CERN

Hybrid Cloud for CERN

Dr Helge Meinhard / CERN-IT

CxP Forum du Numérique

23-Jun-2016

CERN

- International organisation close to Geneva, straddling Swiss-French border, founded 1954
- Facilities for fundamental research in particle physics
- 21 member states, 1 B CHF budget
- 3’197 staff, fellows, apprentices, …
- 13’128 associates

1954: 12 Member States

Members: Austria, Belgium, Bulgaria, Czech republic, Denmark
Finland, France, Germany, Greece, Hungary, Israel, Italy,
Netherlands, Norway, Poland, Portugal, Slovak Republic, Spain,
Sweden, Switzerland, United Kingdom
Candidate for membership: Romania
Associate member: Serbia
Observers: European Commission, India, Japan, Russia,
Turkey, UNESCO, United States of America
Numerous **non-member states with collaboration agreement**

2’531 staff members, 645 fellows,
21 apprentices

7’000 member states, 1’800 USA,
900 Russia, 270 Japan, …

“Science for peace”

Birth place of World-wide Web

Tools (1): LHC

Exploration of a new energy frontier
in p-p and Pb-Pb collisions

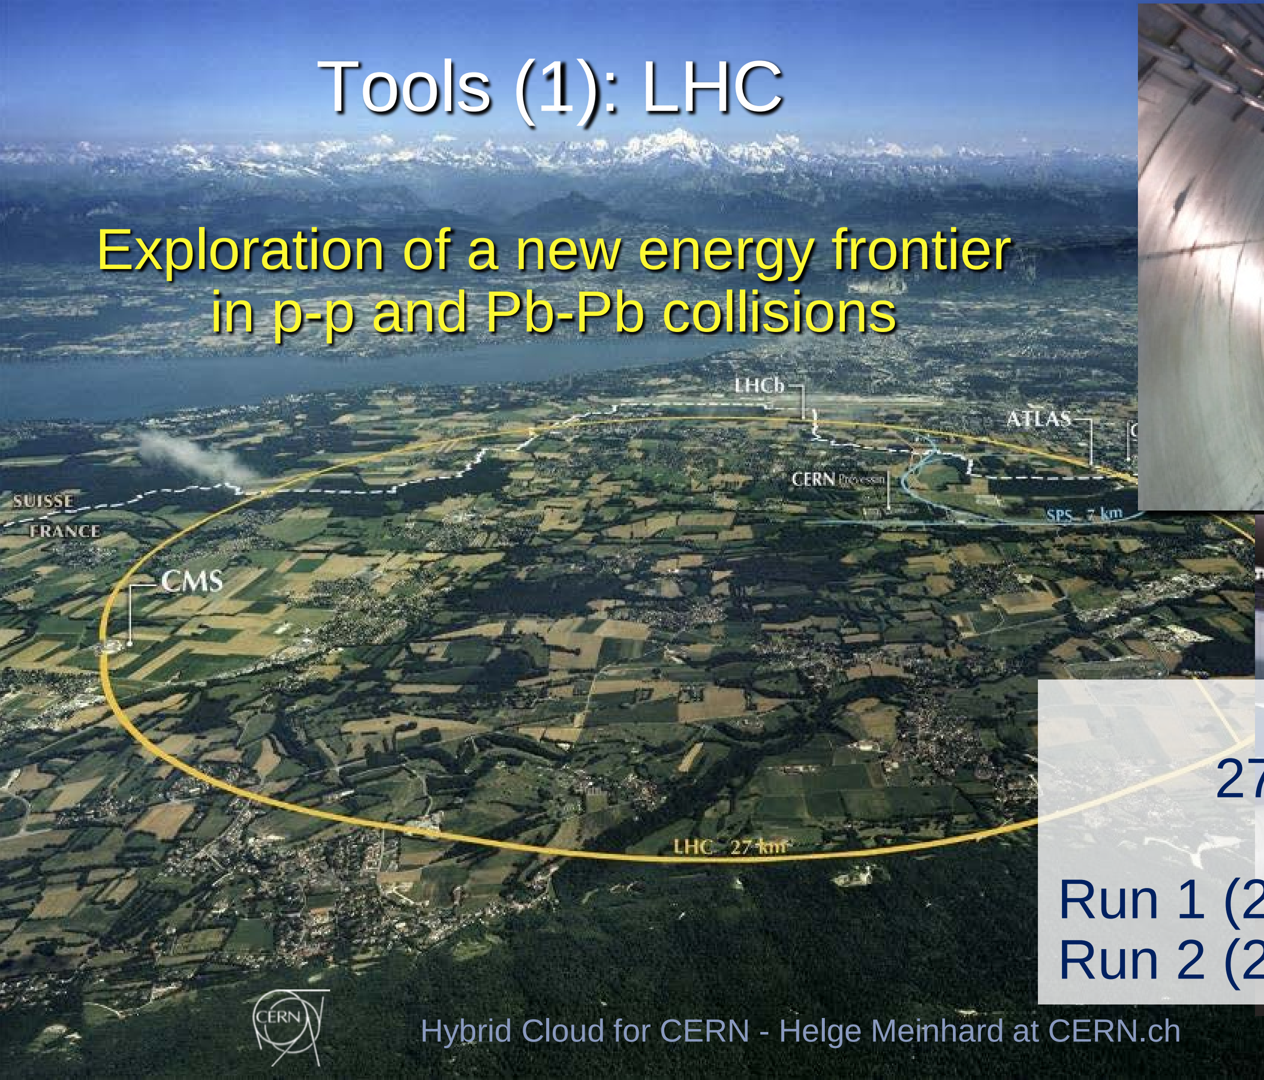

LHC ring:
27 km circumference

Run 1 (2010-2013): 4+4 TeV
Run 2 (2015-2018): 6.5 + 6.5 TeV

Tools (2): Detectors

Exploration of a new energy frontier
in p-p and Pb-Pb collisions

Tools (2): Detectors

ATLAS (A Toroidal Lhc ApparatuS)

- 25 m diameter, 46 m length, 7’000 tons
- 3’000 scientists (including 1’000 grad students)
- 150 million channels
- 40 MHz collision rate
- Event rate after filtering: 300 Hz in Run 1; up to 1’000 Hz in Run 2

Results so far

- Many… the most spectacular one being
- 04 July 2012: Discovery of a “Higgs-like particle”
- March 2013: The particle is indeed the **Higgs boson**
- 08 Oct 2013 / 10 Dec 2013: **Nobel price** to Peter Higgs and François Englert
    - CERN, ATLAS and CMS explicitly mentioned

