
CERN

Cloud Trends and Challenges

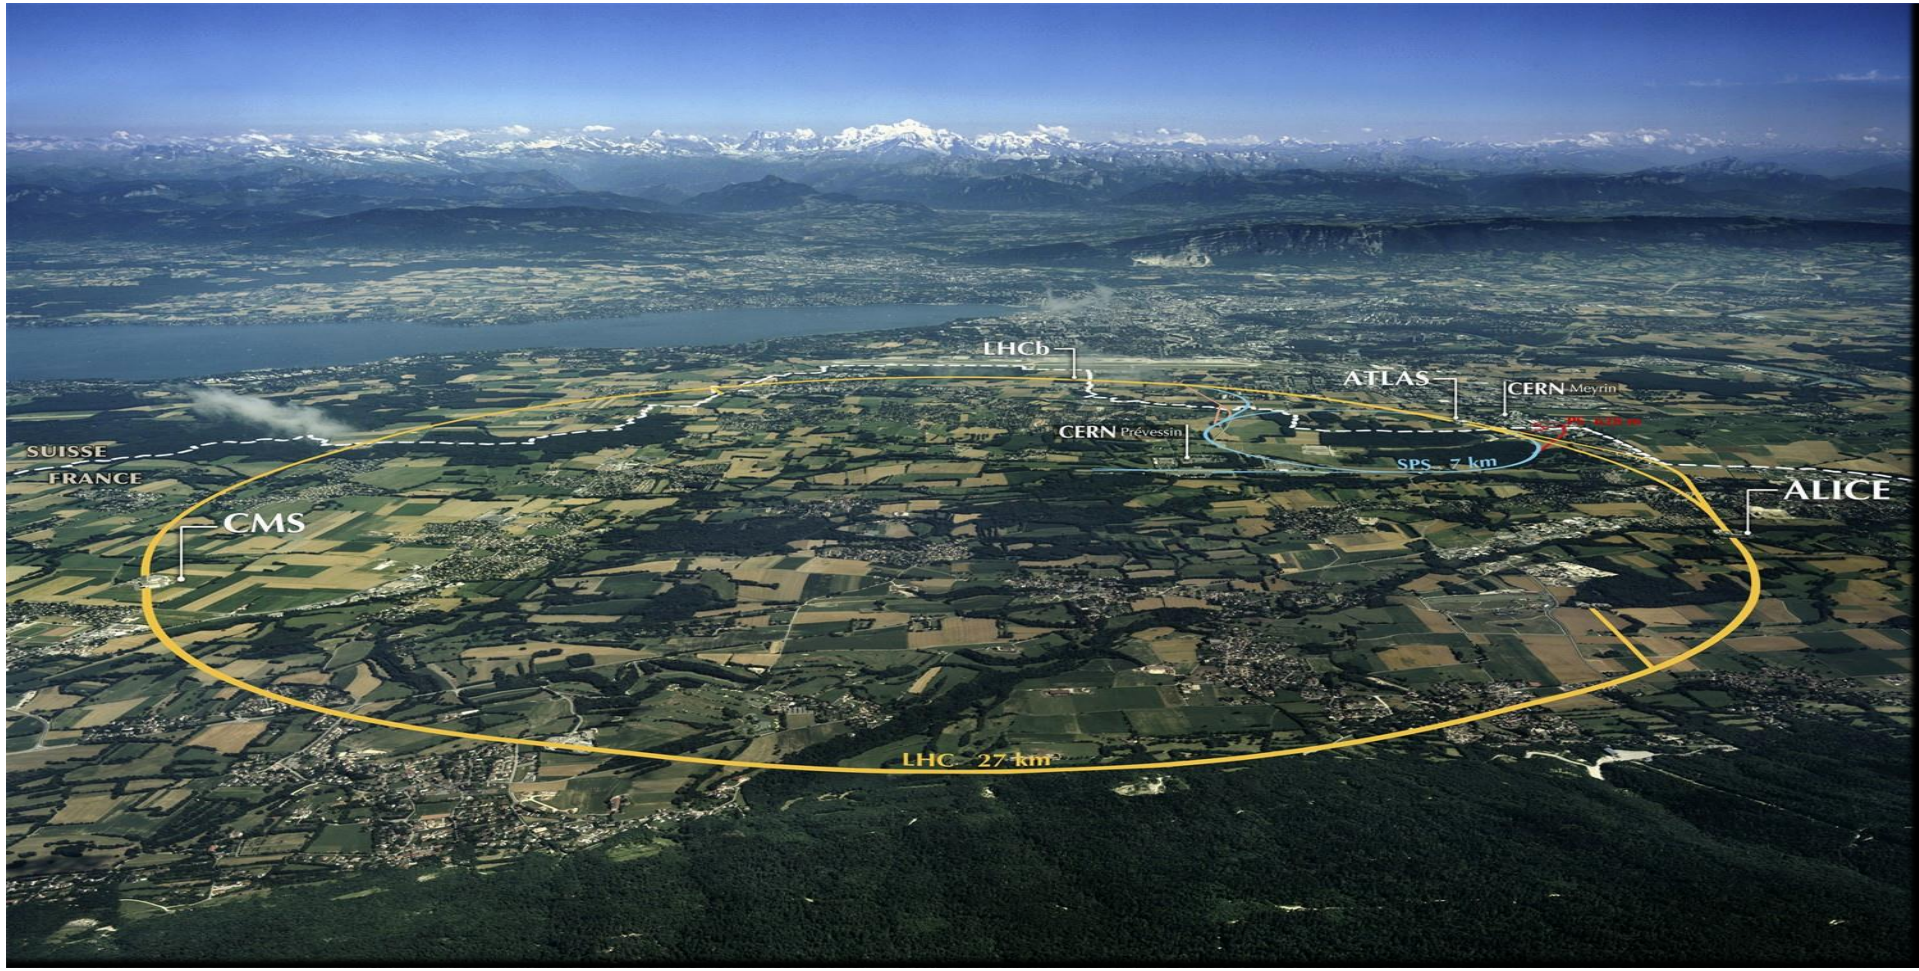

Tim Bell – tim.bell@cern.ch

Compute Growth Outlook

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

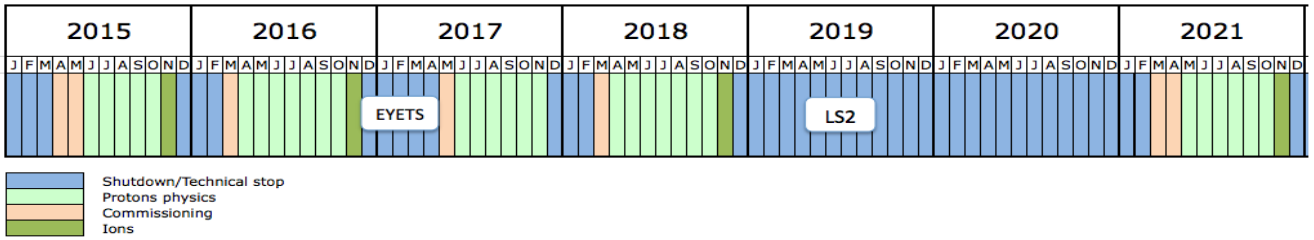

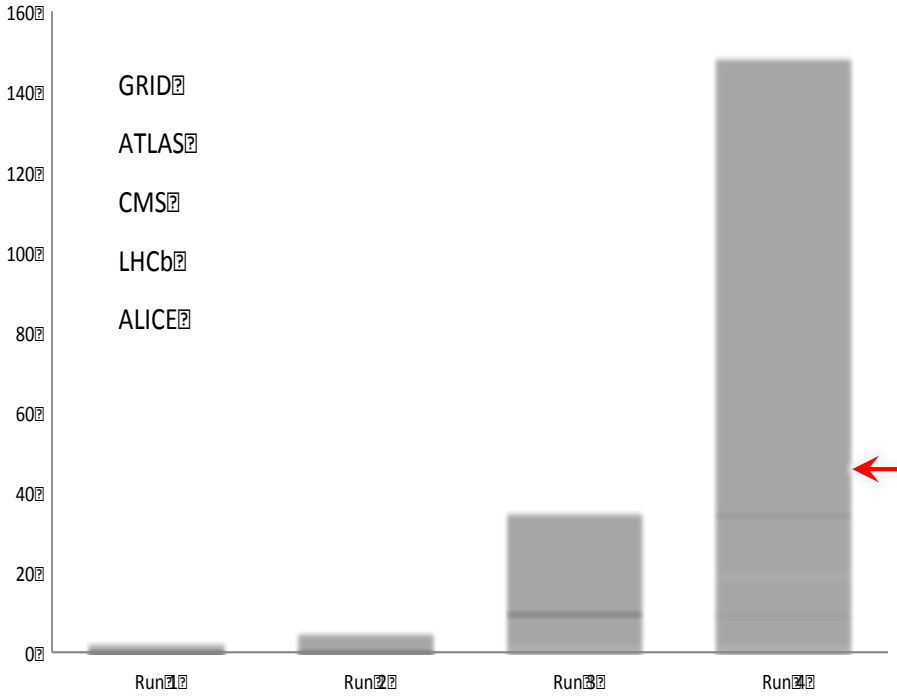

