


Computing at CERN



Lorena Lobato Pardavila

International Teacher Weeks Programme 2017

Geneva, 14th August 2017

- ✓ PH-ESE, IT-DB, IT-CM...almost 6 years at CERN.
- ✓ DevOps Manager at IT Department, having fun with BATCH team
- ✓ Computational Neuroscience Systems, Embedded Systems, Cloud Computing and IT challenges
- ✓ Enjoy traveling a bit too much and sport lover! 😊

About me





Science



Computing




Collaboration
& Innovation

**What will I
talk about?**



Science






An Early “Computer”



- Wim Klein
- Calculating the 73rd root of a 500 digit number took less than 3 minutes...
- Not the first CERN Computer! Two “female computers” were already working with mechanical calculators






Bubble Chamber



Spark Chambers



1958, The Ferranti Mercury arrived



Number 1

15 February 1966


Introduction

As computing becomes a more and more widespread and complex activity in the laboratory, the need will increase for a means to have a wider general circulation of background information about different aspects of computing activities than is possible with the present system of Computer Notices. I therefore make no apology for introducing yet another circular which will find its way on to CERN desks. Rather I would express the hope that this newsletter will prove to be a useful source of general information on computer use and performance, programming developments and the requirements of different kinds of computer users, as well as on future plans for computers, programming and computer uses in the laboratory. The newsletter will be edited by the Computer Manager and any comments, criticism or suggestion

CDC 6600



1972, Super Computer Installed




Tapes being sent up from B513 basement



Mainframes



Commodity Computing



Today's Data centre





Computing



Networking

➤ Science without borders

- Data exchange across the iron curtain
- China and scientific world (IHEP to CERN)

➤ Truly international Internet

- 1989 first external TCP/IP connection
- 1990 Principle link US-EU from CERN
- 1991 80% of the internet capacity installed in Europe for international traffic was terminated at CERN

➤ CERN Internet eXchange Point (CIXP)

Vague but exciting ...

CERN DD/OC

Tim Berners-Lee, CERN/DD

Information Management: A Proposal

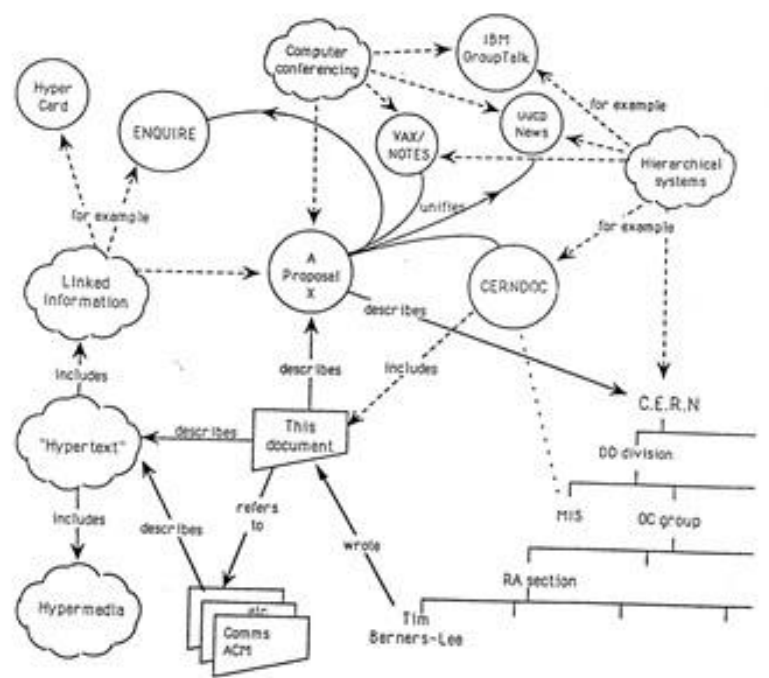
March 1989

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



World Wide Web



World Wide Web

← → ↻ ⓘ info.cern.ch/hypertext/WWW/

World Wide Web

The WorldWideWeb (W3) is a wide-area [hyper](#)

Everything there is online about W3 is linked di

[What's out there?](#)

Pointers to the world's online information

[Help](#)

on the browser you are using

[Software Products](#)

A list of W3 project components and their

[Technical](#)

Details of protocols, formats, program int

[Bibliography](#)

Paper documentation on W3 and referenc

[People](#)

A list of some people involved in the proj

[History](#)

A summary of the history of the project.

[How can I help ?](#)

If you would like to support the web..

[Getting code](#)


Getting the code by [anonymous FTP](#) , etc




er's [W3 news](#) , [Frequently Asked Questions](#) .