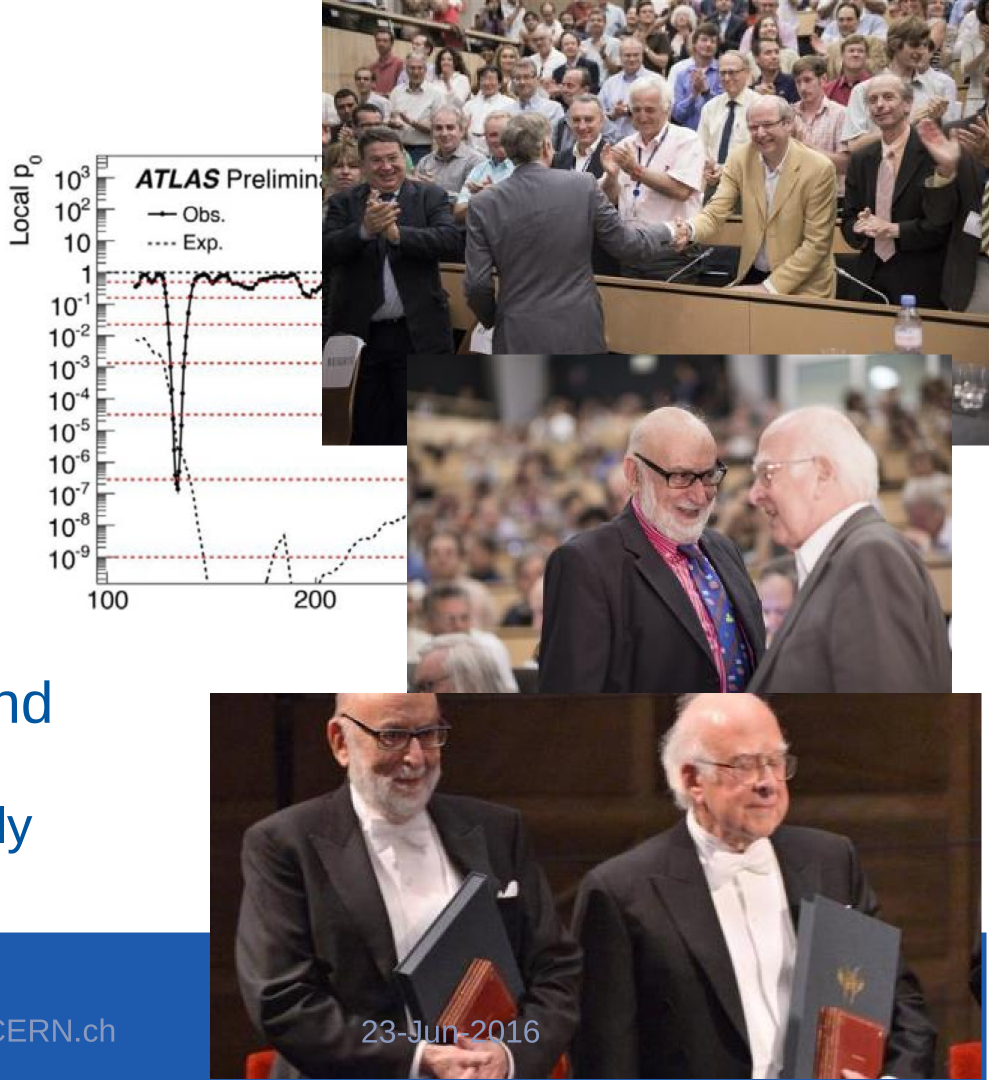

- Up to 6 GB/s to be permanently stored after filtering
- Almost 30 PB/y in Run 1
- Expect ~50 PB/y in Run 2
- 2023: 400 PB/y(?)

