HL-LHC, WLCG and EOSC

2nd ASTERICS-OBELICS Workshop

16-19 October 2017, Barcelona, Spain.

H2020-Astronomy ESFRI and Research Infrastructure Cluster
(Grant Agreement number: 653477).

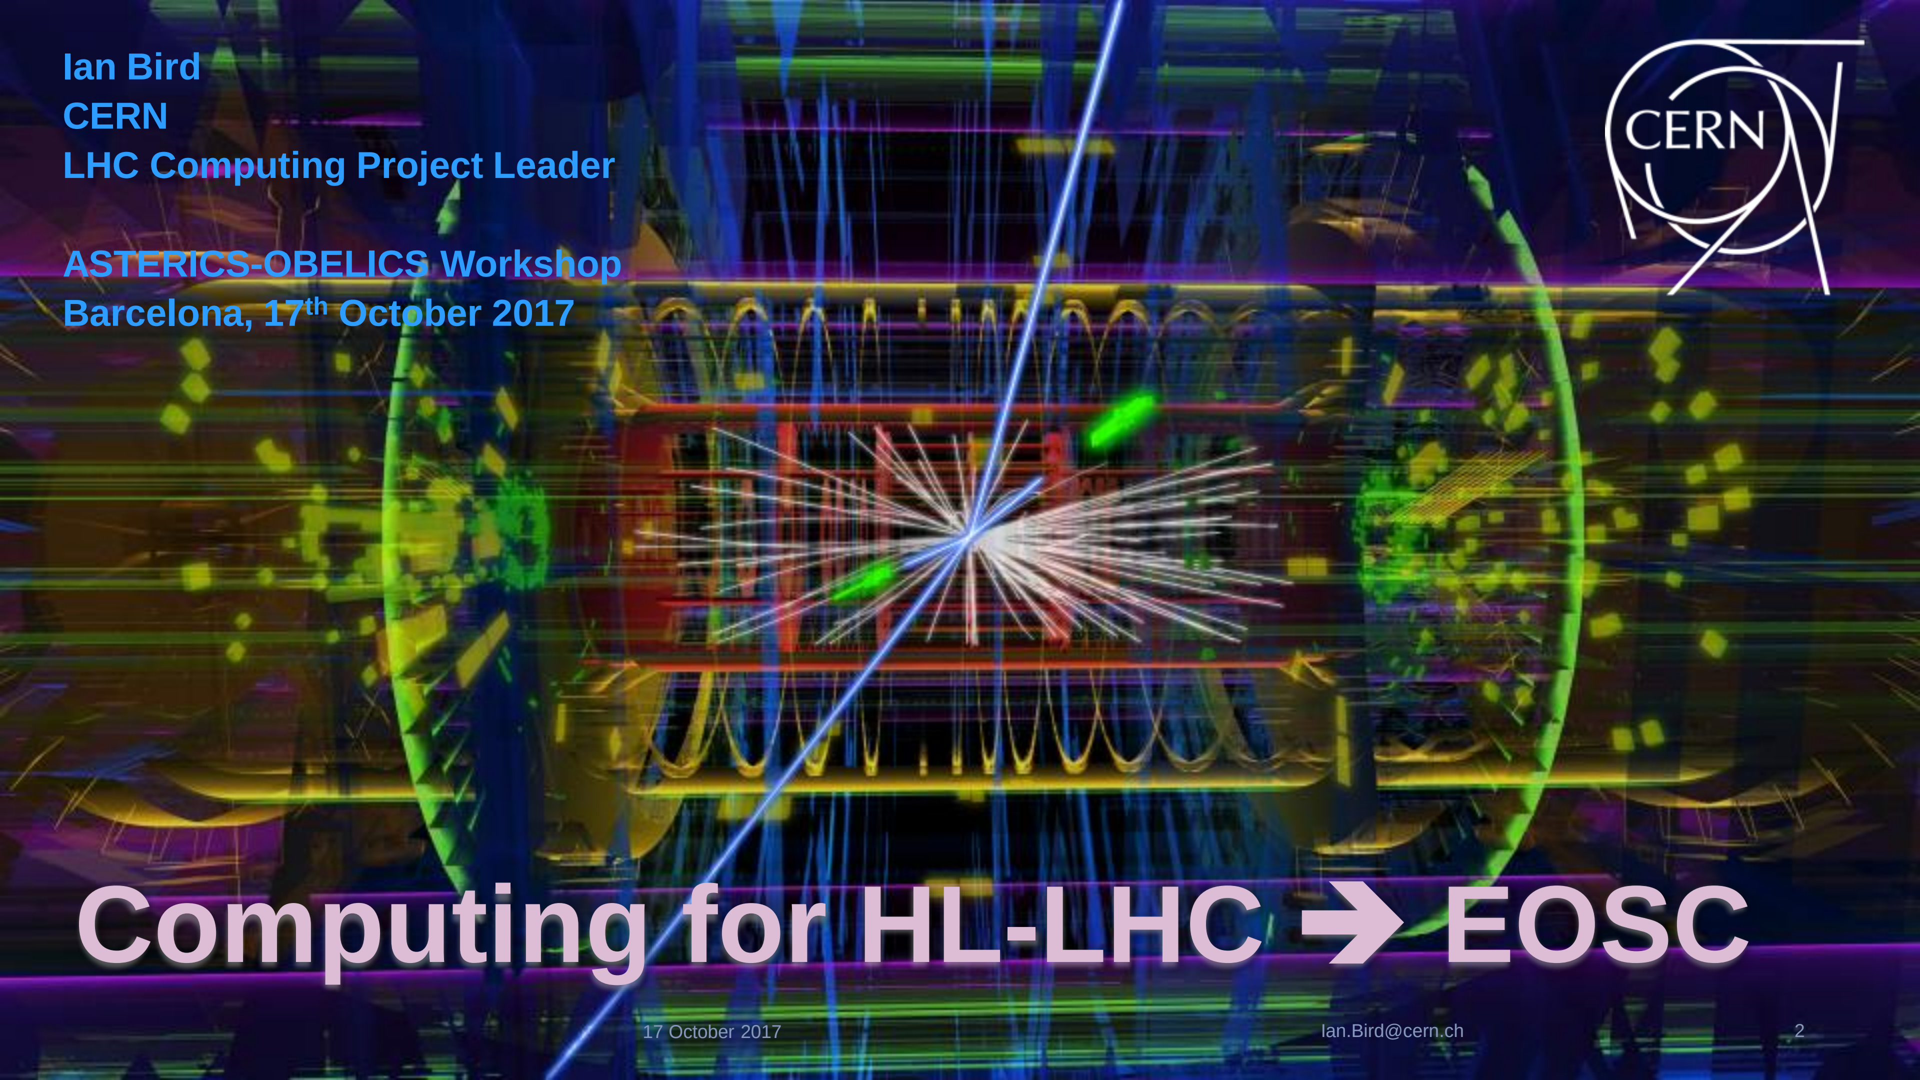

Ian Bird
CERN
LHC Computing Project Leader

ASTERICS-OBELICS Workshop
Barcelona, 17th October 2017

Computing for HL-LHC ➔ EOSC

The Large Hadron Collider (LHC)