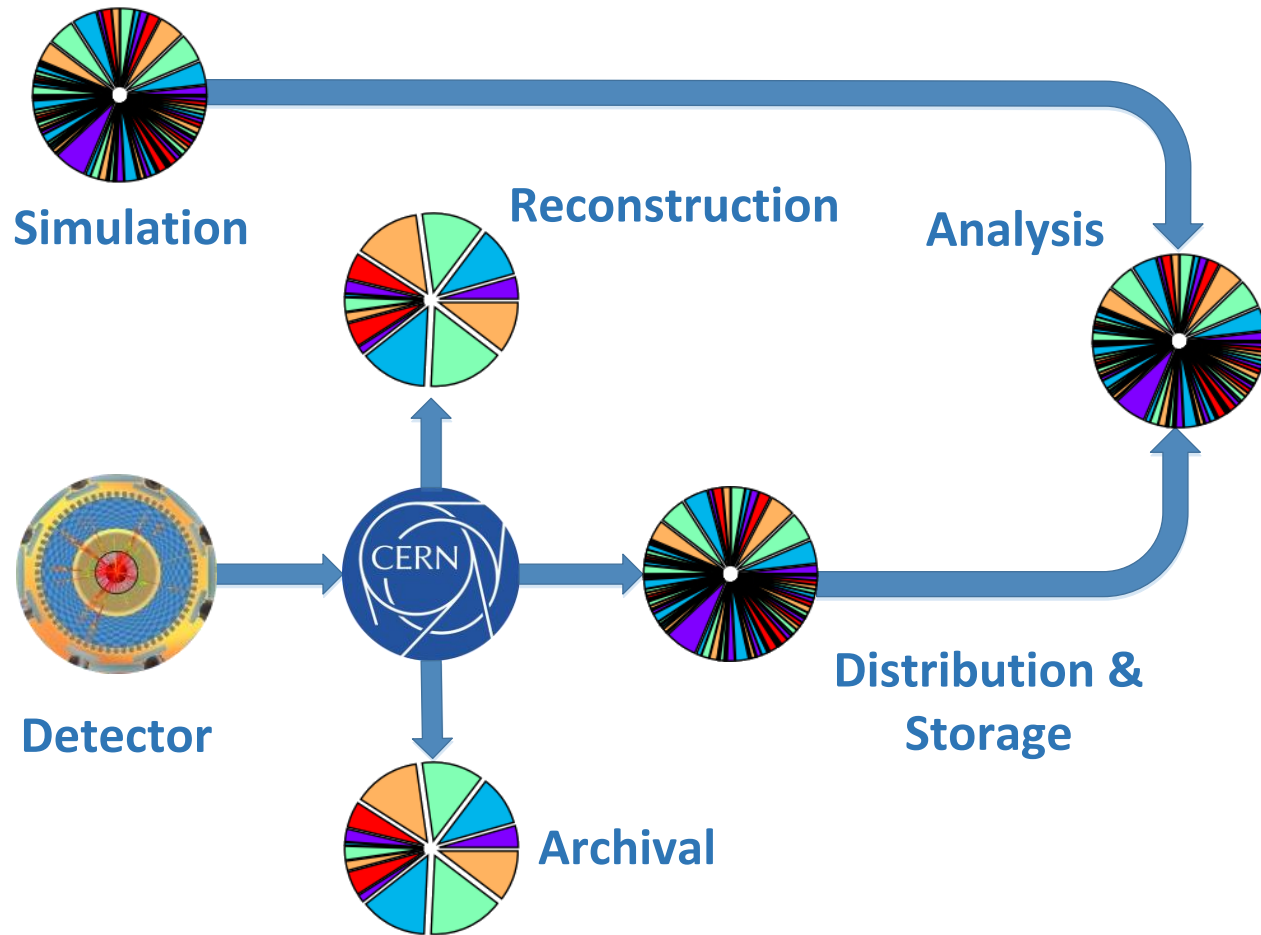
Where does the data come from?



- 4 experiments. 99.99 % data filtered at the experiment
- 1 PB/second!

