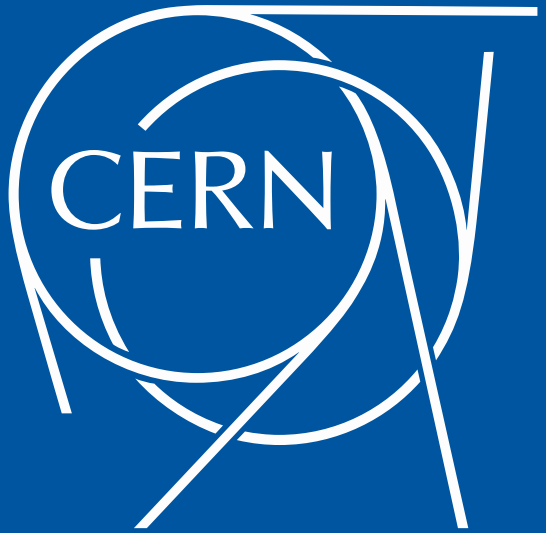
CERN

CERN Computing Infrastructure

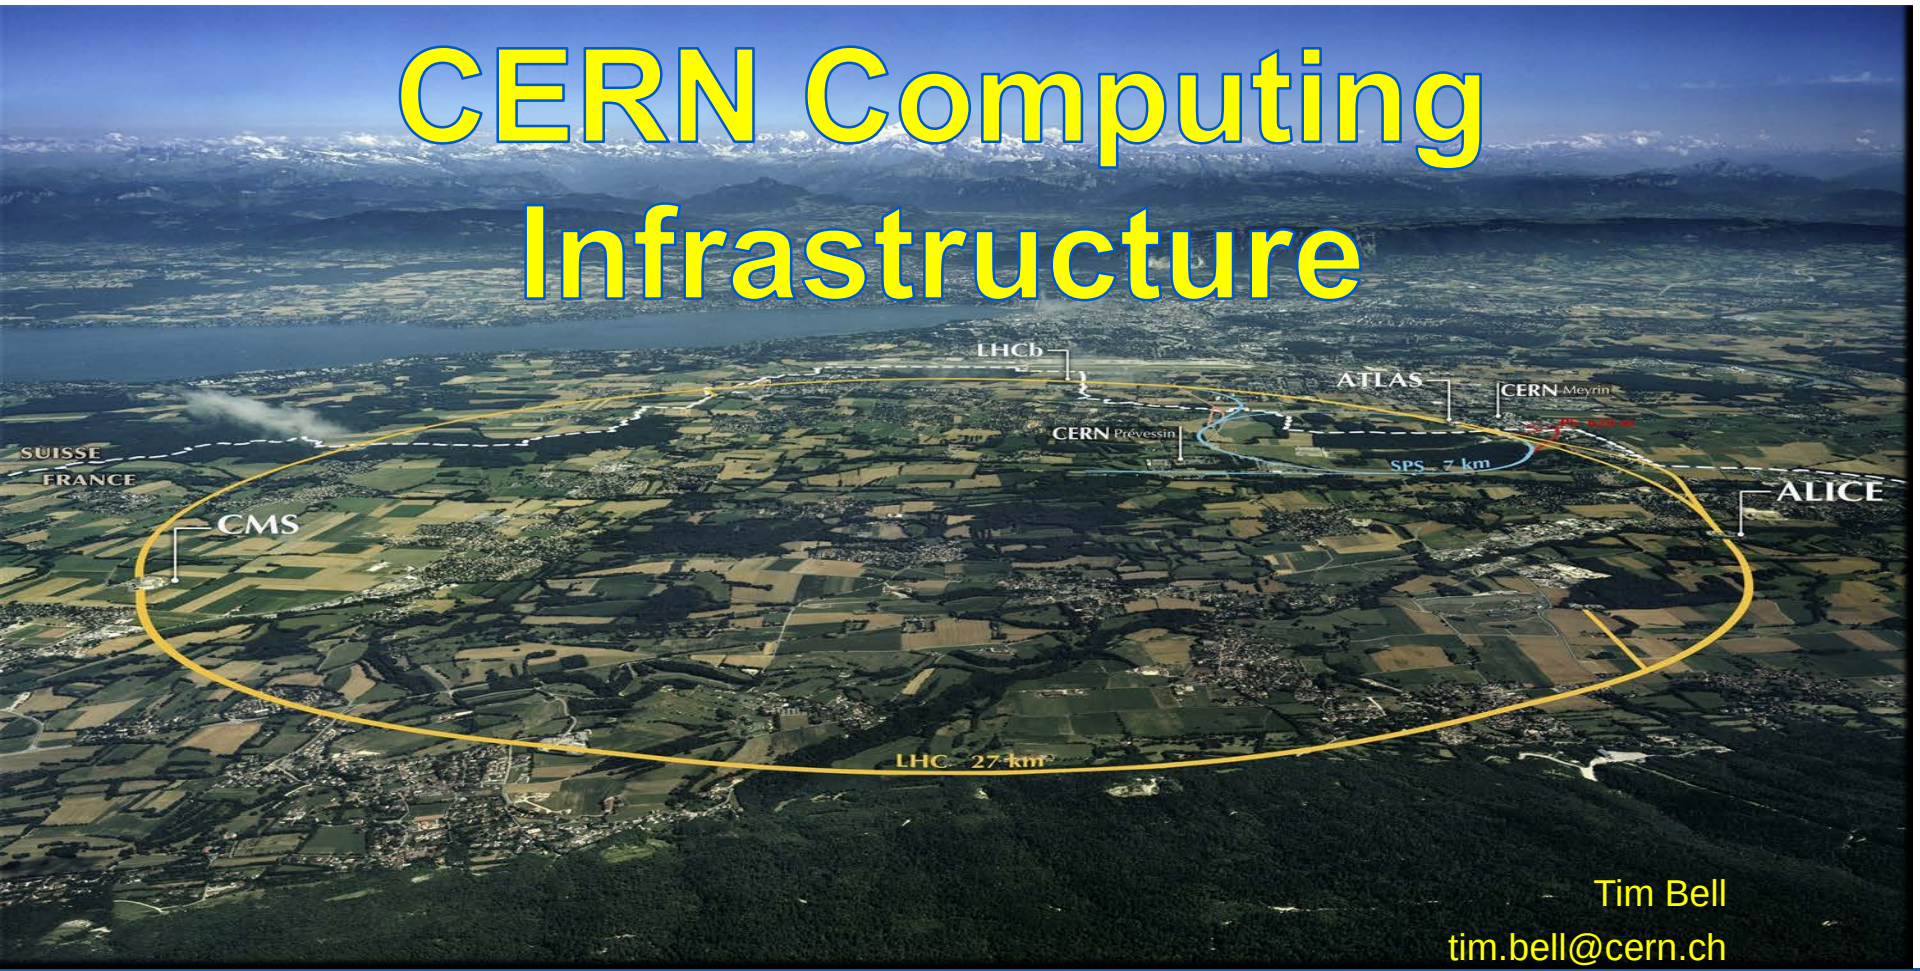

Tim Bell
tim.bell@cern.ch

Computing Infrastructure

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

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**

CERN Tool Chain

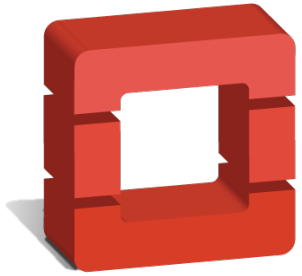

elasticsearch.

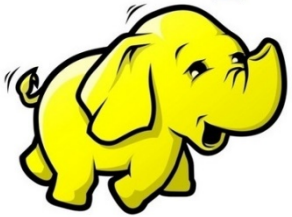

OpenStack Status

- 4 OpenStack clouds at CERN
  - Largest is ~155,000 cores in ~5,000 servers with 60,000 to be installed for 1H 2016