Compute: Growth > x50
Moore's law only x16

← What we can afford

… and 400PB/year by 2023

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

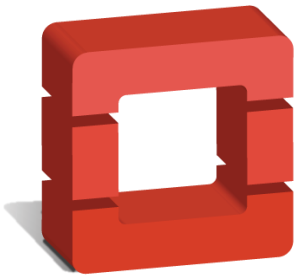

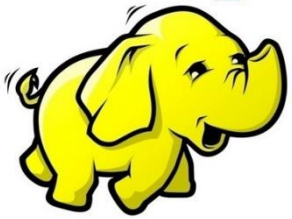

CERN OpenStack Project

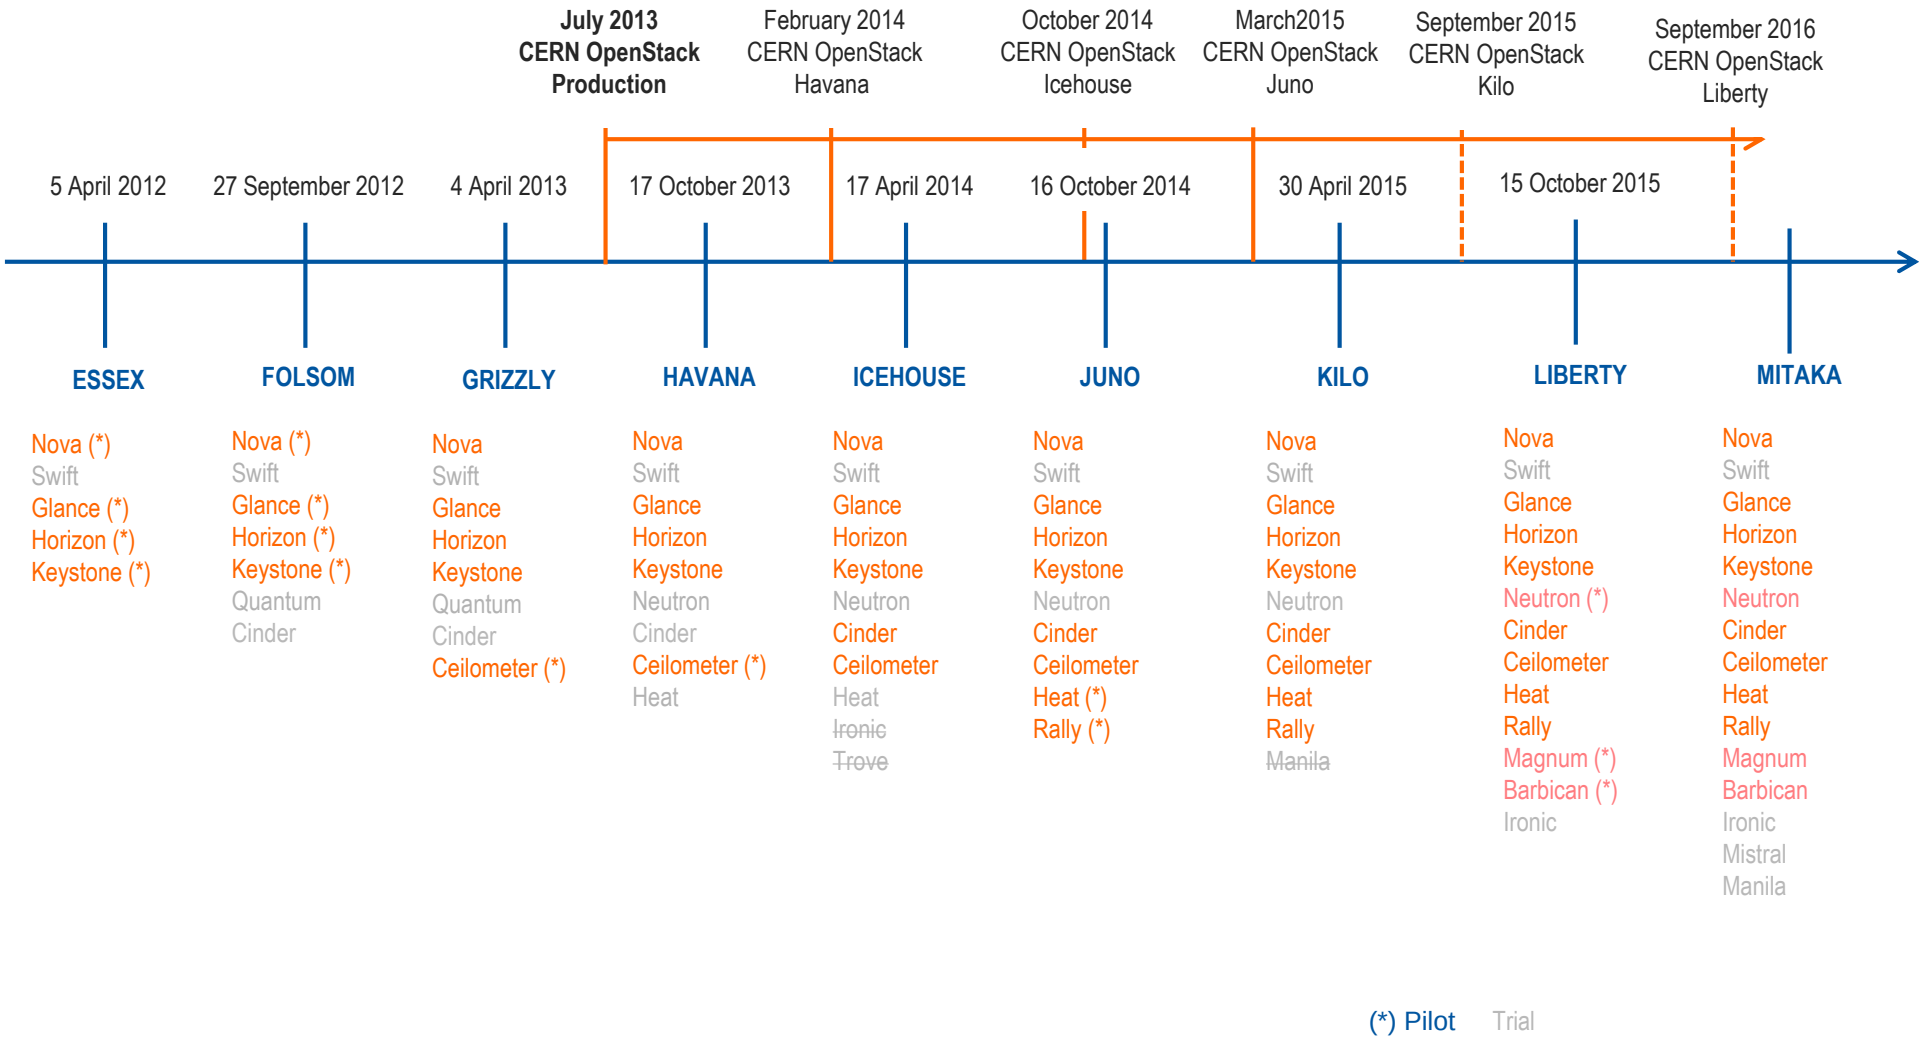

(*) Pilot Trial

Not Just The Software

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

e.g.

- 200 people arrive and leave CERN / month
- User skill levels vary widely
- Application range from server consolidation, dev/test to production compute

