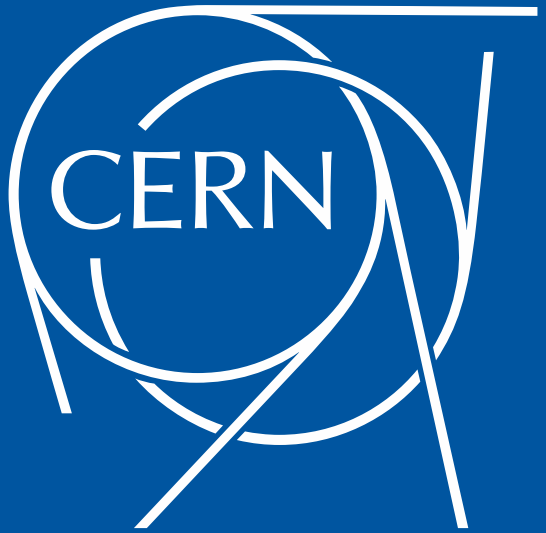
CERN

CERN Computing Infrastructure

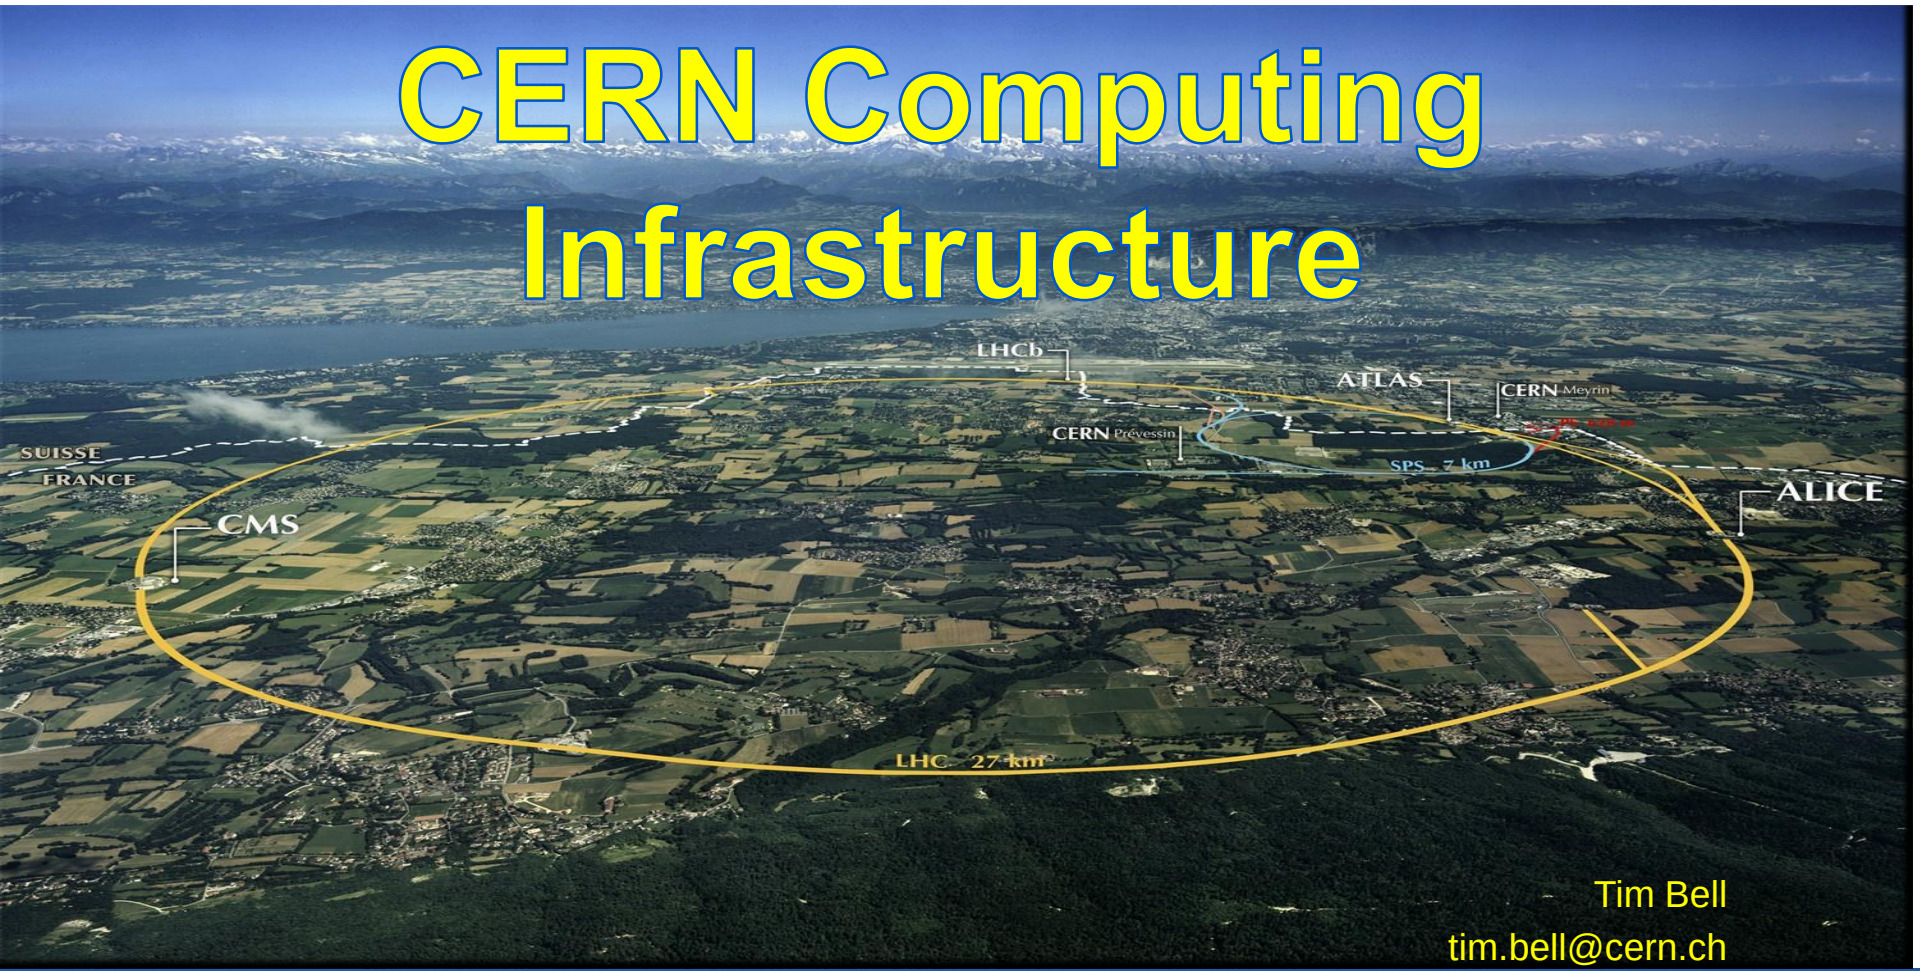

Tim Bell
tim.bell@cern.ch

Computing Infrastructure

- Diverse computing services
  - Physics computing
  - IT and Experiment Services
  - Administrative Computing
- Target is for
  - Standardised procedures
  - Bulk purchasing

