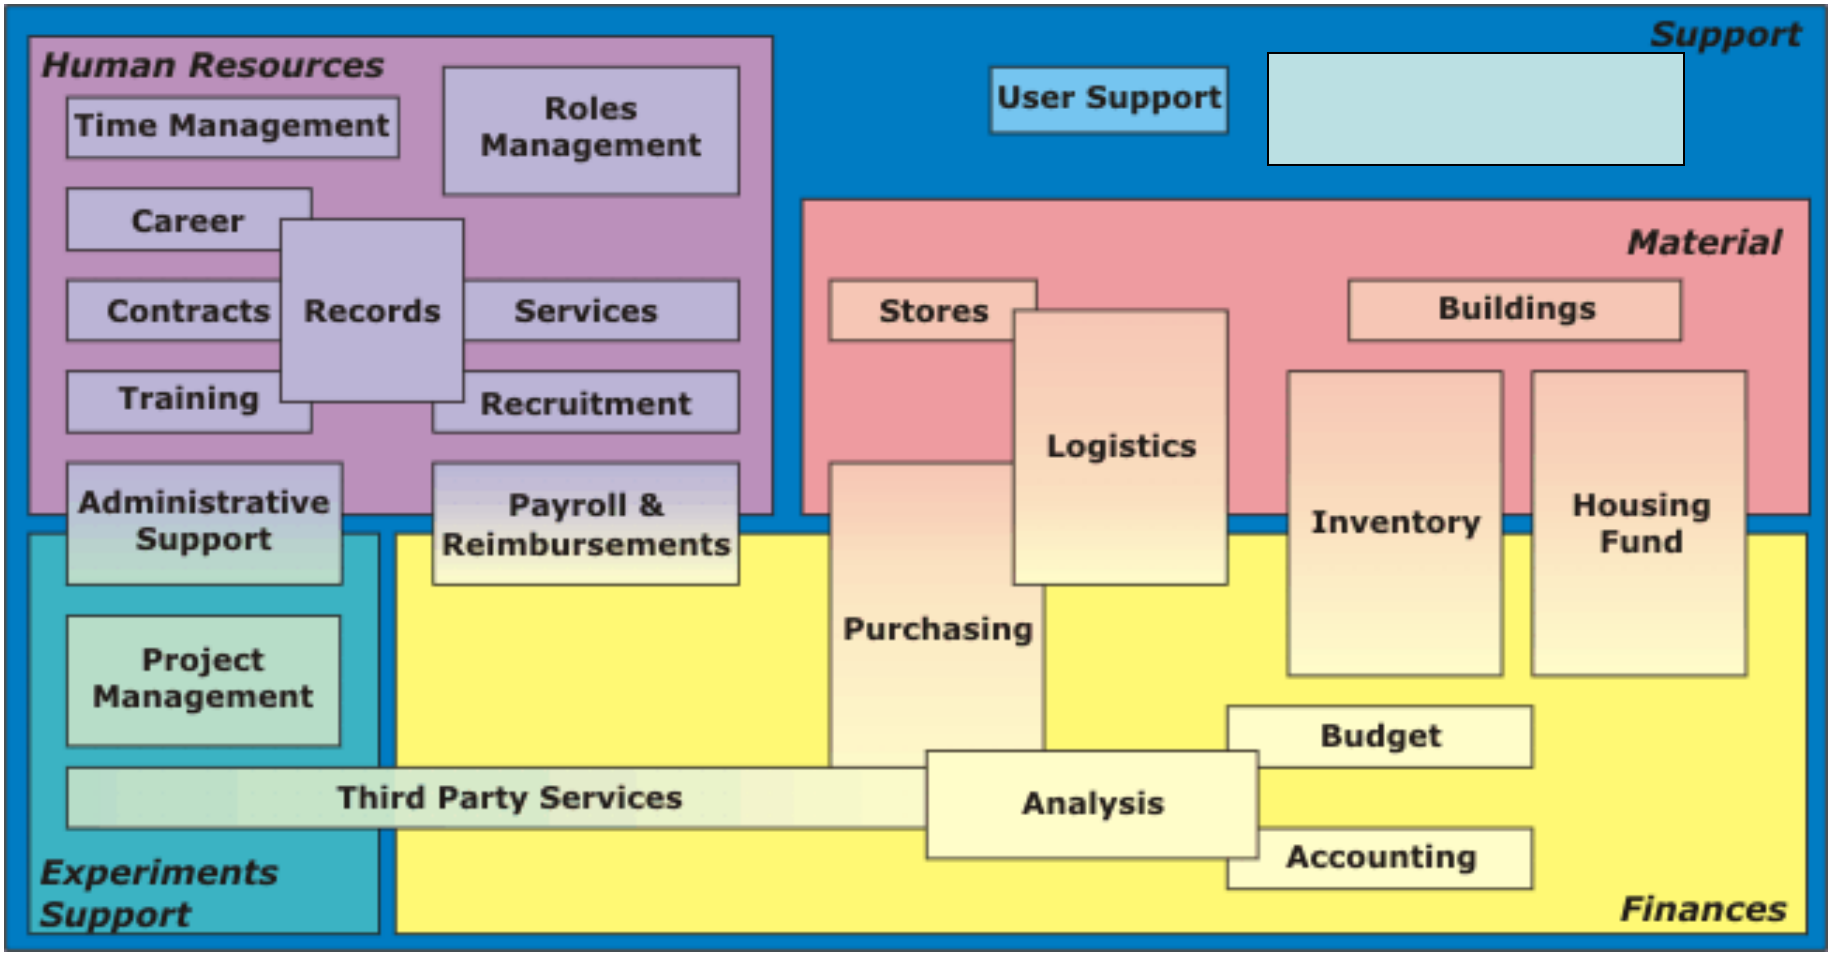

Administrative Information Services

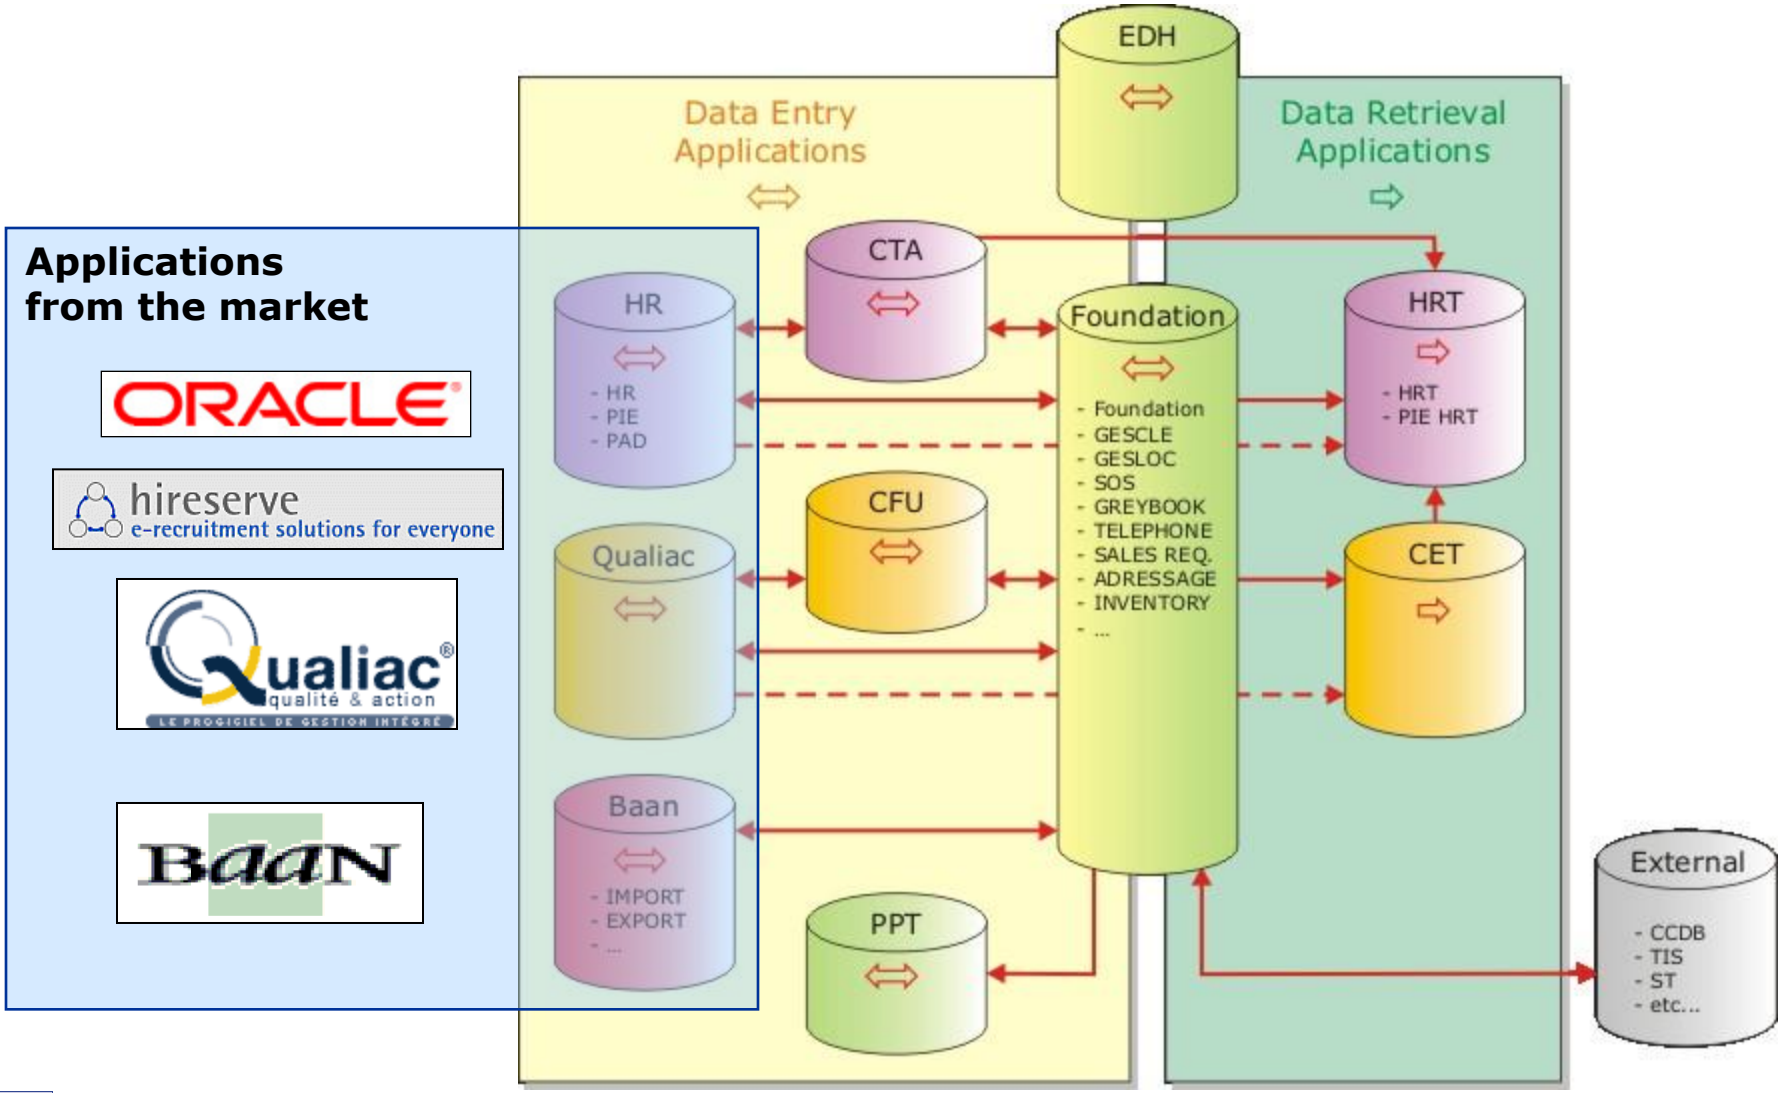

Software for Open Science