Subbu Allamaraju @ eBay

OpenStack@CERN Status

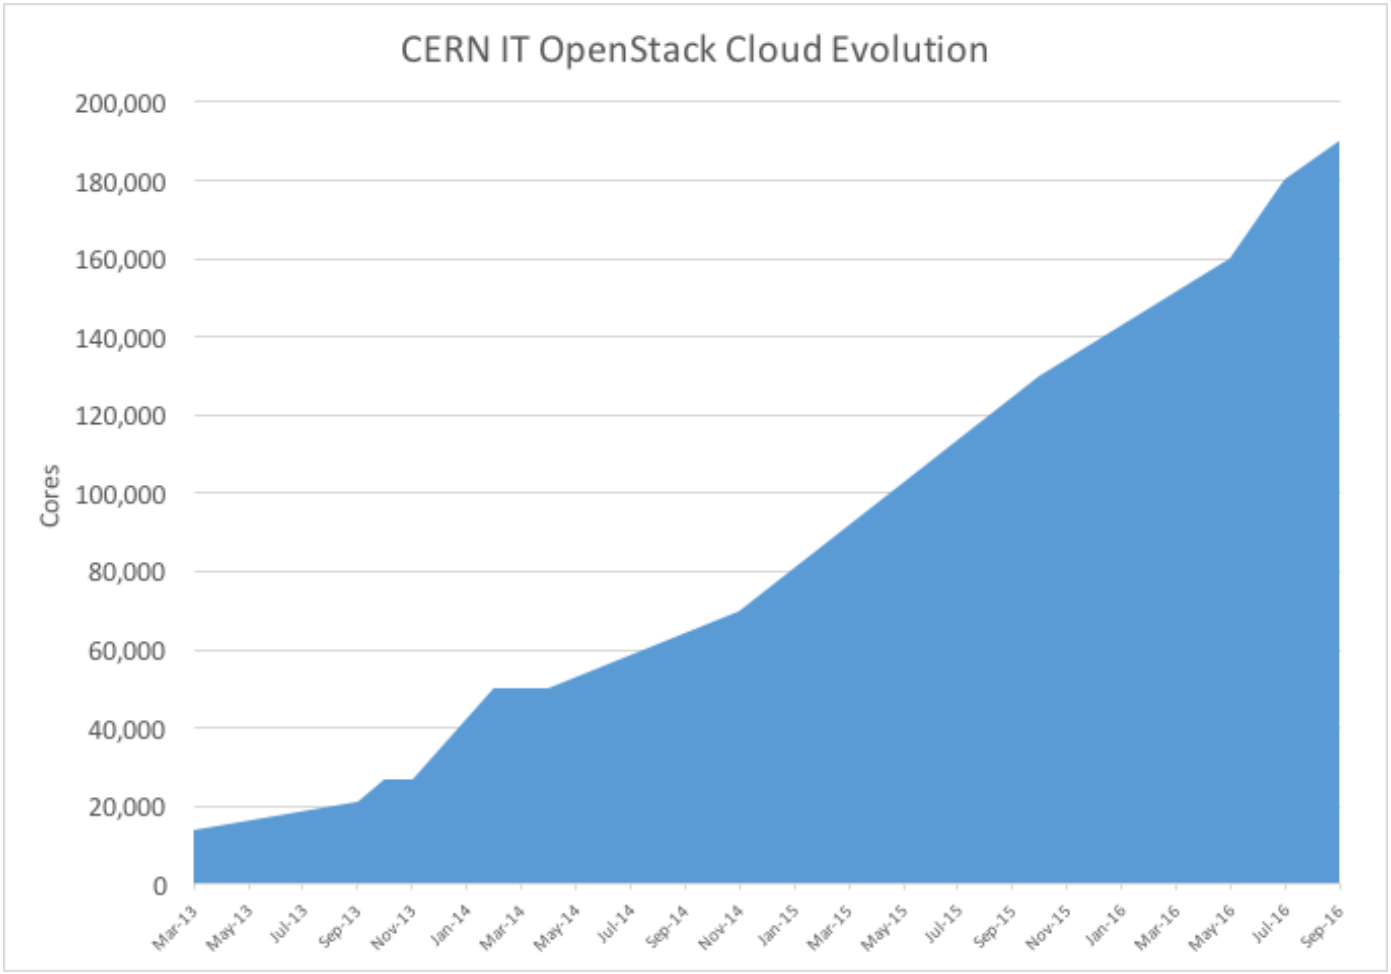

In production:

- >190K cores
- >7000 hypervisors

~100,000 additional cores being installed in next 6 months

90% of CERN's compute resources are now delivered on top of OpenStack

Cultural Change Impact

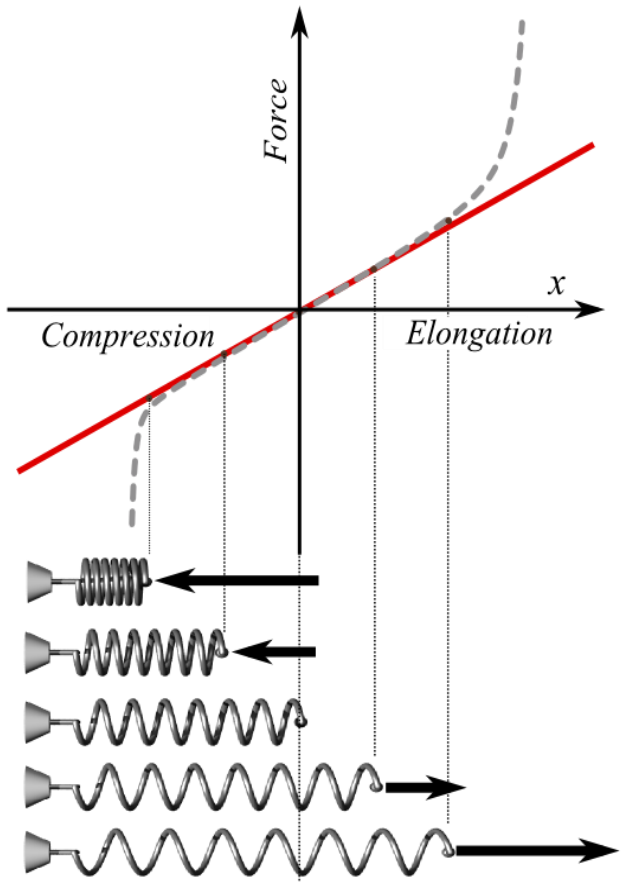

- Move to Agile technologies brings great benefits
  - Reduced deployment time
  - Continuous integration validation
  - Flexibility
- Don’t forget Hooke’s Law (adapted)
  - Under load, an organisation can extend proportional to external force
  - Too much load leads to permanent deformation
- Ensure the tail is moving fast as well as the head
