

Computing in High-Energy Physics

Dr Helge Meinhard / CERN-IT
CERN openlab summer student lecture
06 July 2016



CERN



CERN



CERN

- International organisation close to Geneva, straddling Swiss-French border, founded 1954



CERN

- International organisation close to Geneva, straddling Swiss-French border, founded 1954
- Facilities for fundamental research in particle physics



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



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



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, ...



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

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, ...



CERN

“Science for peace”

- 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, ...



CERN – Where the Web was Born



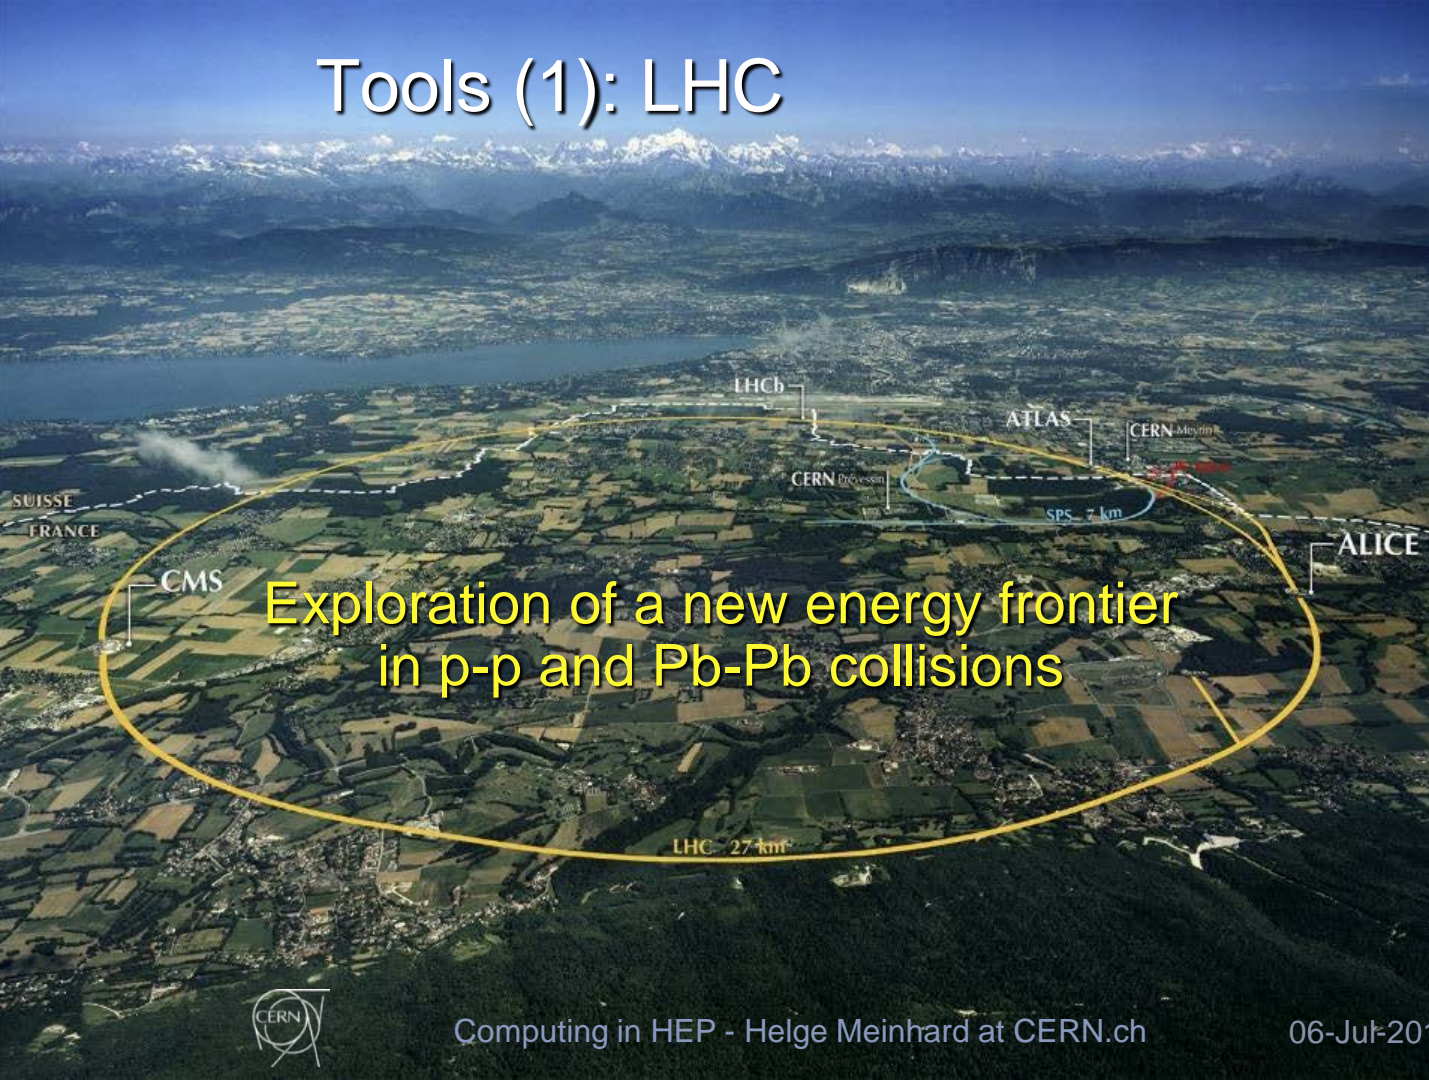
CERN – Where the Web was Born



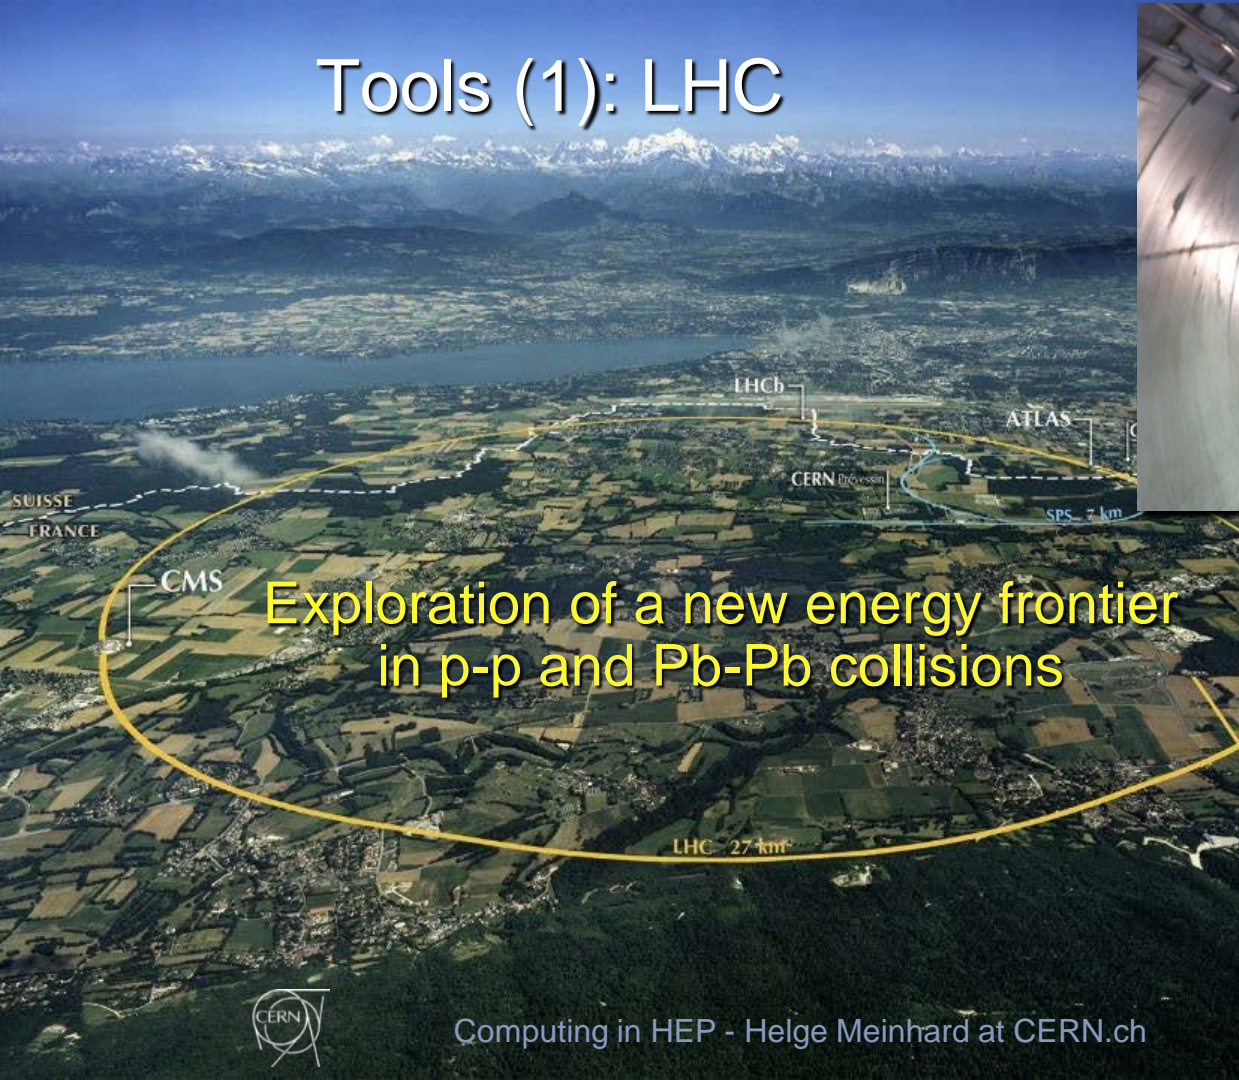
CERN – Where the Web was Born



Tools (1): LHC



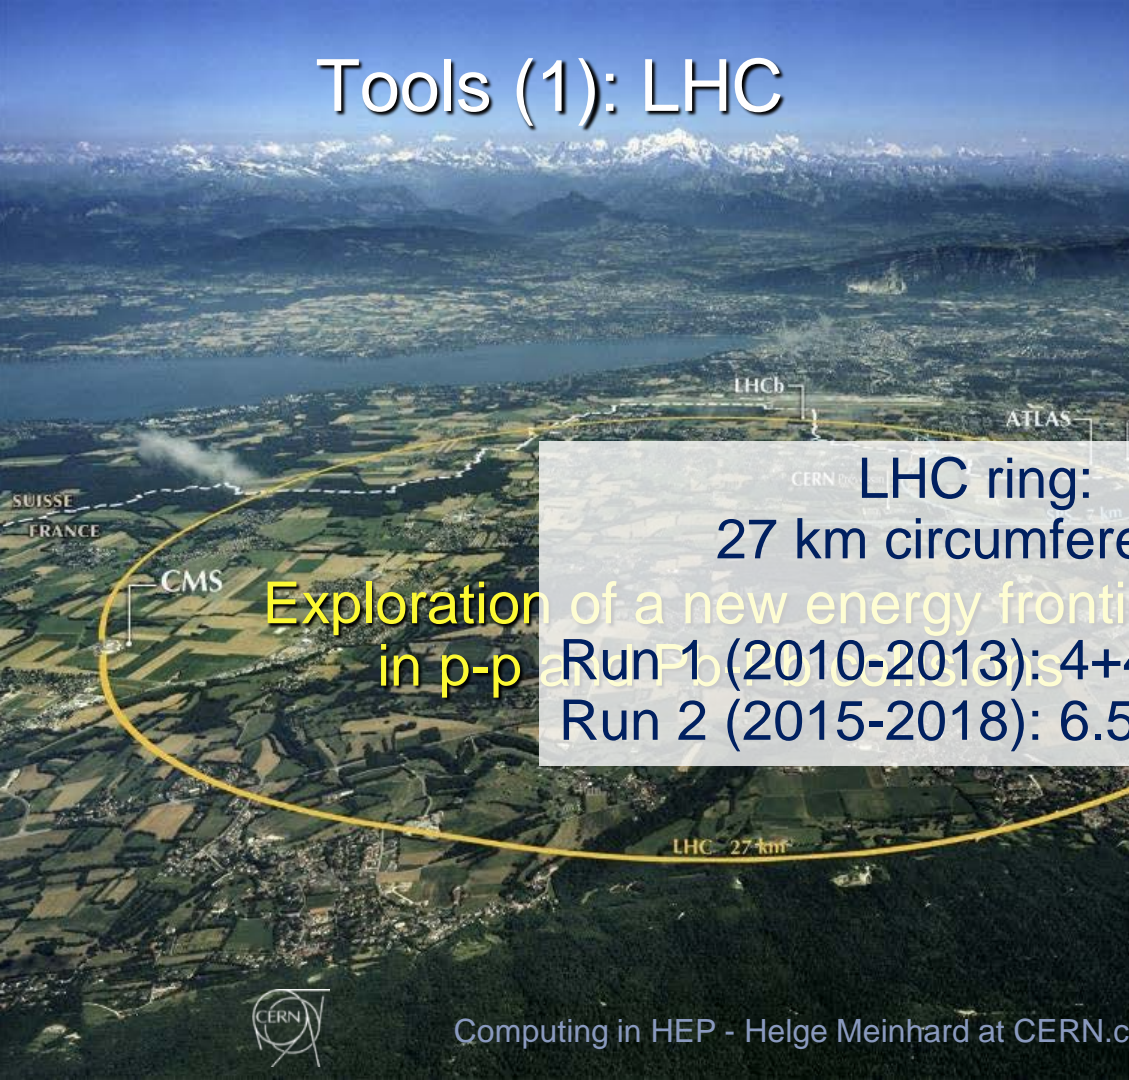
Tools (1): LHC



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



Tools (1): LHC



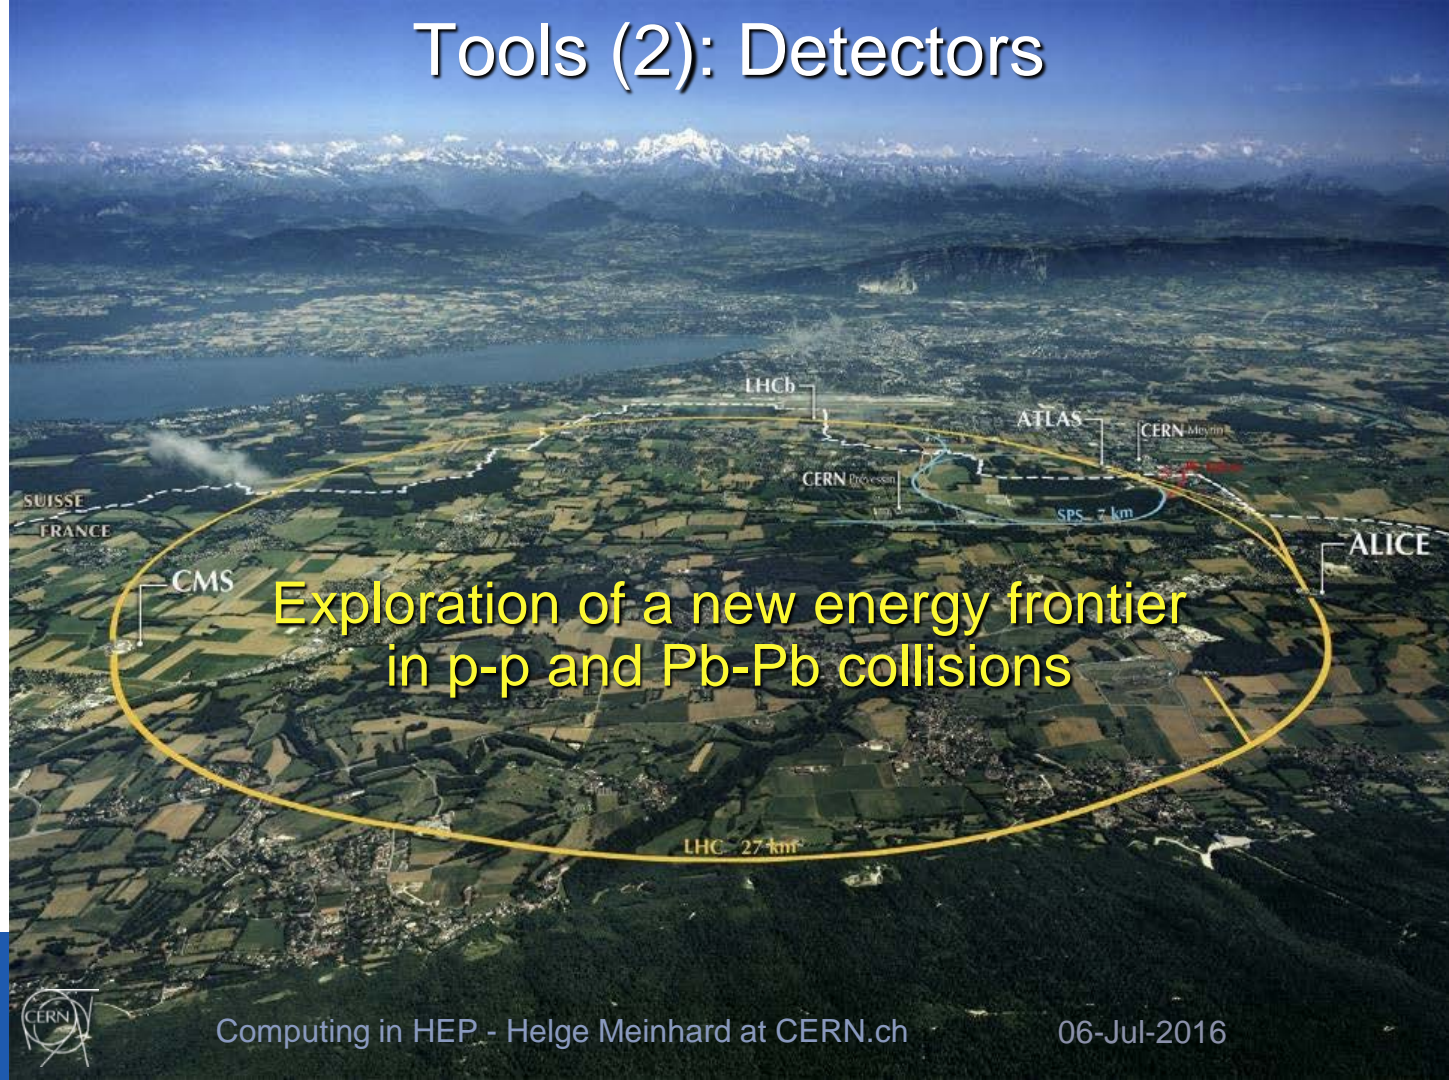
Exploration of a new energy frontier
in p-p collisions

Run 1 (2010-2013): 4+4 TeV

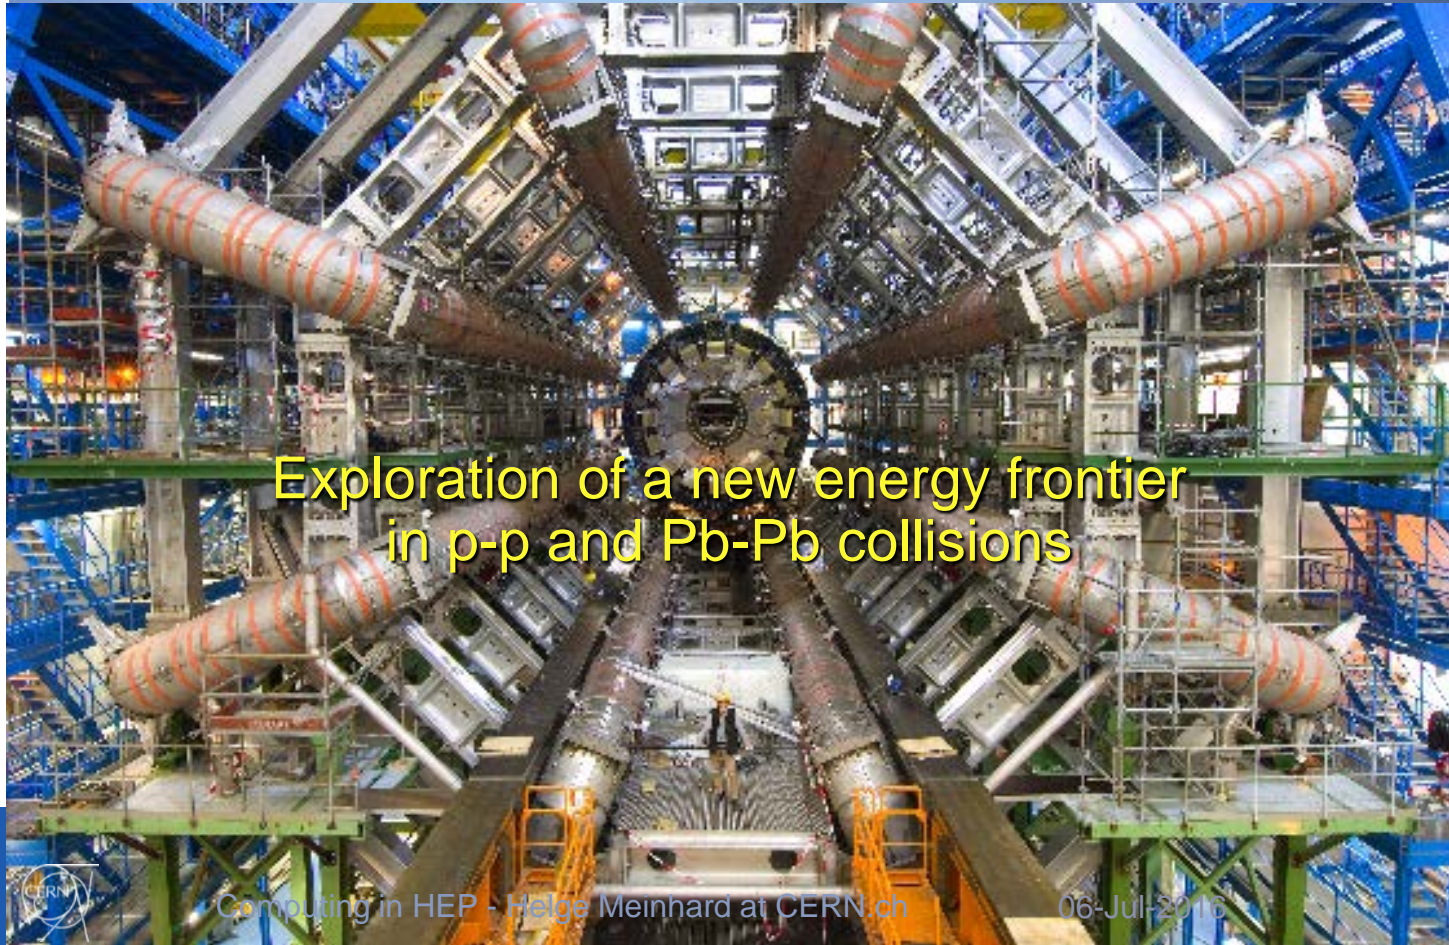
Run 2 (2015-2018): 6.5 + 6.5 TeV



Tools (2): Detectors



Tools (2): Detectors



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

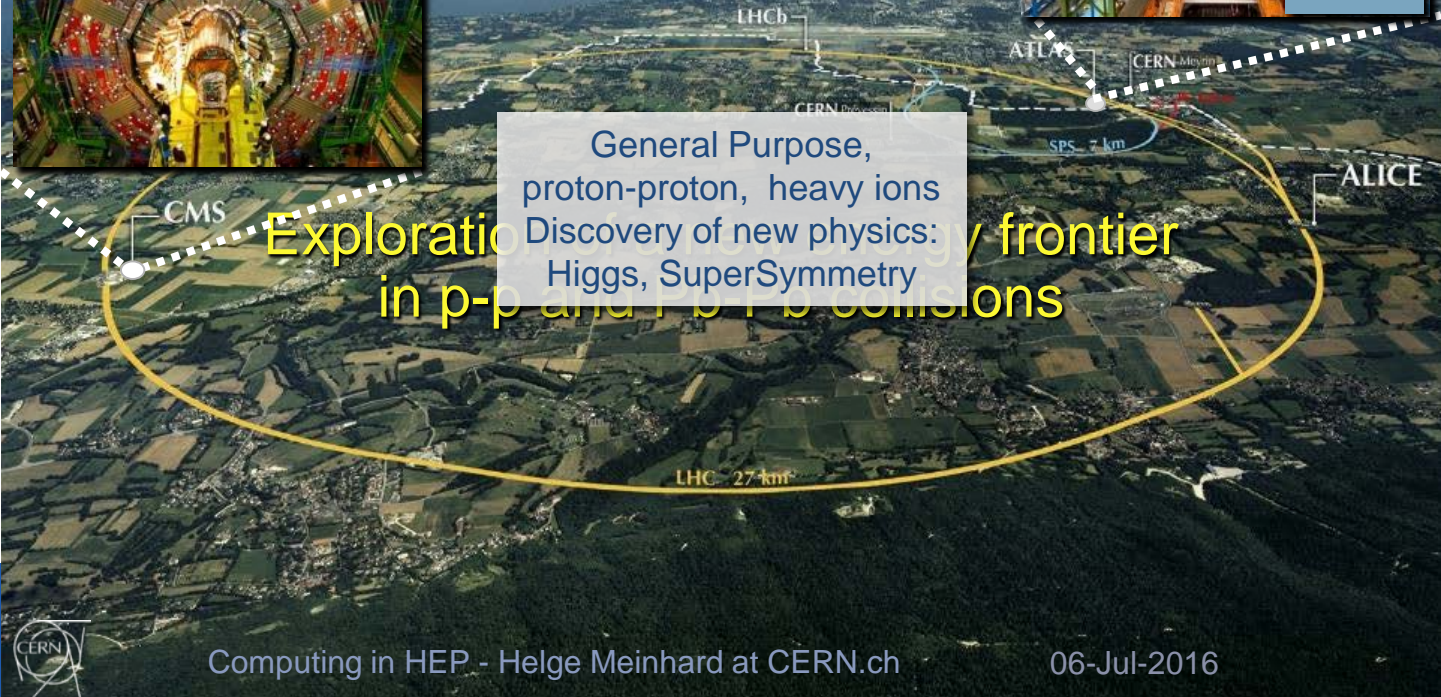
Tools (2): Detectors

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

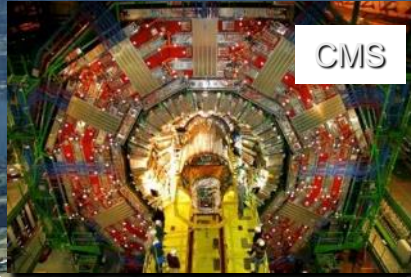
Exploration of a new energy frontier
in p-p and Pb-Pb collisions