The Worldwide LHC Computing Grid

An International collaboration to distribute and analyse LHC data

Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists

Tier-0 (CERN):
data recording, reconstruction and distribution

Tier-1:
permanent storage, re-processing, analysis

Tier-2:
Simulation, end-user analysis

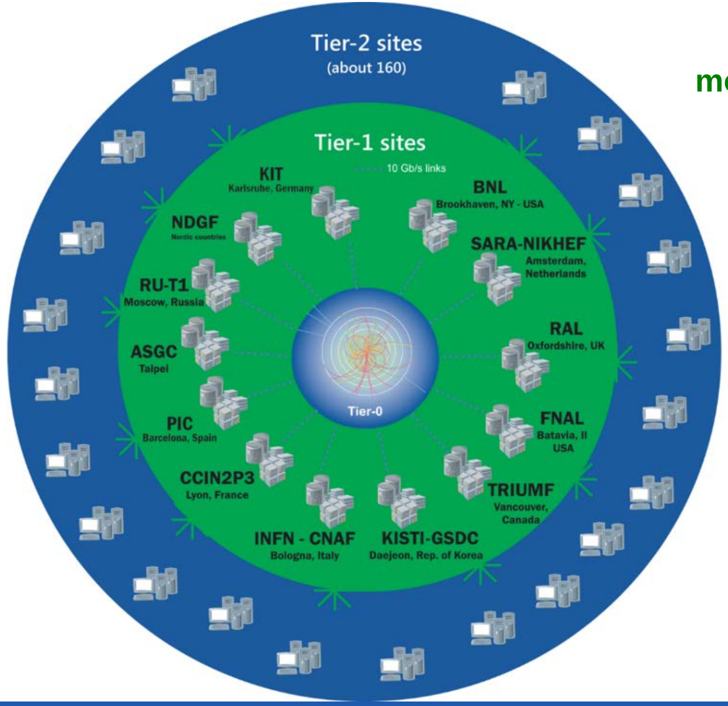

more than 170 sites, 40 countries

~350’000 cores

500 PB of storage

> 2 million jobs/day

10-100 Gb links

Challenges

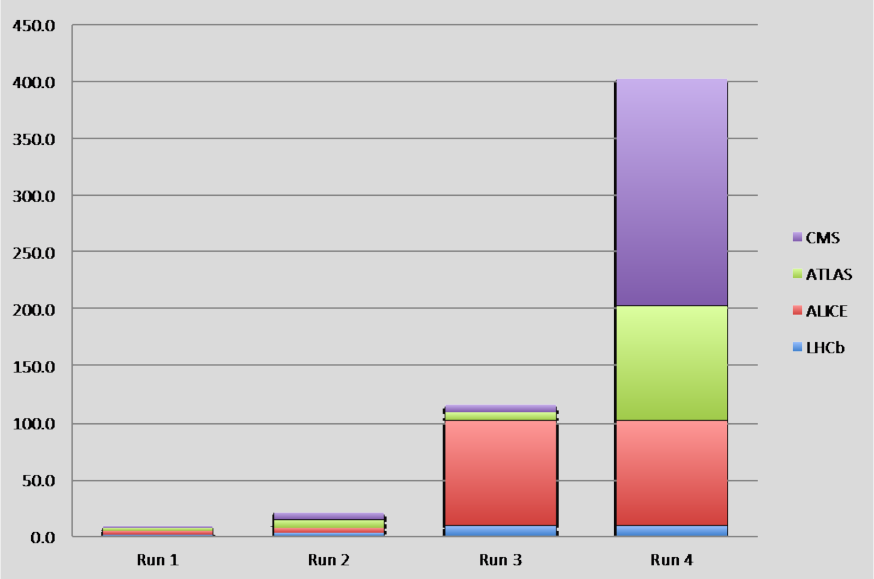

Run 2:

- Moore's law helps, but not sufficient
- Large effort spent to improve software efficiency
- Exploit multi-threading, new instruction sets, …
- Still need factor 2 in terms of cores, storage etc.

Tier-0: 15% of WLCG

MEYRIN DATA CENTRE

()	last_value ()
Number of Cores in Meyrin	151,159
Number of Drives in Meyrin	83,709
Number of 10G NIC in Meyrin	9,307
Number of 1G NIC in Meyrin	23,647
Number of Processors in Meyrin	25,215
Number of Servers in Meyrin	13,377
Total Disk Space in Meyrin (TB)	175,900
Total Memory Capacity in Meyrin (TB)	613

WIGNER DATA CENTRE

()	last_value ()
Number of Cores in Wigner	43,328
Number of Drives in Wigner	23,180
Number of 10G NIC in Wigner	1,399
Numer of 1G NIC in Wigner	5,067
Number of Processors in Wigner	5,418
Number of Servers in Wigner	2,712
Total Disk Space in Wigner (TB)	71,738
Total Memory Capacity in Wigner (TB)	172

NETWORK AND STORAGE

()	last_value ()
Tape Drives	104
Tape Cartridges	20,517
Data Volume on Tape (TB)	144,038
Free Space on Tape (TB)	41,023
Routers (GPN)	140
Routers (TN)	30
Routers (Others)	108
Switches	3,712

Transforming In-House Resources

We now have

- Full support for physical and virtual servers
- Full support for remote machines
- Horizontal view
  - Responsibilities by layers of service deployment
- Large fraction of resources run as **private cloud under OpenStack**
- Scaling to large numbers
  (> 15’000 physical, several 100’000s virtual)
- Support for dynamic host creation/deletion
  - Deploy new services/servers in hours rather than weeks/months
  - Optimise operational and resource efficiency

Scaling up Further: Public Clouds (1)

- Additional resources, perhaps later replacing on-premise capacity
- Potential benefits:
  - Economy of scale
  - More elastic, adapts to changing demands
  - Somebody else worries about machines and infrastructure

Scaling up Further: Public Clouds (2)

- Potential issues:
  - Cloud provider’s business models not well adapted to procurement rules and procedures of public organisations
  - Lack of skills for and experience with procurements
  - Market largely not targeting compute-heavy tasks
    - Performance metrics/benchmarks not established
  - Legal impediments
  - Not integrated with on-premise resources and/or publicly funded e-infrastructures