  - Application support
  - Cultural challenges
  - Process change

Nova Cells

Top level cell

- Runs API service
- Top cell scheduler

Child cells run

- Compute nodes
- Scheduler
- Conductor
- ~40 cells

Version 2 coming

- Default for all

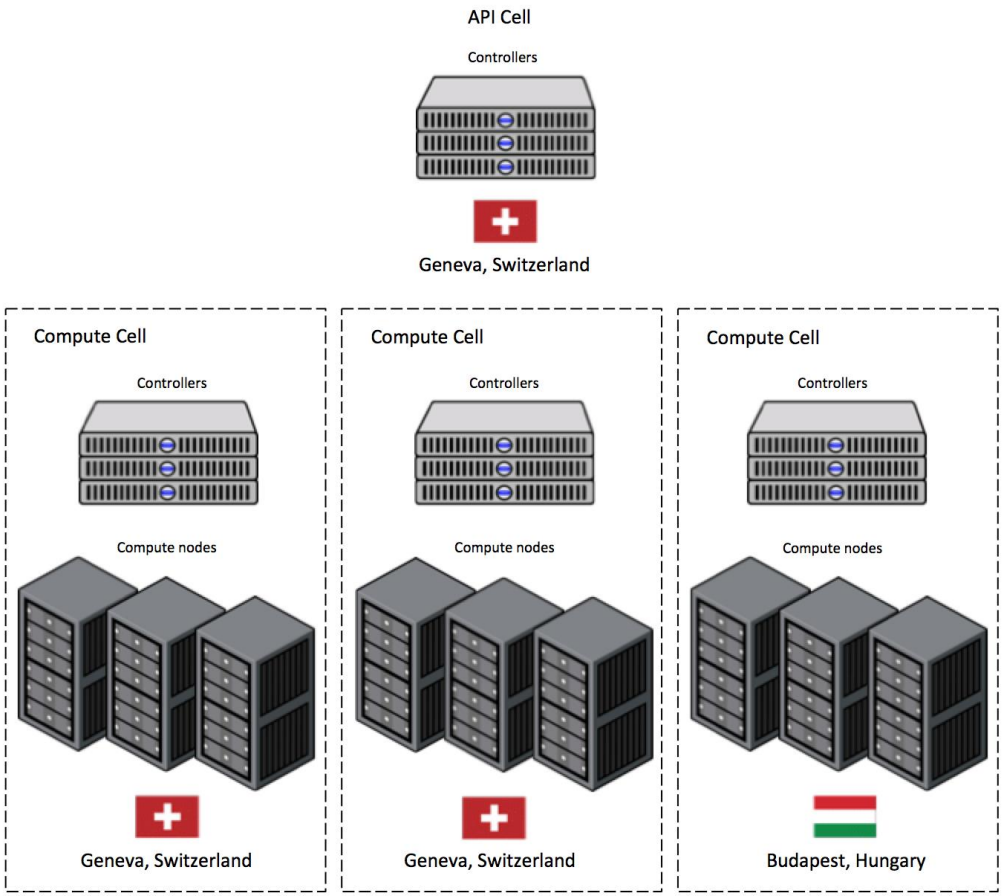

On-Premise Challenges

- Keeping up
  - Releases every 6 months
- Tetris scheduling of different flavors of VMs
  - Rackscale architectures may help in future
- Consistency of configuration
  - > 20 different hardware configurations
- Simulating elasticity on premise
  - Want to have good utilisation of resources
  - Shrink opportunistic batch computing