  - 3 other instances with 45,000 cores total
- Collaborating with companies at every 6 month open design summits
  - Austin summit in 2016 with ~7,800 attendees
  - Share experiences and design next release
  - Operations meetups such as large deployments and scientific working groups
  - All CERN code of interest is contributed upstream

Onwards the Federated Clouds

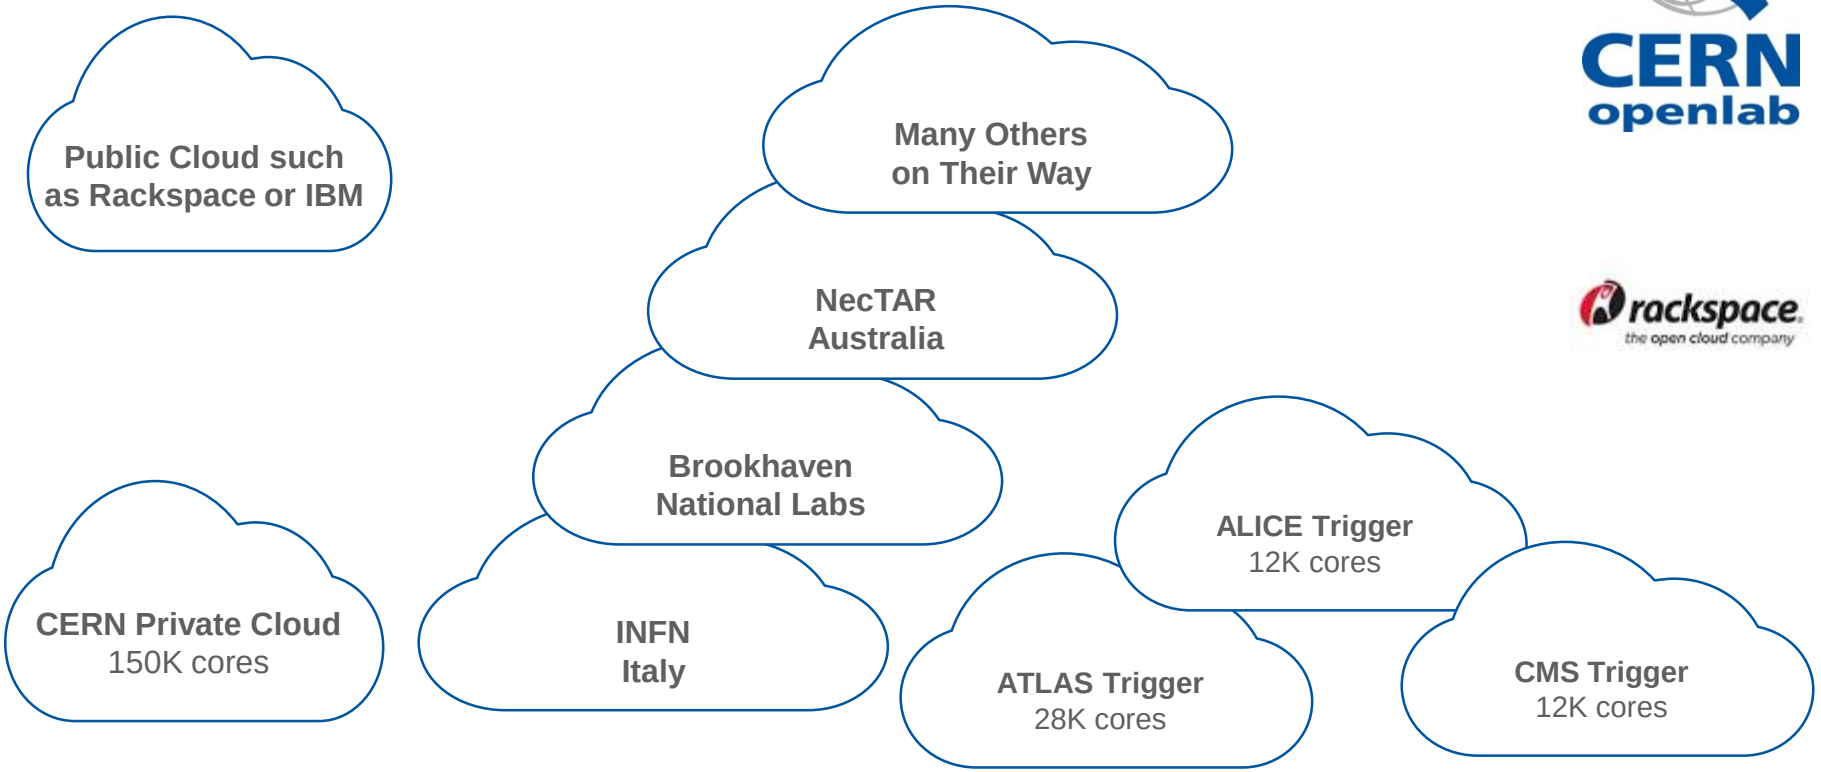

HELIX NEBULA THESCIENCECLOUD

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, Prologue, sixsq, SWITCH, T-Systems, the SERVER LABS, ultimum technologies