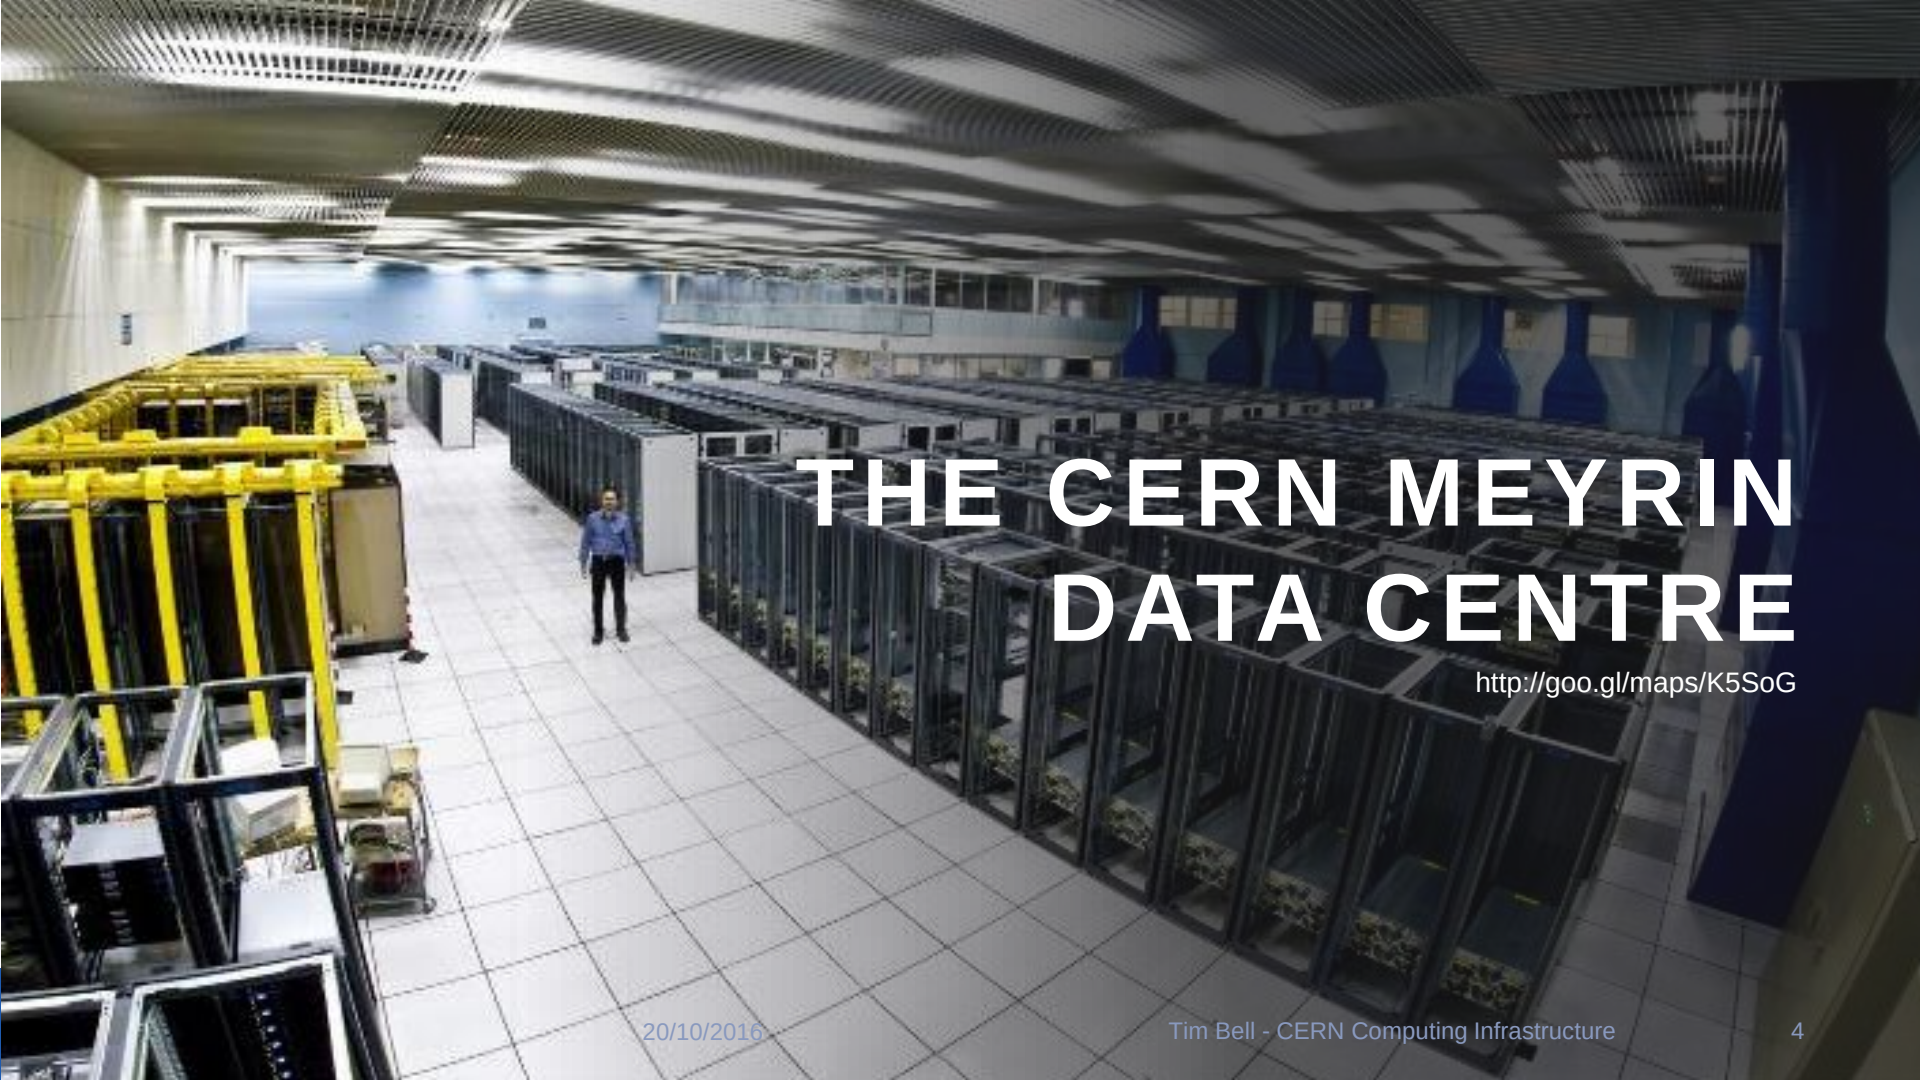

THE CERN MEYRIN DATA CENTRE

http://goo.gl/maps/K5SoG


DANTE
100 GbE
T-Systems
100 GbE
Wigner RCP
CERN
Amsterdam
Netherlands
Niedersachsen
Berlin
Poznan
Warsaw
Poland
Antwerpen
Brussels
Belgium
Cologne
Germany
Prague
Czech Republic
Krakow
Luxembourg
Luxembourg
Paris
Slovakia
Munich
Vienna
Bratislava
Budapest
Austria
Zurich
Liechtenstein
Bern
Switzerland
France
Hungary
Slovenia
Zagreb
Croatia
Veneto
Milan
Torino
Belgrade

Public Procurement Cycle

Step	Time (Days)	Elapsed (Days)
User expresses requirement		0
Market Survey prepared	15	15
Market Survey for possible vendors	30	45
Specifications prepared	15	60
Vendor responses	30	90
Test systems evaluated	30	120
Offers adjudicated	10	130
Finance committee	30	160
Hardware delivered	90	250
Burn in and acceptance	30 days typical with 380 worst case	280
Total		**280+ Days**