CERN Approach

Series of short procurement projects of increasing size and complexity

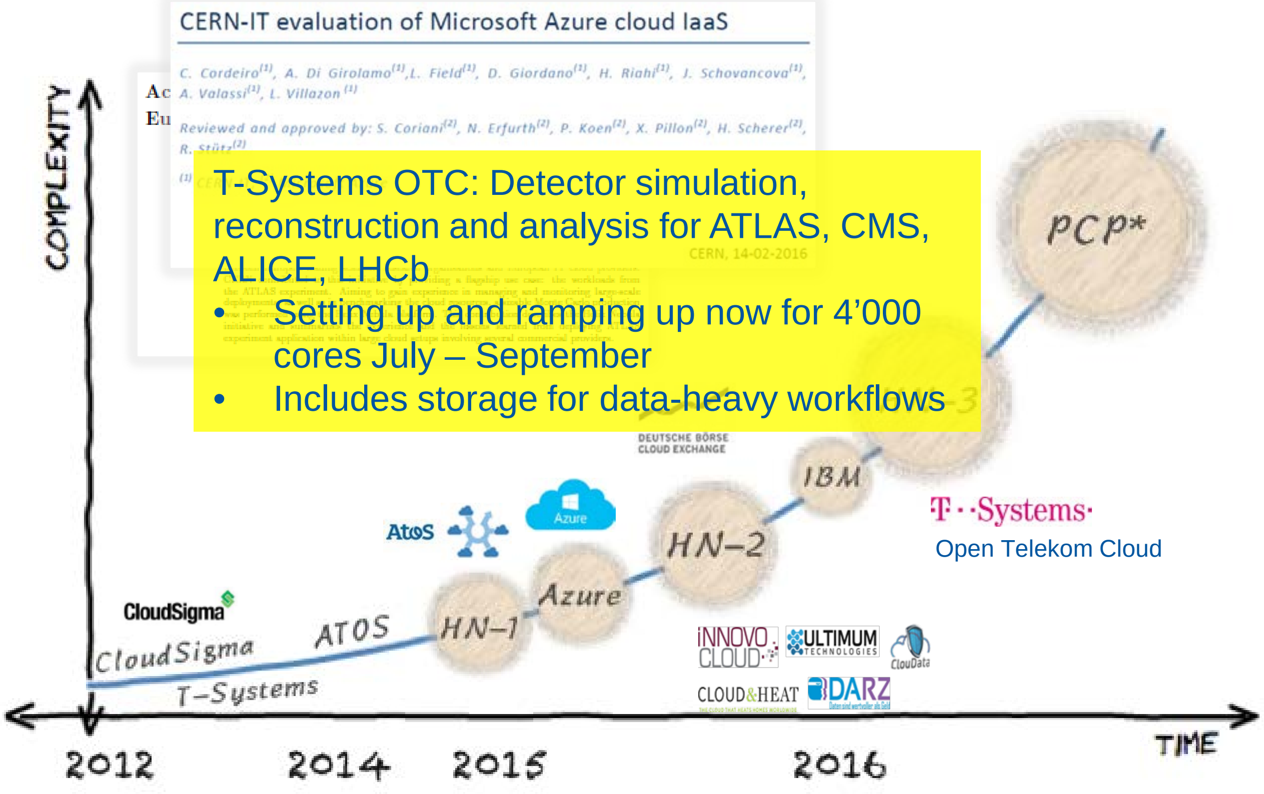

T-Systems OTC: Detector simulation, reconstruction and analysis for ATLAS, CMS, ALICE, LHCb

- Setting up and ramping up now for 4'000 cores July – September
- Includes storage for data-heavy workflows

Some Lessons Learned

- APIs not as stable and well-defined as you would expect them to be
  - Consider requiring support for ecosystem tool such as Terraform, libcloud, jclouds
- Accounting is key – required on user side, too
- Benchmarking/performance metrics required
- Avoid brokers – they risk adding complexity, cost, intransparency
  - YMMV – EMBL and ESA have reported rather positive experience

Future Requirements

- Not only LHC, but a number of particle physics projects with high data rates
  - SuperKEKB, HL-LHC, FCC, LBNF, ILC
- Not only particle physics, but also other physics fields (e.g. astronomy)
  - SKA, LSST, CTA
- Not only physics, but also other sciences (e.g. life sciences, material science)
  - EBI expects data doubling every year (!)