Tools (2): Detectors



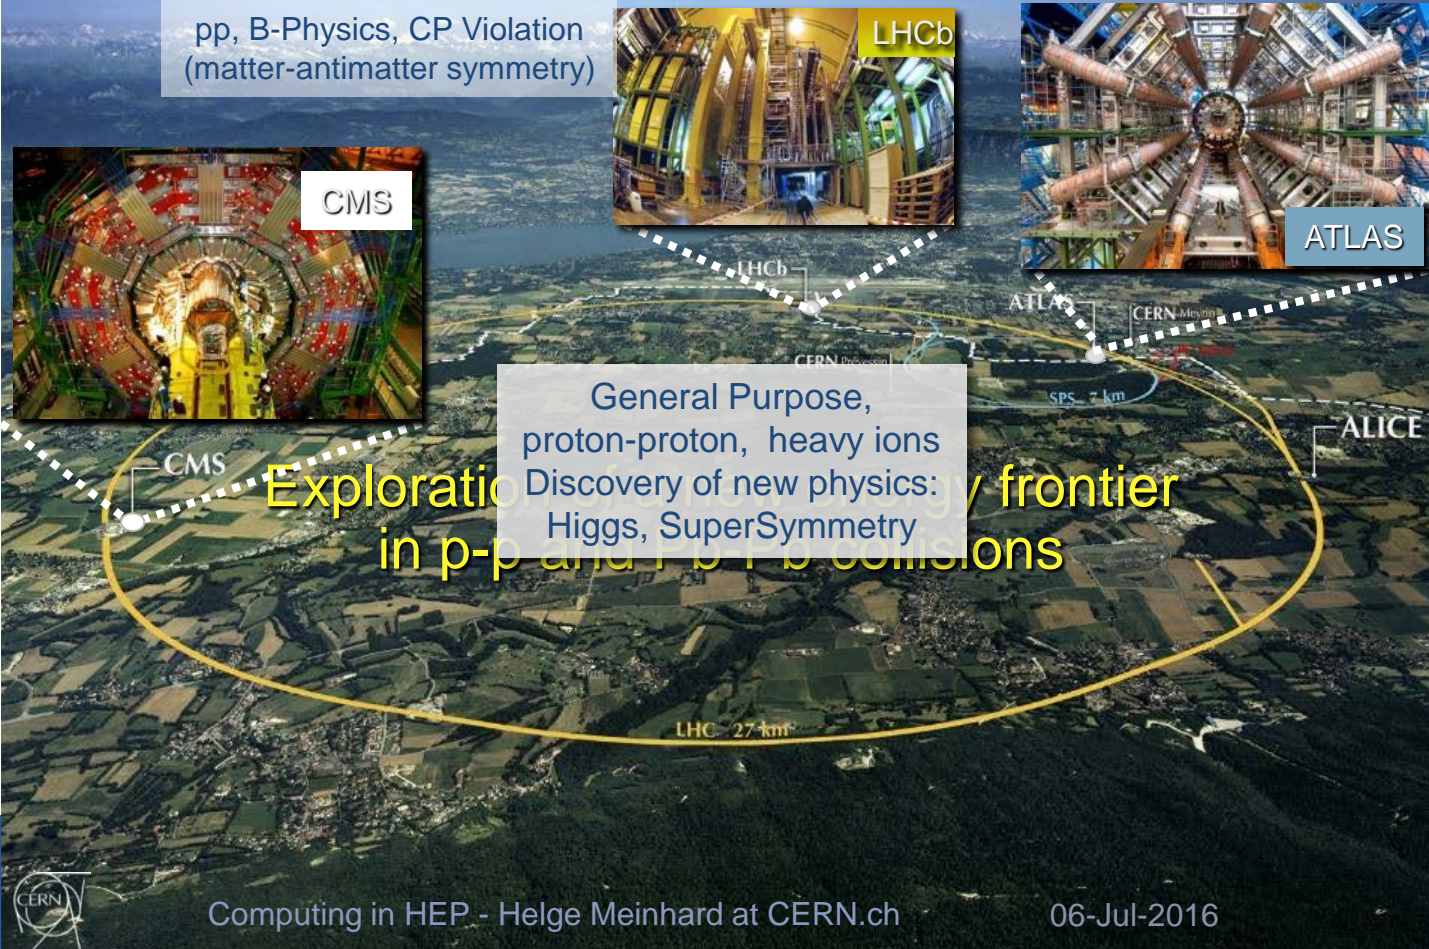
Tools (2): Detectors

pp, B-Physics, CP Violation
(matter-antimatter symmetry)



General Purpose,
proton-proton, heavy ions
Discovery of new physics:
Higgs, SuperSymmetry

Exploration of the energy frontier
in p-p and Pb-Pb collisions



Tools (2): Detectors

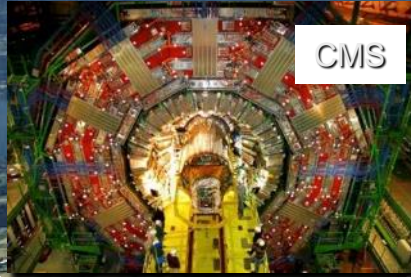
pp, B-Physics, CP Violation
(matter-antimatter symmetry)



LHCb



ATLAS



CMS

General Purpose,
proton-proton, heavy ions
Discovery of new physics:
Higgs, SuperSymmetry

Exploration in p-p and Pb-Pb collisions



ALICE

Heavy ions, pp
(state of matter of early universe)

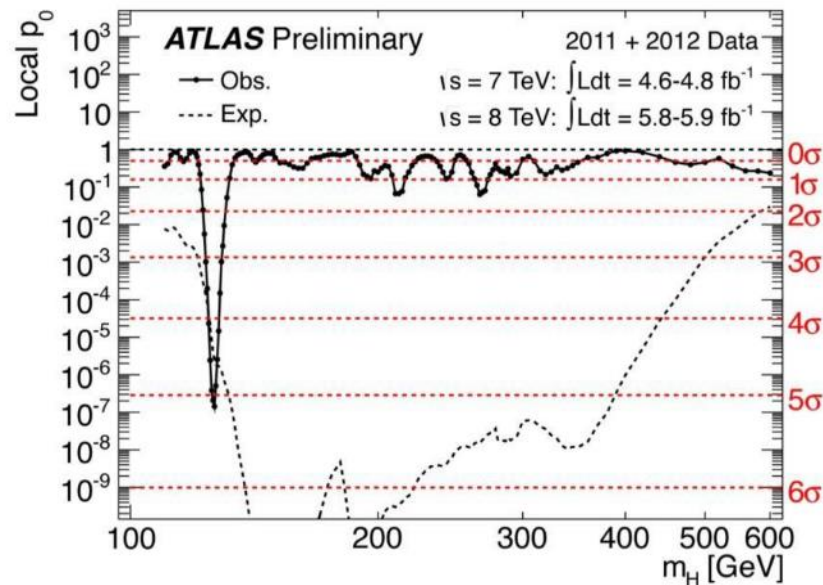


Results so far

- Many... the most spectacular one being

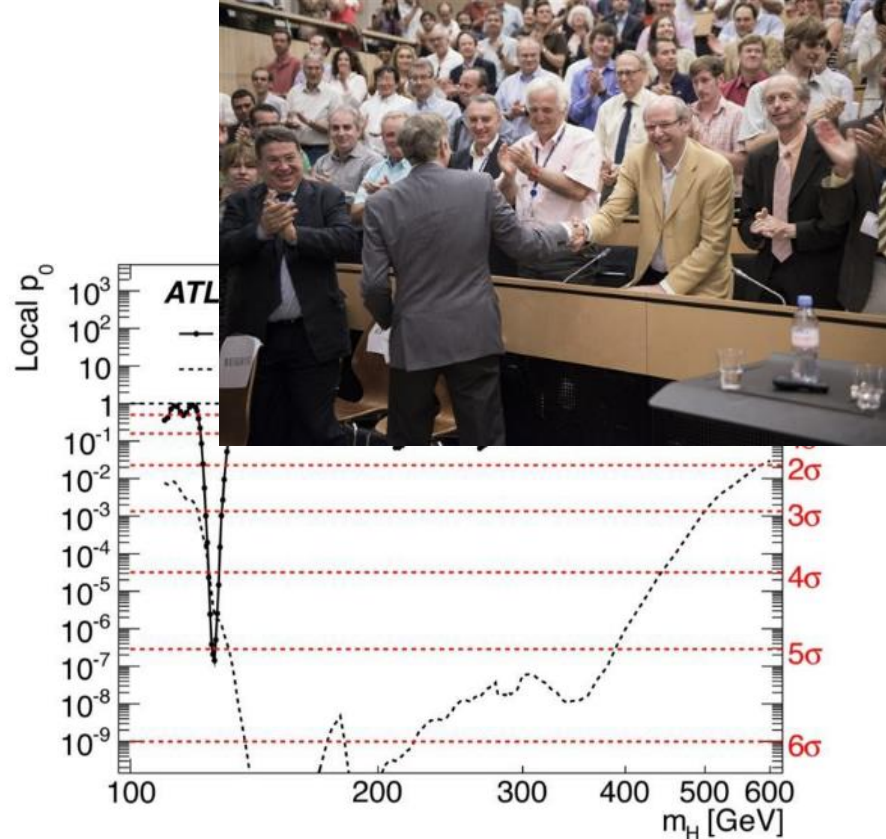
Results so far

- Many... the most spectacular one being
- 04 July 2012: Discovery of a “Higgs-like particle”



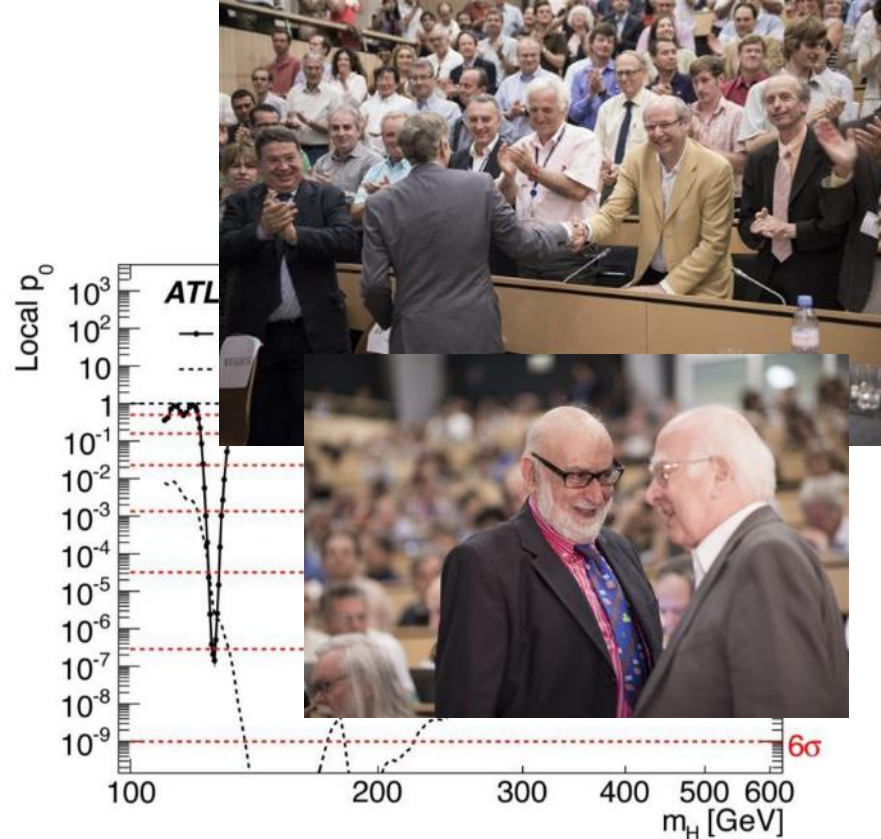
Results so far

- Many... the most spectacular one being
- 04 July 2012: Discovery of a “Higgs-like particle”



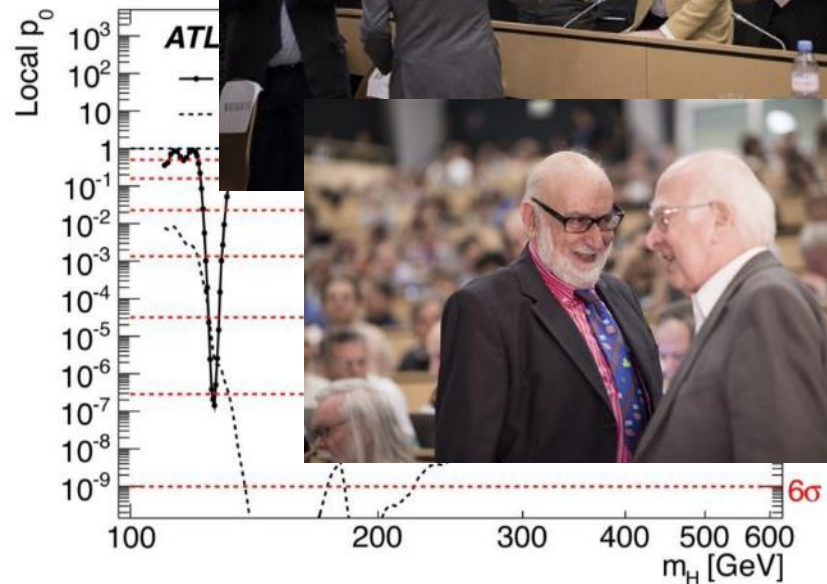
Results so far

- Many... the most spectacular one being
- 04 July 2012: Discovery of a “Higgs-like particle”



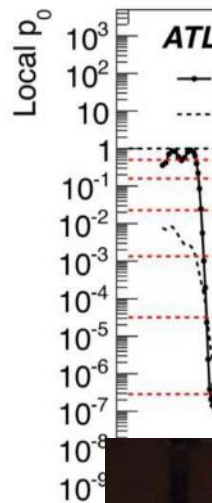
Results so far

- Many... the most spectacular one being
- 04 July 2012: Discovery of a “Higgs-like particle”
- March 2013: The particle is indeed a Higgs boson