O’Reilly Consideration

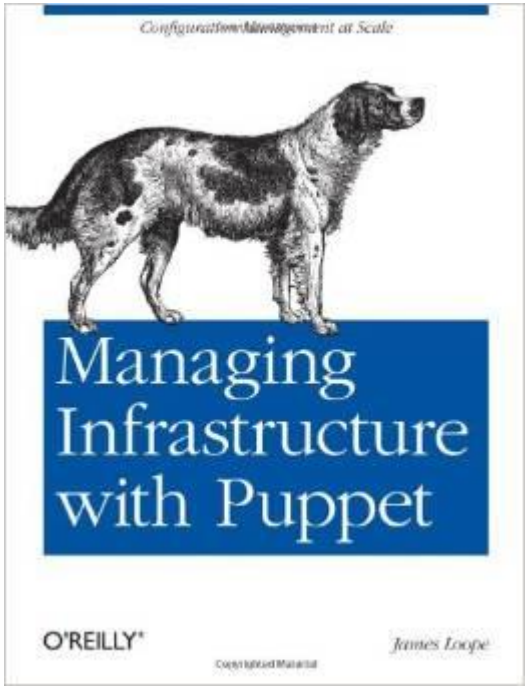

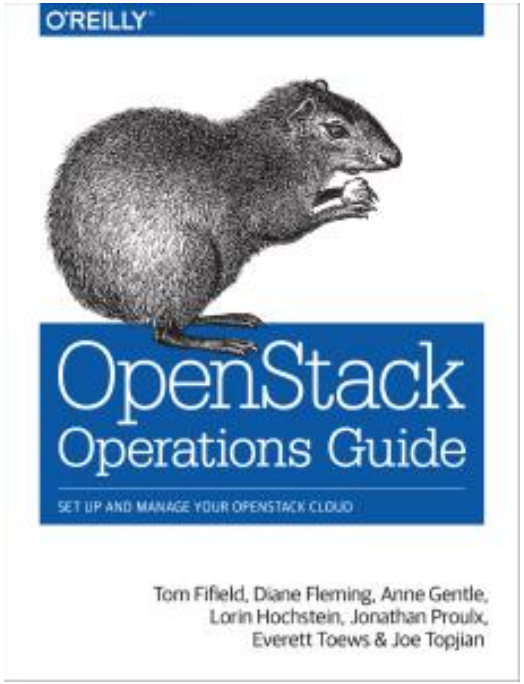

Job Trends Consideration

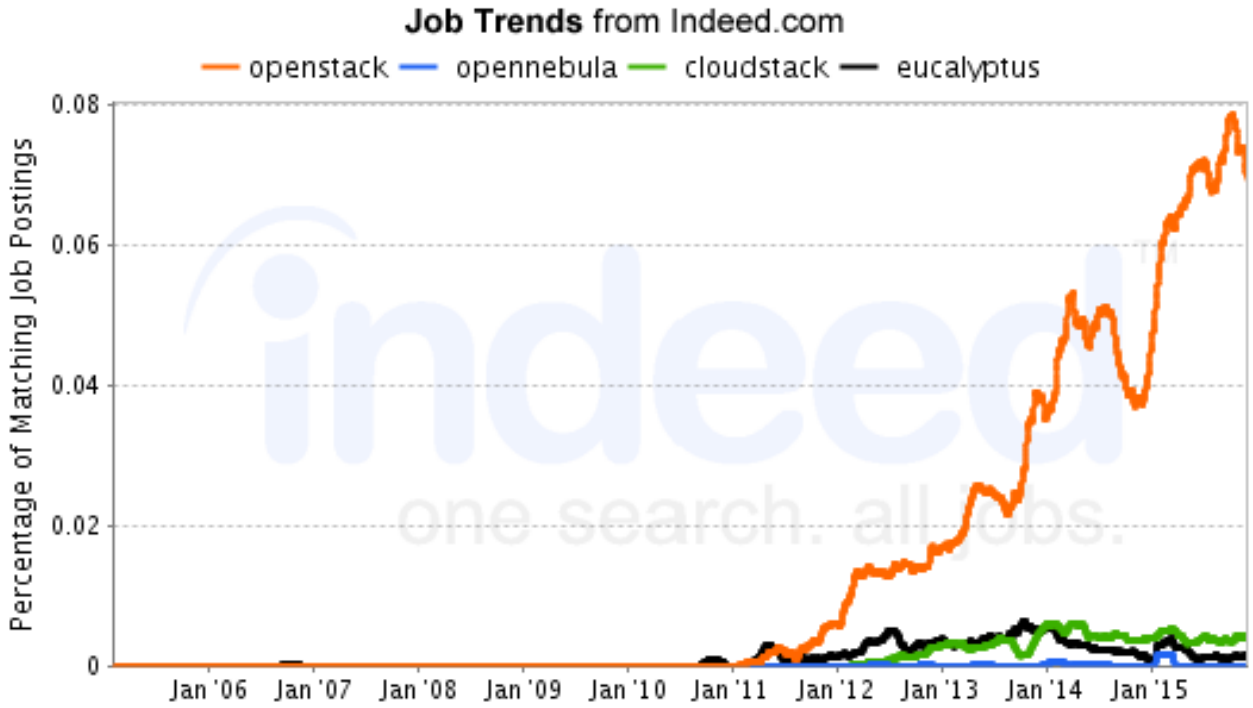

CERN Tool Chain

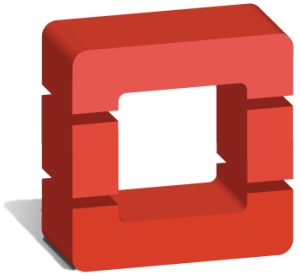

elasticsearch.