Adopters

AWST, Capgemini, CSA cloud security alliance, deac, Ifremer, MEMSET, orange Business Services, RHEA, terradue 2.0, THALES, vito, VU, Yandex

Past, ongoing & future commercial activities@CERN

Support CERN's scientific computing programme

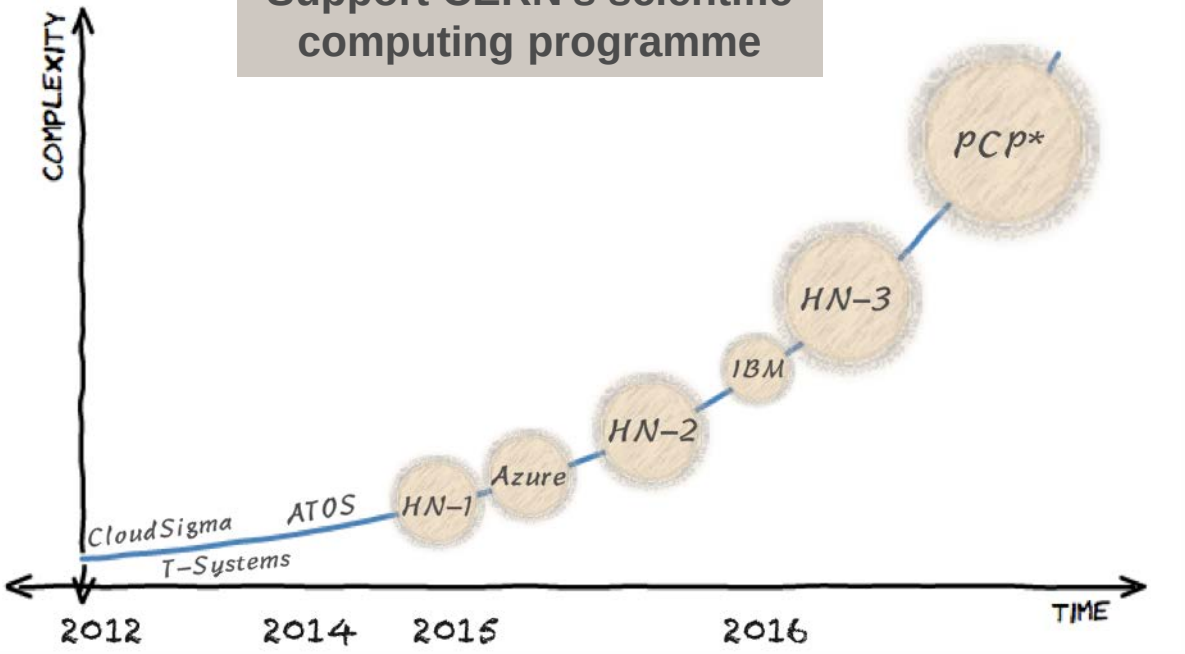

HN - Helix Nebula

- 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

Batch - Present

- Currently running Platform LSF from IBM
  - Around 45,000 running jobs
  - 400,000 jobs/day
  - ~5,000 servers, 90% virtualised
- Grid and local submission

Batch - Future

- Looking to move to HTCondor
  - Improved scalability
  - Open source
  - Good collaboration with development teams
  - Better at handling dynamic resources
- Grid resources now available and local submission pilot starting
- Aim to migrate before end of LS2

Containers

- A potentially disruptive technology
  - Microservices
  - Packaging all software in an object
- New Openlab collaboration project funded by Rackspace
  - Enhance OpenStack container service for scientific use cases
- Evaluating HTCondor container support

Volunteer / Opportunistic

- Non-pledged scavenged resources
  - Volunteers (e.g. home PCs)
  - Institute desktops
  - Small farms with easy deployment
- Unpredictable but significant resources
  - Target CPU bound simulations
- Outreach benefits