Public Cloud Impact

As expensive as computers are, people's time is more

- How we currently schedule processing makes little sense from the perspective of maximize the efficiency of people

One of the biggest improvements in joining to a much larger pool of resources is breaking the idea we need to lay out our resources for average load

- Workflows could be completed as they are defined and not over months

Ian Fisk CHEP2016

Provisioning for Average

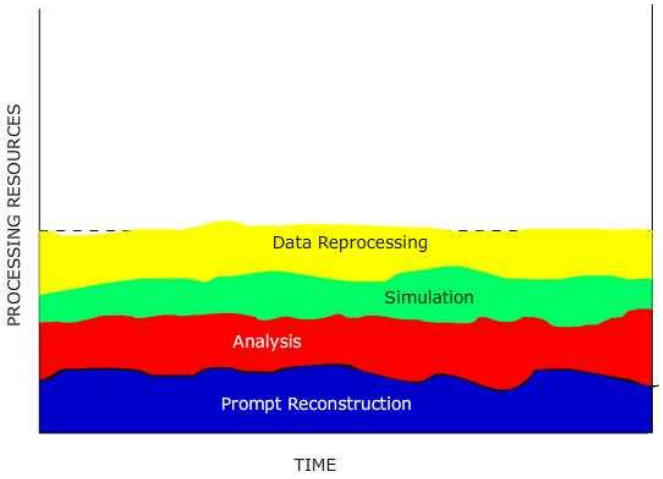

Provisioning for Peak

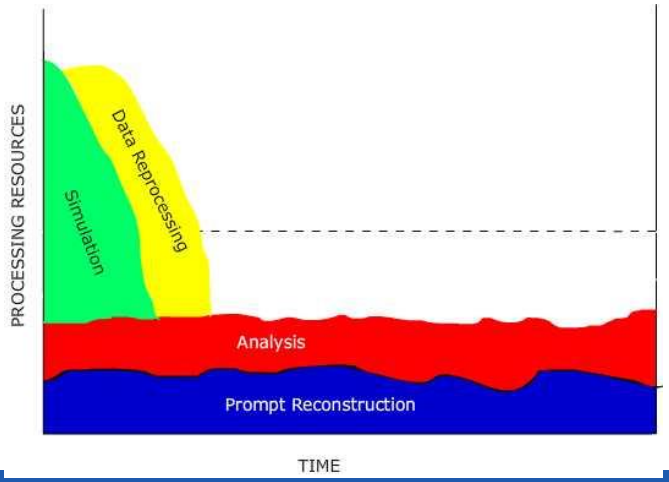

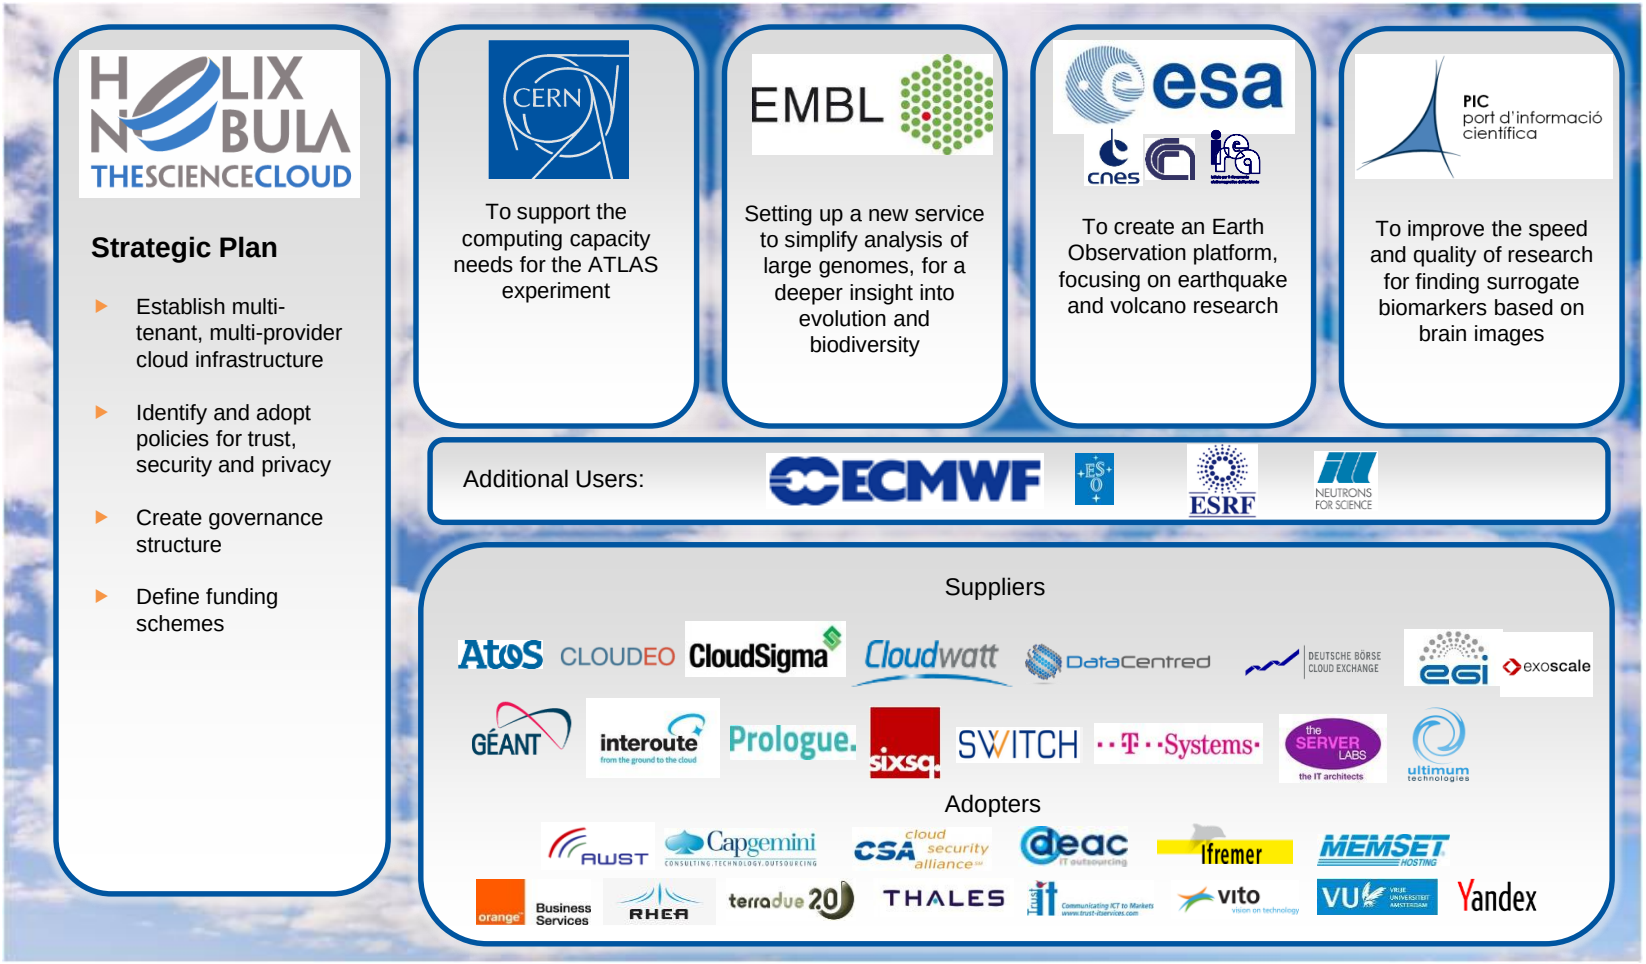
HELIX NEBULA
THESCIENCECLOUD
Strategic Plan
Establish multi-tenant, multi-provider cloud infrastructure
Identify and adopt policies for trust, security and privacy
Create governance structure
Define funding schemes
CERN
To support the computing capacity needs for the ATLAS experiment
EMBL
Setting up 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
THALES
trust-it
vito
VU
Yandex

