

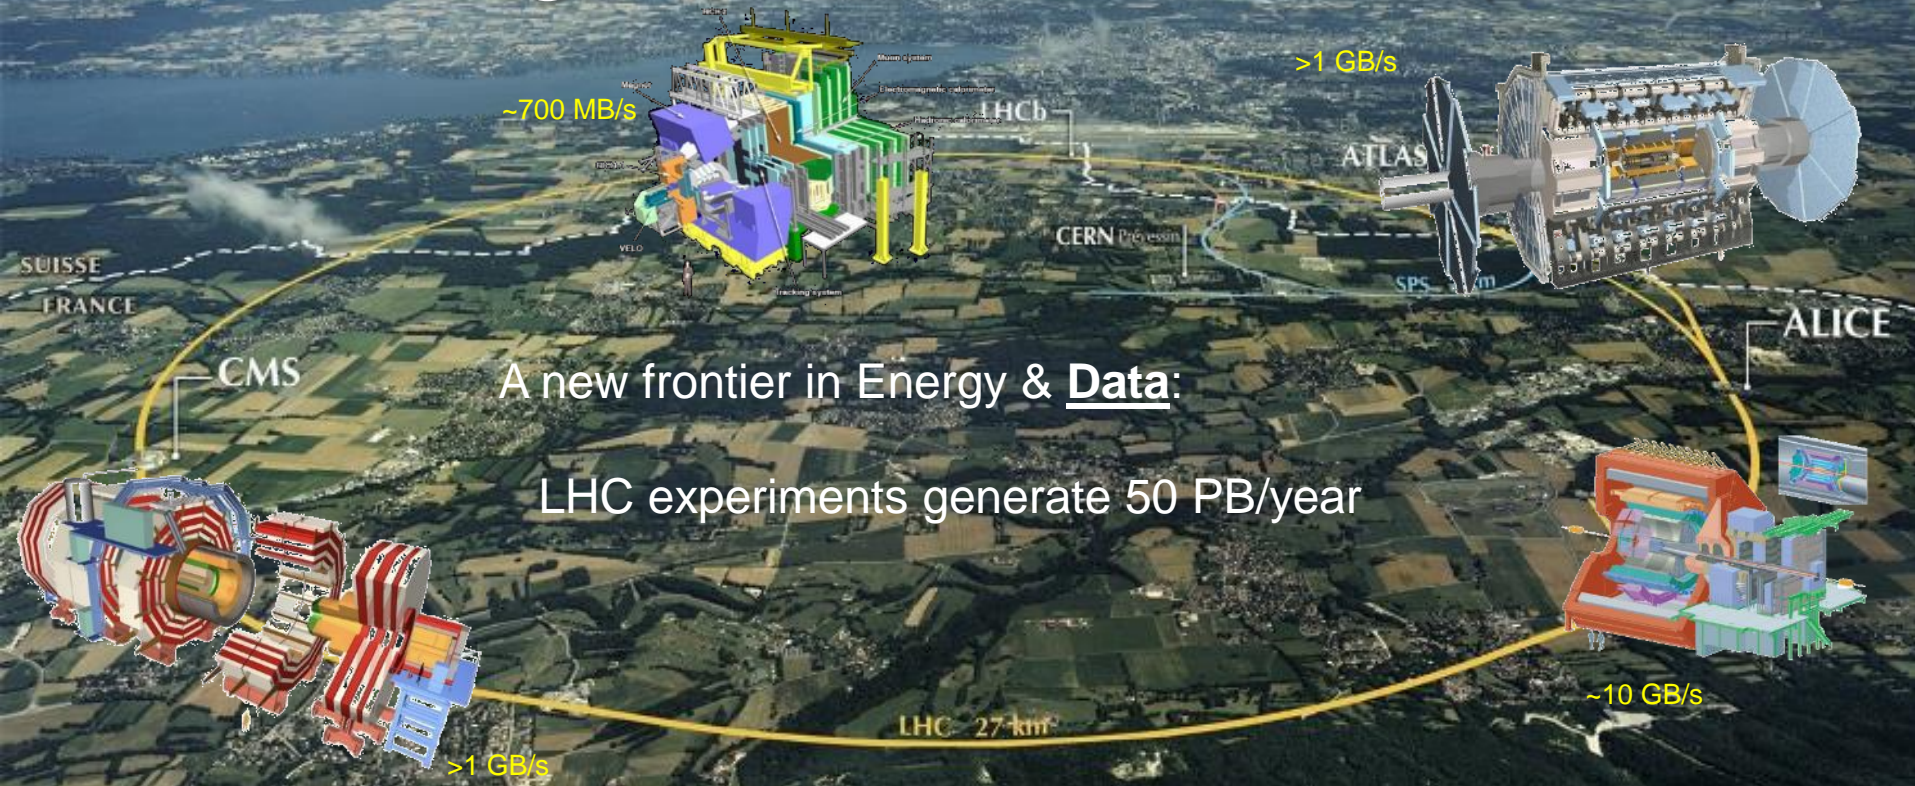
Dr. Ian Bird
CERN
LHC Computing Project Leader

CERN-GW Meeting
CERN, 1st September 2017



Computing for LHC

The Large Hadron Collider (LHC)