ATLAS@HOME

BOINC ATLAS contribution

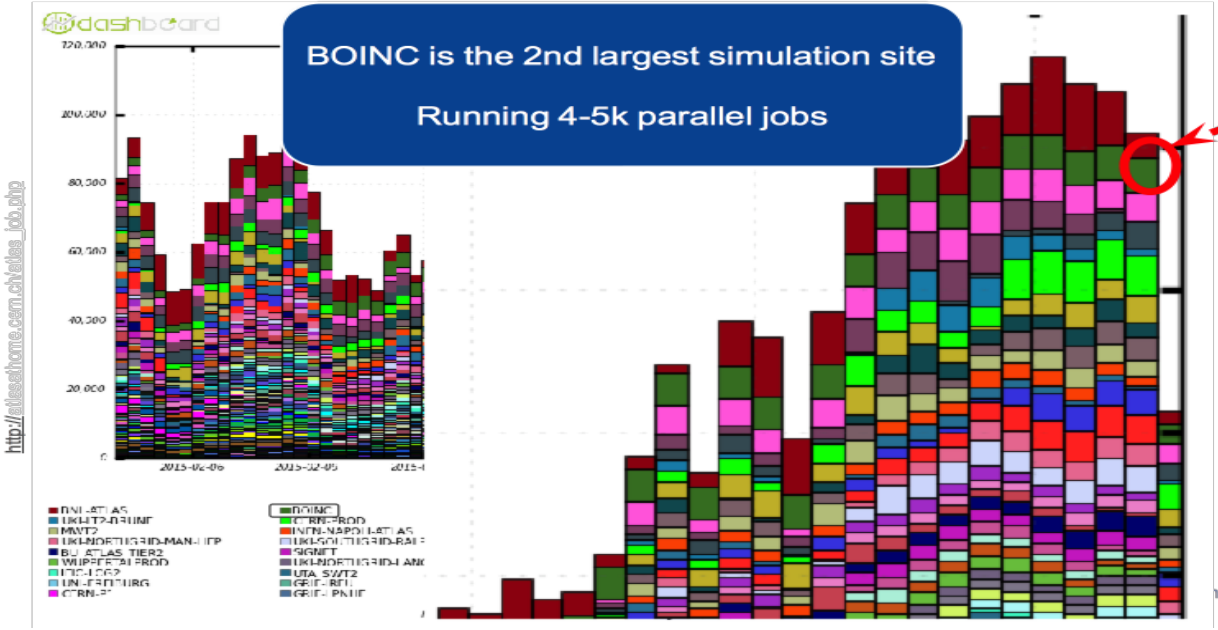

Multiple Projects Running

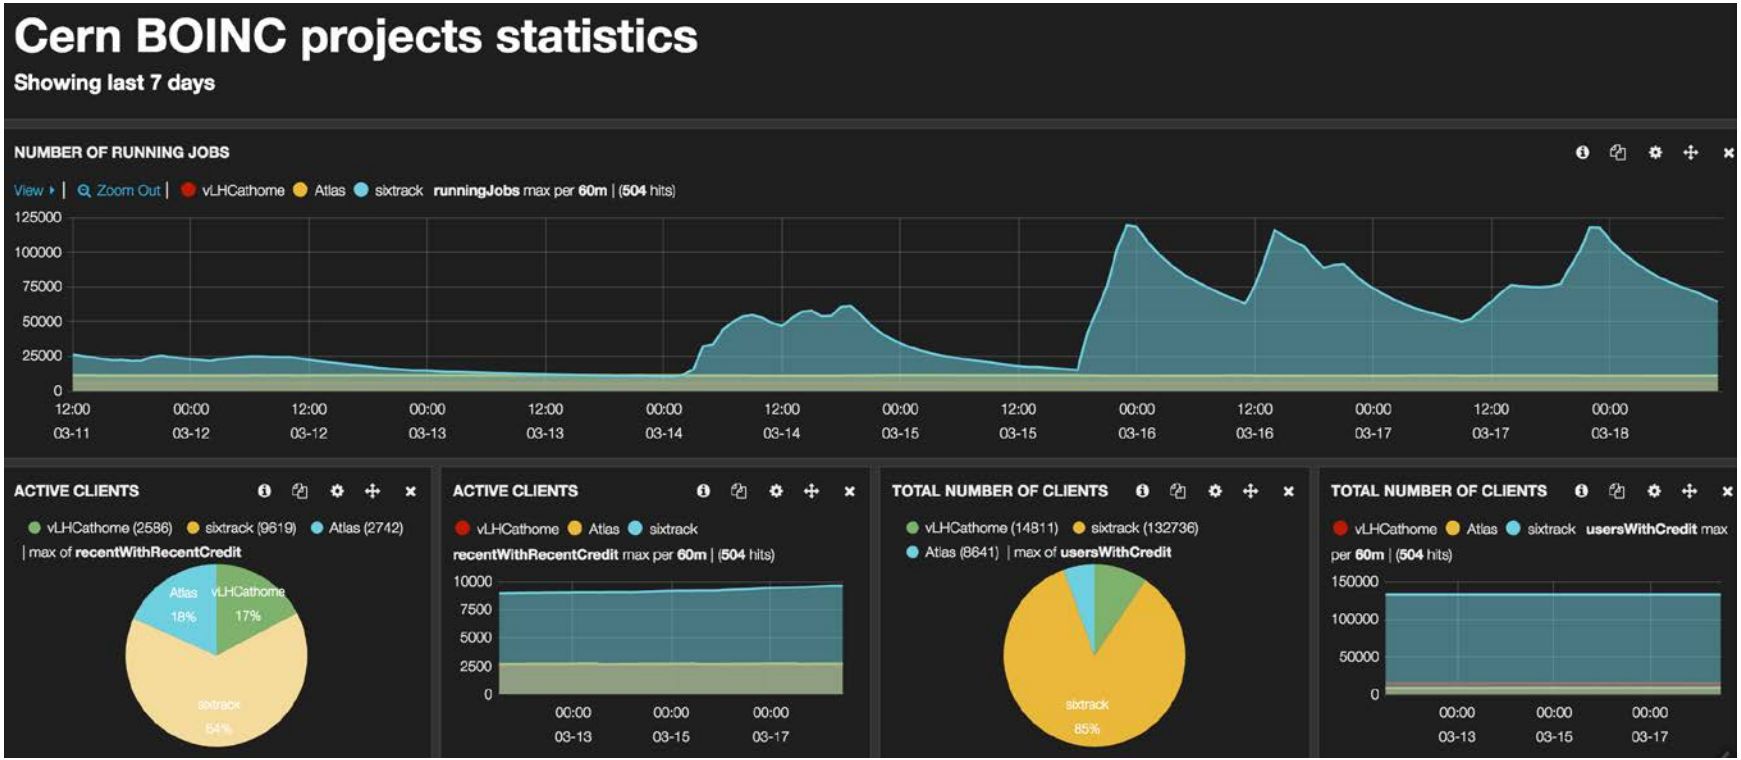

Summary

- Computing infrastructure is rapidly evolving
  - Clouds
  - Batch
  - Opportunistic
- Challenges likely to continue in
  - Sustainability
  - Scale
  - Efficiency