Public Cloud Tests

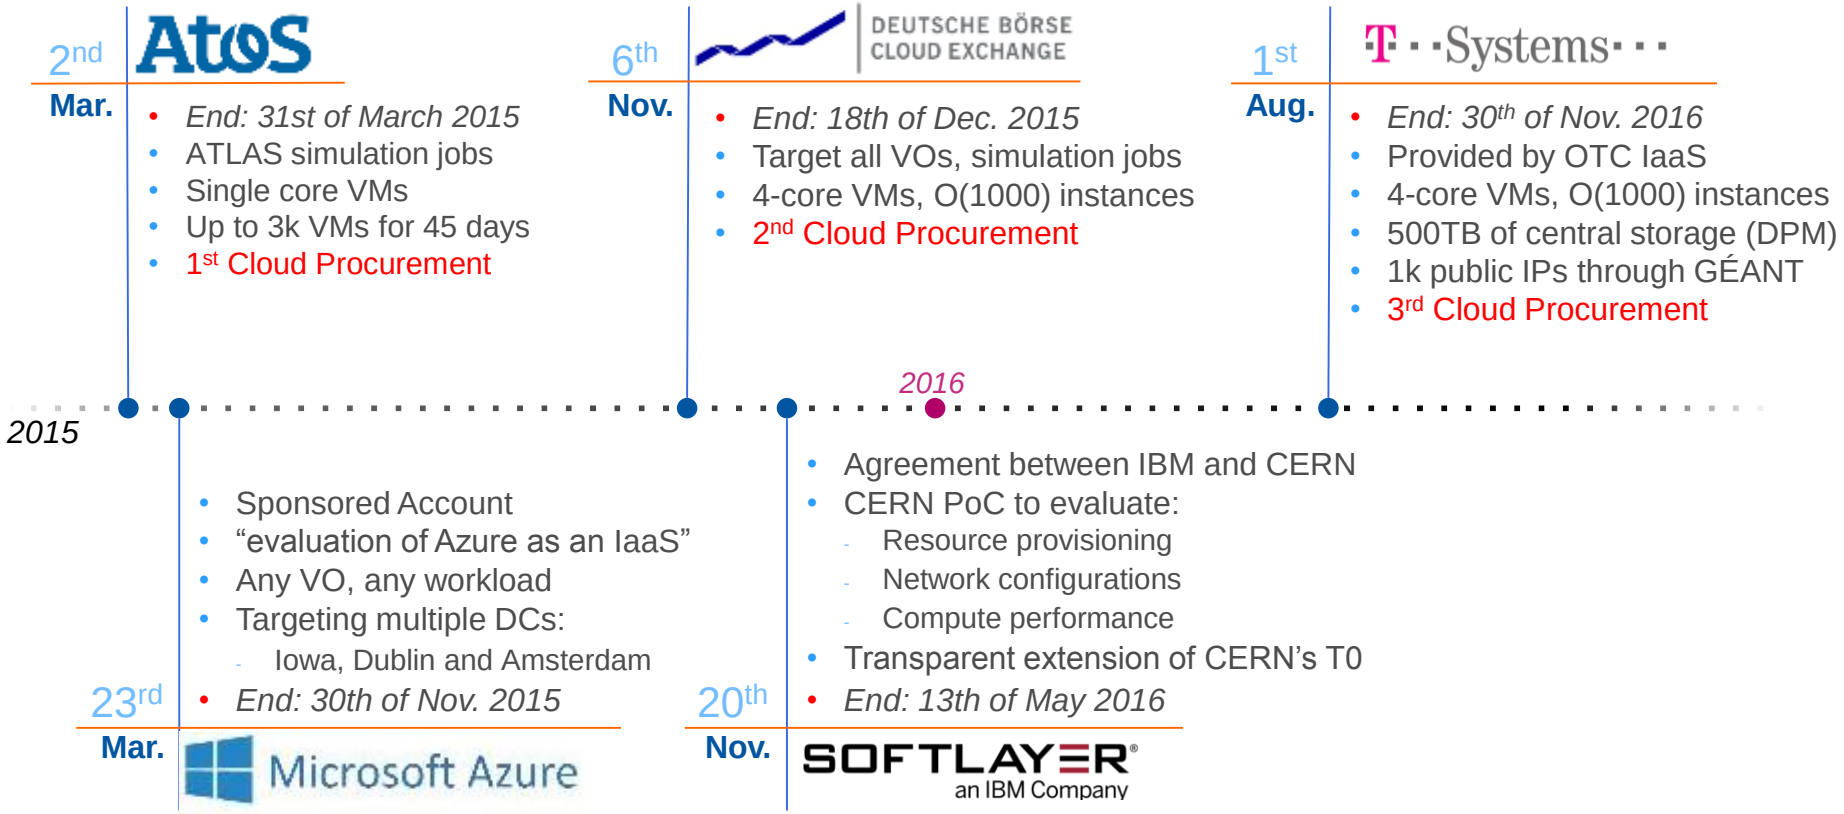

Different Approaches Evaluated

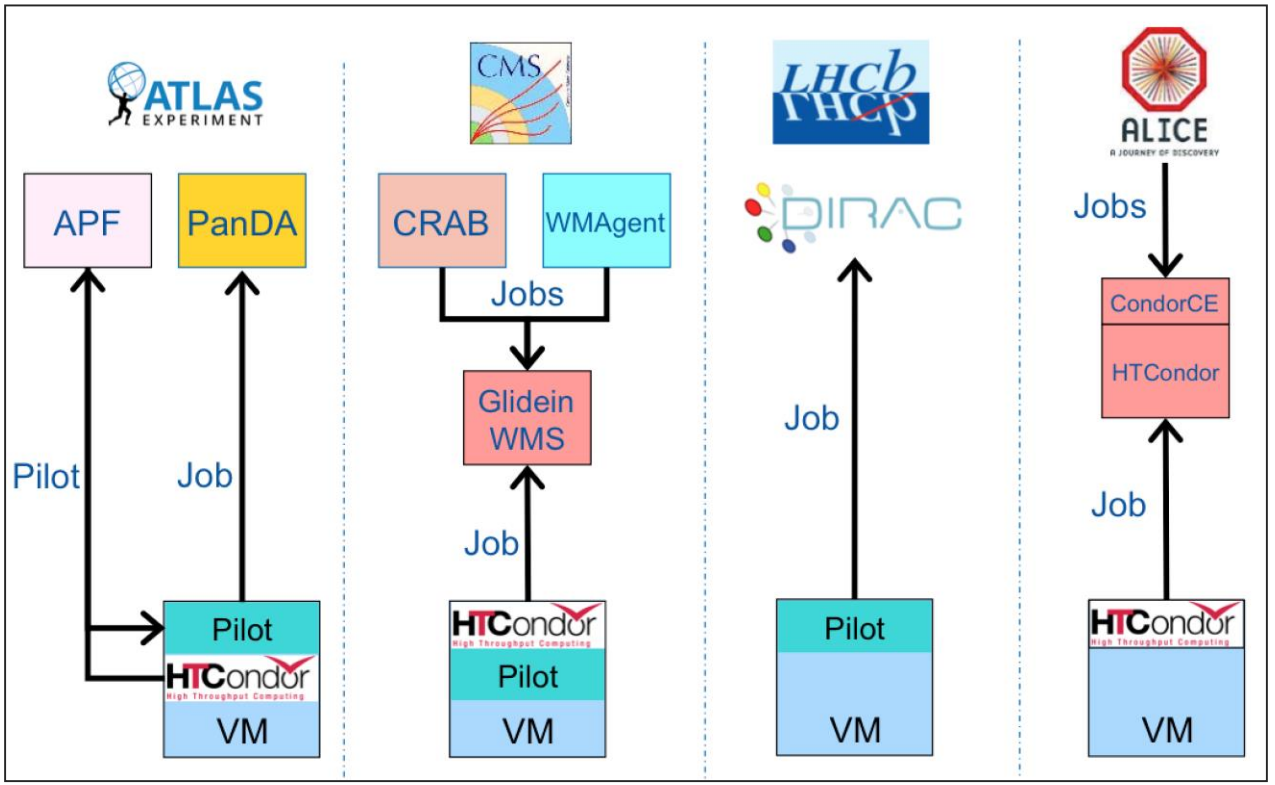

- Standardised on ALICE style for the future
- Extending CERN Batch system with Terraform