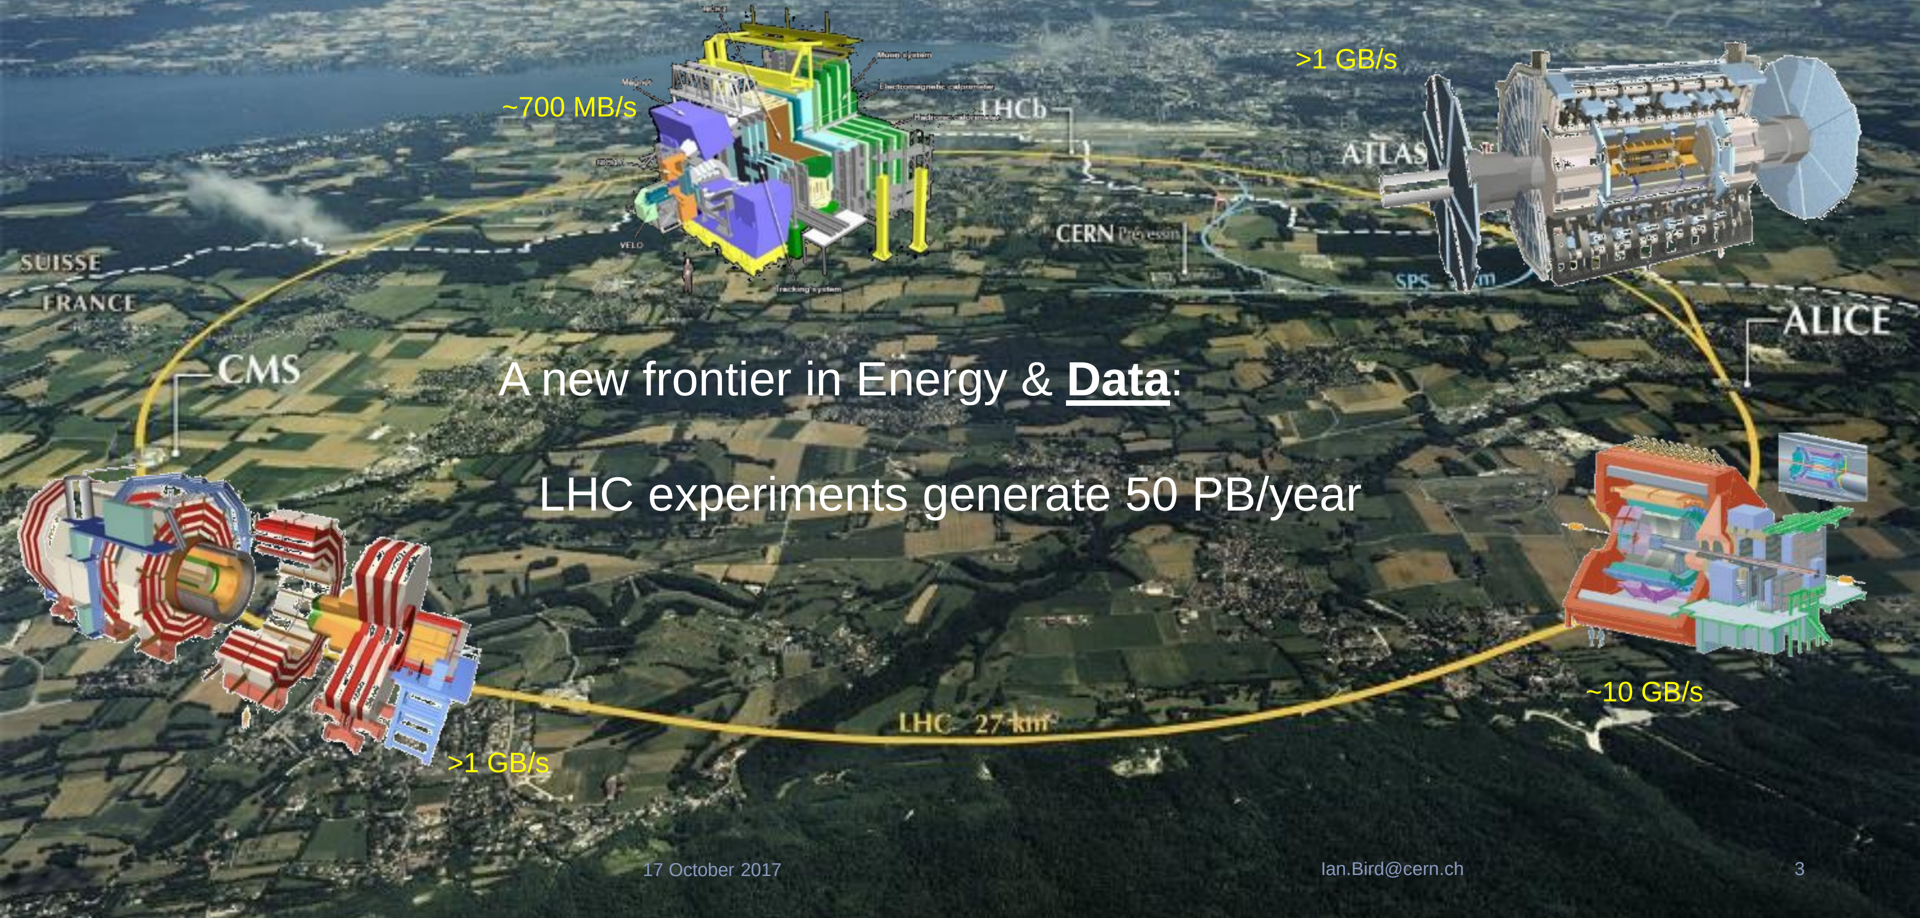

A new frontier in Energy & **Data**:

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)