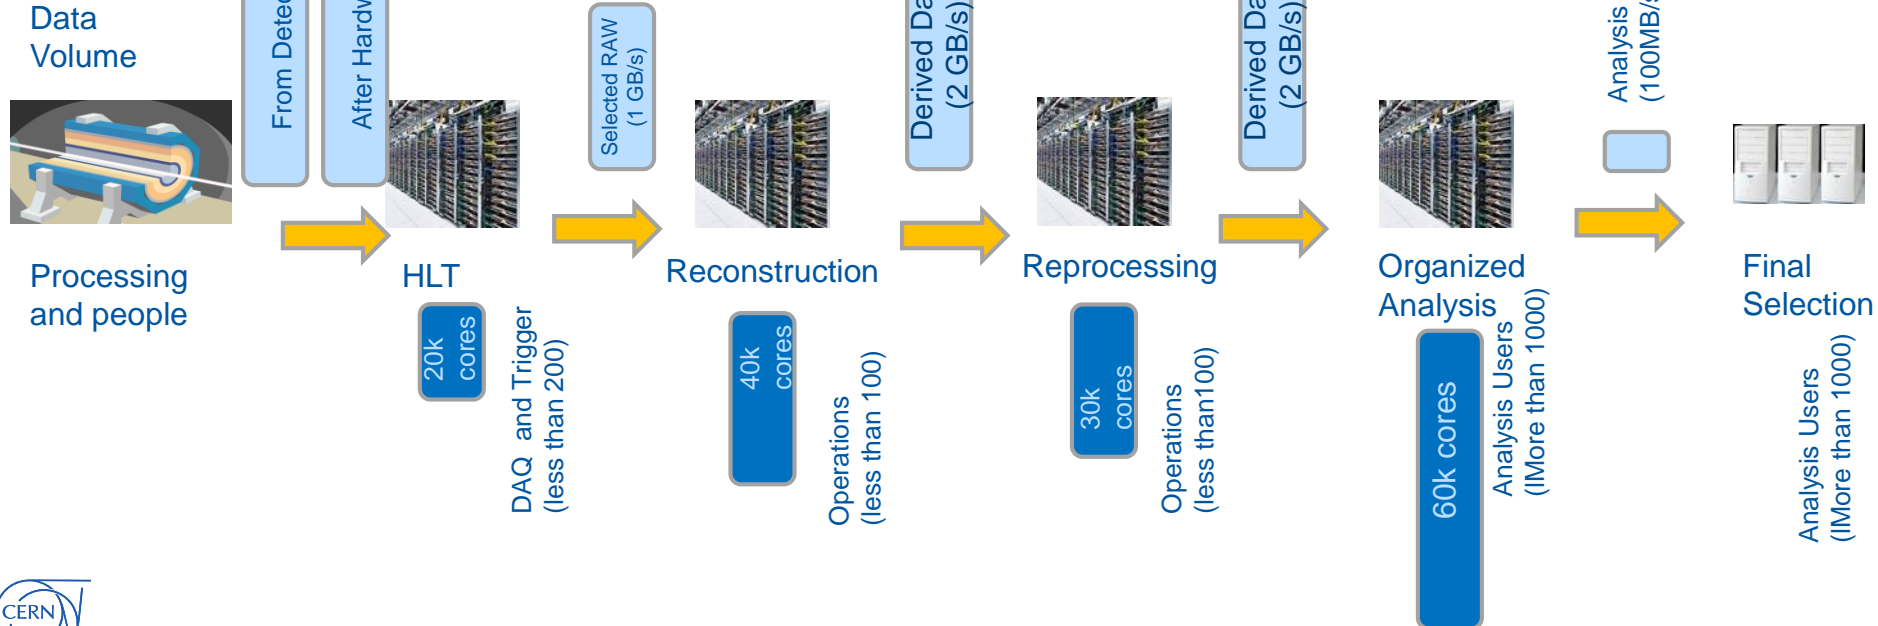
A new frontier in Energy & Data:

LHC experiments generate 50 PB/year

Data Analysis at the LHC

The process to transform raw data into useful physics datasets

- This is a complicated series of steps at the LHC (Run2)




The Worldwide LHC Computing Grid

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

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

Tier-2: Simulation,
end-user analysis



~170 sites,
42 countries

~750k CPU cores

~1 EB 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 MoU Signatures


2017:

- 63 MoU's
- 167 sites; 42 countries



Optical Private Network
Support T0 – T1 transfers
& T1 – T1 traffic
Managed by LHC Tier 0 and
Tier 1 sites

Up to 340 Gbps transatlantic



— T0-T1 and T1-T1 traffic
- - - T1-T1 traffic only
■ = Alice ■ = Atlas ■ = CMS ■ = LHCB
 edoardo.martelli@cern.ch 20161010