Results so far

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

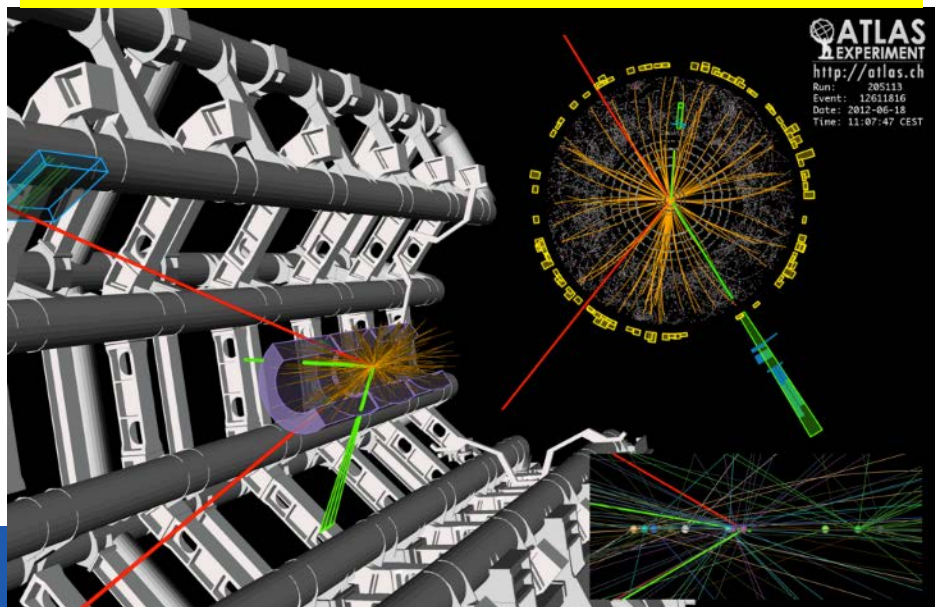


06-Jul-2016

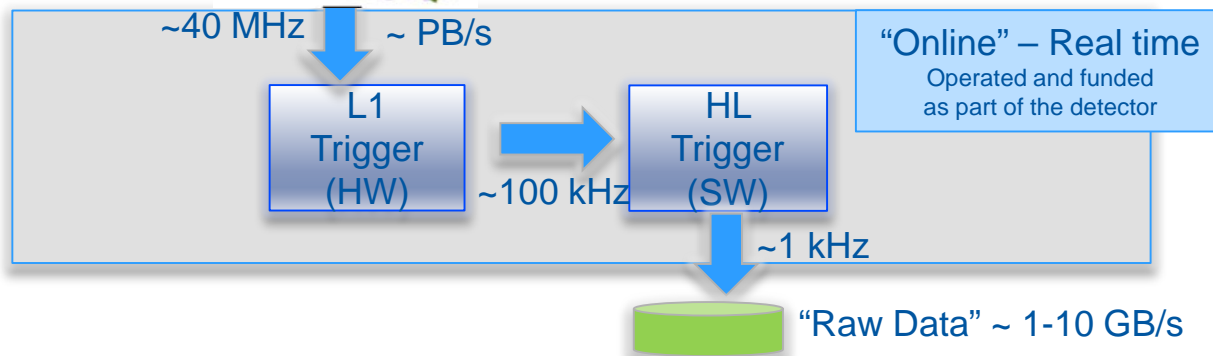
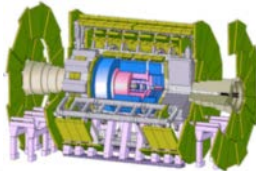
What is the LHC Data?

- 150 million sensors deliver data ... 40 million times per second
- Generates ~ 1 PB per second

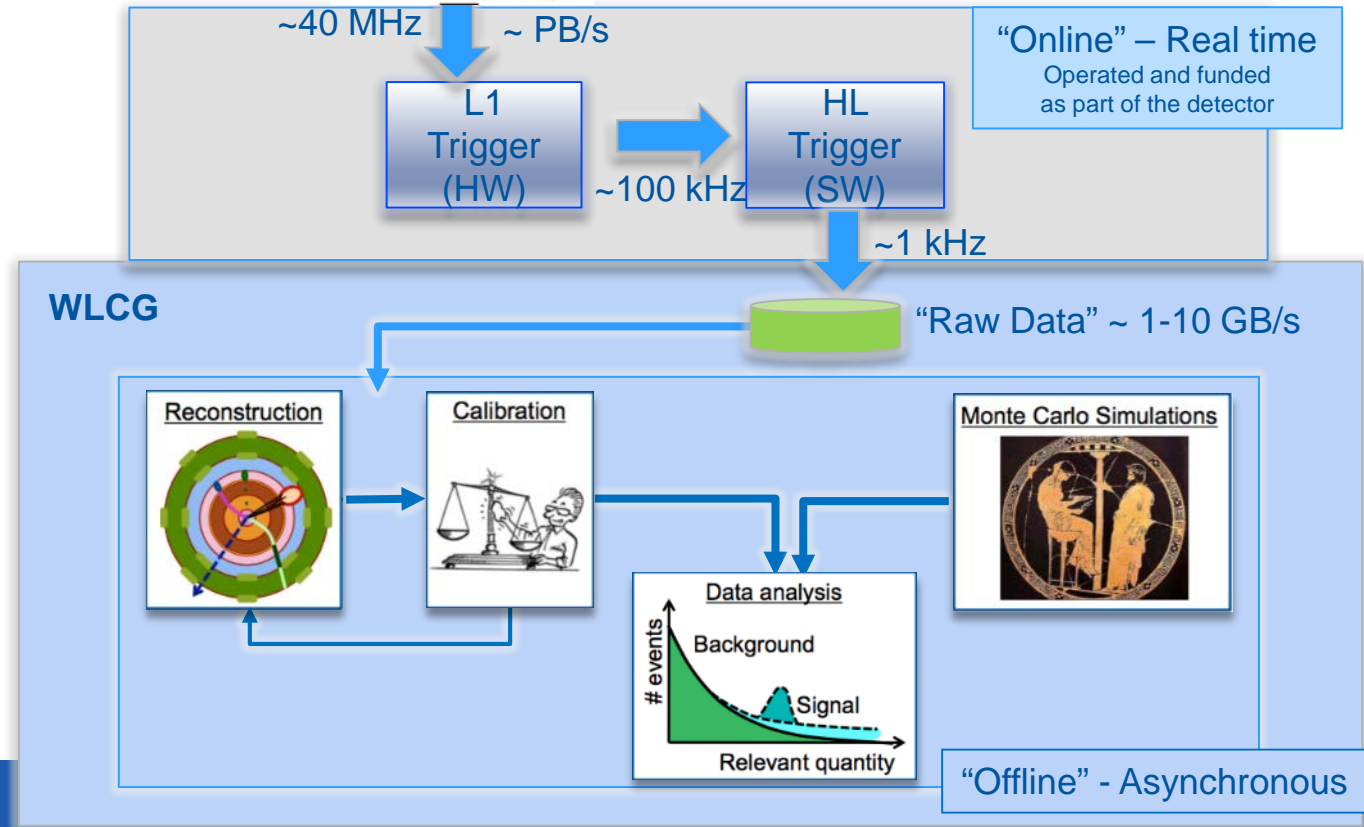
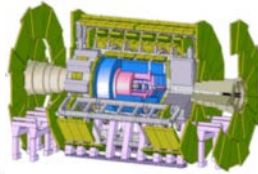
- Raw data:
 - Was a sensor hit?
 - How much energy?
 - What time?
- Reconstructed data:
 - Momentum of tracks (4-vectors)
 - Origin
 - Energy in clusters (jets)
 - Particle type
 - Calibration information
 - ...



HEP Computing



HEP Computing



Nature of the Computing Problem

- Enormous numbers of collisions of proton bunches with each other
 - Data from each collision are small (order 1...10 MB)
 - Each collision independent of all others
- No supercomputers needed
 - Most cost-effective solution is standard PC architecture (x86) servers with 2 sockets, SATA drives, Ethernet network
 - Linux (RHEL variants: Scientific Linux, CentOS) used everywhere
- Calculations are mostly combinatorics – integer (rather than floating-point) intensive

Scale of the Computing Problem

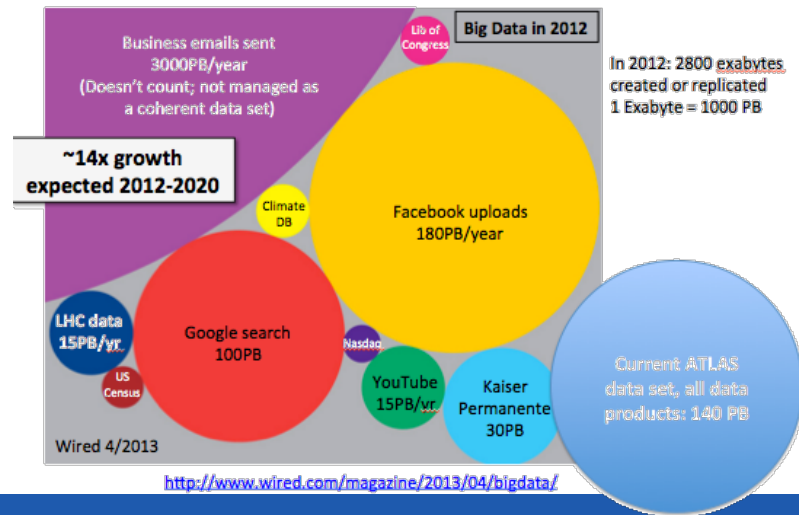
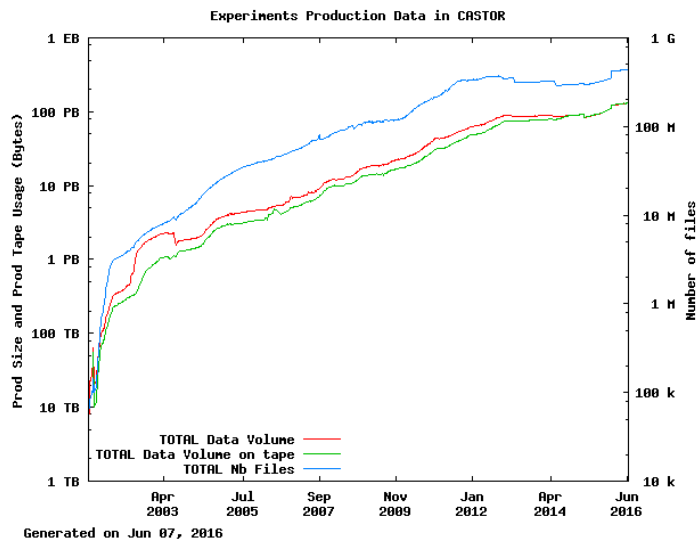
- Raw data: order 1...10 MB per collision event
 - 1 kHz, for $\sim 7 \cdot 10^6$ live seconds / year
 - 7 PB/year per detector
- Several copies, derived data sets, replicated many times for performance, accessibility, etc



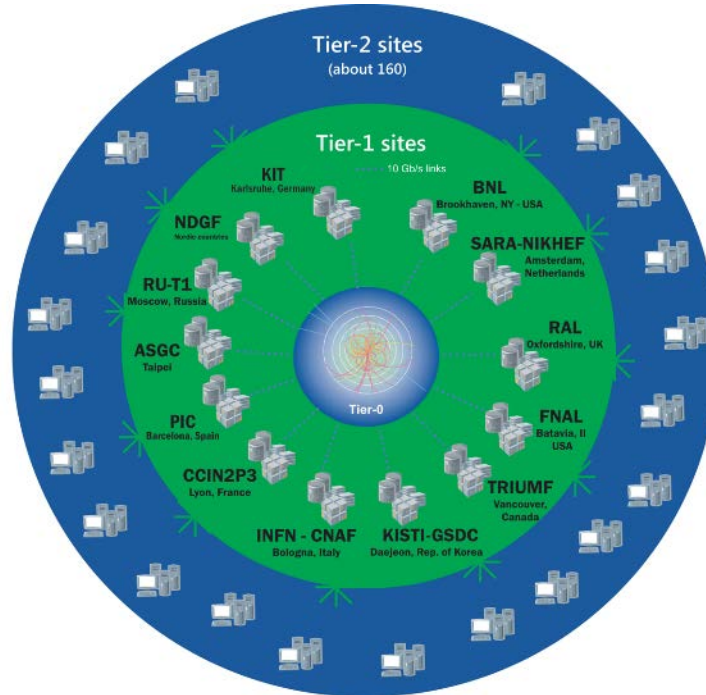
ATLAS (for example) has a managed data set of ~ 200 PB



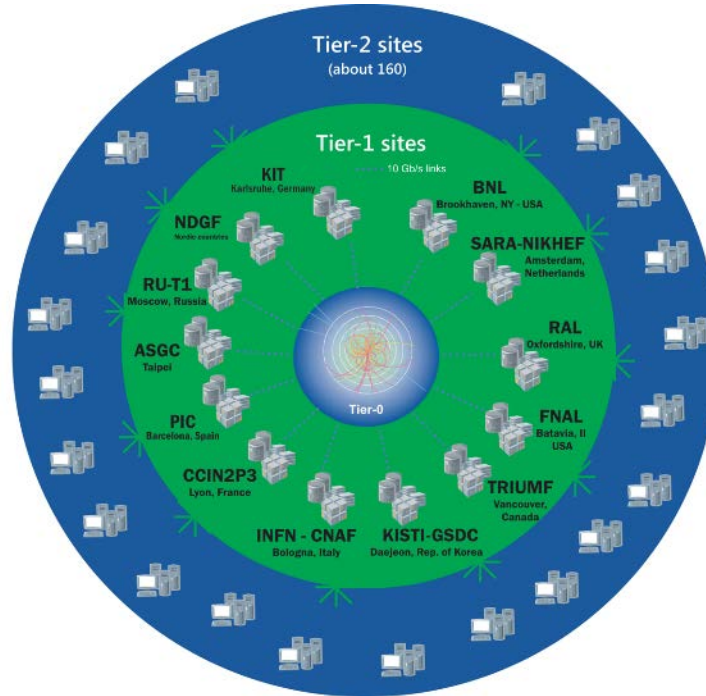
CERN data archive on tape is ~ 130 PB



The Worldwide LHC Computing Grid



The Worldwide LHC Computing Grid



WLCG: An international collaboration to distribute and analyse LHC data

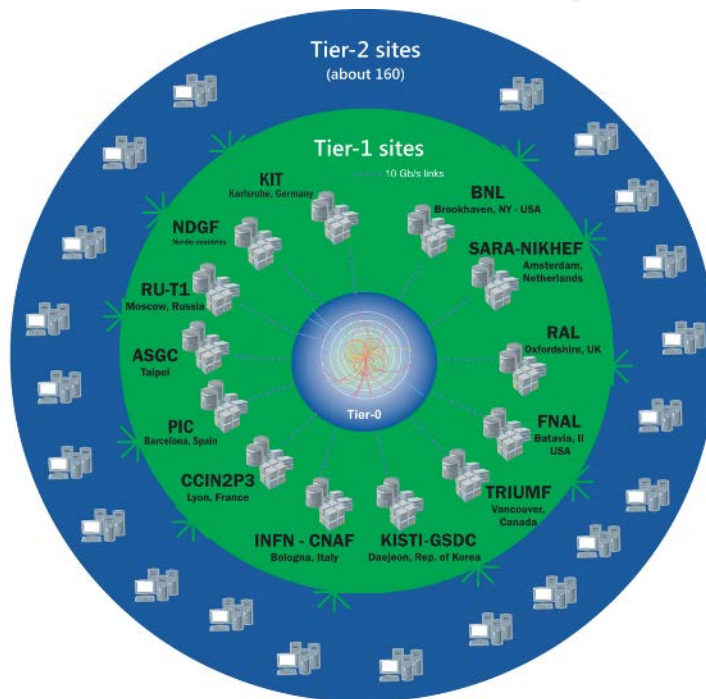
Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists

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

~600'000 cores

500 PB of storage

> 2 million
jobs/day

10-100 Gb links

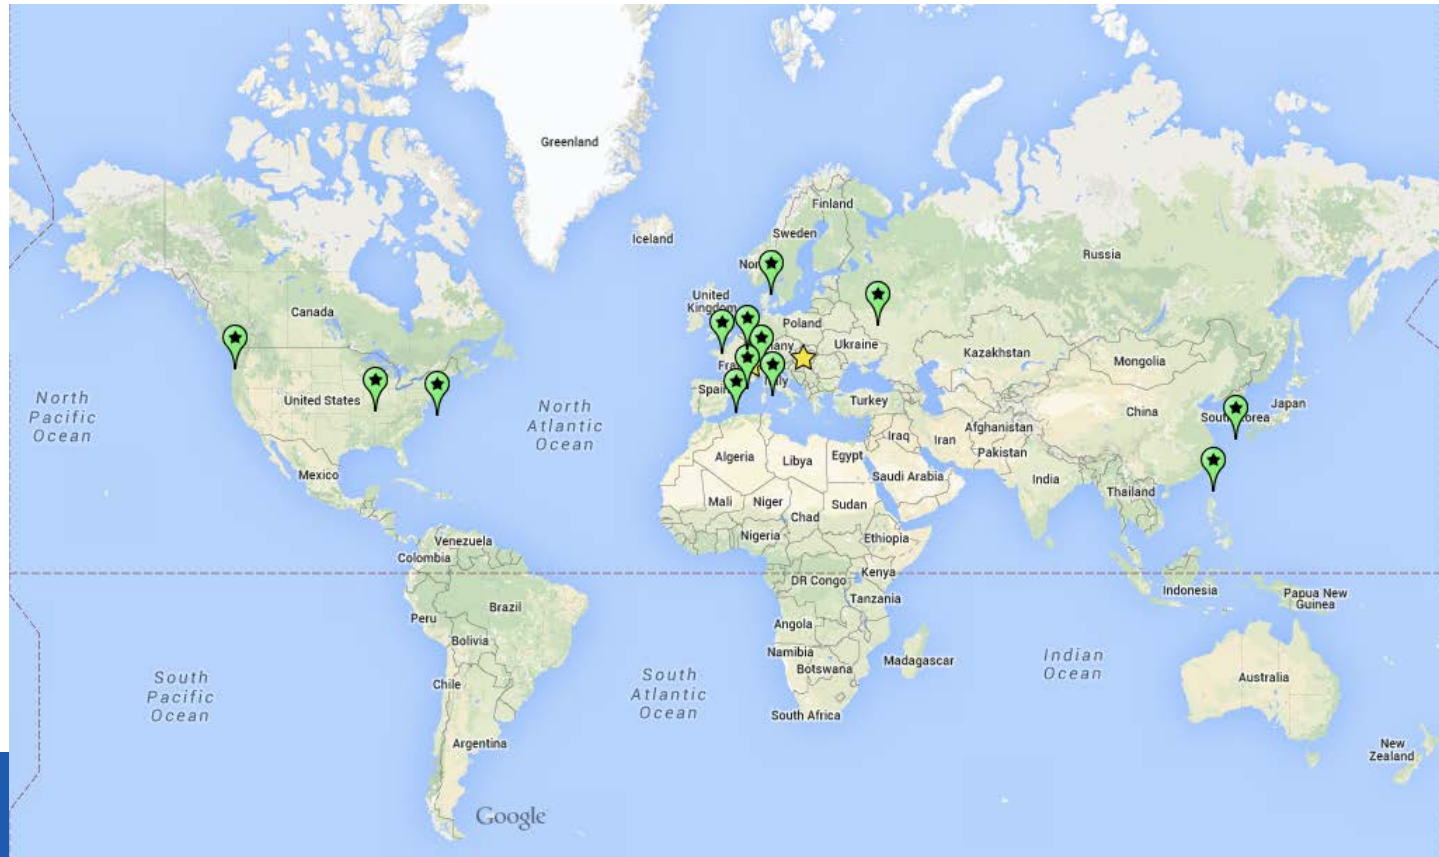
WLCG: An international collaboration to distribute and analyse LHC data

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

WLCG – a World-wide Infrastructure



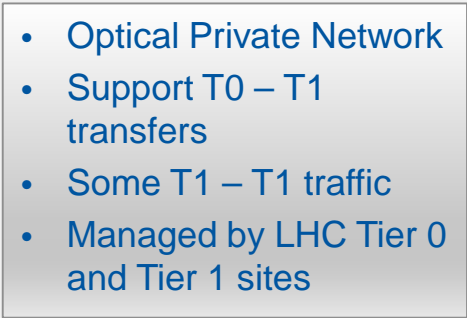
WLCG – a World-wide Infrastructure



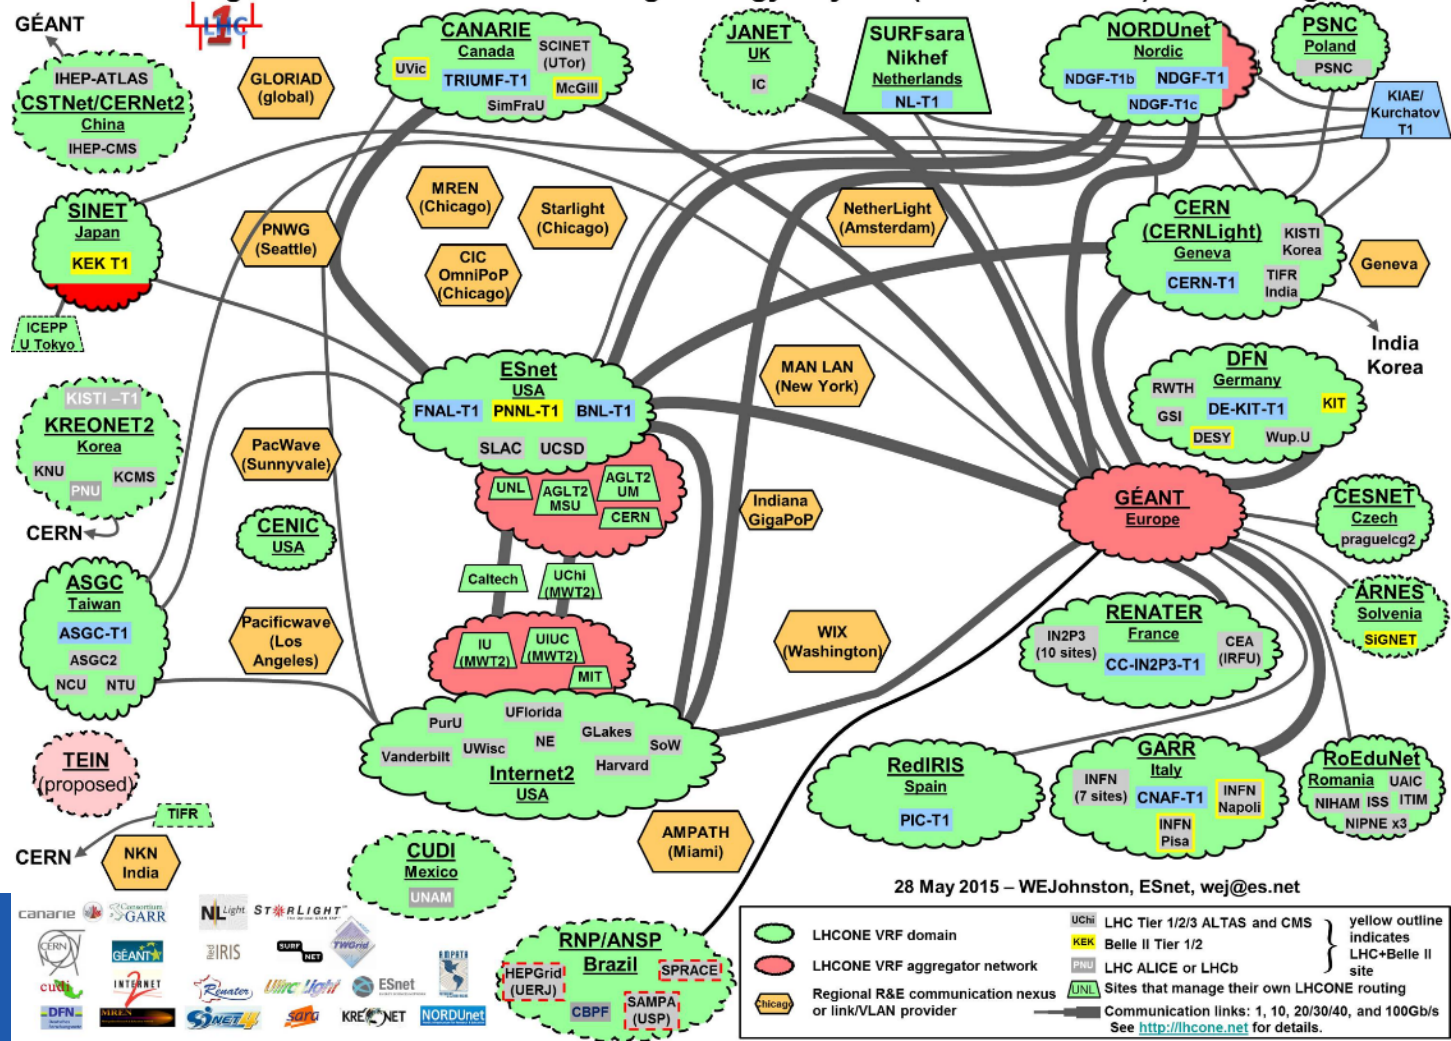
WLCG – a World-wide Infrastructure



2G ASnet



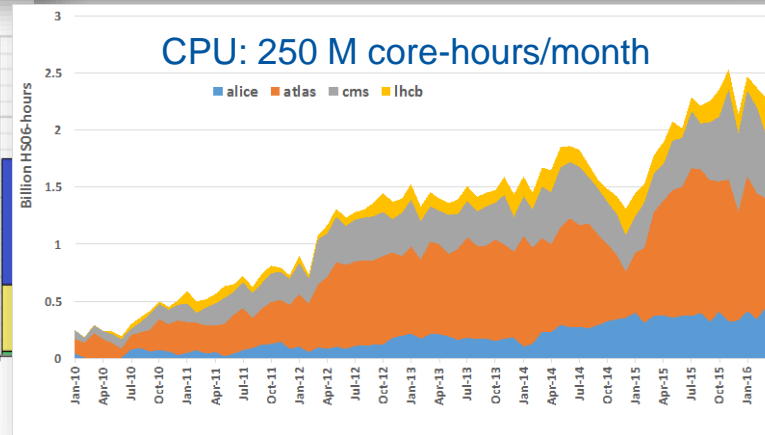
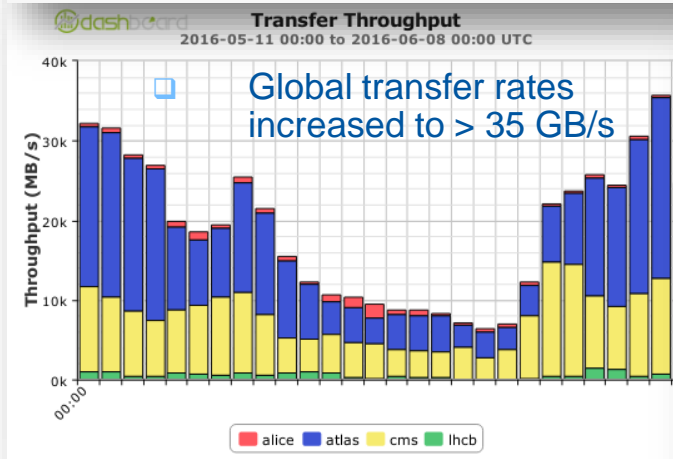
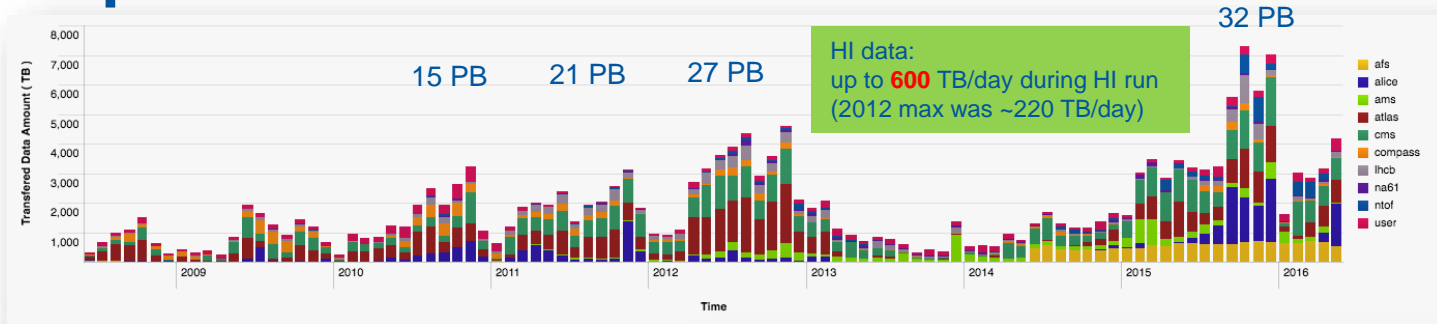
LHCONE: A global infrastructure for the High Energy Physics (LHC and Belle II) data management