For Further Information

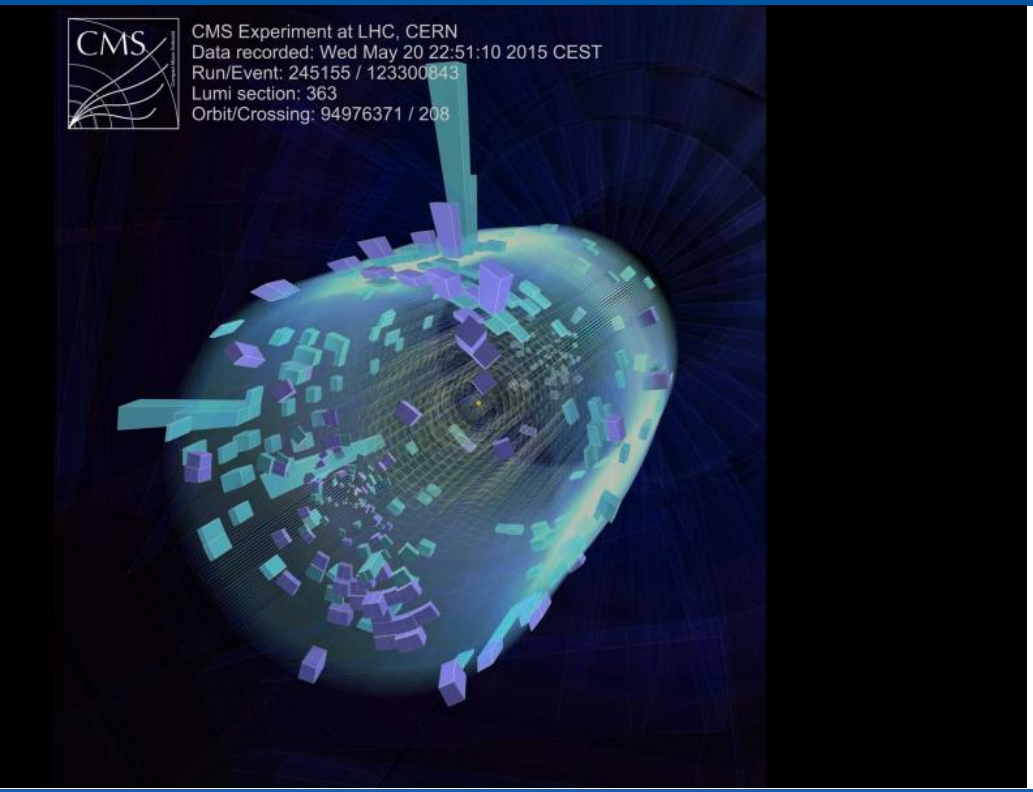

Technical details at http://openstack-in-production.blogspot.fr

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 ?