- Invenio: 514k; Indico: 377k

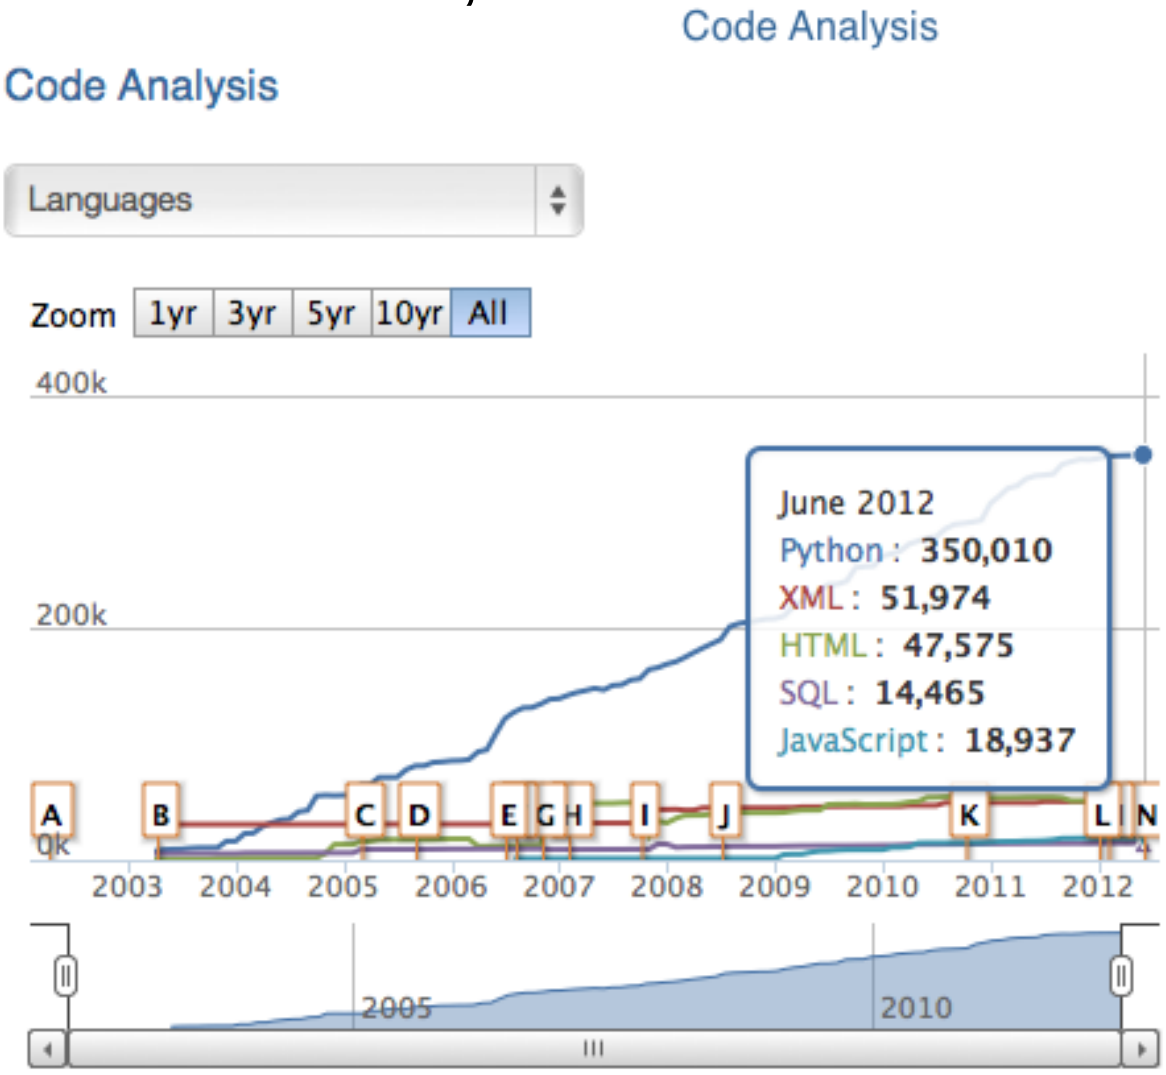

Open Source Software

- Intellectual Property