28 May 2015 – WEJohnston, ESnet, wej@es.net



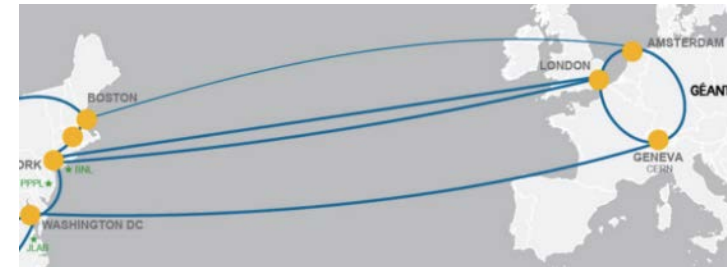
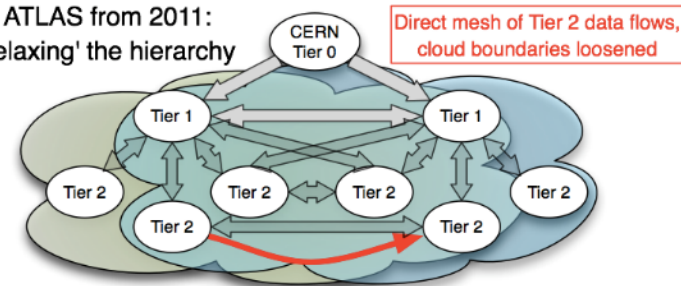
Examples of scale



Distributed model

- Performance & reliability of the networks has exceeded early expectations or fears
 - 10 Gb/s → 100 Gb/s at large centres
 - >100 Gb/s transatlantic links now in place
 - Many Tier 2s connected at 10 Gb/s or better
 - NB. Still concern over connectivity at sites in less-well connected countries
- Strict hierarchical model of Tiers evolved even during Run 1 to make best use of facilities available
 - Move away from the strict roles of the Tiers to more functional and service quality based
 - Better use of the overall distributed system
- Focus on use of resources/capabilities rather than “Tier roles”
 - Data access peer-peer: removal of hierarchical structure

ATLAS from 2011:
'relaxing' the hierarchy



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
● Number 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	13,712

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
● Number 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	13,712

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
● Number 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	13,712

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
● Number 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	13,712

Transforming In-House Resources (1)

Before Wigner deployment:

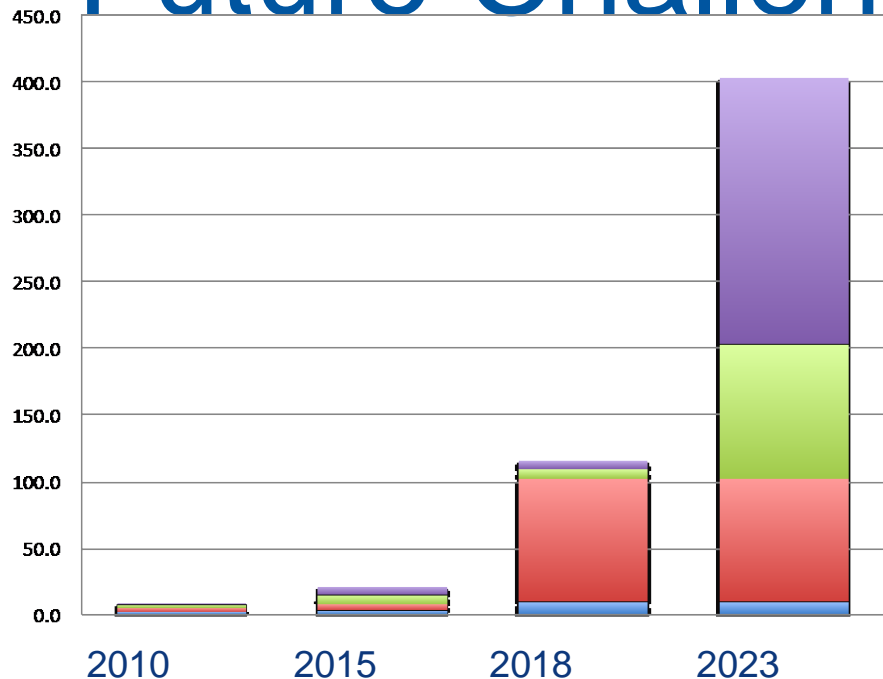
- Physical servers only
 - Inefficient resource usage
 - Strong coupling of services with HW life-cycle
- Vertical view
 - Service managers responsible for entire stack
- Home-made tools of 10 years ago
 - Successful at the time, but Increasingly brittle
 - Lack of support for dynamic host creation/deletion
 - Limited scalability
- Person-power: (at best) constant
 - ... despite many more machines

Transforming In-House Resources (2)

Current situation:

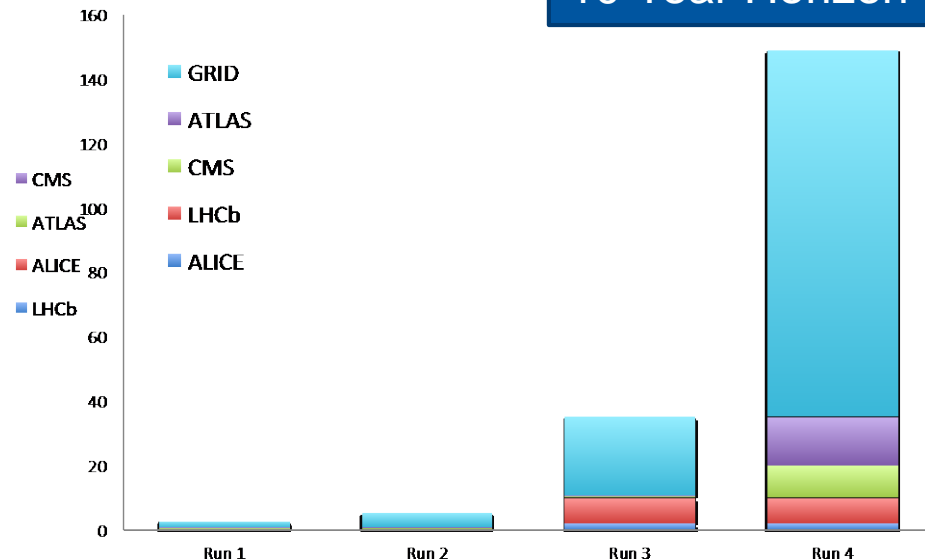
- Full support for physical and virtual servers
- Full support for remote machines
- Horizontal view
 - Responsibilities by layers of service deployment
- Large fraction of resources run as private cloud under OpenStack
- Scaling to large numbers
(> 15'000 physical, several 100'000s virtual)
- Support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