openstack™
CLOUD SOFTWARE

hadoop

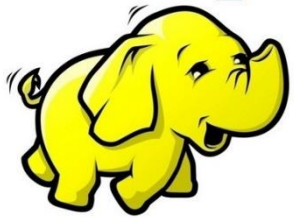

RDO

FOREMAN

RUNDECK

Jenkins

Upstream OpenStack on its own does not give you a cloud service

Packaging

Integration

Burn In

SLA

Monitoring

…

Source: eBay

CERN OpenStack Project

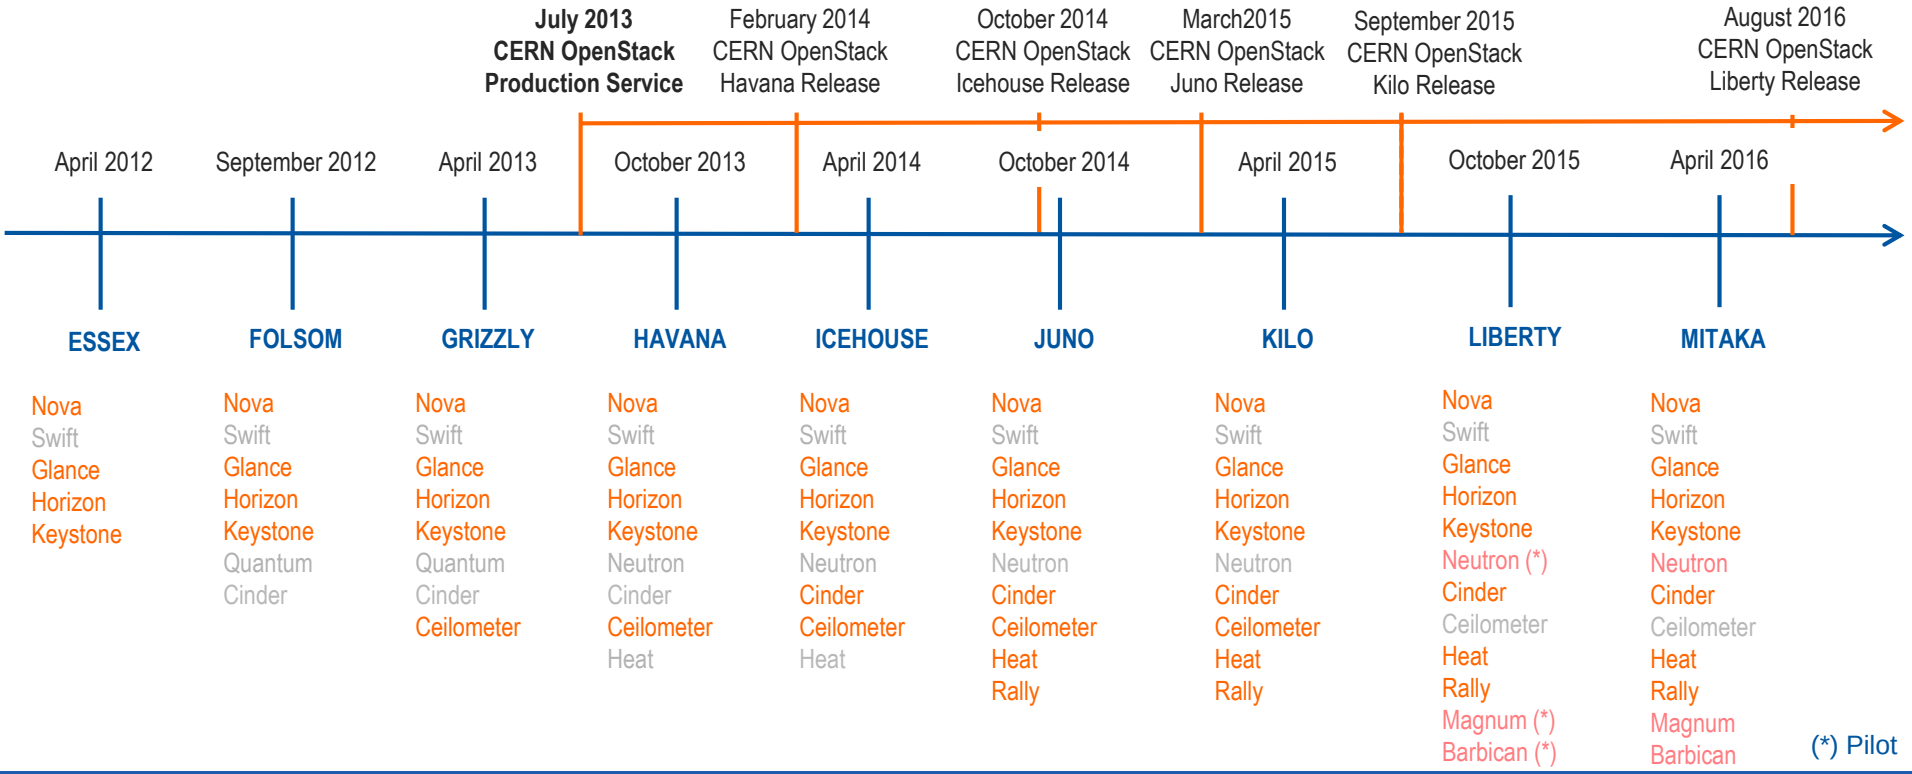

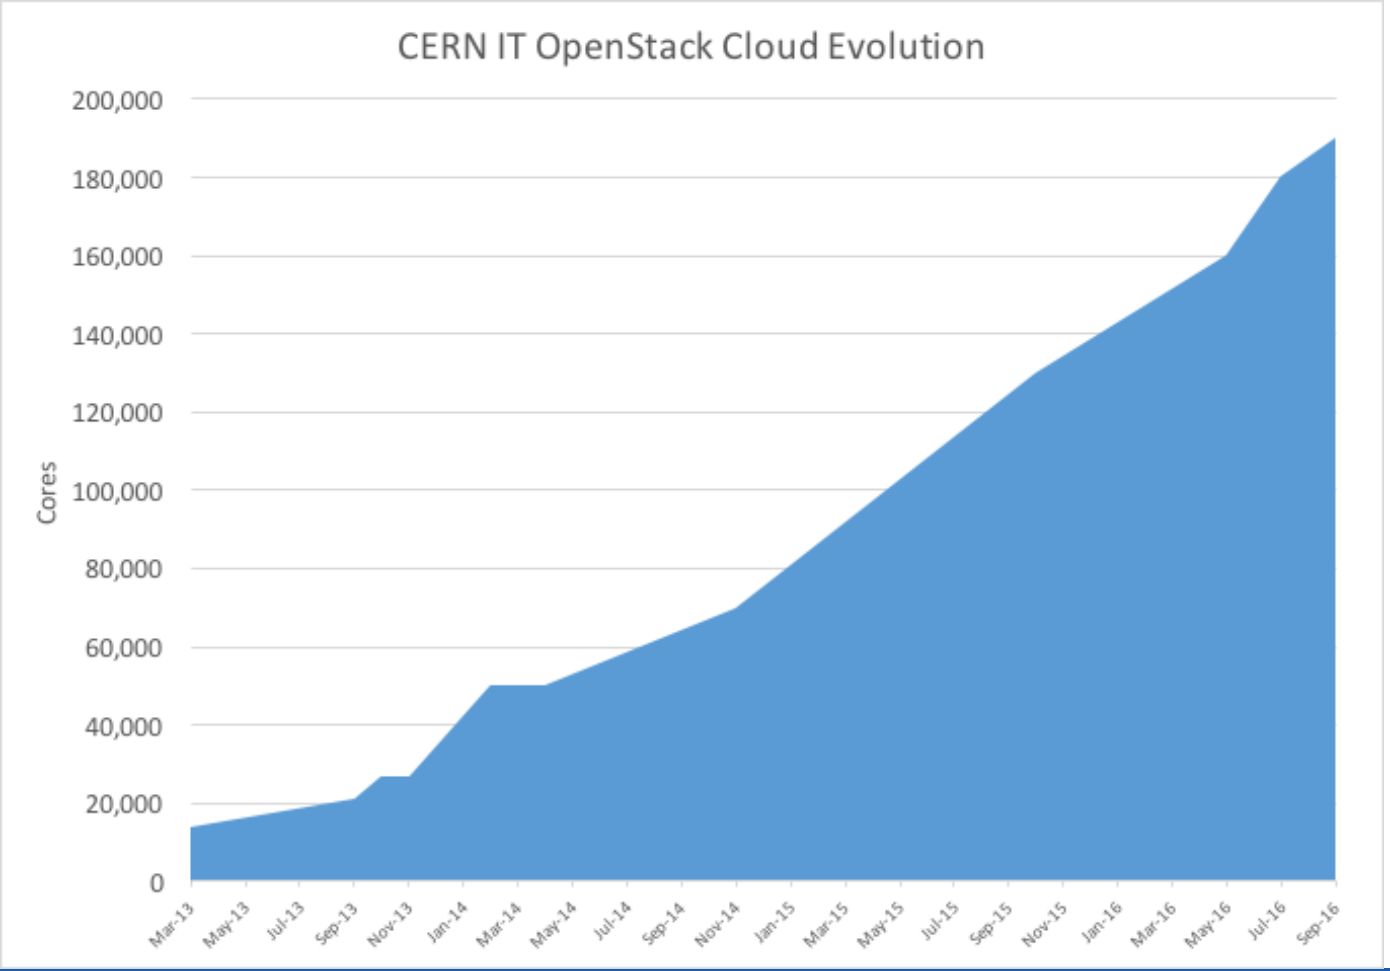

- >190K cores in production under OpenStack
- >90% of CERN compute resources are virtualised
- >5,000 VMs migrated from old hardware in 2016
- >100K cores to be added in the next 6 months