  - Restrictions, restrictions…
- Public Domain
  - A vast common good
  - Use, enjoy, share and build upon
  - IP expired, forfeited or inapplicable

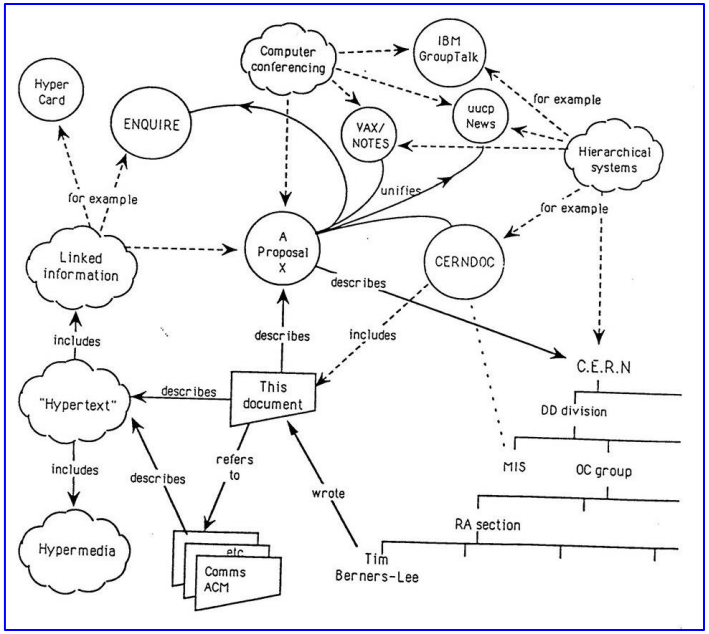

- Not all restrictions are bad!
  - Symmetric collaboration
  - Attribution
  - Standard form licences