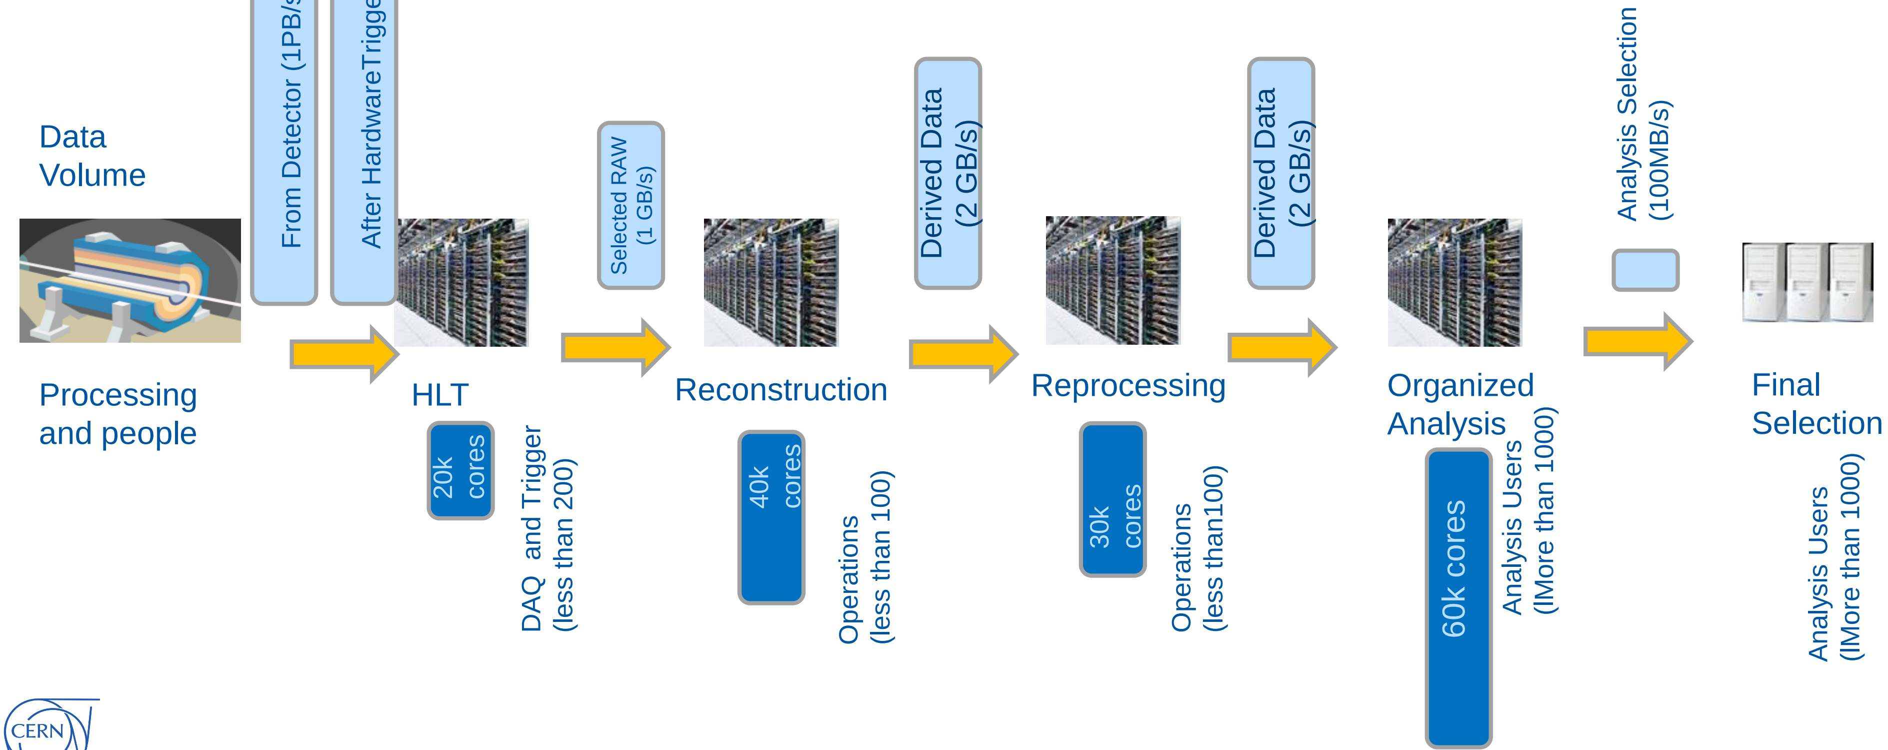

The Worldwide LHC Computing Grid

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

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

Tier-2: Simulation, end-user analysis

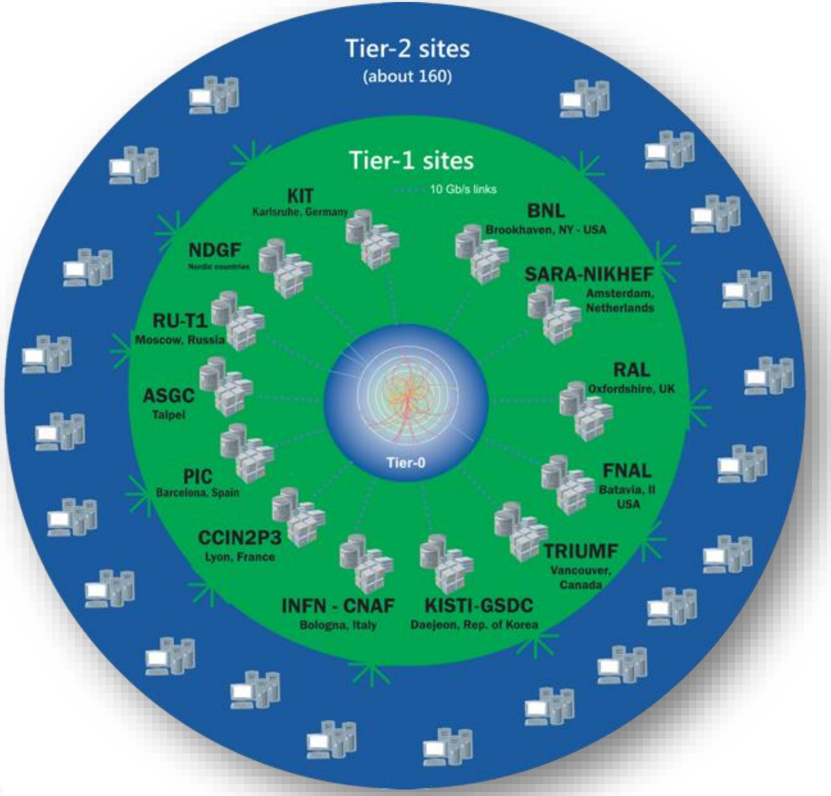