Onwards the Federated Clouds

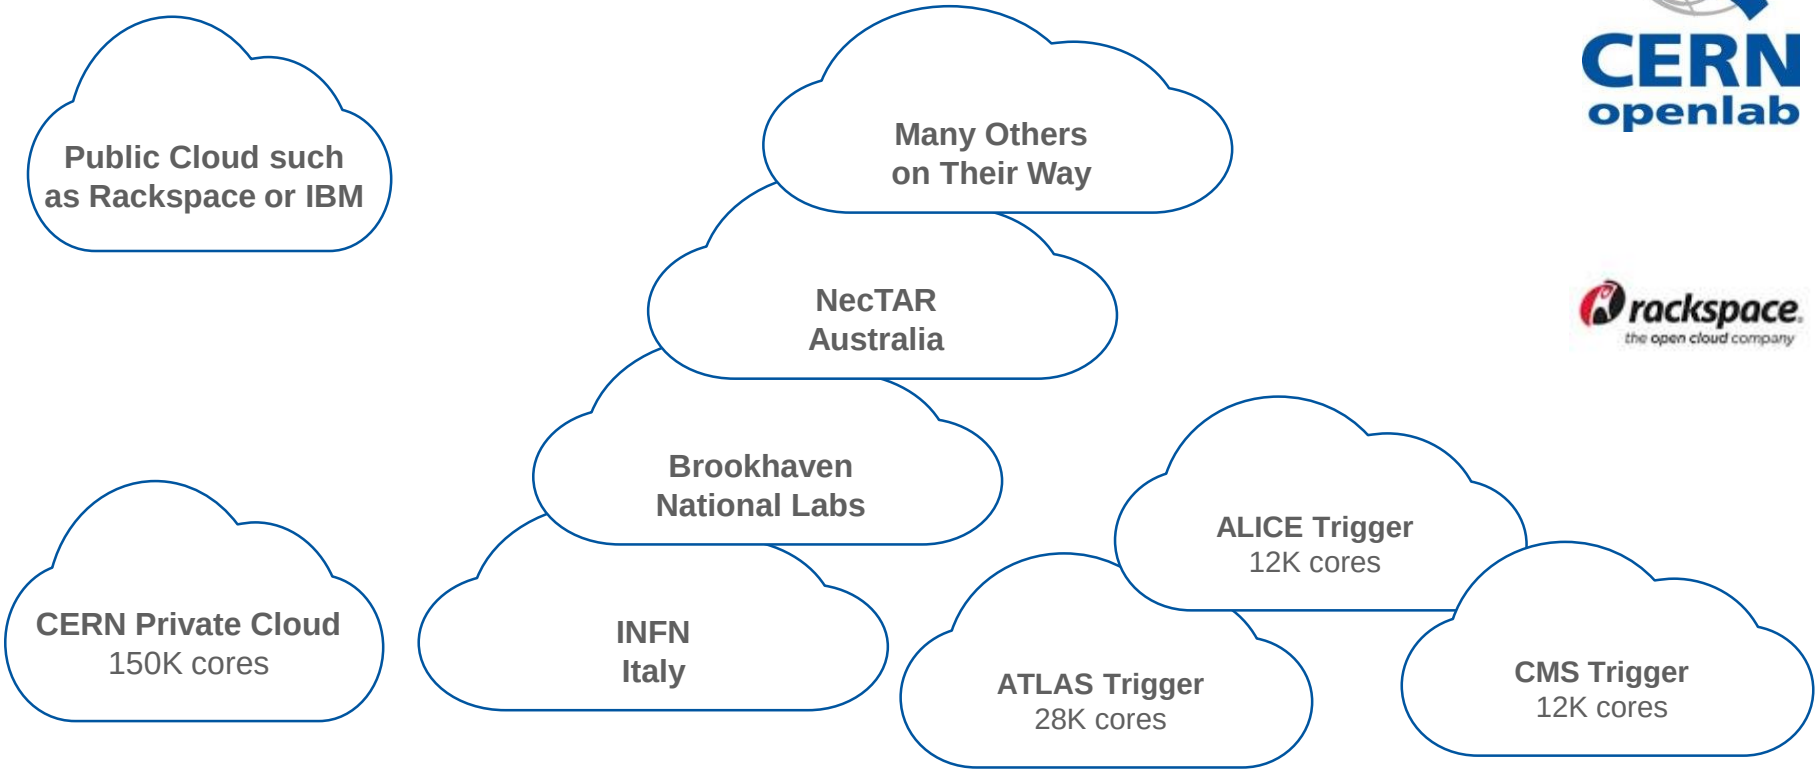

Nova Cells

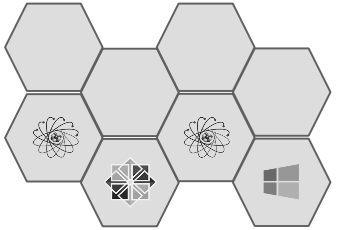

Top level cell

- Runs API service
- Top cell scheduler

Child cells run:

- Compute nodes
- Nova network
- Scheduler
- Conductor

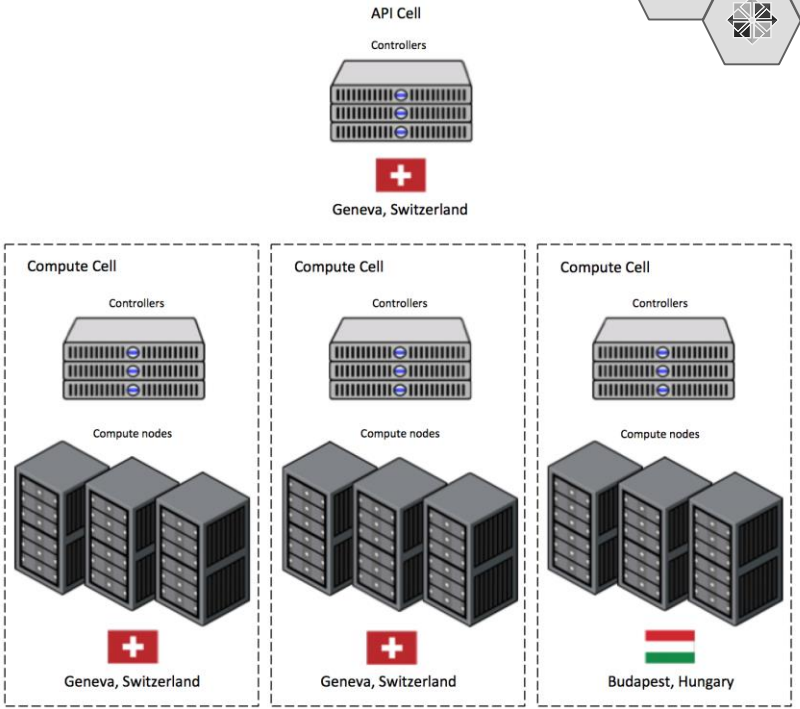

HELIX NEBULA
THESCIENCECLOUD

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

cnes

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

PIC
port d'informació científica

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

Additional Users: ECMWF ESO ESRF ILL NEUTRONS FOR SCIENCE

Suppliers

Atos CLOUDEO CloudSigma Cloudwatt DataCentred DEUTSCHE BÖRSE CLOUD EXCHANGE egi exoscale GÉANT interoute from the ground to the cloud Prologue. sixsq SWITCH T-Systems the SERVER LABS the IT architects ultimum technologies