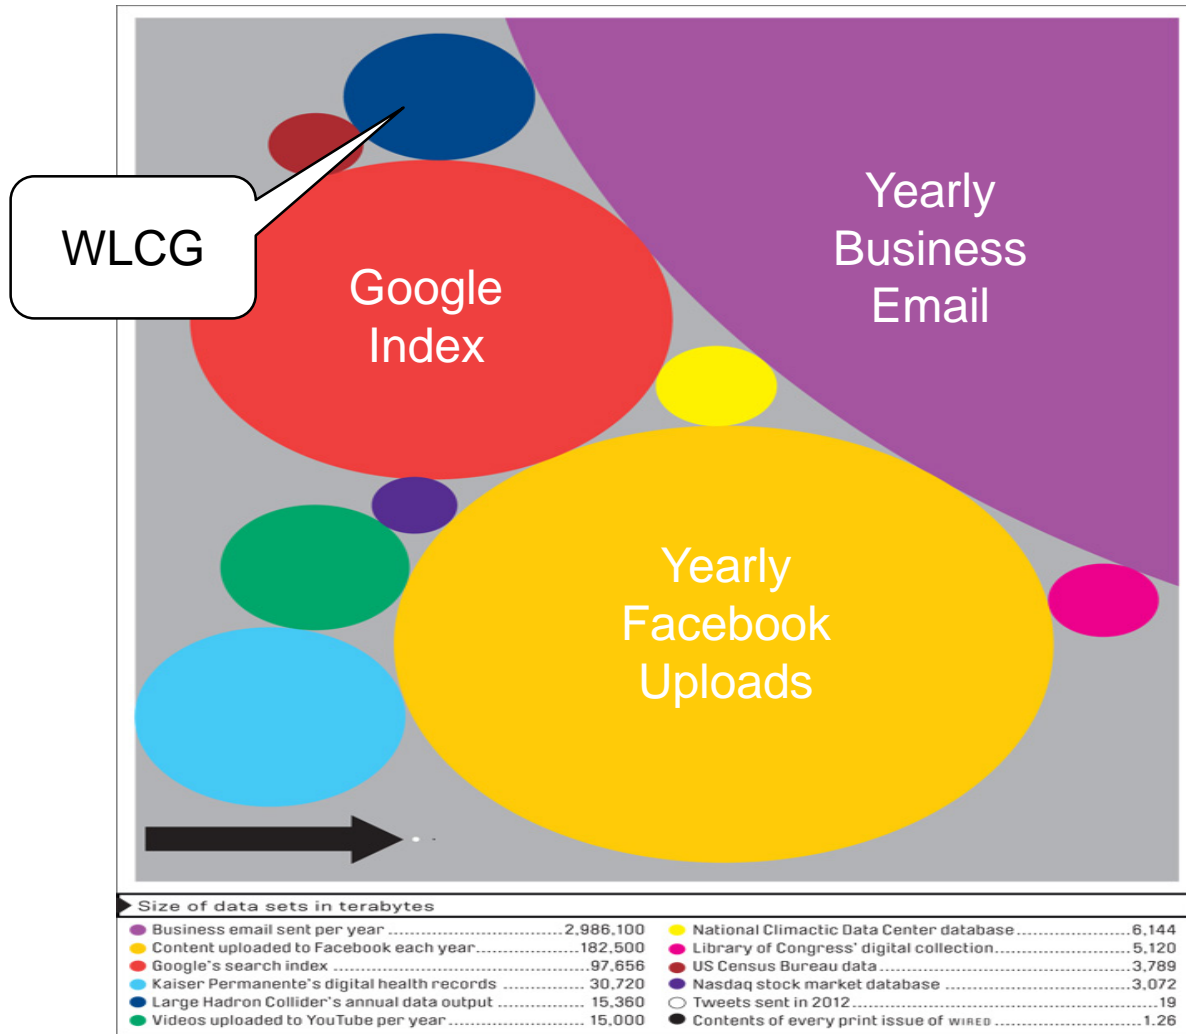


What is the workflow?



- More than half the CPU goes on simulation.
- Most of the rest is reconstruction.
- The remainder is analysis.

How much data are we talking?



- Big Data! 😊
- 2012, 15 PB
- 2017 estimates 50 PB, equivalent to a 12km high stack of DVDs
- CERN can only provide 20%-30% storage and CPU

Data Centre

MEYRIN DATA CENTRE

	last_value
● Number of Cores in Meyrin	121,806
● Number of Drives in Meyrin	56,231
● Number of 10G NIC in Meyrin	11,019
● Number of 1G NIC in Meyrin	12,621
● Number of Processors in Meyrin	16,640
● Number of Servers in Meyrin	9,036
● Total Disk Space in Meyrin (TB)	144,551
● Total Memory Capacity in Meyrin (TB)	595


WIGNER DATA CENTRE

	last_value
● Number of Cores in Wigner	56,000
● Number of Drives in Wigner	29,686
● Number of 10G NIC in Wigner	2,981
● Number of 1G NIC in Wigner	6,579
● Number of Processors in Wigner	7,002
● Number of Servers in Wigner	3,504
● Total Disk Space in Wigner (TB)	97,286
● Total Memory Capacity in Wigner (TB)	221


NETWORK AND STORAGE

	last_value
● Tape Drives	104
● Tape Cartridges	24,729
● Data Volume on Tape (TB)	188,124
● Free Space on Tape (TB)	35,213
● Routers (GPN)	148
● Routers (TN)	31
● Routers (Others)	101
● Switches	3,833