~170 sites, 42 countries

~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

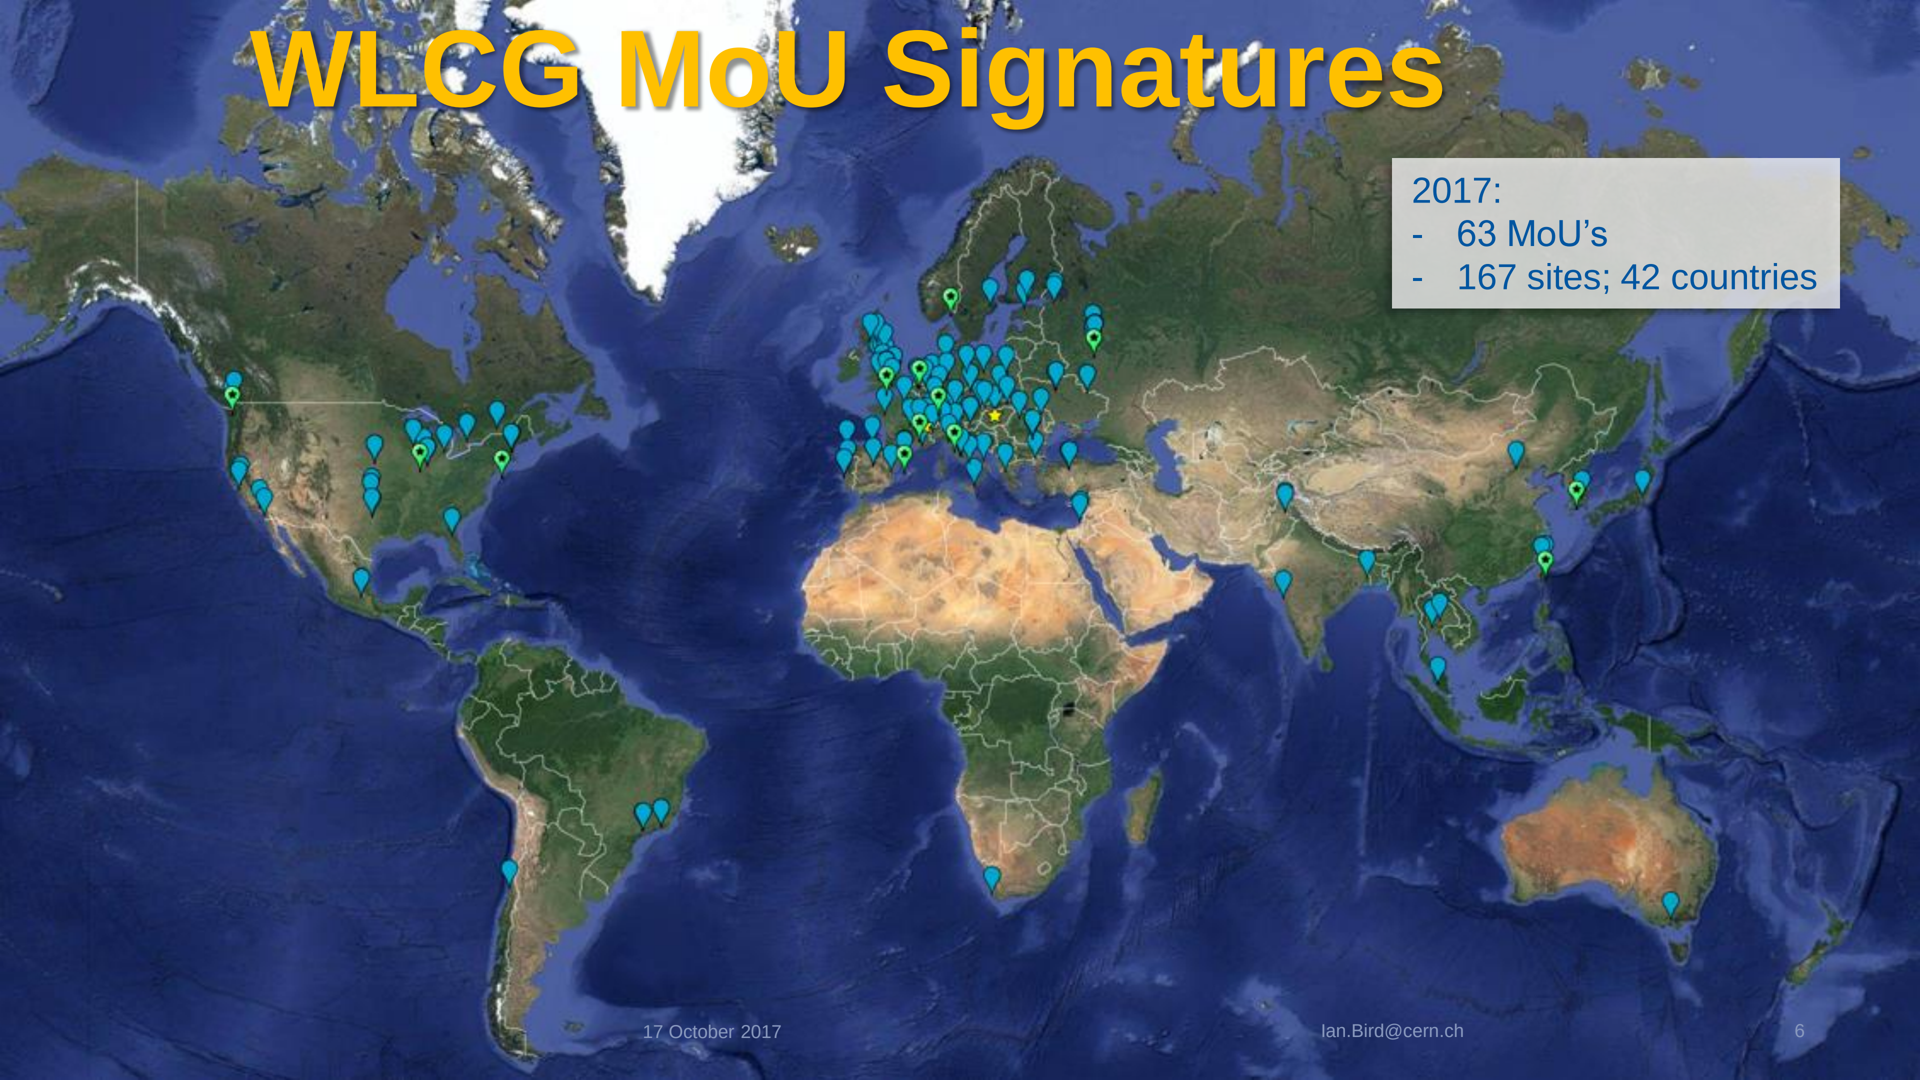

2017:
- 63 MoU's
- 167 sites; 42 countries

Networks

LHCOPN

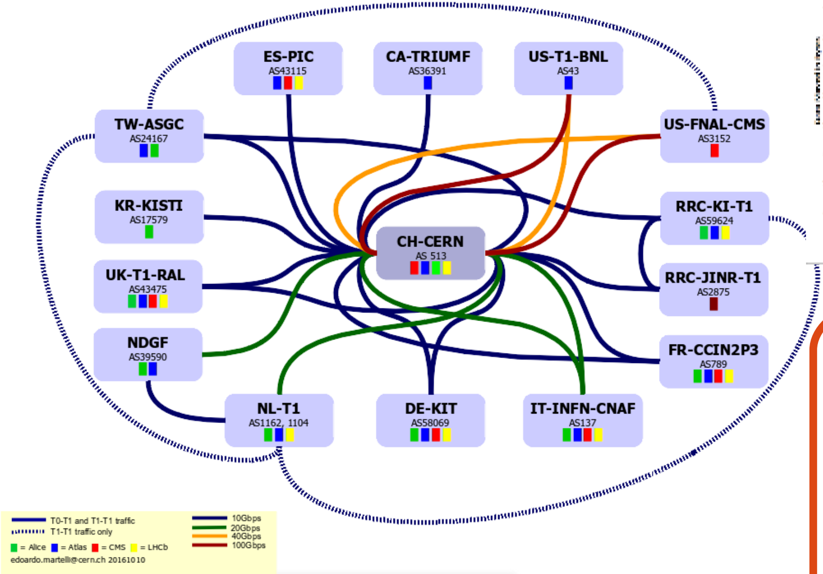

Optical Private Network
Support T0 – T1 transfers
& T1 – T1 traffic
Managed by LHC Tier 0 and
Tier 1 sites

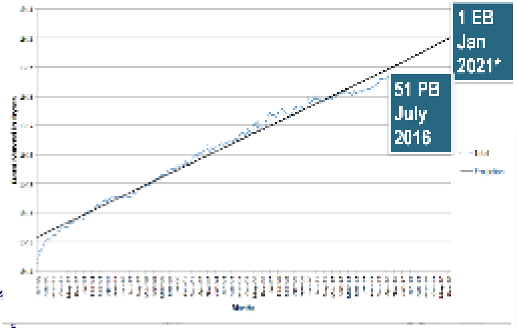

LHCOne: Overlay network
Allows NREN’s to manage HEP traffic on general purpose network
Managed by NREN collaboration