Future Challenges for LHC



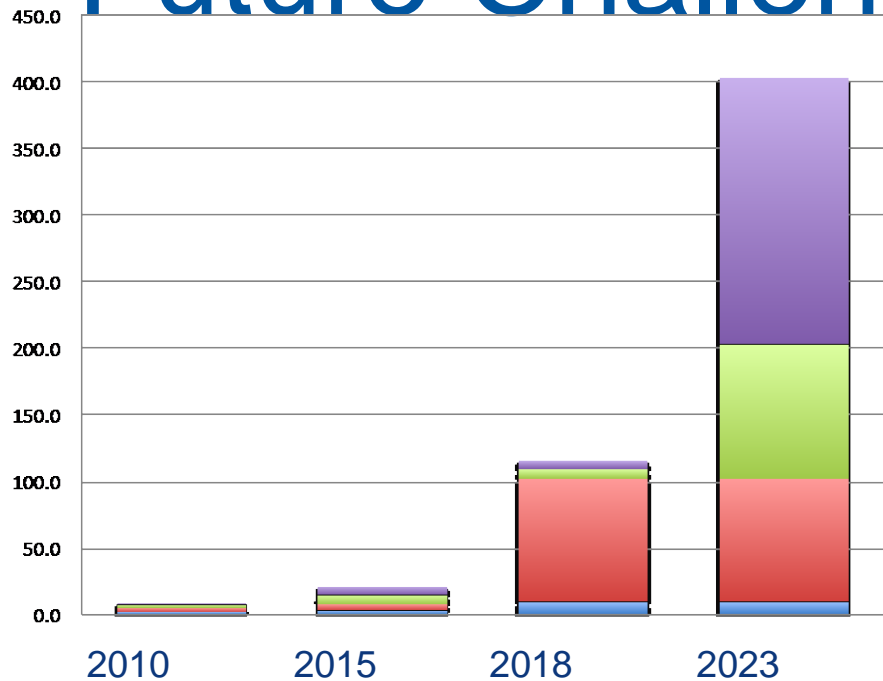
Data: ~25 PB/yr → 400 PB/yr

10 Year Horizon

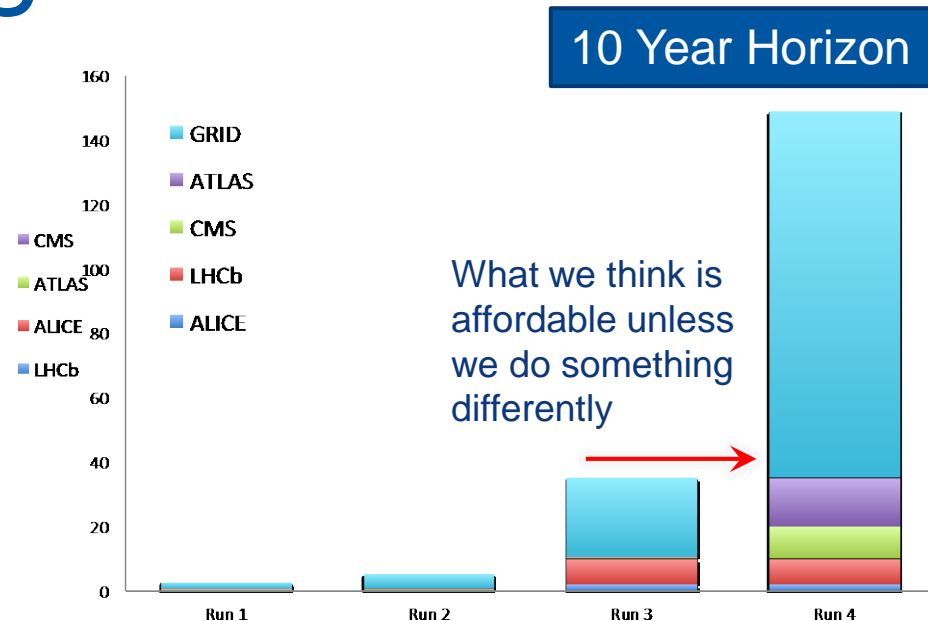


Compute: Growth > x50

Future Challenges for LHC



Data: ~25 PB/yr → 400 PB/yr

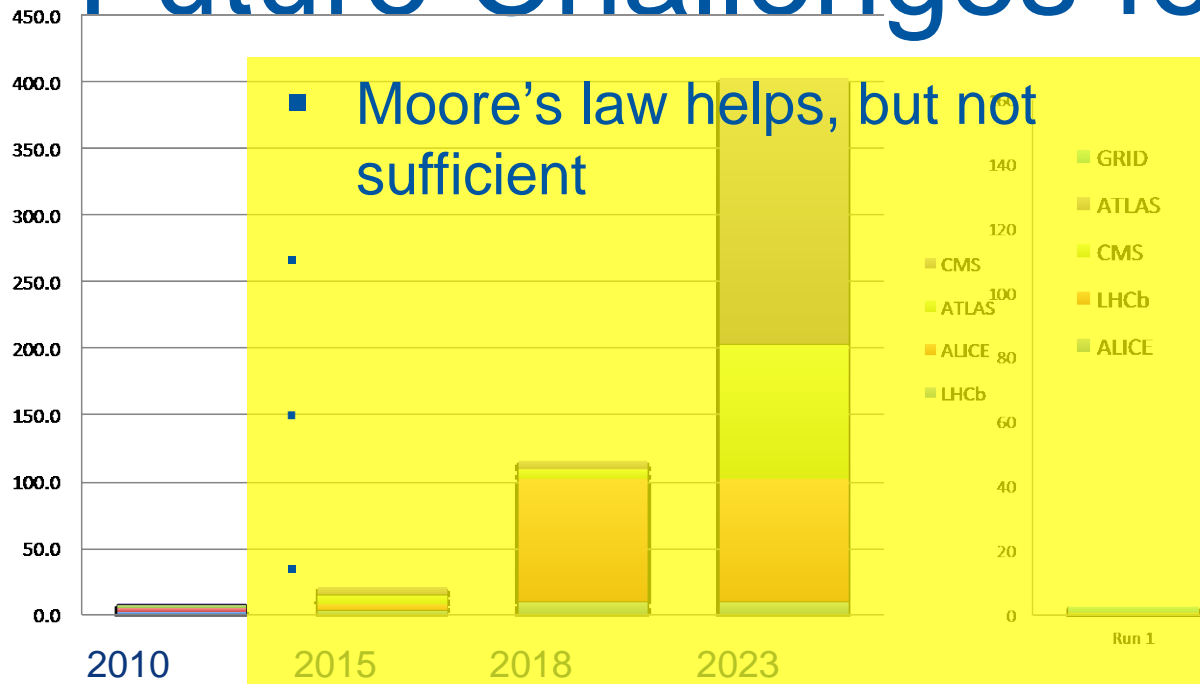


10 Year Horizon

What we think is affordable unless we do something differently

Compute: Growth > x50

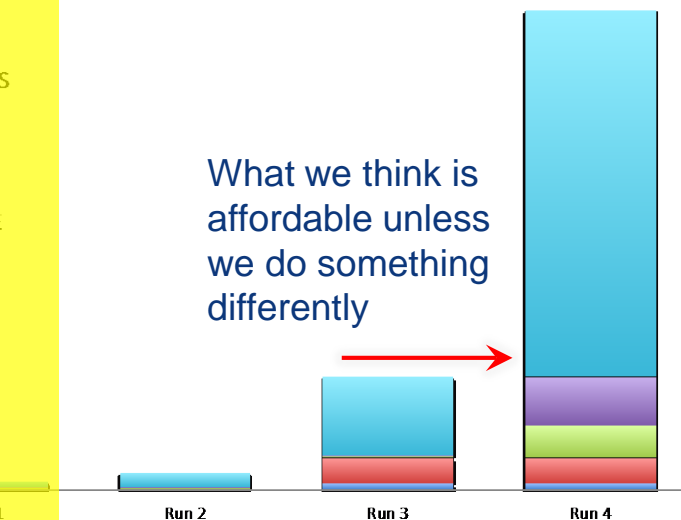
Future Challenges for LHC



Data: ~25 PB/yr → 400 PB/yr

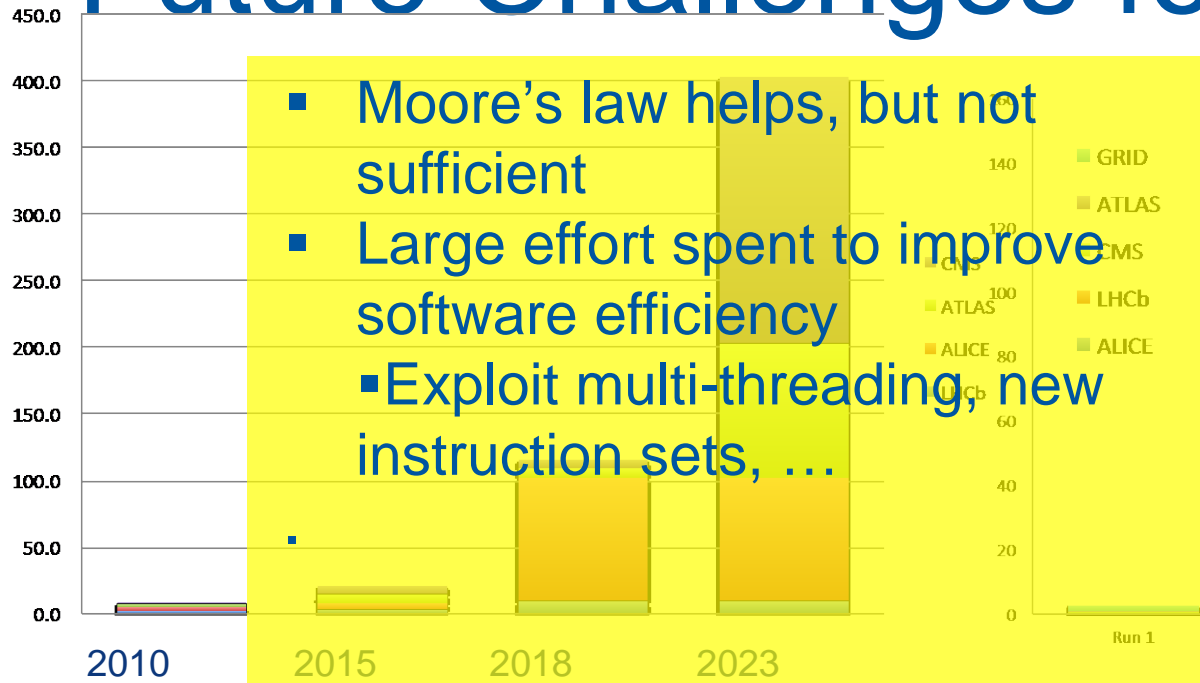
10 Year Horizon

What we think is affordable unless we do something differently



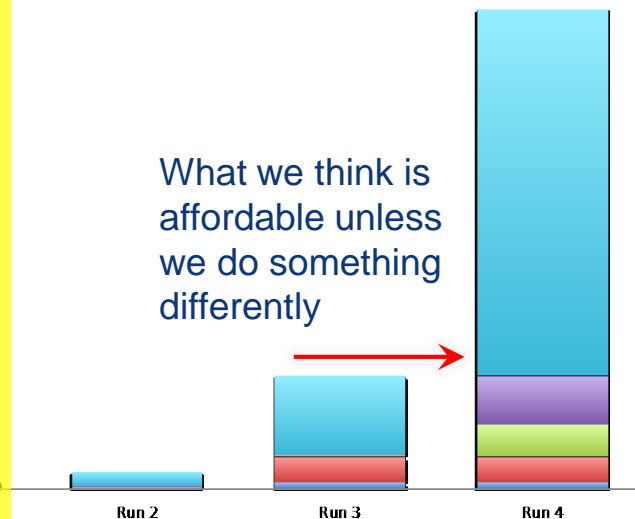
Compute: Growth > x50

Future Challenges for LHC



10 Year Horizon

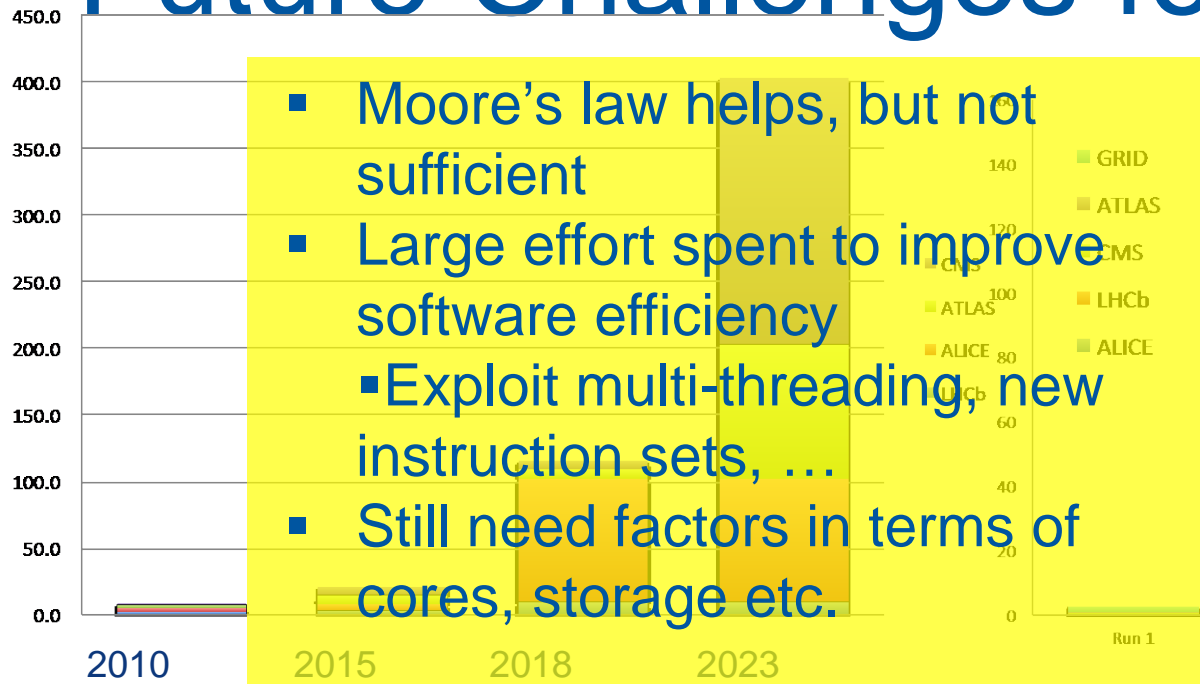
What we think is affordable unless we do something differently



Data: ~25 PB/yr → 400 PB/yr

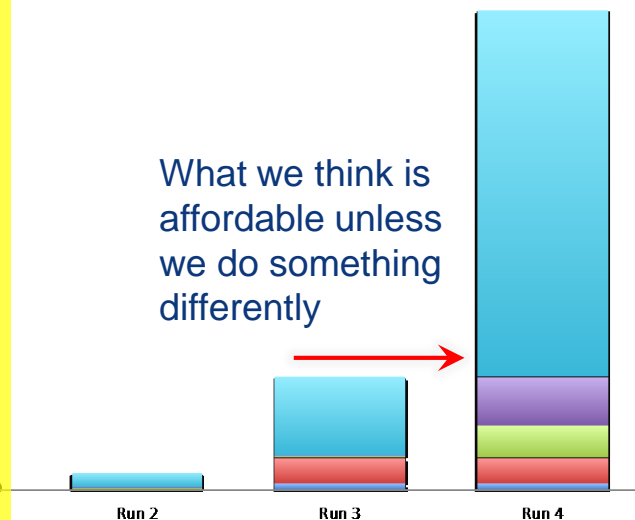
Compute: Growth > x50

Future Challenges for LHC



10 Year Horizon

What we think is affordable unless we do something differently



Data: ~25 PB/yr → 400 PB/yr

Compute: Growth > x50

Trends – Software

- Recognizing the need to re-engineer HEP software
 - New architectures, parallelism everywhere, vectorisation, data structures, ...
- HEP Software Foundation (HSF) set up (<http://hepsoftwarefoundation.org>)
 - Community wide – buy-in from major labs, experiments, projects
 - Goals:
 - Address rapidly growing needs for simulation, reconstruction and analysis of current and future HEP experiments
 - Promote the maintenance and development of common software projects and components for use in current and future HEP experiments
 - Enable the emergence of new projects that aim to adapt to new technologies, improve the performance, provide innovative capabilities or reduce the maintenance effort
 - Enable potential new collaborators to become involved
 - Identify priorities and roadmaps
 - Promote collaboration with other scientific and software domains

Opportunistic resources

- Today this has become more important
 - Opportunistic use of:
 - HPC facilities
 - Large cloud providers
 - Other offers for “off-peak” or short periods
 - ...
 - All at very low or no cost (for hardware)
 - But scale and cost are unpredictable
- Also growing in importance:
 - Volunteer computing (citizen science)
 - BOINC-like (LHC@home, ATLAS/CMS/LHCb@home, etc)
 - Now can be used for many workloads – as well as the outreach opportunities

Drivers of Change

- Must reduce the (distributed) provisioning layer of compute to something simple, we need a hybrid and be able to use:
 - Our own resources
 - Commercial resources
 - Opportunistic use of clouds, grids, HPC, volunteer resources, etc.
- Move towards simpler site management
 - Reduce operational costs at grid sites
 - Reduce “special” grid middleware support cost
- Today (2015) it is cheaper for us to operate our own data centres
 - We use 100% of our resources 24x365
- We also get a large synergistic set of resources in many Tier 2s – essentially for “free” – over and above the pledged resources
- However, commercial pricing is now getting more competitive
 - Large scale hosting contracts, commercial cloud provisioning

Scaling up Further: Commercial 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: Commercial Clouds (2)

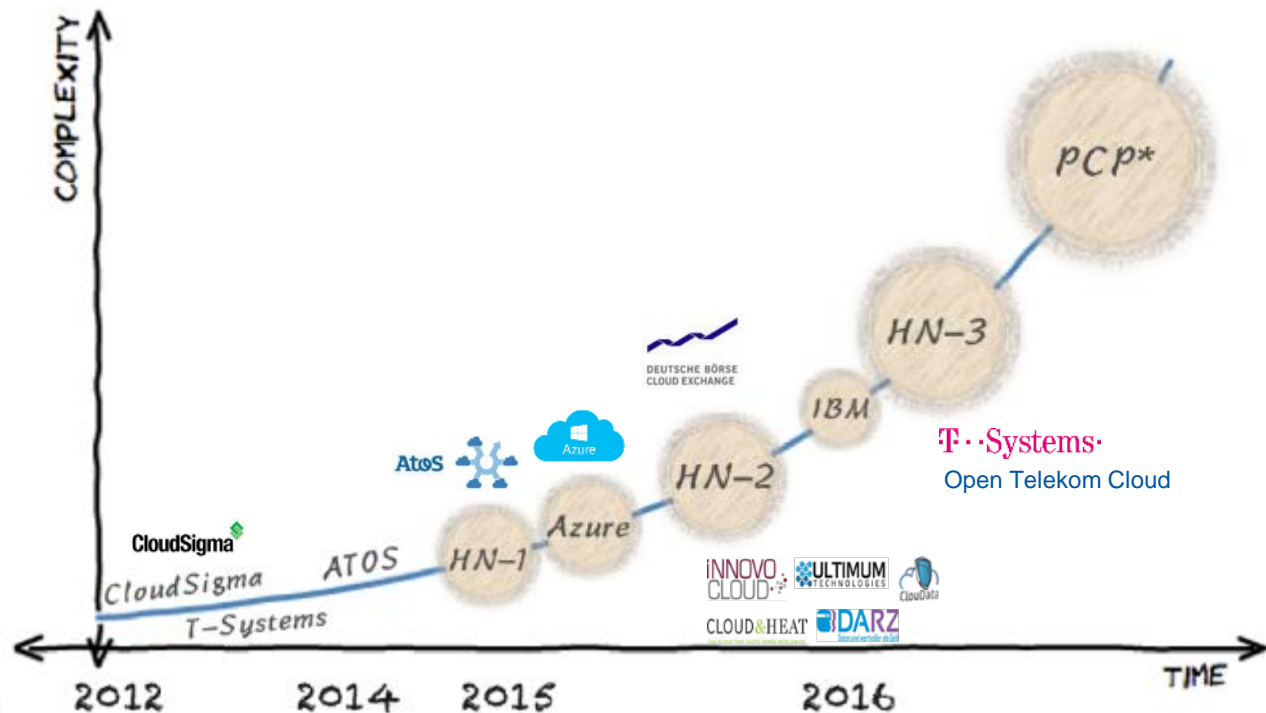
- Potential issues:
 - Cloud provider's business models not well adapted to procurement rules and procedures of public organisations
 - Lack of skills for and experience with procurements
 - Market largely not targeting compute-heavy tasks
 - Performance metrics/benchmarks not established
 - Legal impediments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

Science Clouds

- Experiments and sites have made many explorations of use of private and commercial clouds
 - Cloud infrastructures at many sites
 - Use of AWS, Google, Rackspace etc. by FNAL, BNL, CERN, experiments, others
 - Helix Nebula EC project in Europe (together with other sciences)
 - Also testing real commercial procurements to understand cost
 - So far most use has been simulation, only now looking at data-intensive use cases