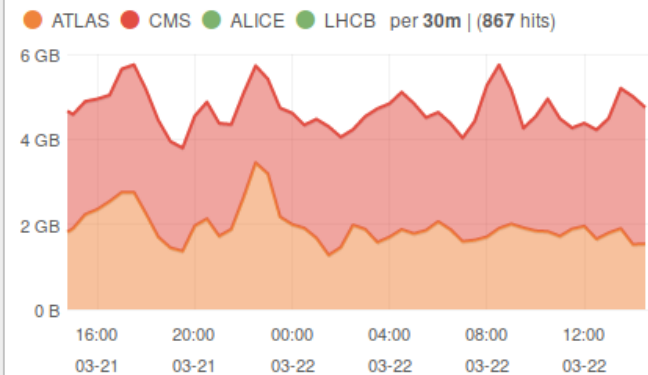
BATCH JOBS (#)



EOS ACTIVE DATA TRANSFERS (#)




FILE TRANSFER THROUGHPUT (GB/S)



Data Challenge

- Few places can store it
- HEP community distributed
- Resource Management
 - Computing power and Data Storage capacity
 - Resources of CERN and its experiments
- Remote access...? - Distributed solution...;)



Distributed Systems

- Collection of independent computers
- No shared-memory
- Continuous availability
- Easily expandable



Grid Computing



- Collection of a large number of computer resources from multiple locations connected together in a network to reach a common goal.
- Sharing computing resources and storage resources
- Many computers acting as a single one

Cloud?