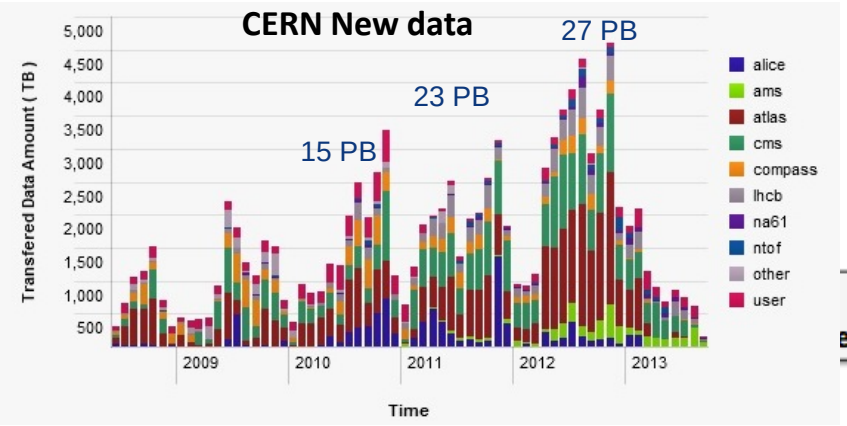

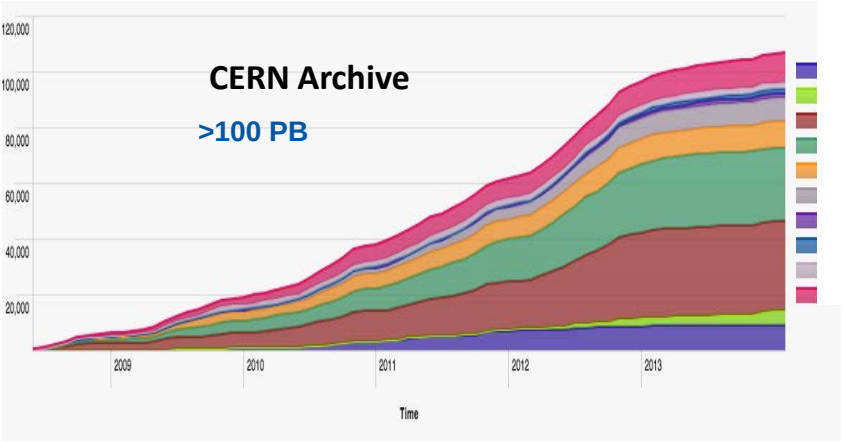

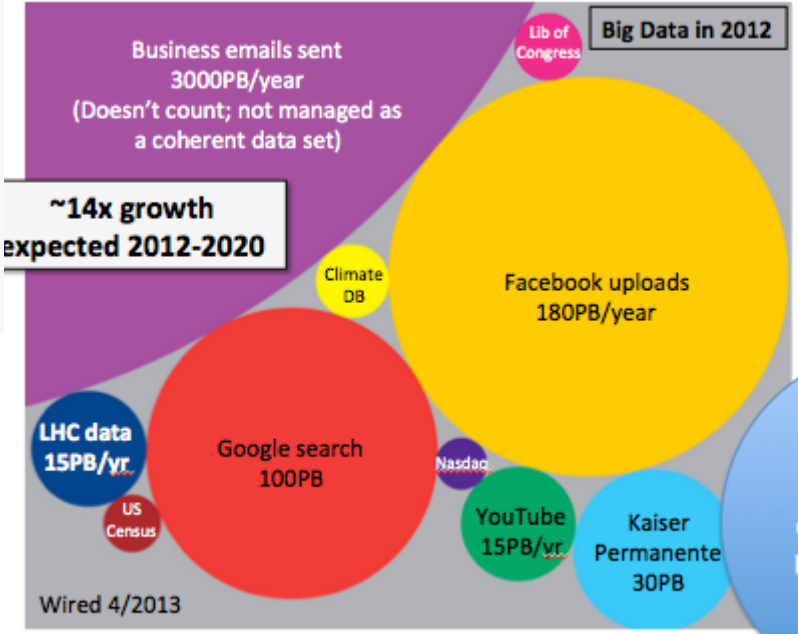

In 2012: 2800 exabytes created or replicated
1 Exabyte = 1000 PB

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

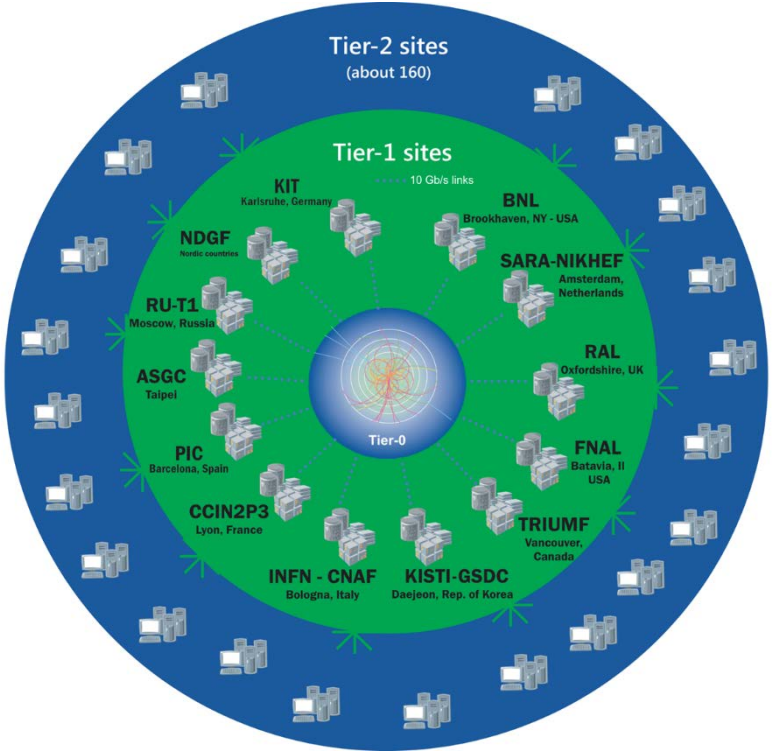

nearly 170 sites, 40 countries

~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

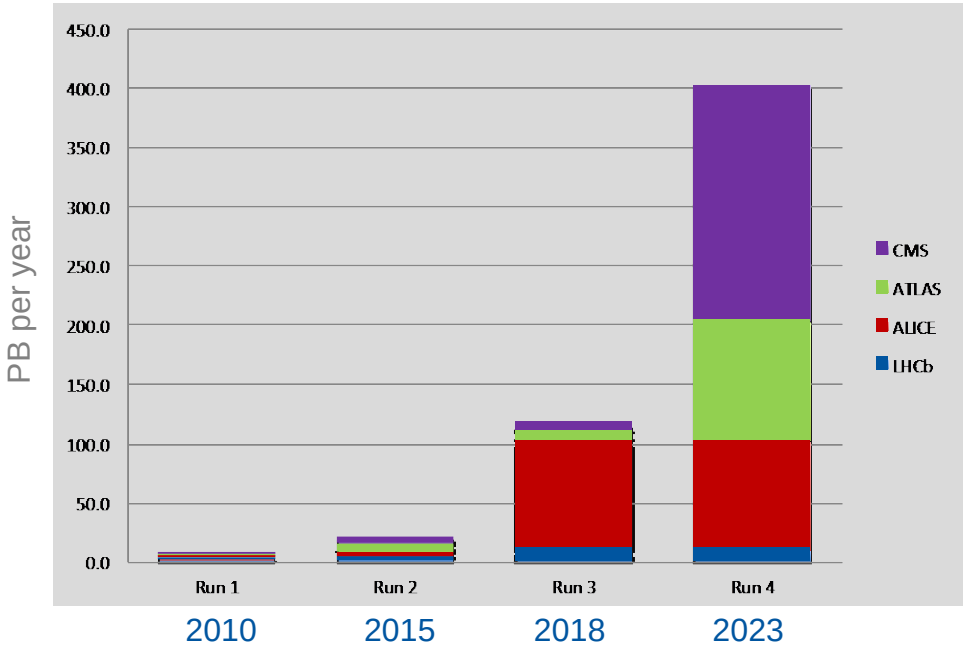

Where is x3 improvement ?