© CERN with free licence

Computing

- Collaborative

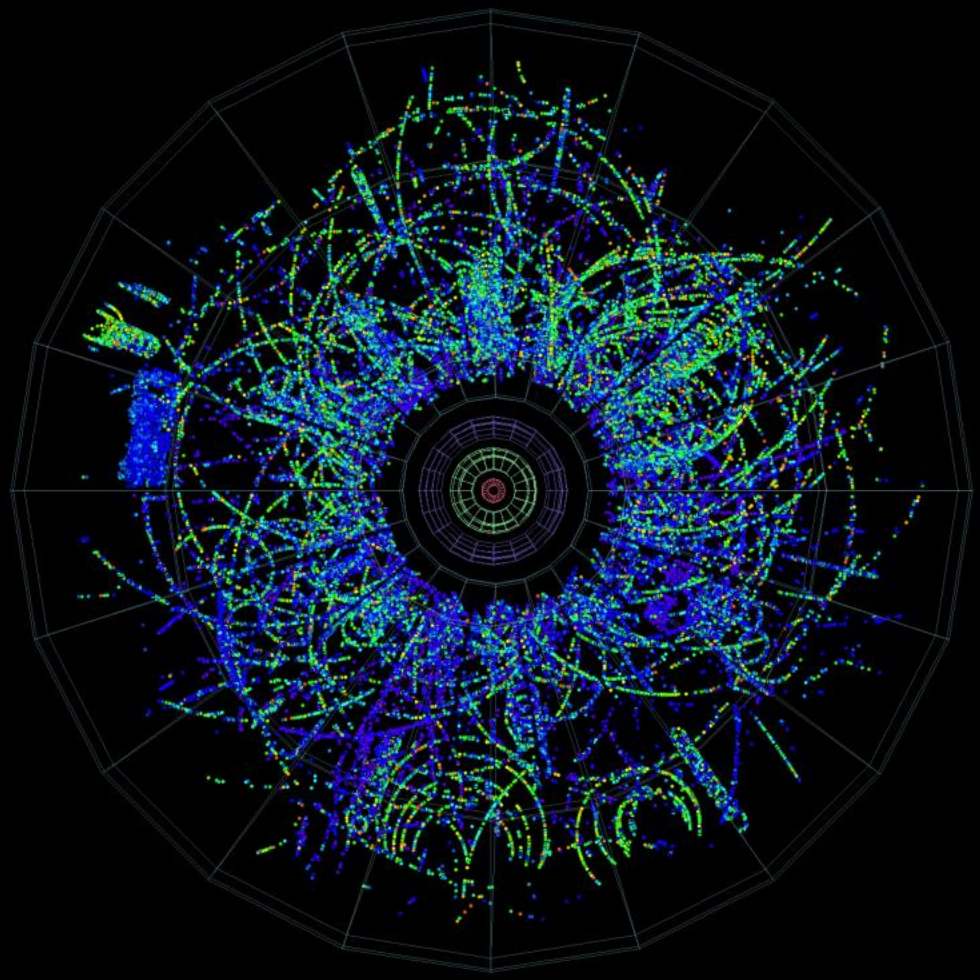

Videoconference

- 250 meeting rooms of all sizes on site
  - 90 equipped for video conference
    - Legacy + VidyoPanorama
  - 13 equipped for VC + Webcast
- 500 legacy endpoints worldwide
  - Non centrally managed