HNSciCloud – Project Phases

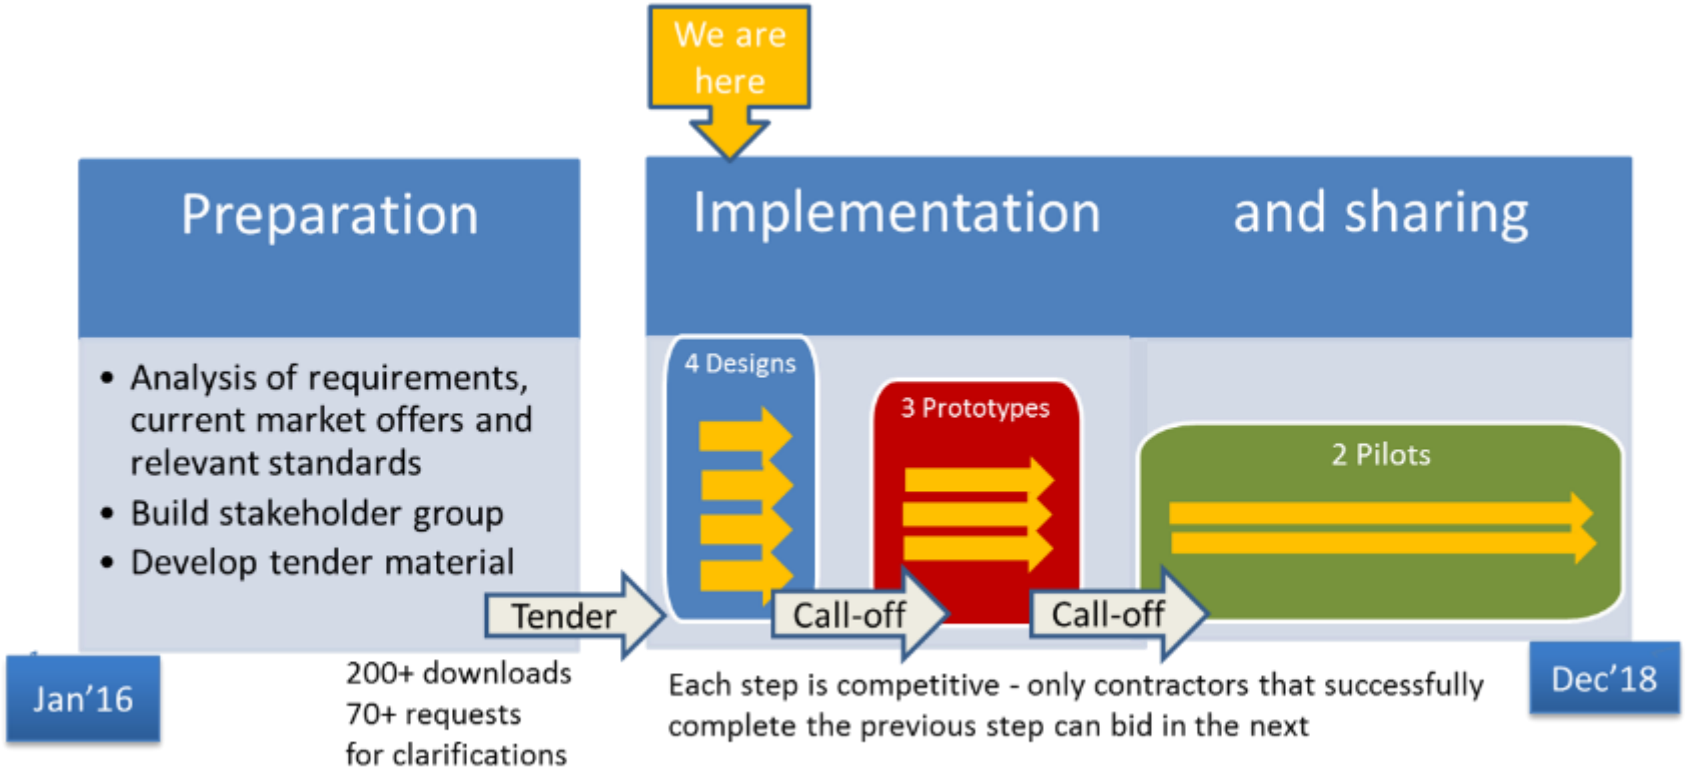

Public Cloud Challenges I

- Service Model
  - Spot vs dedicated instances
  - SLAs are very diverse
- APIs
  - No well adopted IaaS standard
- Storage
  - Experiment frameworks not fully exploiting cloud storage
  - Block and Object Store compared to HEP technologies

Public Cloud Challenges II

- Networking
  - VPNs, IPv4, DNS, routing to the LHC/GEANT networks
  - Cost is defined by the application
- Accounting
  - How to cross-check the bill?
  - Benchmarking, availability, pledge model
- Higher level functions are very unique
  - Lambda, Batch, Machine Learning
  - Significant value but risk of lock in
- Skills profiles changes significantly
  - Disk changing to contract management

Summary

- Clouds have been effectively used for 3 years to deliver LHC compute capacity
- Major cultural and technology changes have been successfully addressed
- More work required to understand usage, costs and procurement of public cloud resources

Discussions…

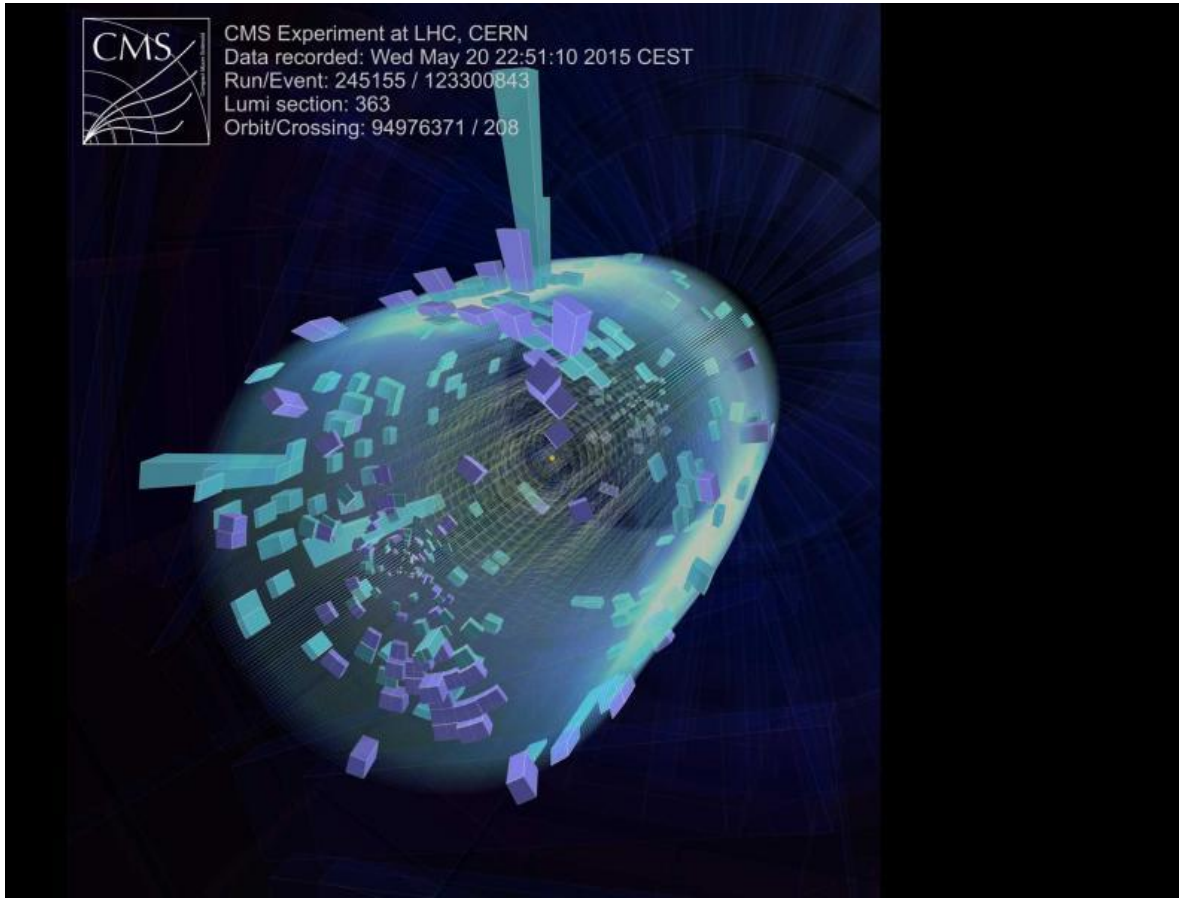

OpenStack Technical details at http://openstack-in-production.blogspot.fr

Helix Nebula information at http://hnscicloud.eu

Expanded Facilities in Hungary