Adopters

AWST Capgemini CONSULTING.TECHNOLOGY.OUTSOURCING CSA cloud security alliance deac IT outsourcing Ifremer MEMSET HOSTING orange Business Services RHEA terradue 2.0 THALES it Communicating ICT to Markets vito vision on technology VU VRIJE UNIVERSITEIT AMSTERDAM Yandex

Past, ongoing & future commercial activities@CERN

Support CERN's scientific computing programme

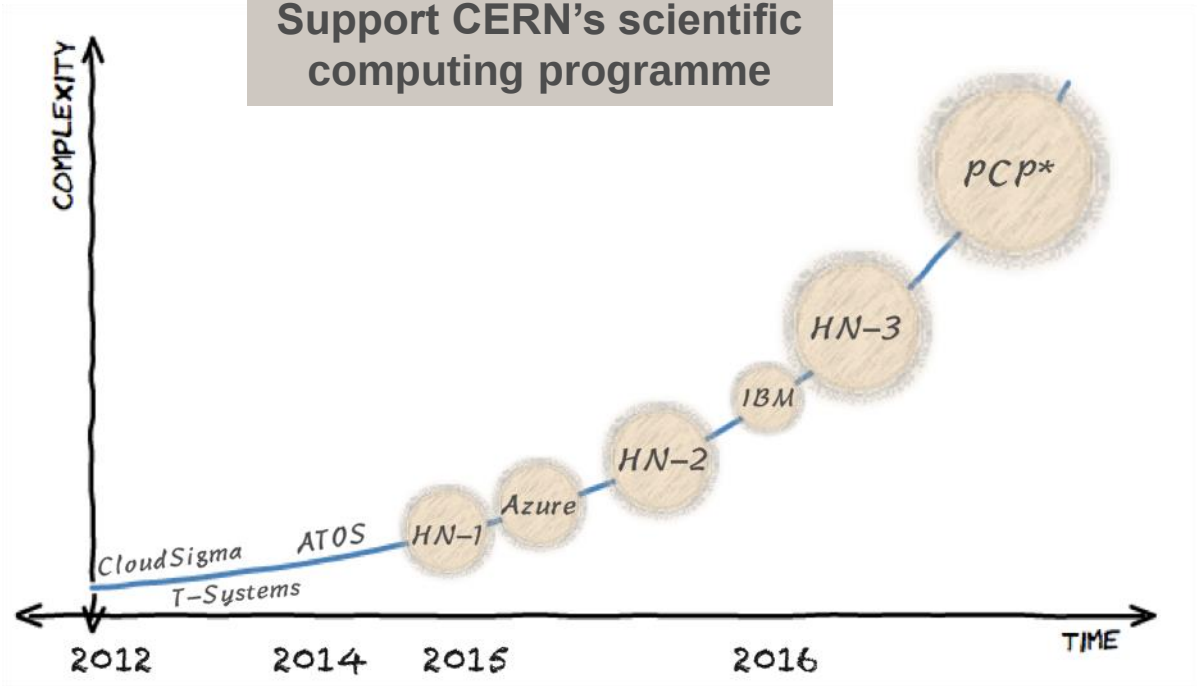

HN - Helix Nebula

- Partnership between research organization and European commercial cloud providers

* EC co-funded joint Pre-Commercial Procurement (PCP) project: https://indico.cern.ch/event/319753