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

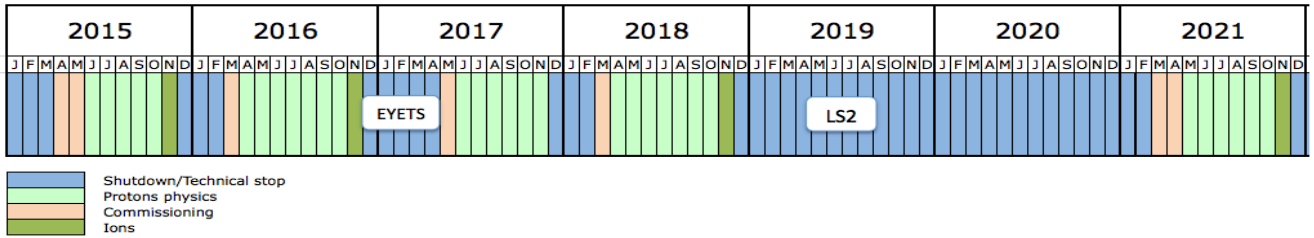

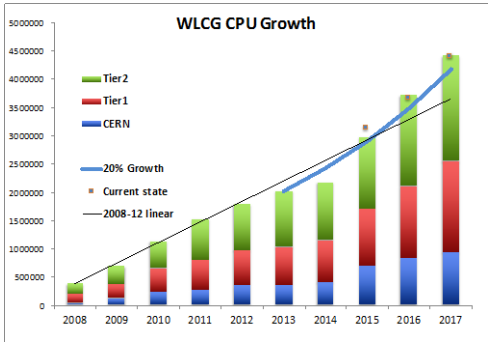

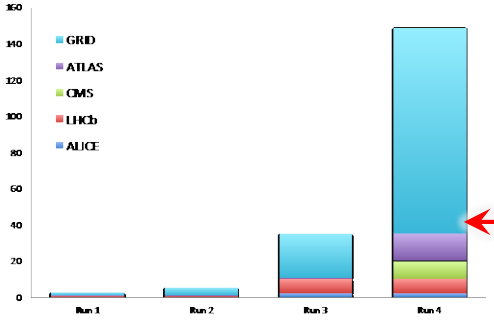

Compute: Growth > x50

What we think is affordable unless we do something differently

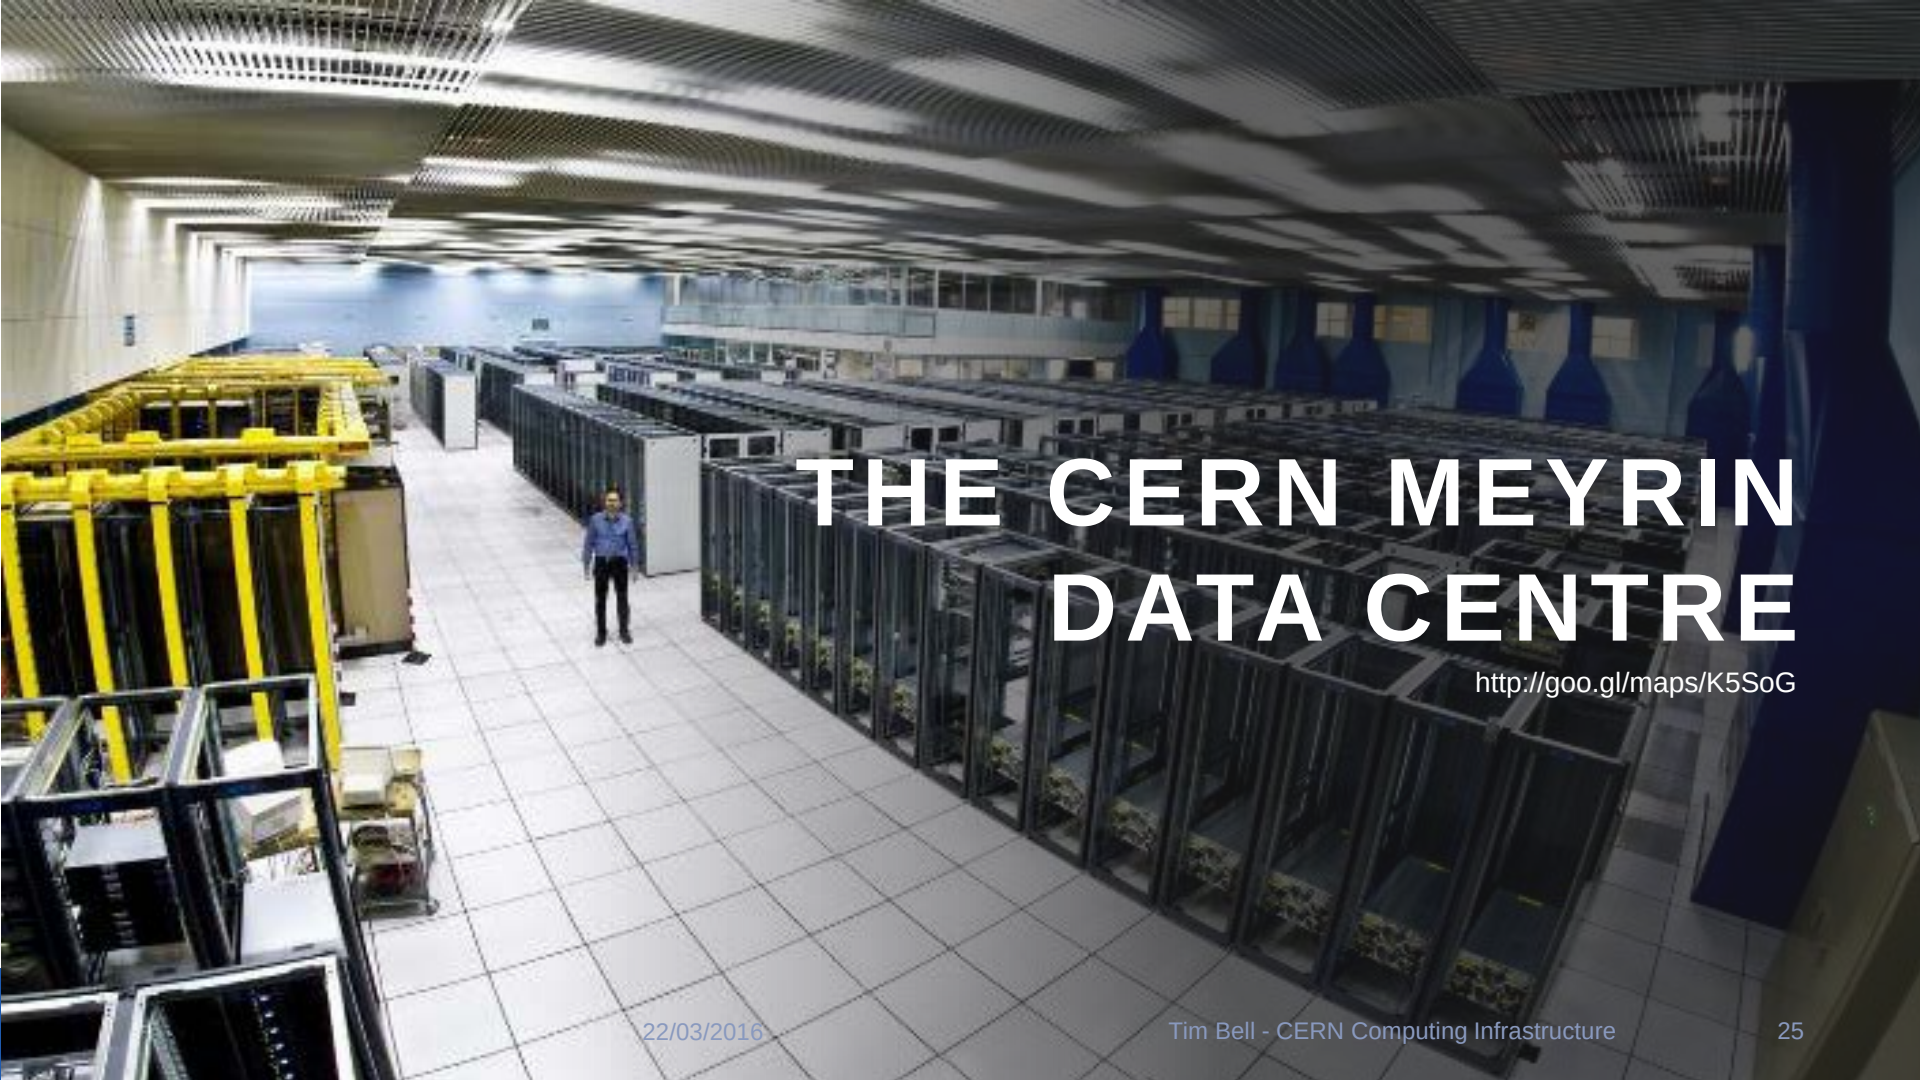

THE CERN MEYRIN DATA CENTRE

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
Luxembourg
Luxembourg
Prague
Czech Republic
Krakow
Paris
Slovakia
Munich
Vienna
Bratislava
Budapest
Austria
Hungary
Zurich
Liechtenstein
Berne
Switzerland
France
Slovenia
Zagreb
Croatia
Vojvodina
Veneto
Milan
Torino
Belgrade