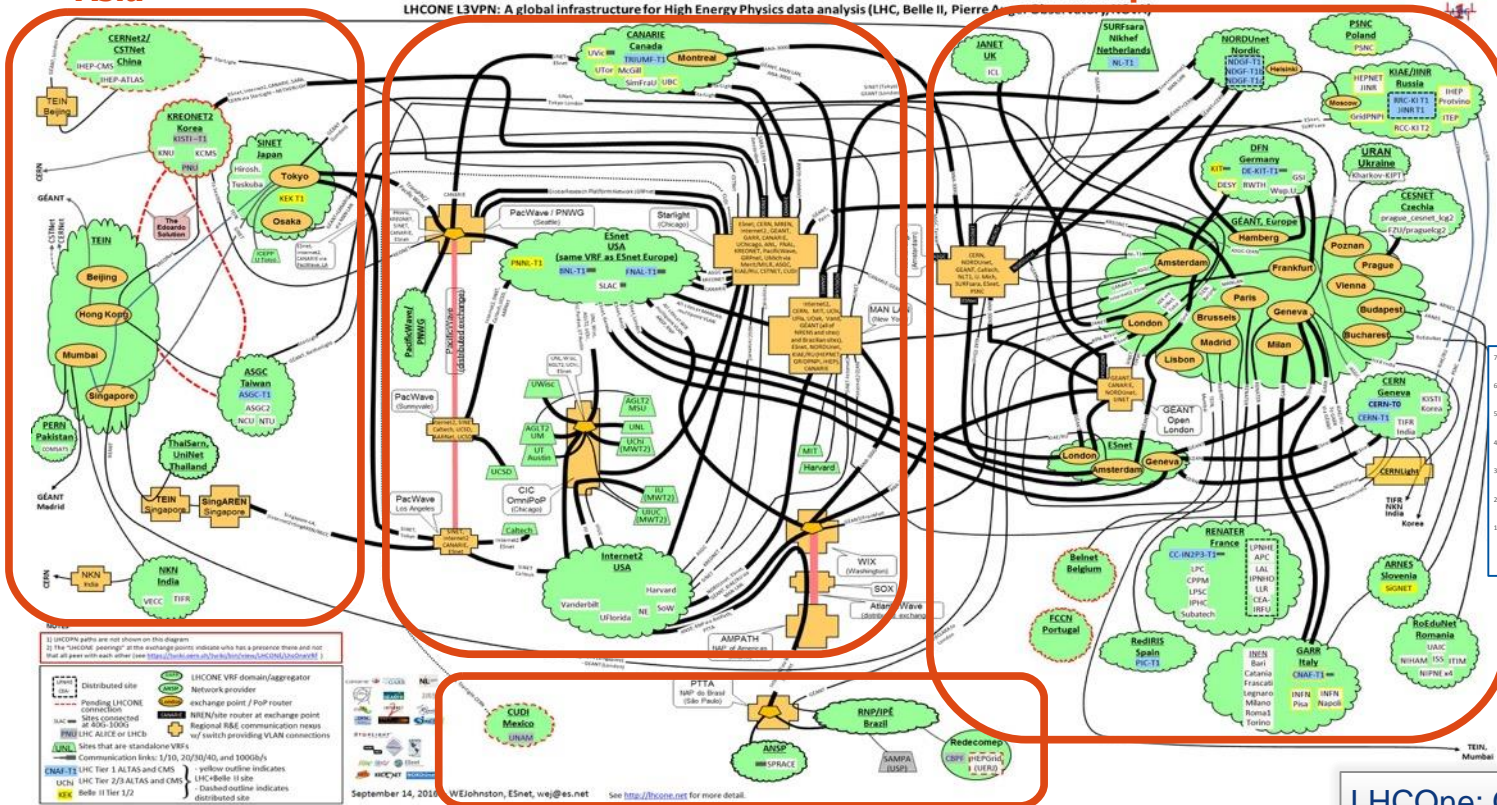
Asia

North America

Europe



LHCONE L3VPN: A global infrastructure for High Energy Physics data analysis (LHC, Belle II, Pierre Auger)



1) LHCONE paths are not shown on this diagram
 2) The "LHCONE peering" at the exchange points indicates who has a presence there and not that all peers with each other (see <https://wiki.lhc-cone.net/peering>)

- Distributed site
- LHCONE VRF domain/aggregator
- Network provider
- Exchange point / POP router
- Pending LHCONE connection
- Sites connected at 100G/100G
- NREN/site router at exchange point
- Regional R&E communication peers
- LHC L3VPN or LHCED
- uf switch providing VLAN connections
- Sites that are standalone VRFs
- Communication links: 1/10, 20/30/40, and 100Gb/s
- Yellow outline indicates LHC Tier 1 ATLAS and CMS
- Yellow outline indicates LHC Belle II site
- Yellow outline indicates Belle II Tier 1/2
- Yellow outline indicates distributed site


South America

LHCOne: Overlay network
 Allows NREN's to manage HEP traffic on general purpose network
 Managed by NREN collaboration