** Other work has been conducted outside CERN, such as the Amazon Pilot project at BNL for ATLAS

Containers on Clouds

For the user

- Interactive
- Dynamic

For IT - Magnum

- Secure
- Managed
- Integrated
- Scalable

For industry and EU

- Prepare for the next disruptive technology at scale

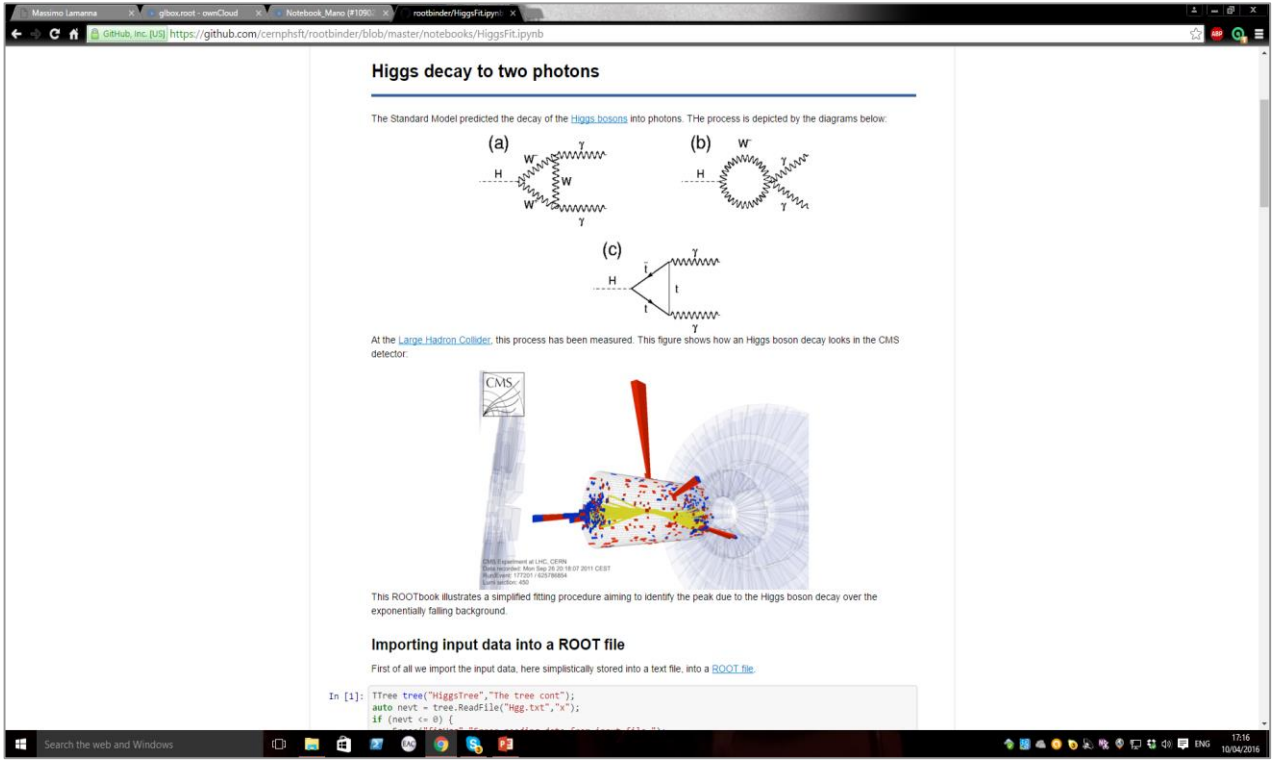

For Further Information

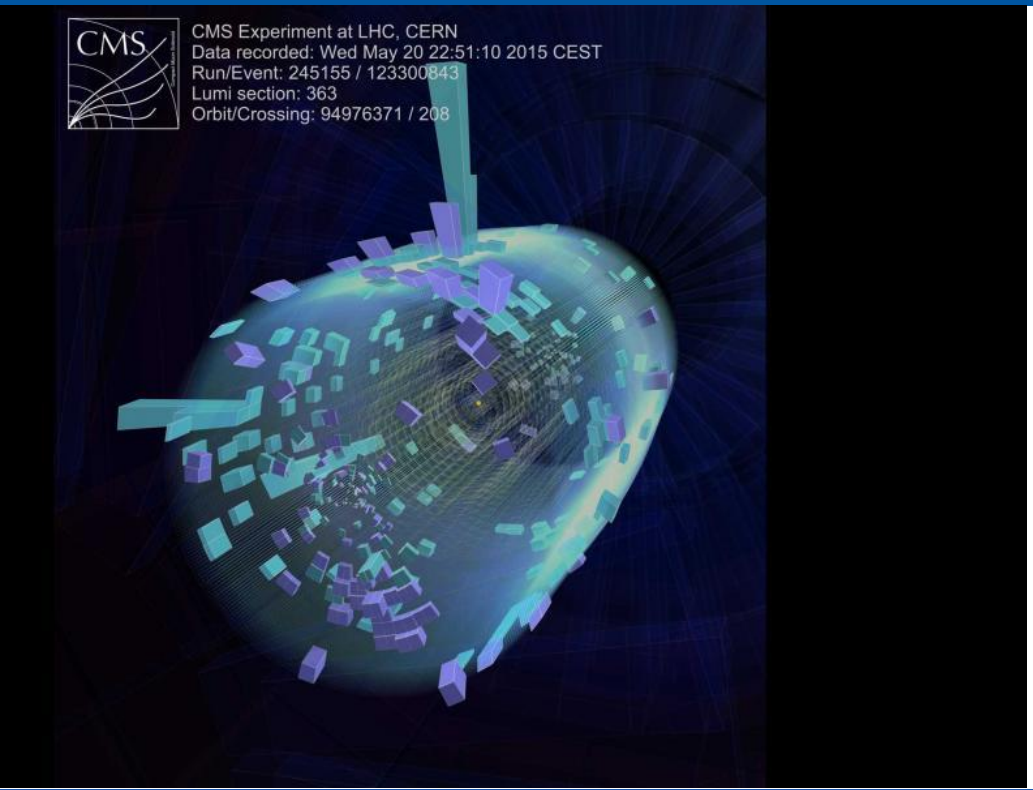

Technical details at http://openstack-in-production.blogspot.fr

Helix Nebula Initiative at http://www.helix-nebula.eu/

Scientific Working Group at https://wiki.openstack.org/wiki/Scientific_working_group

Some history of scale…

Date	Collaboration sizes	Data volume, archive technology
Late 1950's	2-3	Kilobits, notebooks
1960's	10-15	kB, punchcards
1970's	~35	MB, tape
1980's	~100	GB, tape, disk
1990's	~750	TB, tape, disk
2010's	~3000	PB, tape, disk

For comparison:
1990's: Total LEP data set ~few TB
Would fit on 1 tape today

Today: 1 year of LHC data ~27 PB

Innovation Dilemma

- How can we avoid the sustainability trap ?
  - Define requirements
  - No solution available that meets those requirements
  - Develop our own new solution
  - Accumulate technical debt
- How can we learn from others and share ?
  - Find compatible open source communities
  - Contribute back where there is missing functionality
  - Stay mainstream

Are CERN computing needs really special ?

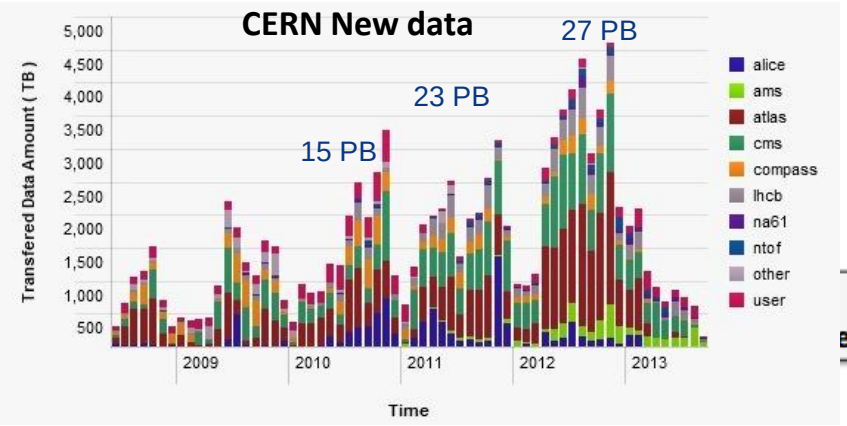

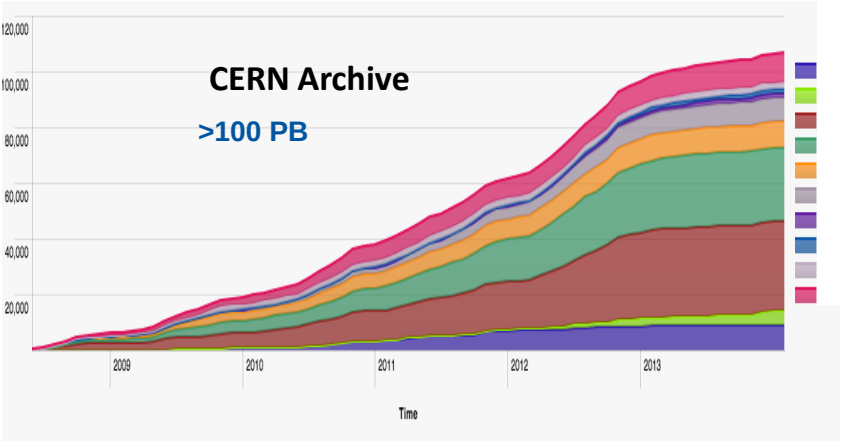

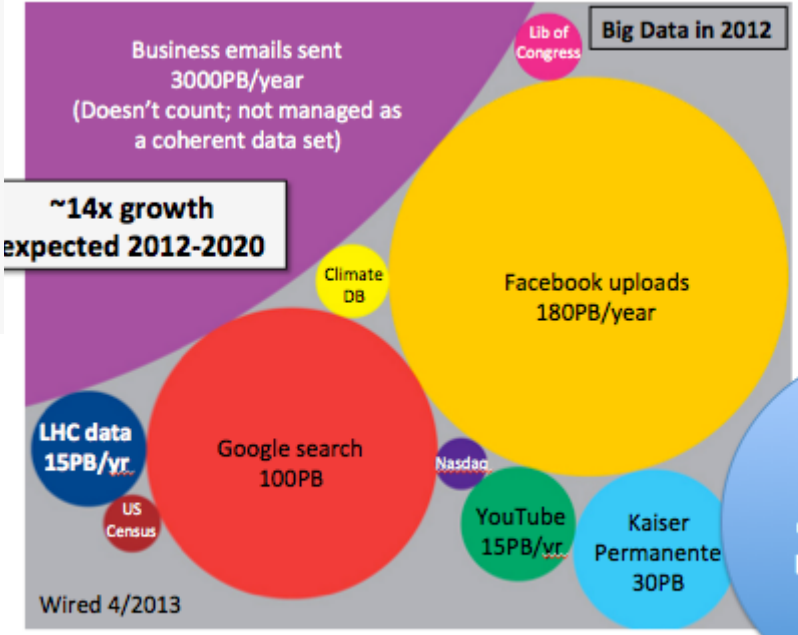

In 2012: 2800 exabytes created or replicated
1 Exabyte = 1000 PB

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

OpenStack Collaborations

- Large Deployment Team
  - Walmart, Yahoo!, Rackspace, eBay, Paypal, …
- Containers
  - Rackspace, Red Hat
- OpenStack Scientific Working Group
  - Not just academic
  - High Performance and High Throughput