Cloud

- On Demand
- Dynamically provisioned & metered by e.g. Amazon, Microsoft Azure


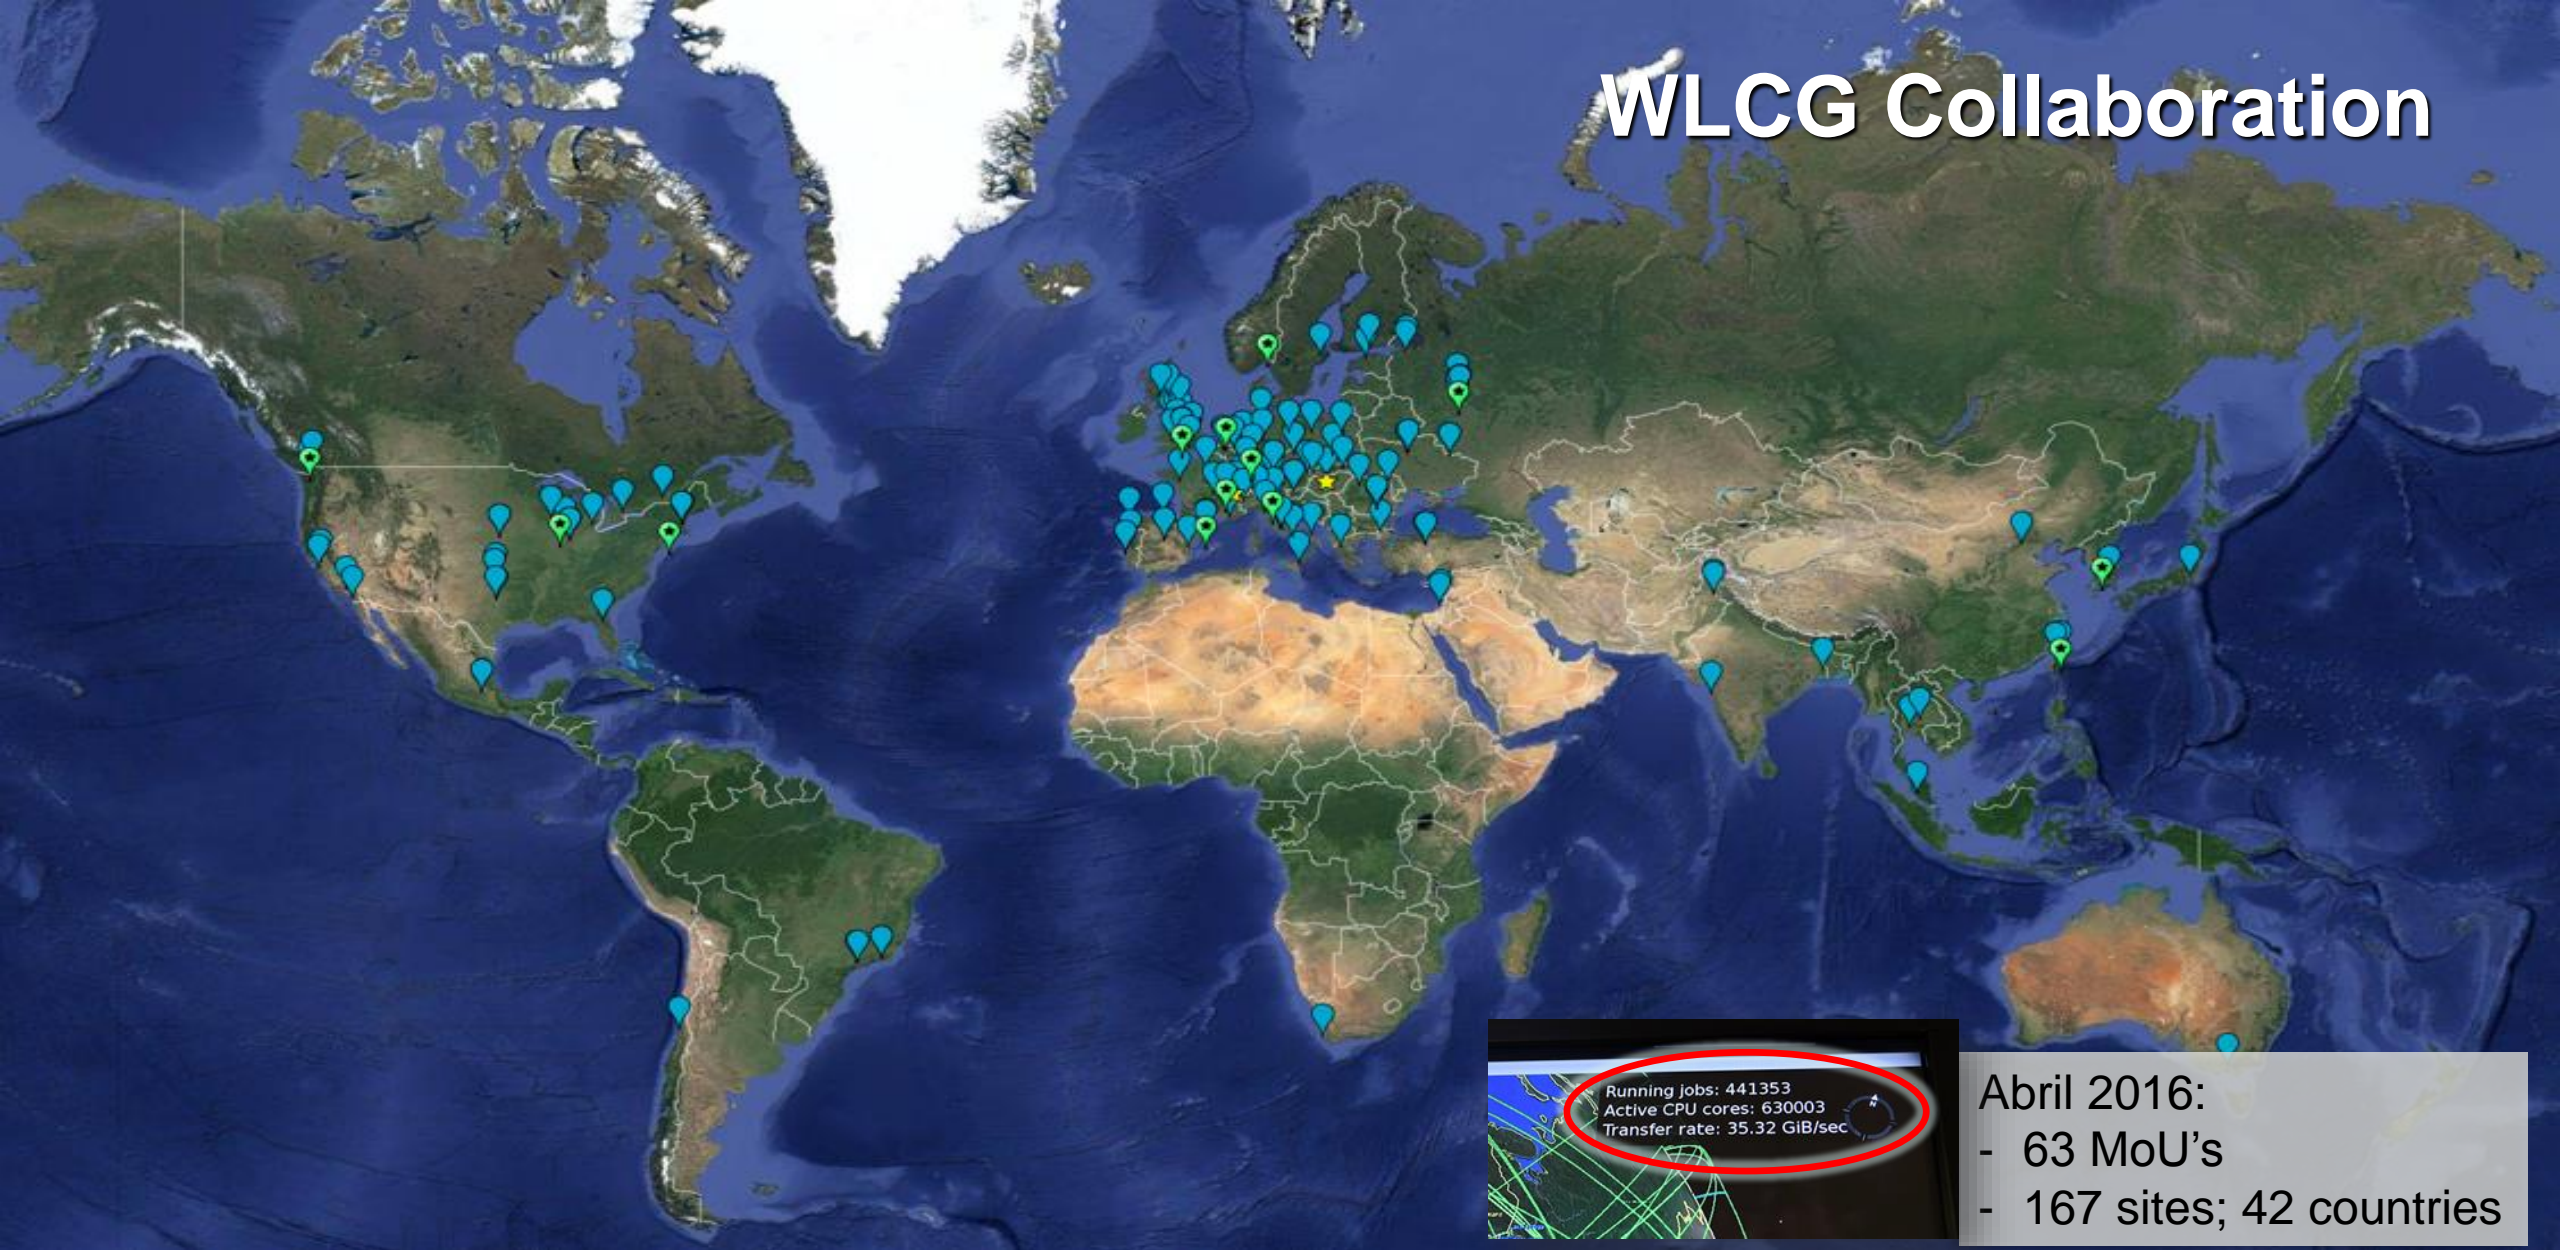


Grid

- Fixed size
- Collaborative, run by community



WLCG Collaboration




Abril 2016:
- 63 MoU's
- 167 sites; 42 countries

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

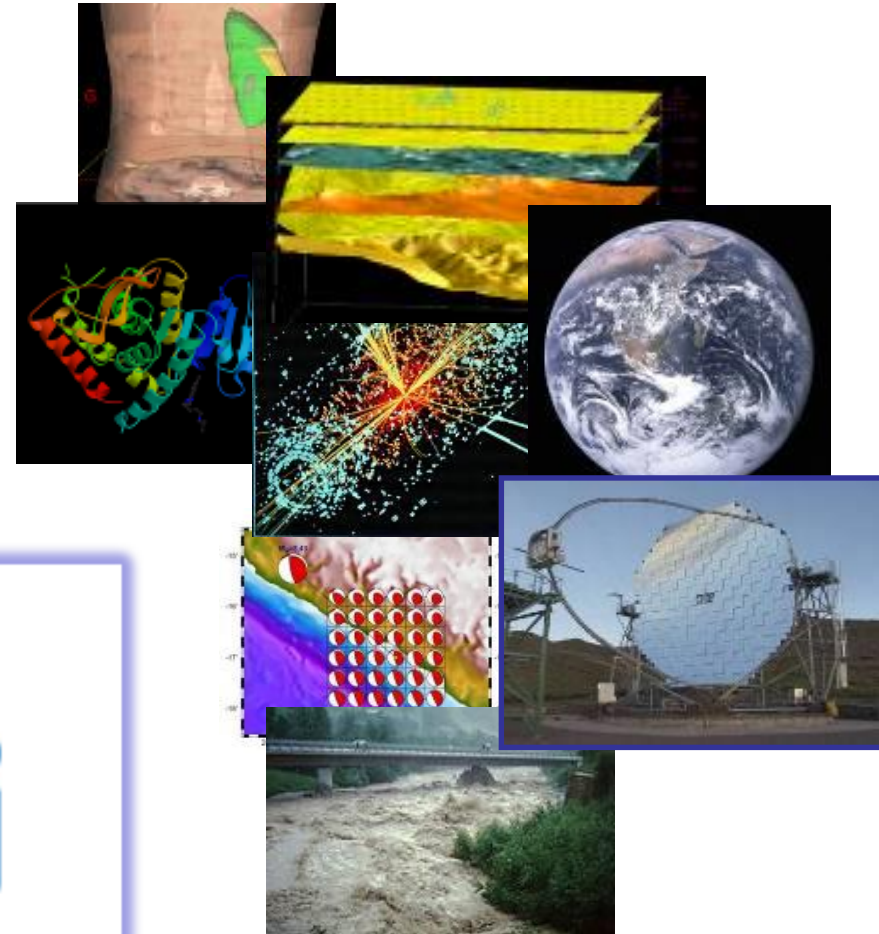
Volunteer/Oppportunistic

- Scavenged resources
 - Volunteers (e.g. home PCs)
 - Institute desktops
 - Supercomputing backfill
 - Small farms with easy deployment!
- Unpredictable but significant resources
 - Target CPU bound simulations (not data intensive)
 - Over 50% of LHC compute is simulation!
- Outreach benefits