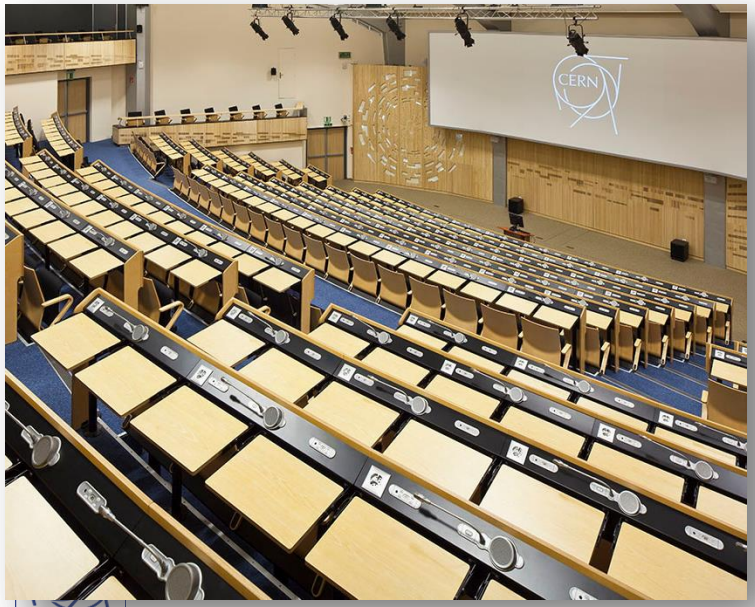

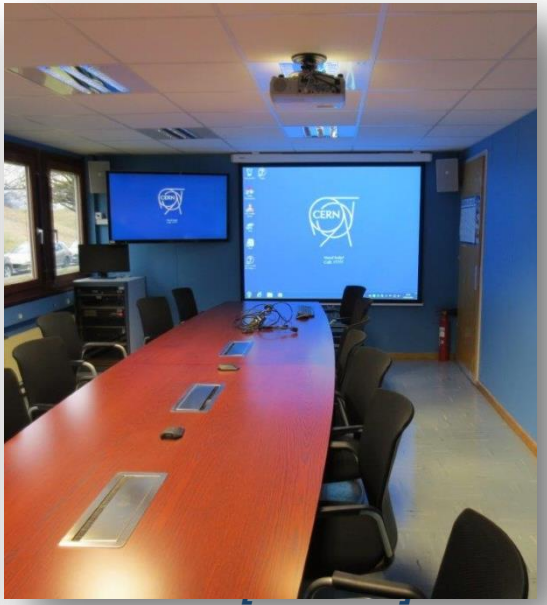

CERN Vidyo Worldwide Service Topology

- 3200 meetings/month
- 852 simultaneous connections
- 252 in one meeting
- 50M minutes last year / 40k downloads

Recruitment

- Asynchronous video screening

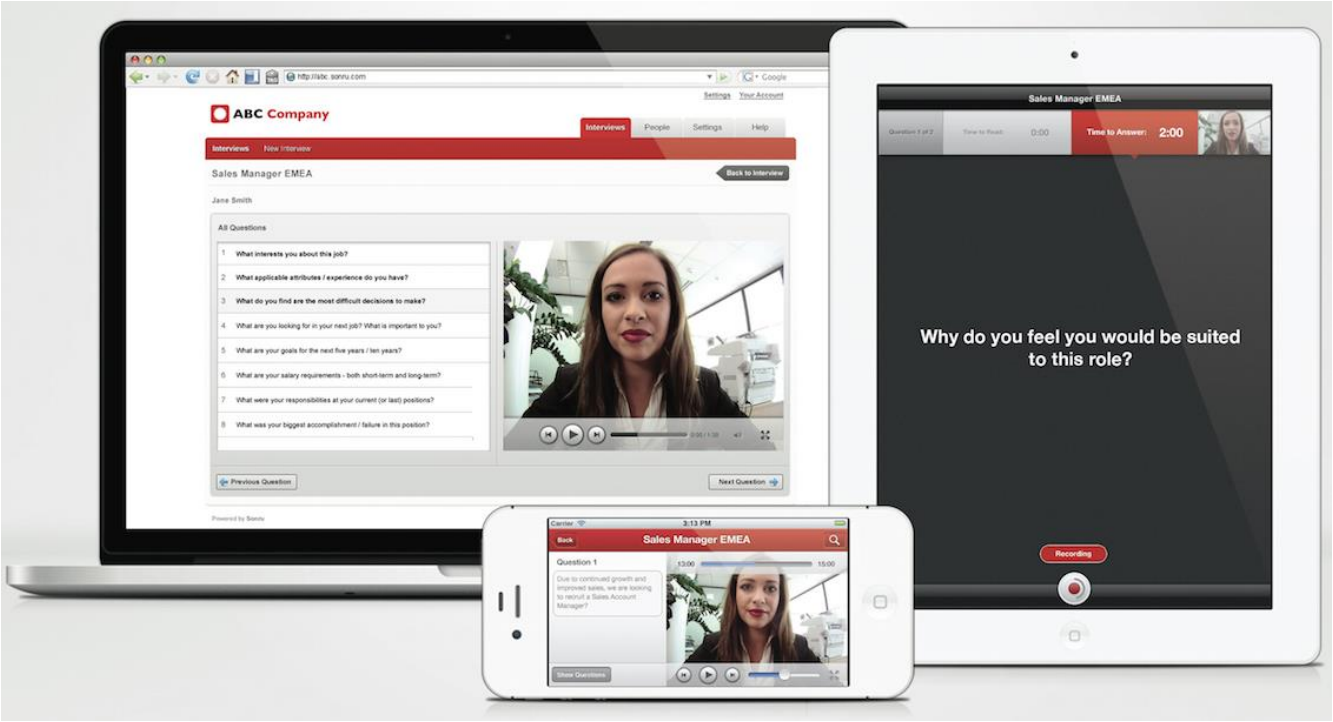

- Cost savings in bringing people to interview
- Multi-lingual – recruit from over 20 countries