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

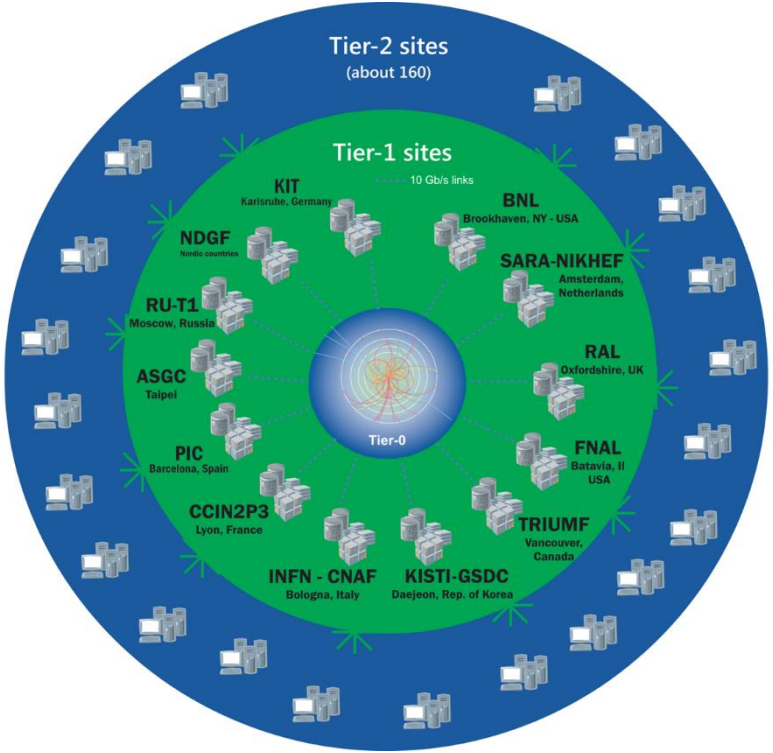

nearly 170 sites, 40 countries

~350’000 cores

500 PB of storage

> 2 million jobs/day

10-100 Gb links

LHC Data Growth

Expecting to record 400PB/year by 2023 with the High Luminosity LHC upgrade

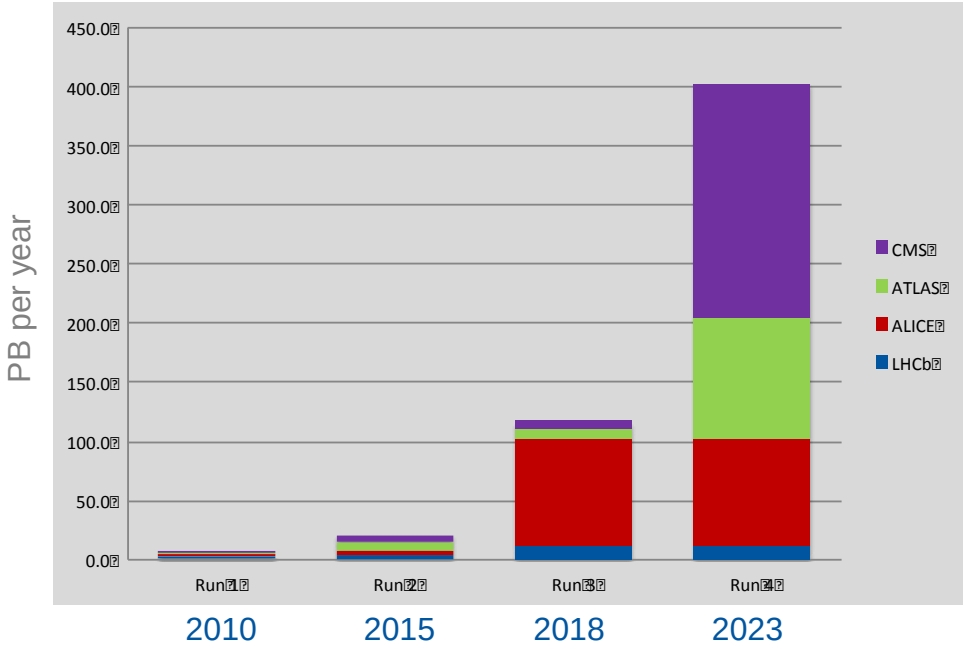

Where is x3 improvement ?

The outline LHC schedule out to 2035 presented by Frederick Bordry to the SPC and FC June 2015 can be found here

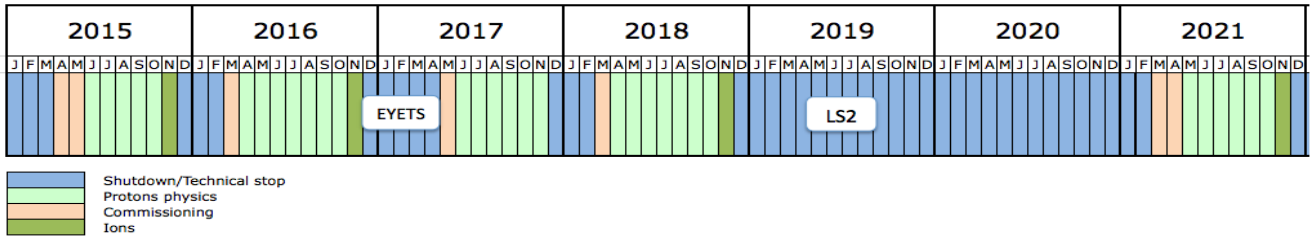

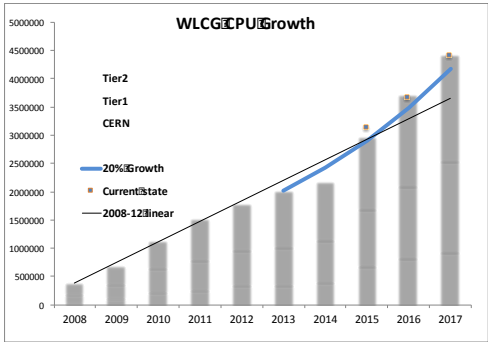

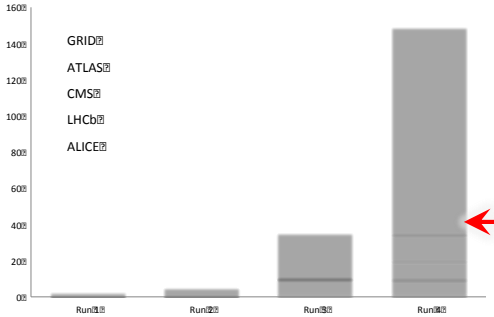

Compute: Growth > x50

What we think is affordable unless we do something differently