CERN Approach to Commercial Clouds

Series of short procurement projects of increasing size and complexity

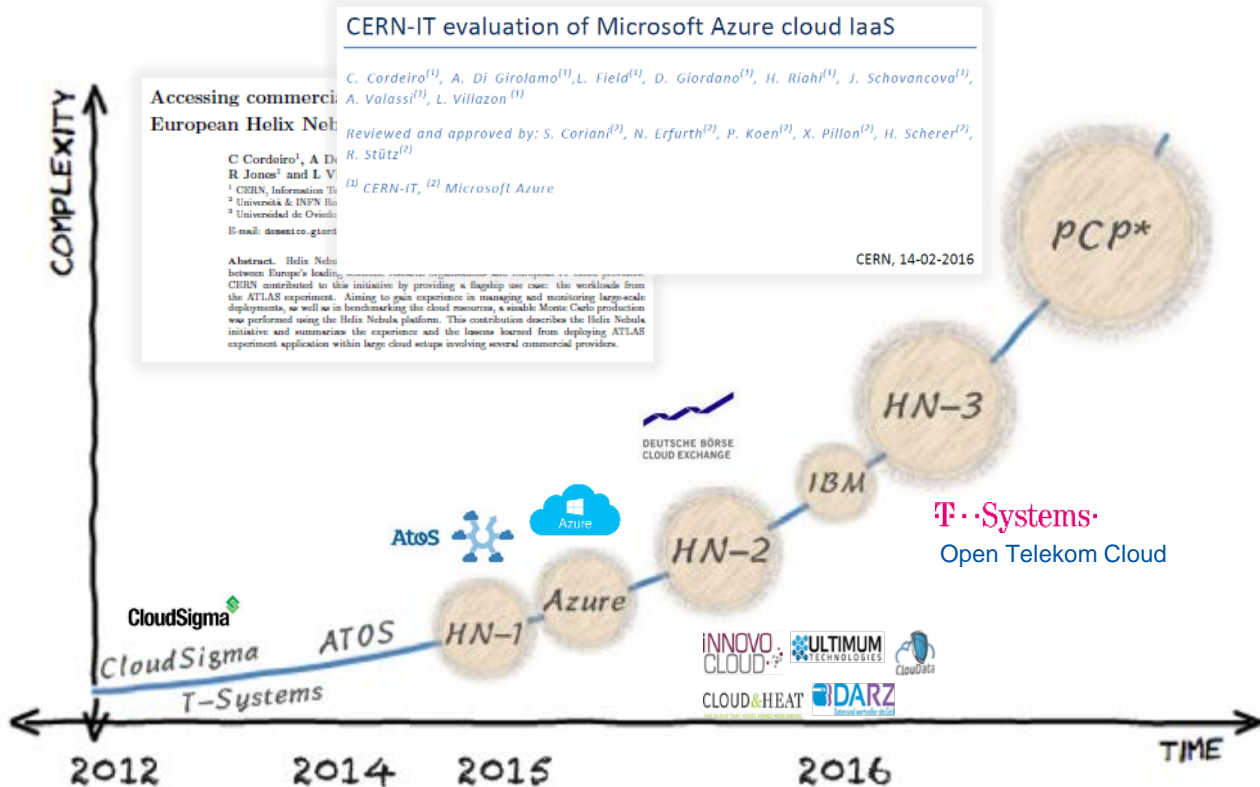


Series of short procurement projects of increasing size and complexity



CERN Approach to Commercial Clouds

Series of short procurement projects of increasing size and complexity

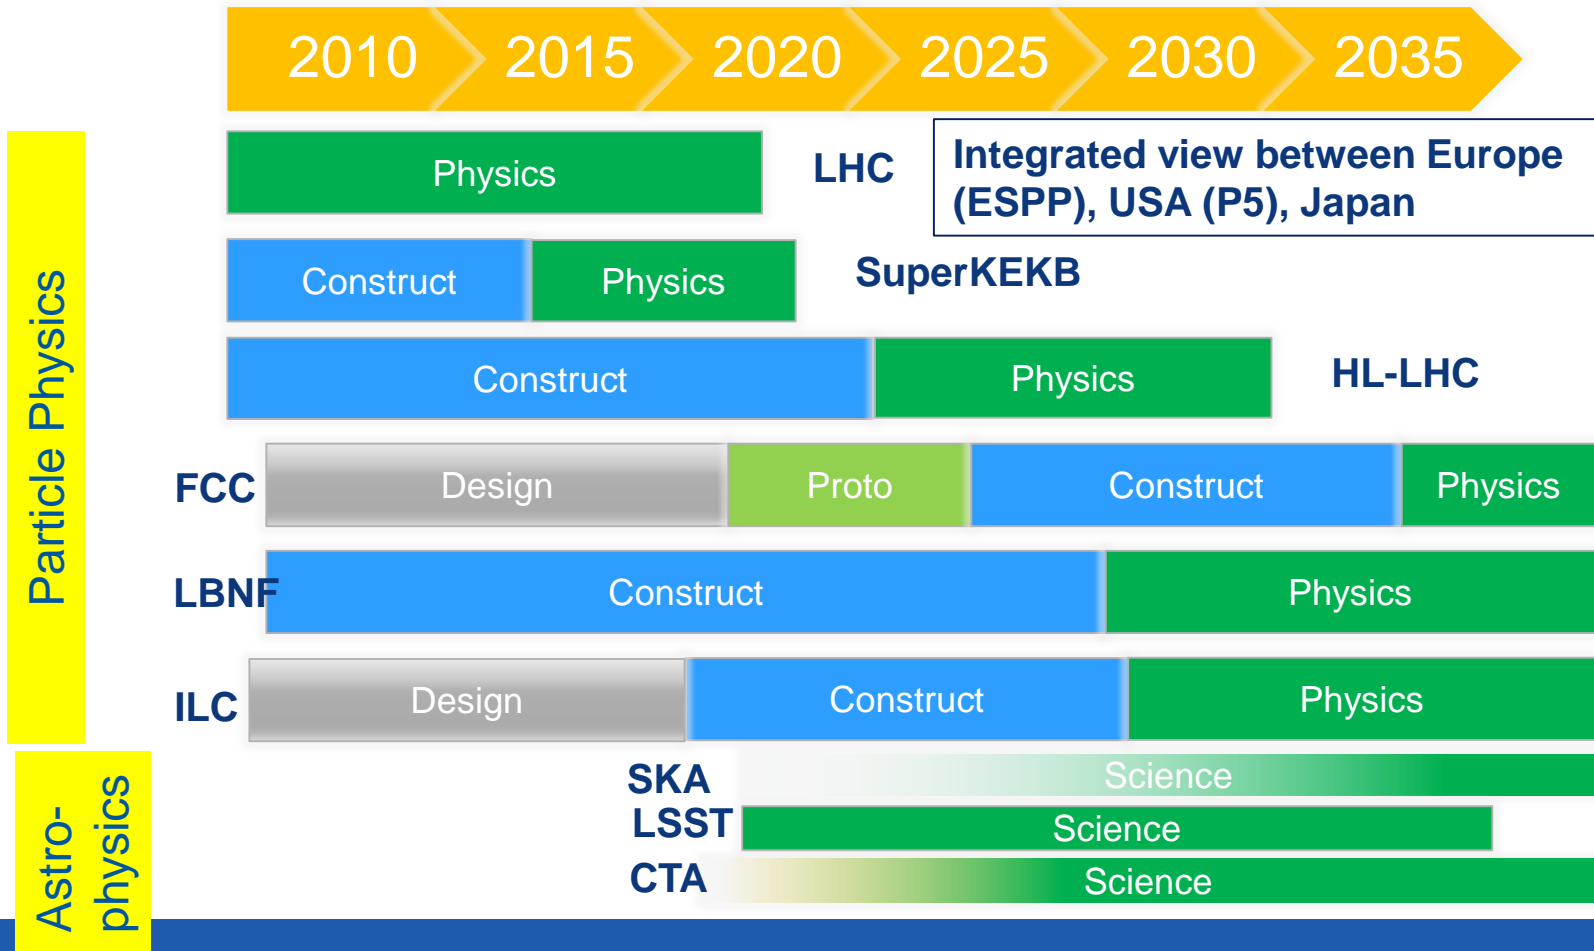


CERN Procurements

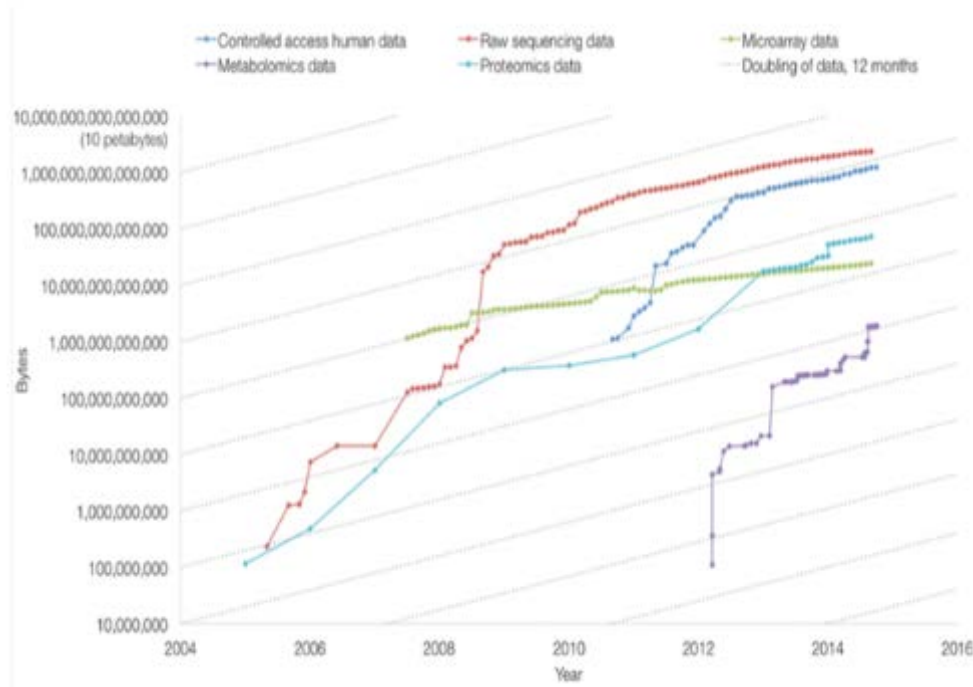
- HN-1 (ATOS): Detector simulation for ATLAS
- (Microsoft Azure evaluation)
- HN-2 (DBCE): Detector simulation for ATLAS, CMS, ALICE, LHCb
- (IBM SoftLayer evaluation)
- HN-3 (T-Systems OTC): Detector simulation, reconstruction and analysis for ATLAS, CMS, ALICE, LHCb
 - Using 4'000 cores July – September
 - Includes storage for data-heavy workflows

Future Challenges beyond LHC

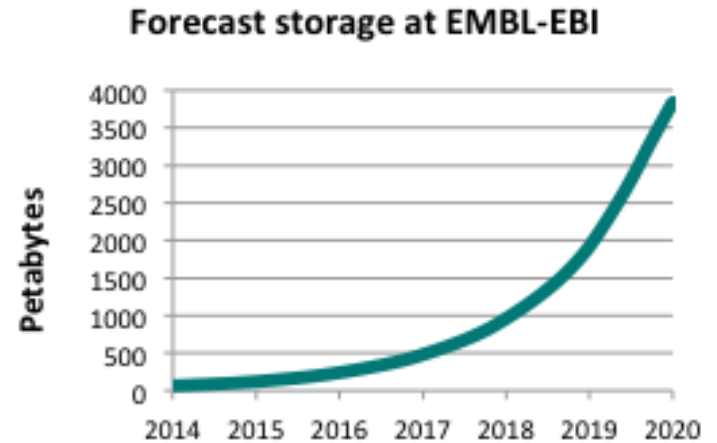
- Not only LHC, but a number of particle physics projects with high data rates
- Not only particle physics, but also other physics fields (e.g. astronomy)
- Not only physics, but also other sciences (e.g. life sciences, material science)



Not only Physics



Growth of EBI (European Bioinformatics Institute) repositories, lines are 12 month doubling



HELIX NEBULA Science Cloud

Joint Pre-Commercial Procurement

Procurers: CERN, CNRS, DESY, EMBL-EBI, ESRF,
IFAE, INFN, KIT, SURFSara, STFC

Experts: Trust-IT & EGI.eu

The group of procurers have committed

- >1.6M€ of procurement funds
- Manpower for testing/evaluation
- Use-cases with applications & data
- In-house IT resources

To procure innovative IaaS level cloud services integrated into a hybrid cloud model

- Commercial cloud services
- European e-Infrastructures

Services will be made available to end-users from many research communities

Co-funded via H2020 (Jan'16-Jun'18)

- Grant Agreement 687614



Total procurement commitment >5M€

Computing in HEP - Helge Meinhard at CERN.ch

06-Jul-2016

33



User groups to be supported

- **High Energy Physics**

- LHC experiments
- Belle II
- COMPASS



- **Astronomy**

- CTA – Cherenkov Telescope Array
- MAGIC
- Pierre Auger Observatory



- **Life Sciences**

- ELIXIR
- Euro-BioImaging
- Pan-Cancer
- BBMRI
- WeNMR



- **Photon/Neutron science**

- PETRA III, European XFEL, 3DIX, OCEAN, OSIRIS



- **Long tail of science**



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

- Technical and administrative work well advanced
 - Consortium agreement signed
 - Subcommittees established and working
- Tender announced in Jan 2016
- Open Market consultation successfully held on 17 March 2016
- Tender material in final phase of preparation
 - To be published very soon

Conclusions

- LHC computing has successfully managed to use unprecedented data volumes in science
- Initially used purpose-built tools, many of more general use for data-intensive science
 - Adaptation/generalisation may be needed
- More and more open-source tools and new technologies are being adopted/tested
- Future expectations of data volumes will still need further innovations
- Commercial clouds have a large potential of addressing the requirements of public research organisations for ever more resources and of dealing with peak demands
- A full integration of commercial clouds with on-premise resources and public e-infrastructures is required

Thank you for your attention