Data at CERN

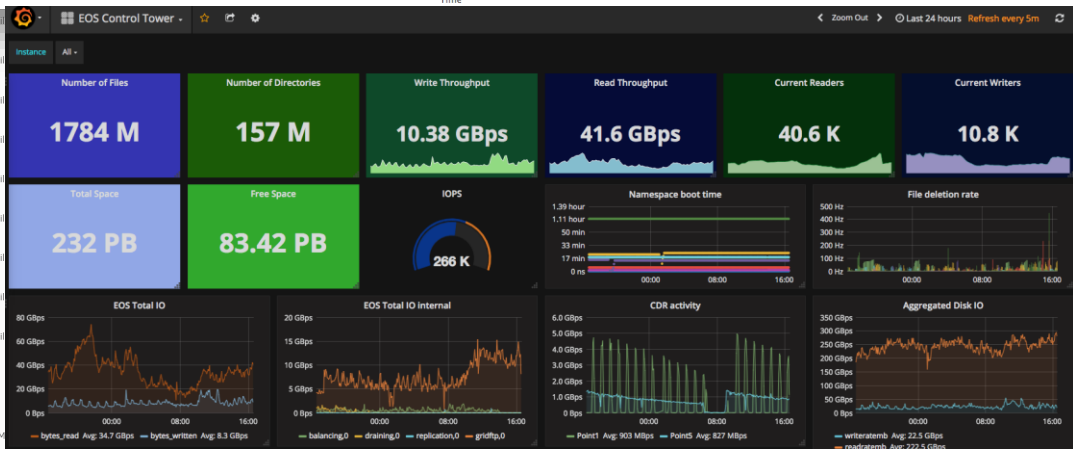

Transferred Data Amount per Virtual Organization for WRITE Requests



Transferred Data Amount per Virtual Organization for WRITE Requests




Physics Data in CASTOR



Data distribution

- Global transfer rates increased to 30-40 GB/s (>2 x Run1)



Increased performance everywhere:

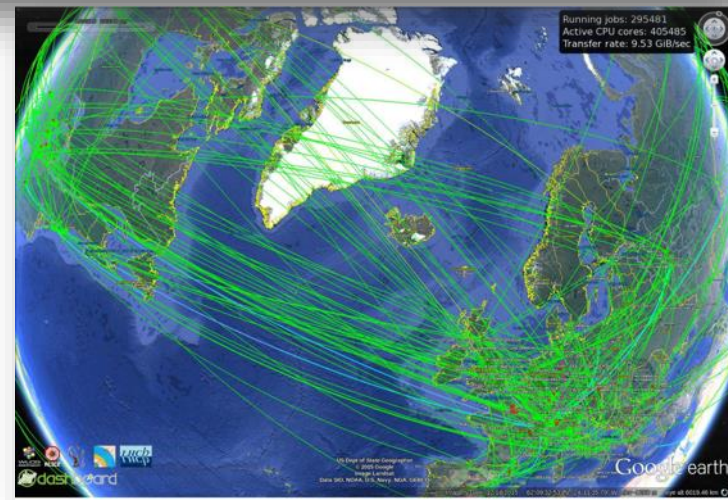
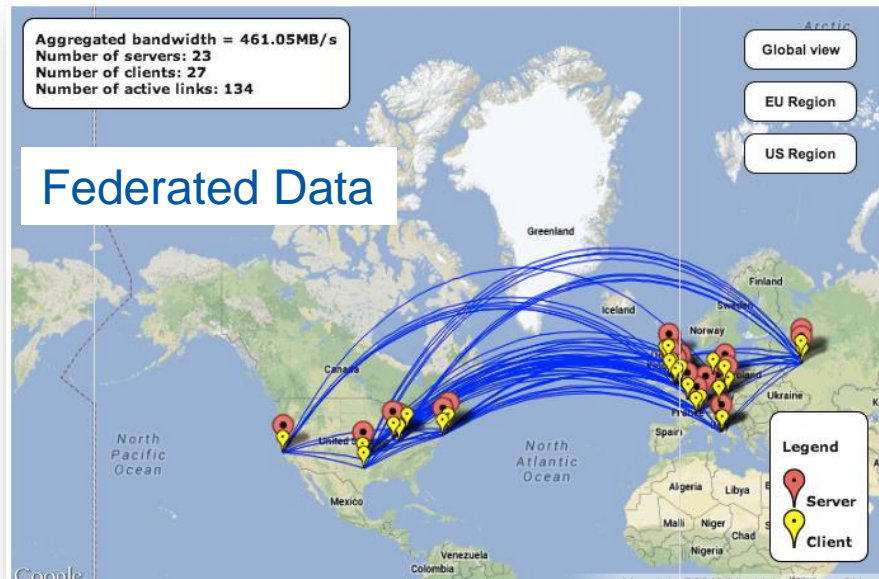
- Data acquisition >10PB / month
- Data transfer rates > 35 GB/s globally

Regular transfers of >80 PB/month with ~100 PB/month during July-October
(many billions of files)


Y

Aggregated bandwidth = 461.05MB/s
Number of servers: 23
Number of clients: 27
Number of active links: 134


Federated Data



Worldwide computing




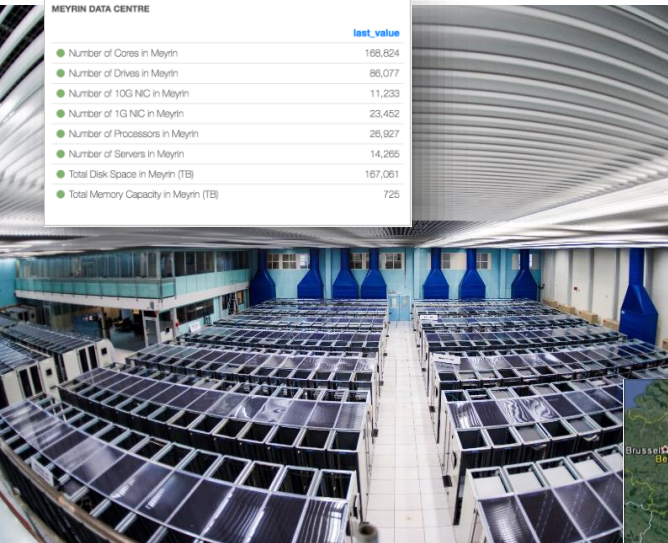
CPU Delivered: HS06-hours/month



Peak delivery:
190M core-days/month
(~ 600k cores permanently)