O’Reilly Consideration

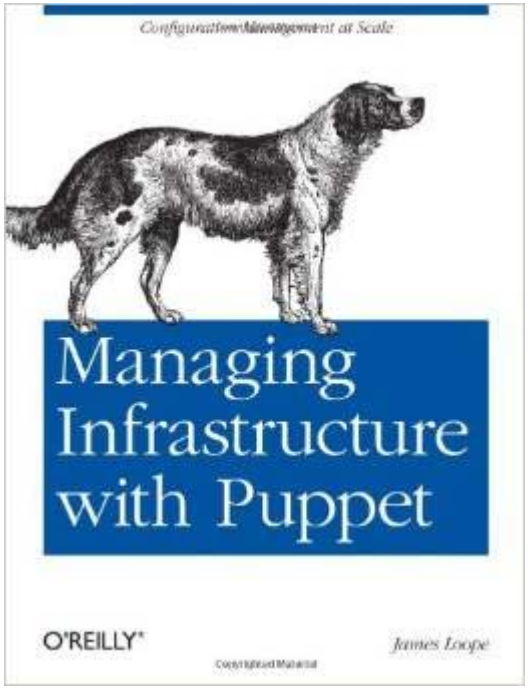

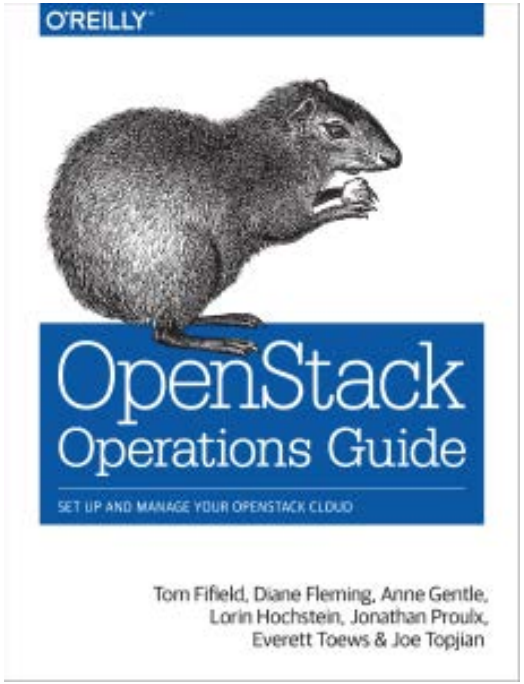

Job Trends Consideration

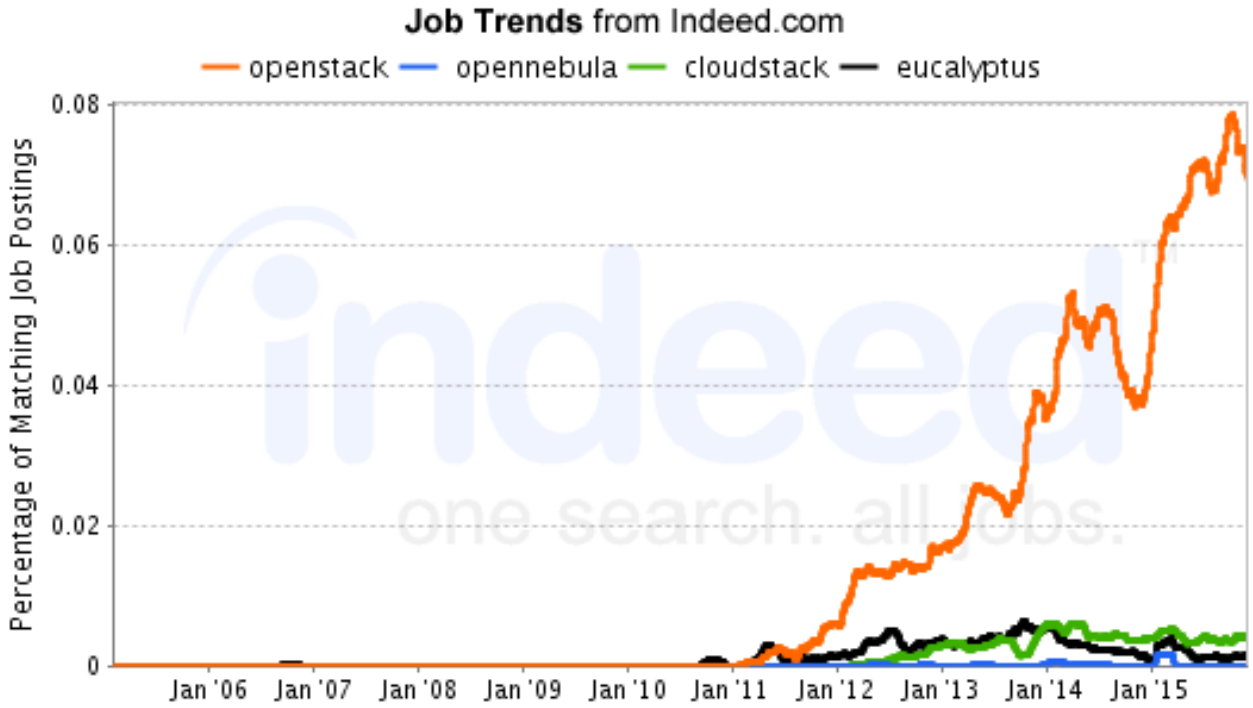

Upstream OpenStack on its own does not give you a cloud service

Packaging

Integration

Burn In

SLA

Monitoring

…

Source: eBay