Shared Infrastructures: EGI

- Astronomy & Astrophysics
- Civil Protection
- Computational Chemistry
- Comp. Fluid Dynamics
- Computer Science/Tools
- Condensed Matter Physics
- Earth Sciences
- Fusion
- High Energy Physics
- Life Sciences

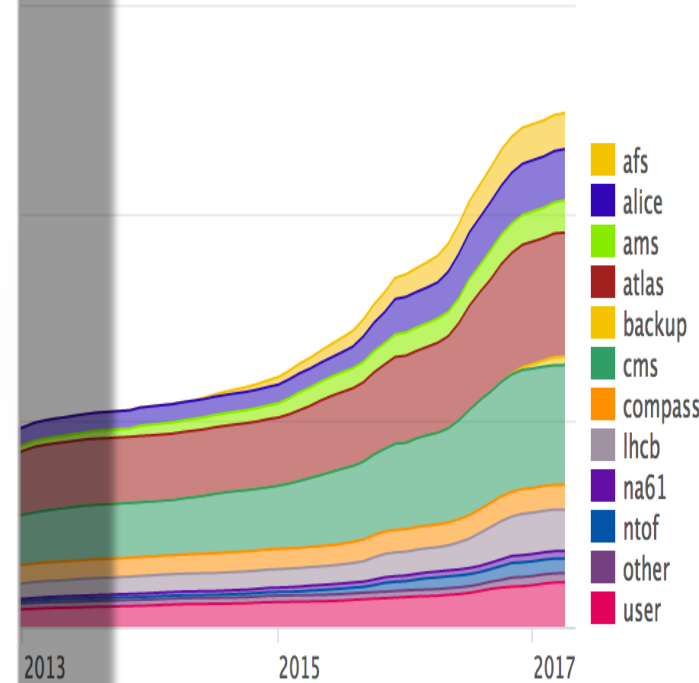



Data Centres



Data Transferred to...

- 203 PB physics
- Now 50 PB/year (1000 Pb/year)
- Preserving long-term data
- Exa-scale Resources



Evolution of computing model

- Consolidation
- WAN access
- Resource Diversity
- Standard Solutions
- Cohabitation (big astro projects coming online)




Collaboration & Innovation




Seminars, training courses, academic training


CERN openlab Summer Students Programme



Year	Candidates	Selected
2007	25	15
2008	15	15
2009	15	15
2010	92	15
2011	330	15
2012	490	15
2013	750	22
2014	850	23
2015	1540	40
2016	1480	39

- 200+ press cuttings
- 150,000 visits to our website
- 50+ events, visits, lectures
- 100+ presentations
- 50+ news articles, press releases, case studies




15 years CERN openlab


Applications to cross-disciplinary research

Credit to Hannah Short

European Commission projects



Bridging the gap between schools and universities through informal education

<https://up2university.eu>



Strategic Plan

- ▶ Establish multi-tenant, multi-provider cloud infrastructure
- ▶ Identify and adopt policies for trust, security and privacy
- ▶ Create governance structure
- ▶ Define funding schemes



To support the computing capacity needs for the ATLAS experiment



Setting up a new service to simplify analysis of large genomes, for a deeper insight into evolution and biodiversity



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



To improve the speed and quality of research for finding surrogate biomarkers based on brain images

Additional Users:



Suppliers



Adopters



Visitors from ~ all Countries

Including

- Antarctica
- Vatican City

56% from Europe

57k Records

- 11k Software
- 3k Datasets

700 Communities

- Projects
- Institutes
- Subjects
- Conferences
- Publishers

Open Science: Zenodo

11 111 111 111 111 111



- Infrastructure
- Impact
- Supports LTOS & large groups



COLLABORATIVE TOOLS



International Teacher Weeks 2017 – Computing at CERN



Videoconference

250 meeting rooms of all sizes on site

- 100 equipped for video conference
- Legacy + VidyoPanorama
- 16 equipped for VC + Webcast