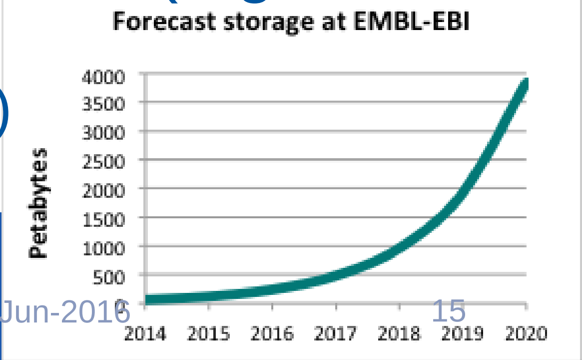

HELIX NEBULA The Science Cloud
Joint Pre-Commercial Procurement

Procurers: CERN, CNRS, DESY, EMBL-EBI, ESRF,
IFAE, INFN, KIT, SURFSara, STFC
Experts: Trust-IT & EGI.eu

Procurers have committed funds (>1.6M€), manpower, use-cases with applications & data and in-house IT resources

Objective: procure innovative IaaS level cloud services
- Fully and seamlessly integrating commercial cloud (Iaas) resources with in-house resources and European e-Infrastructures
- To form a hybrid cloud platform for science

Services will be made available to end-users from many research communities: High-energy physics, astronomy, life sciences, neutron/photon sciences, long tail of science

Co-funded via H2020 (Jan'16-Jun'18)
- Grant Agreement 687614
- Total procurement volume: >5M€

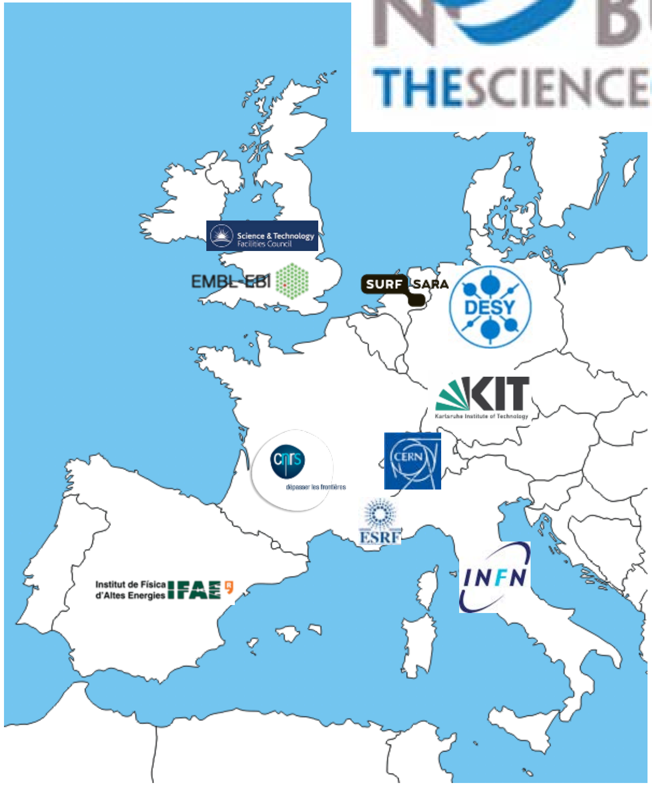

Total procurement commitment >5M€

Technical Challenges

- Compute
  - Integration of some HPC requirements
- Storage
  - Caching at provider's site, if possible automatically (avoid managed storage)
- Network
  - Connection via GÉANT
  - Support of identity federation (eduGAIN) for IT managers
- Procurement
  - Match of cloud providers' business model with public procurement rules

HNSciCloud – Current Status

- Project started in January 2016
- Tender announced in Jan 2016
- Open Market consultation successfully held on 17 March 2016
- Tender material in final phase of preparation
  - To be published this summer

HNSciCloud – Contacts

- Interested?
  - See http://www.hnscicloud.eu/
  - Subscribe to hnscicloud-announce@cern.ch

Summary

- Public clouds have a large potential of addressing the requirements of public research organisations for ever more resources and of dealing with peak demands
- Using public clouds isn't as easy as you would like it to be
- A full integration of public clouds with on-premise resources and public e-infrastructures is required, and is technically and administratively challenging
- Commercial cloud services are expected to play an increasing role in the computing models of scientific Research Infrastructures as part of a hybrid cloud platform

Merci pour votre attention

http://cern.ch

http://cern.ch/it-dep

http://cern.ch/wlcg

http://www.hnscicloud.eu/

Accelerating Science and Innovation