CERN Facilities today

	last_value
Number of Cores in Meyrin	168,824
Number of Drives in Meyrin	86,077
Number of 10G NIC in Meyrin	11,233
Number of 1G NIC in Meyrin	23,452
Number of Processors in Meyrin	26,927
Number of Servers in Meyrin	14,265
Total Disk Space in Meyrin (TB)	167,061
Total Memory Capacity in Meyrin (TB)	725

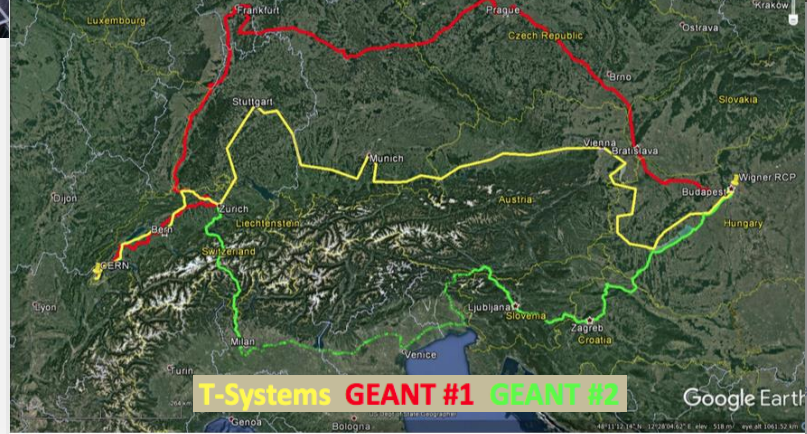


	last_value
Number of Cores in Wigner	56,000
Number of Drives in Wigner	29,694
Number of 10G NIC in Wigner	2,981
Number of 1G NIC in Wigner	6,579
Number of Processors in Wigner	7,002
Number of Servers in Wigner	3,504
Total Disk Space in Wigner (TB)	97,315
Total Memory Capacity in Wigner (TB)	221



2017:

- 325 k cores
- 250 PB raw disk

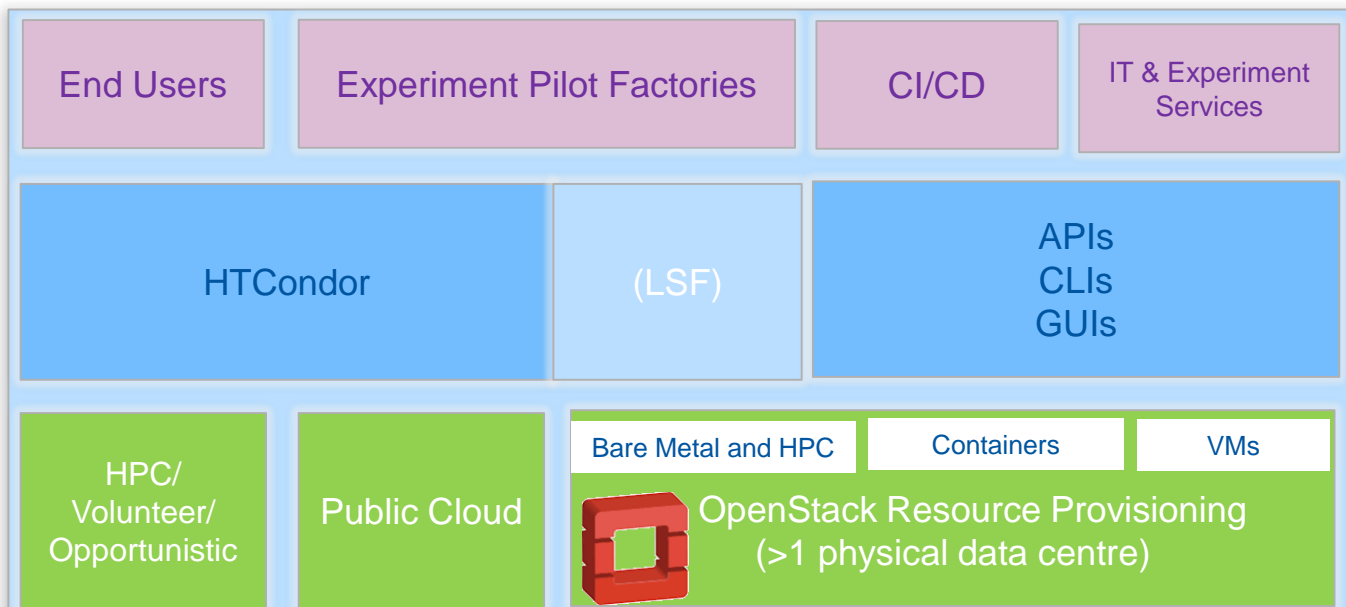


2017-18/19

- Upgrade internal networking capacity
- Refresh tape infrastructure

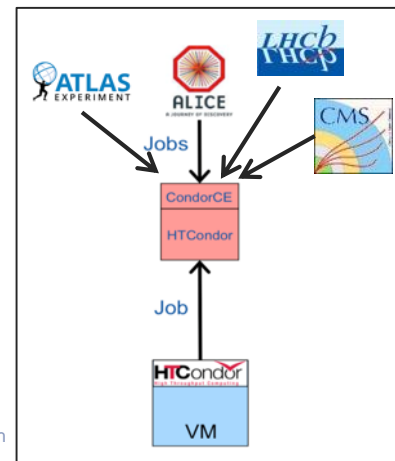


Provisioning services



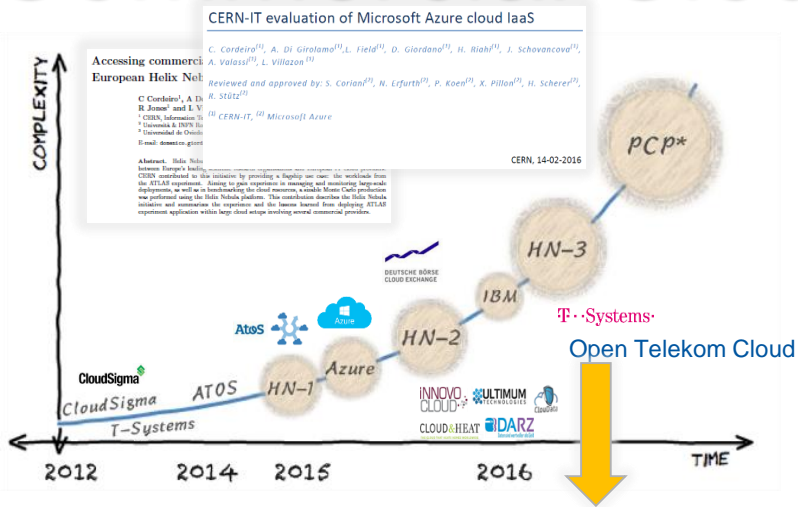


Moving towards Elastic Hybrid IaaS model:

- In house resources at full occupation
- Elastic use of commercial & public clouds
 - Assume “spot-market” style pricing




Commercial Clouds




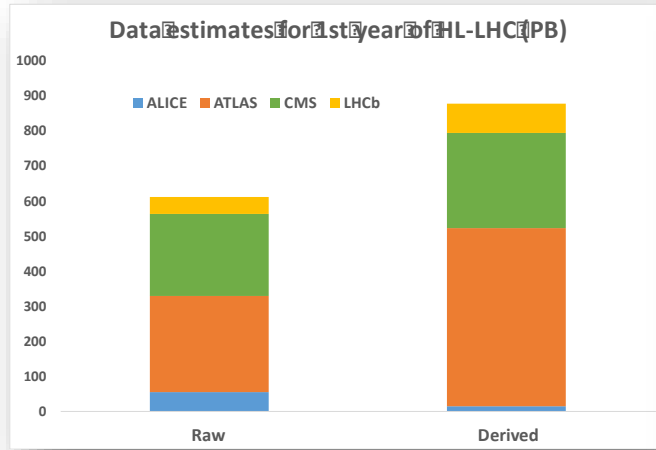


LHC Schedule

• Peak luminosity — Integrated luminosity



Future Challenges

2010 2011 2012 2013 2014 2015 2016 2017 2018 2019 2020 2021 2022 2023 2024 2025 2030?



Data:

- Raw 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB

CPU:

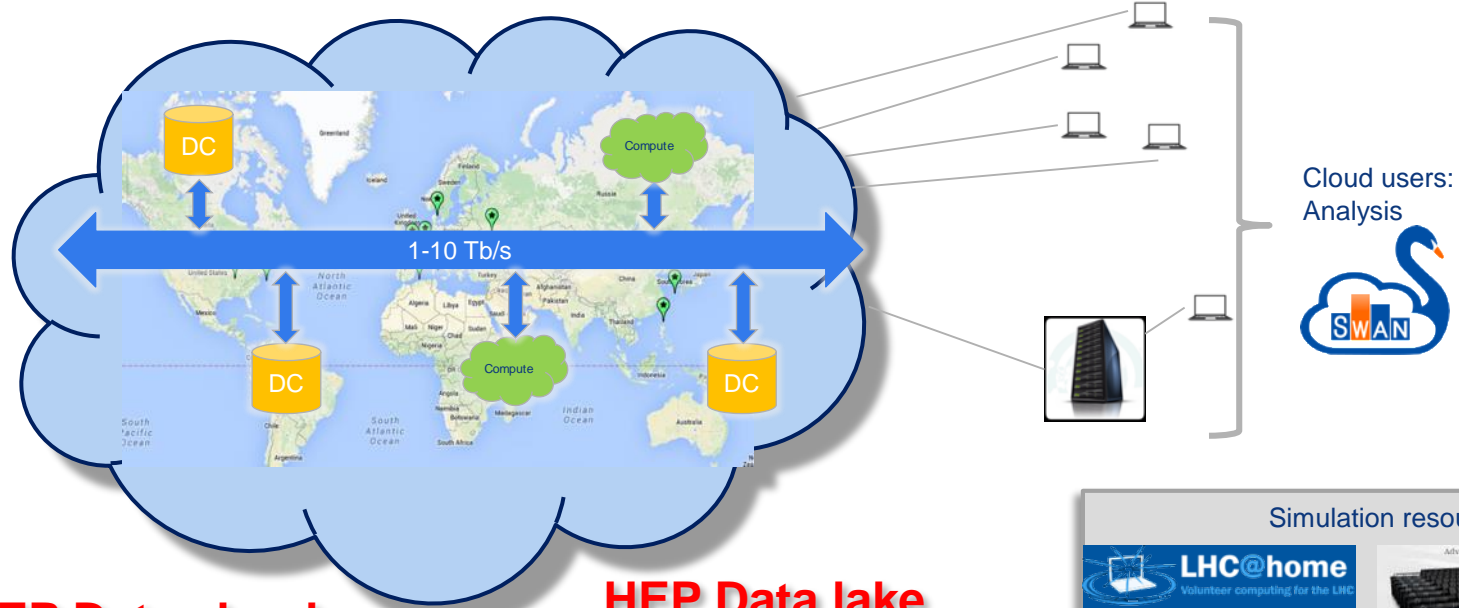
- x60 from 2016

- ❑ Raw data volume for LHC increases exponentially and with it processing and analysis load
- ❑ Technology at ~20%/year will bring x6-10 in 10-11 years
- ❑ Estimates of resource needs at HL-LHC x10 above what is realistic to expect from technology with reasonably constant cost

10-year challenges

- ❑ HL-LHC will be a multi-Exabyte challenge
 - Storage and compute needs x10 above what naïve technology extrapolation will bring
 - Need to drive down costs: focus on performance, efficiency, operations, etc. → changes in computing and infrastructure models are necessary
- ❑ But there is experience:
 - ~15 years of grid development and successful operation for science
 - CERN has been operating a distributed DC for >5 years
 - Large internet companies provide tools and experience that did not exist when we started WLCG
 - Tools for managing interconnected DCs, cloud provisioning, etc.
 - Starting to prototype federated structures for the future

Possible Model for future HEP computing infrastructure



HEP Data cloud
Storage and compute

HEP Data lake
Storage and compute

A data lake is a place to put all the data enterprises (may) want to gather, store, analyze and turn into insights and action, including structured, semi-structured and unstructured data

Simulation resources

LHC@home
Volunteer computing for the LHC

Advancing the Era of Accelerated Computing

A world map showing numerous blue location pins across all continents, representing simulation resources. The map includes labels for major oceans and continents.

Software

HSF Set up in response to recognition that software will be key to success for HL-LHC and the future



The HEP Software Foundation (HSF) facilitates coordination and common efforts in high energy physics (HEP) software and computing internationally.

The HSF is now beginning community process to develop a consensus roadmap for HEP Software and Computing R&D for the 2020s. More information about this can be found on the [Community White Paper \(CWP\)](#) page on the HSF site.

Meetings

All our activities and ideas are discussed weekly in our HSF meeting. Feel free to participate!

- [HSF Weekly Meeting #71, November 3, 2016](#)
- [HSF Weekly Meeting #69, September 15, 2016](#)
- [HSF Weekly Meeting #68, September 8, 2016](#)

[Full list of meetings »](#)

Newsletter

If you would like to stay updated, please subscribe to our newsletter:

- [Third HSF Workshop](#)
- [Sharing ideas and code](#)
- [HSF Newsletter - Logo Contest and Packaging Working Group](#)

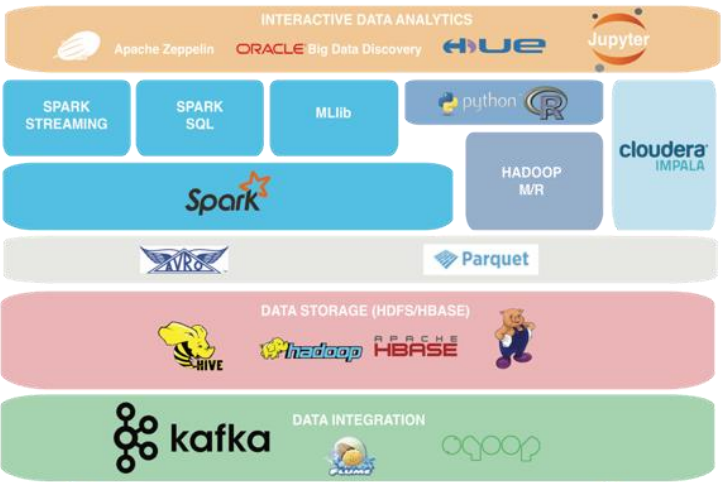
[Older newsletters »](#)

Activities

Our plenty of activities span from our [working groups](#), organizing [events](#) to supporting projects as [HSF projects](#), and channeling communication within the community with [discussion forums](#), [technical notes](#) and a [knowledge base](#).

[How to get involved »](#)

New analytics



- ❑ Big Data tools
- ❑ Machine/Deep learning

Machine Learning

ML in Atlas

- ❑ Machine Learning (or rather Multi Variate Analysis as we used to call it) used almost since first data taking (2010) for reconstruction and analysis
- ❑ In most cases, Boosted Decision Tree with Root-TMVA, but recent explosion of usage and studies (see later)

- Data taking
 - Real time event categorization
 - Data monitoring & certification robot
- Data Reconstruction
 - Calorimeter reconstruction
 - Boosted object jet tagging
- Data Processing
 - Computing Resource Optimization
 - Predicting data popularity
 - Intelligent networking
- Data Analysis
 - CMS assistance service
 - Big data reduction and analysis
 - Model independent search



Machine learning and data analytics are hot topics at CERN openlab workshop

Wednesday, 4 May 2016

Last week, CERN openlab held a workshop on machine learning and data analytics. The event, which took place on Friday 29 April, saw experts from both research and industry gather in the CERN IT Department for a full-day of presentations and lively discussion.