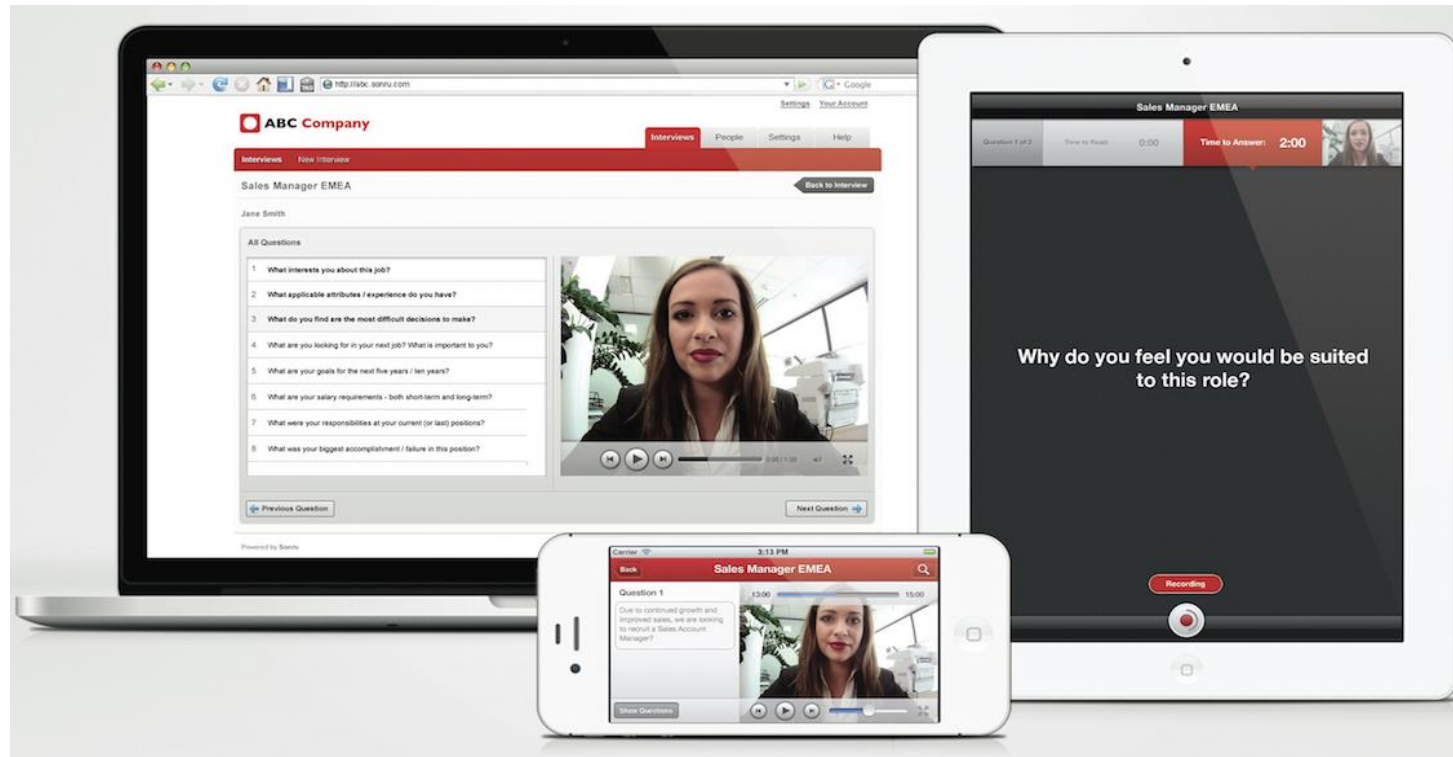
500 legacy endpoints worldwide

- Non centrally managed
- Licenses + Software
- Configuration of the computing resources




CERN Vidyo Worldwide Service Topology

- 8184 meetings/month
- 941 simultaneous connections
- 252 in one meeting
- 50M minutes last year / 40k downloads



- Asynchronous video screening
- Cost savings in bringing people to interview
- Multi-lingual – recruits from over 20 countries

Recruitment



CERN's social media

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

THANK you! Questions?



SOURCE: <http://www.slidegenius.com/blog/dilbert-on-powerpoint-serious-powerpoint-lessons-silly-comic-strip/>

You can find me at:

Email: lorena.lobato@cern.ch

Twitter: @lobatopardavila

CREDITS

Tim Smith, Hannah Short and Germán Cancio for their content and ideas
Resources for free: Photographs by [Pexels](#) and [Pixabay](#)
All the colleagues who have spent time telling me stories about IT 😊



I want to know more about...

- **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>



I want to know more about...

Social Media at CERN

- **Twitter:** <http://twitter.com/CERN>
- **Facebook:** <http://facebook.com/cern>
- **Google+:** <http://google.com/+CERN>
- **Youtube:** <http://youtube.com/CERN>
- **Linkdin:** <http://linkedin.com/company/cern>





www.cern.ch