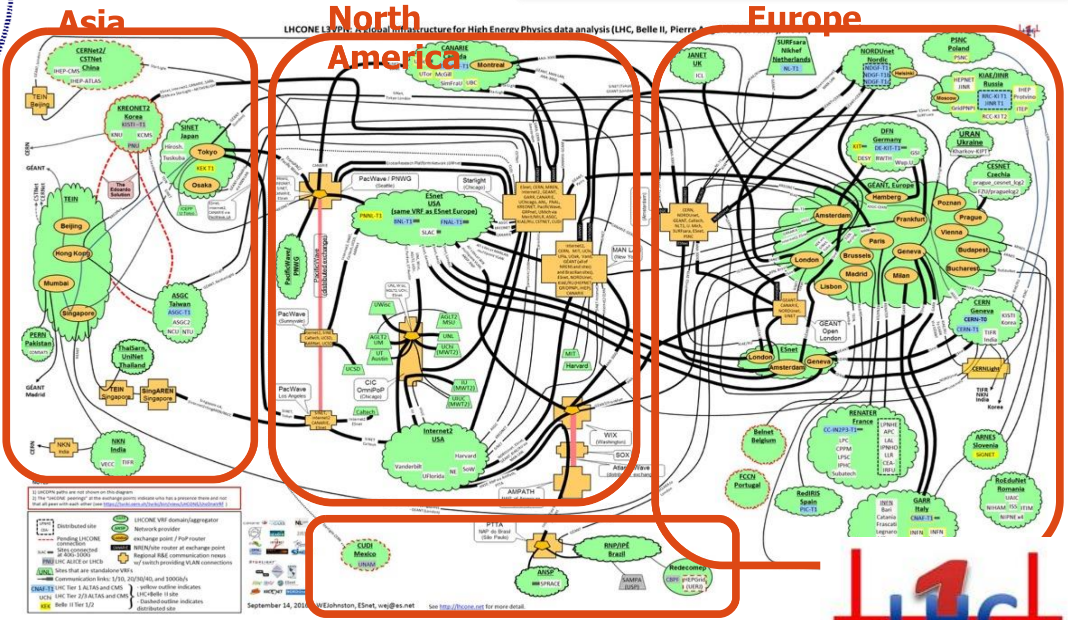

South America

Up to 340 Gbps transatlantic

Ian.Bird@cern.ch

LHC Data

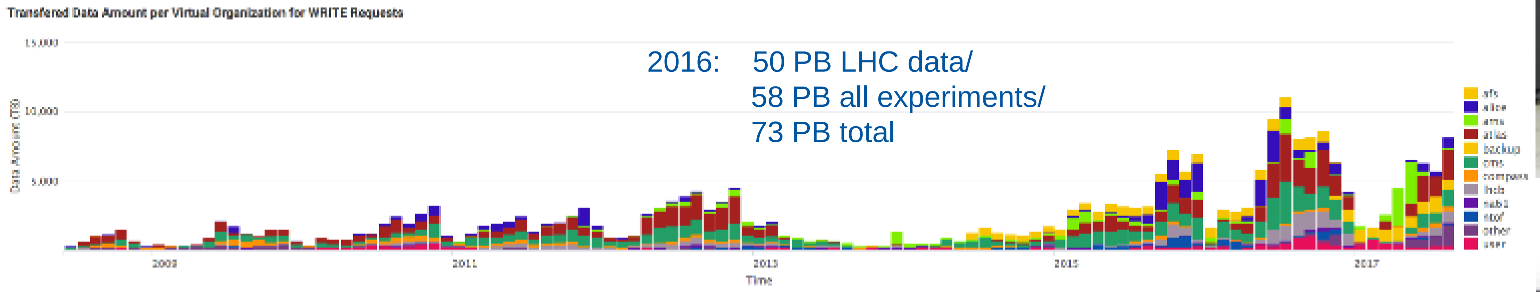

2016: 50 PB LHC data/
58 PB all experiments/
73 PB total

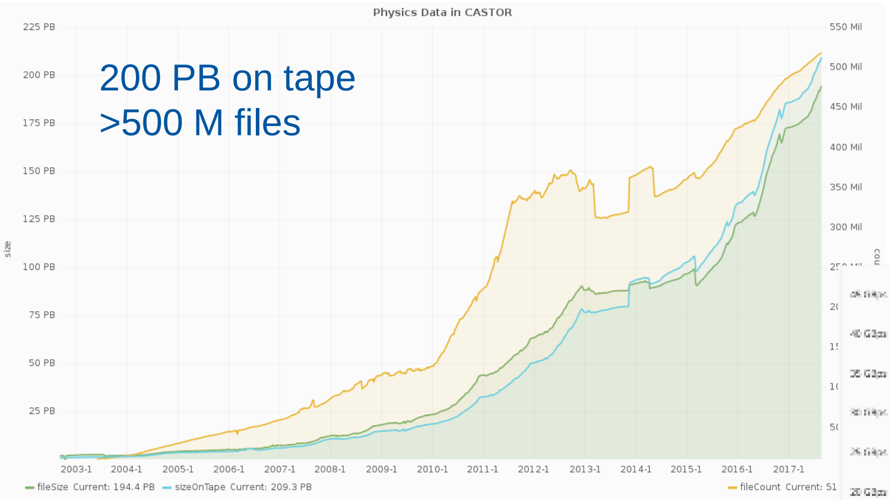

200 PB on tape
>500 M files

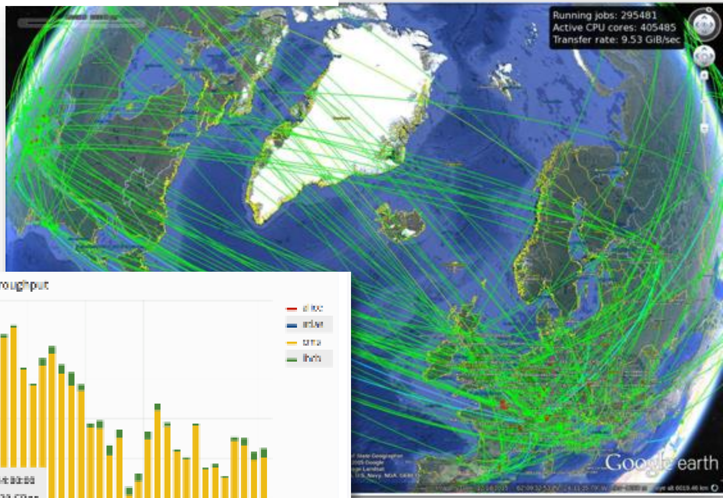