The morning featured presentations from representatives of the four large LHC experiments – ALICE, ATLAS, CMS, and LHCb – on their current projects and future challenges in these areas.

During the afternoon, representatives from industry were also invited to give their perspective. This included presentations from CERN's external partner companies Intel and Samsung, as well as contributors Cisco and associates NetScout. The other companies to present at the event were Cloudera, Google, IBM, Microsoft, and Nvidia.

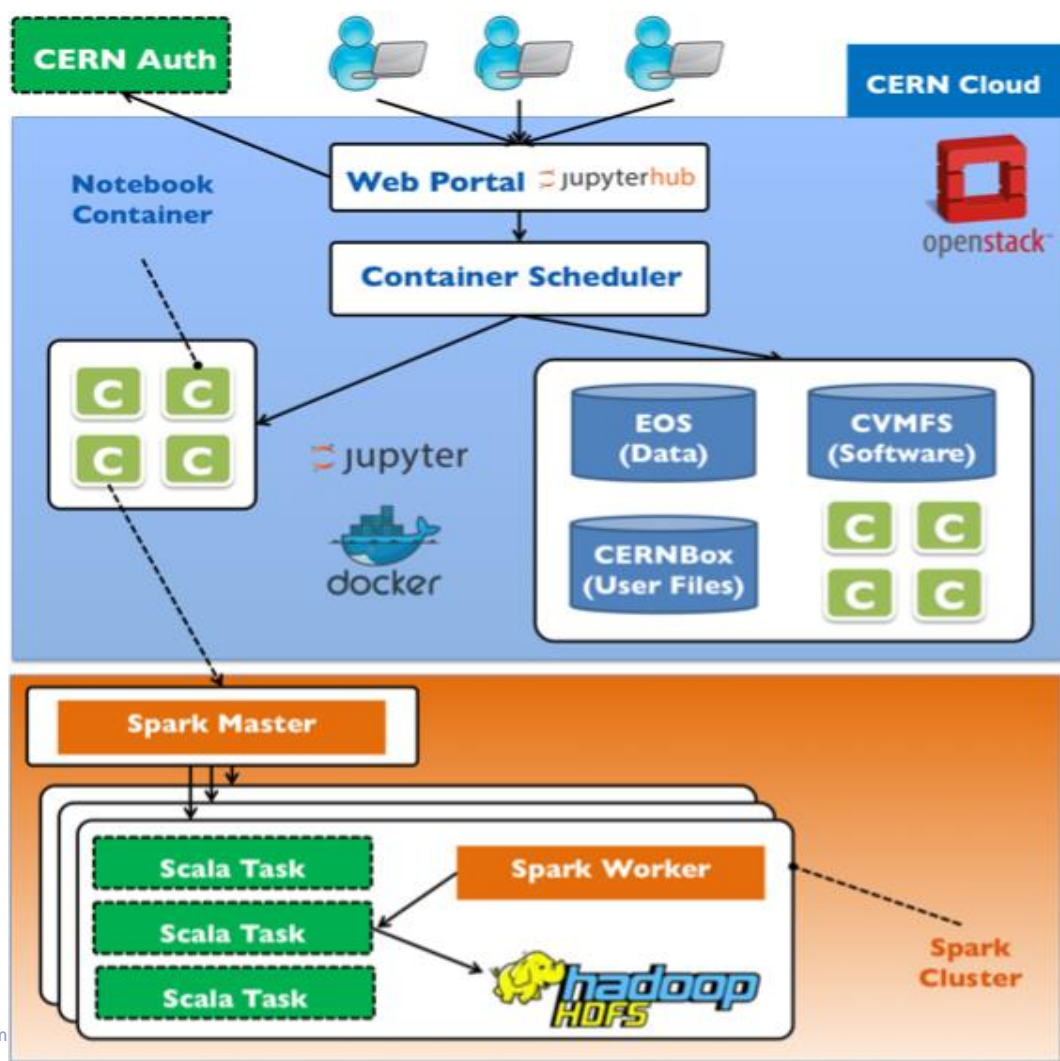
"The event provided a great opportunity for experts from both industry and the LHC experiments to discuss their activities – as well as the challenges they face – in the exciting area of machine learning and data analytics," says Maria Girone, CERN spokeswoman. She particularly advocated the engagement of all participants in lively and constructive discussions. Over the course of the day many commonalities and areas of potential collaboration emerged."



SWAN



- ❑ Provides a web-based analysis facility – via notebooks
- ❑ Transparent access to scalable back-end analysis infrastructure
 - Clouds, Spark, Hadoop, ML, etc.
- ❑ Performance is defined by the infrastructure
- ❑ Provides the analysis portal in a “data cloud” or “data lake” model



Conclusions

- ❑ WLCG has been very successful in providing the global computing environment for physics at the LHC
- ❑ Engagement and contributions of the worldwide community have been essential for that
- ❑ LHC upgrades over the coming decade will give new challenges and opportunities
 - Technology will change our computing models
- ❑ We see a lot of potential synergy with other sciences: leadership in scientific big-data in the next decade