CERN's social media

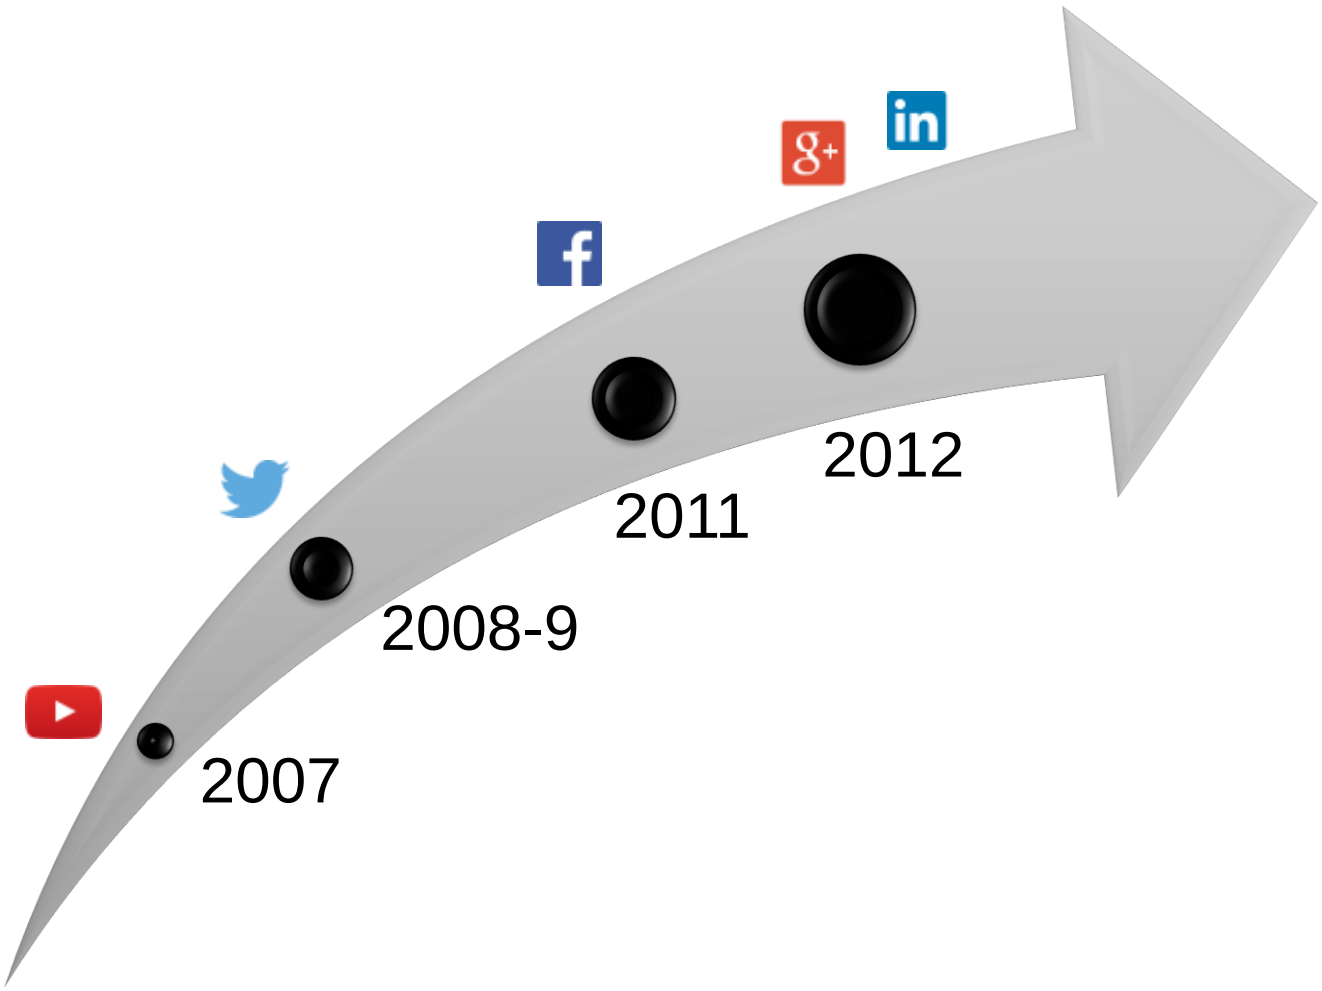

975K
309K
92K
40K
20K
12K

Impact

Courtesy of Twiplomacy, November 2013

Behind Open Science

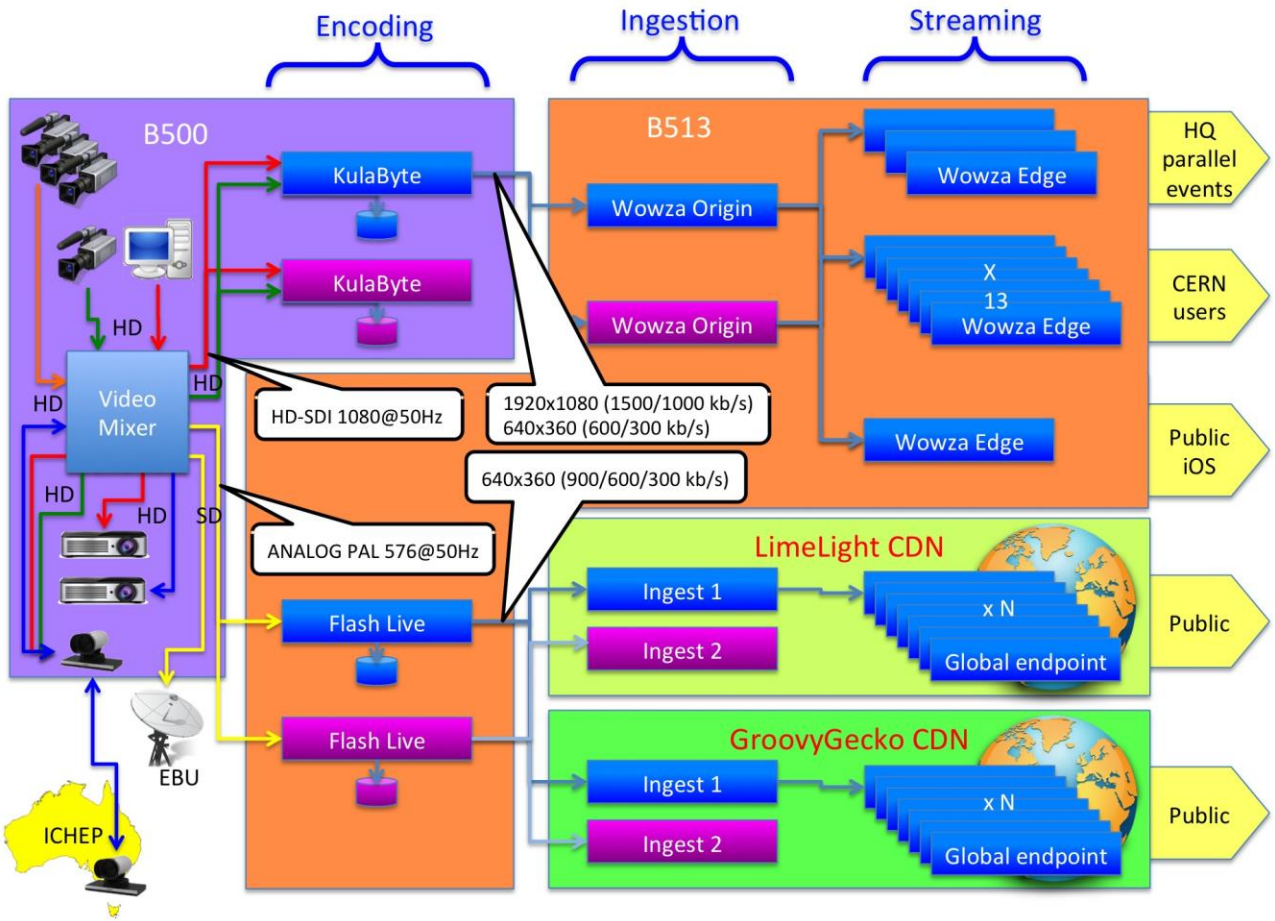

…Biggest Scientific Event ever

WebCast ~Million

TV ~Billion


THANK YOU

www.cern.ch

Links (I)

- Contact:
  - Tim.Smith@cern.ch

- More information:
  - IT Department: http://information-technology.web.cern.ch
  - The LHC Grid: http://wlcg.web.cern.ch
  - Google Street view in CC:
    - https://www.google.ch/maps/@46.232624,6.045747,3a,75y,162.48h,90t/data=!3m5!1e1!3m3!1sBU7JKhoaY_H9JVPFHcH8JA!2e0!3e5?hl=en
    - http://lego-scavenger-hunt.web.cern.ch
  - IT Archives: https://it-archives.web.cern.ch

Links (II)

- Social Media at CERN
  - http://twitter.com/CERN
  - http://twitter.com/CERN_FR
  - http://facebook.com/cern
  - http://google.com/+CERN
  - http://youtube.com/CERN
  - http://linkedin.com/company/cern