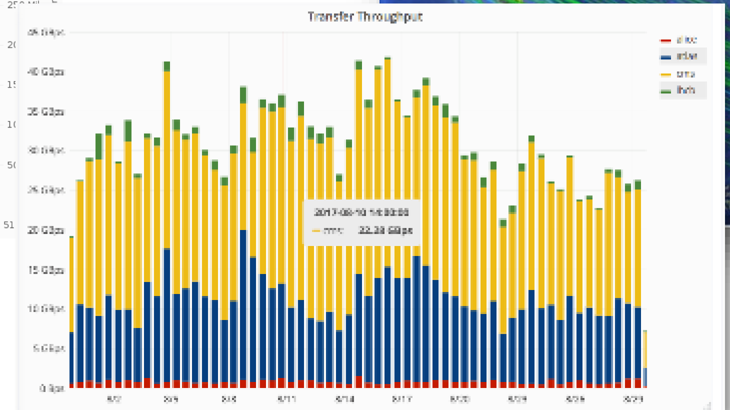

Worldwide computing

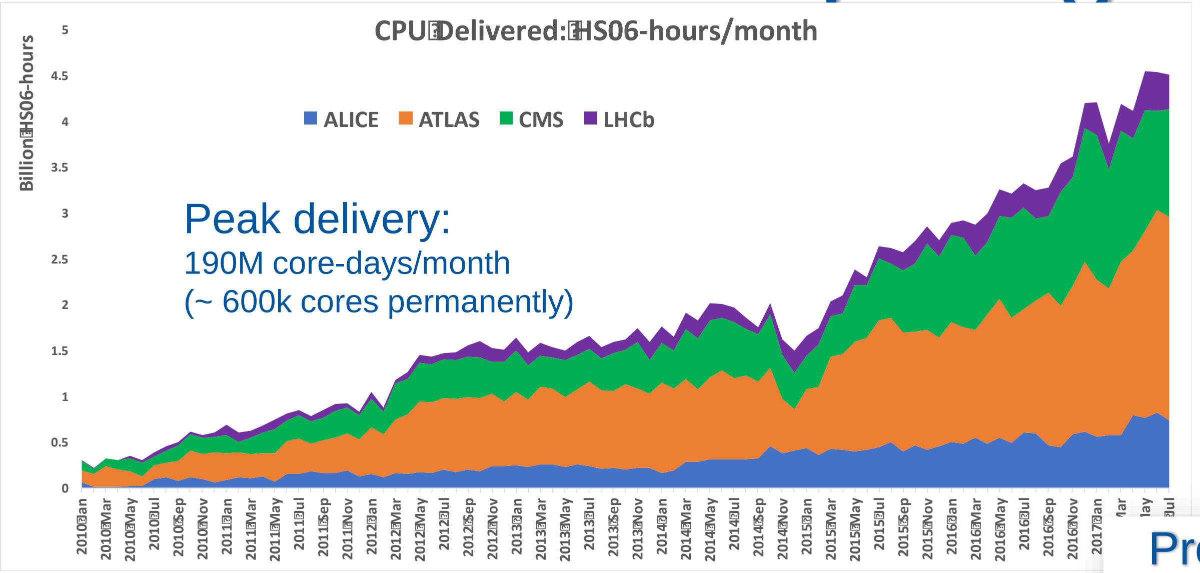

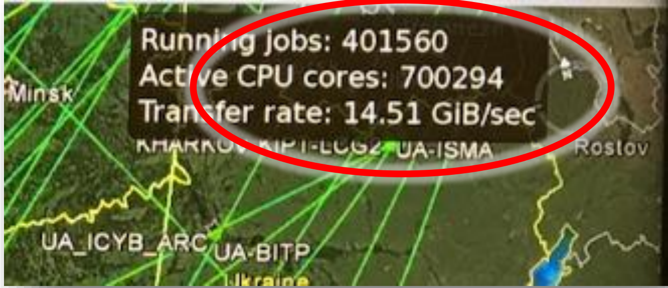

17 October 2017

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

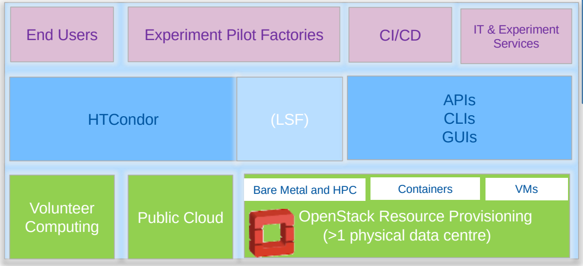

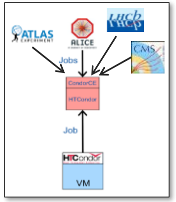

CERN Facilities today

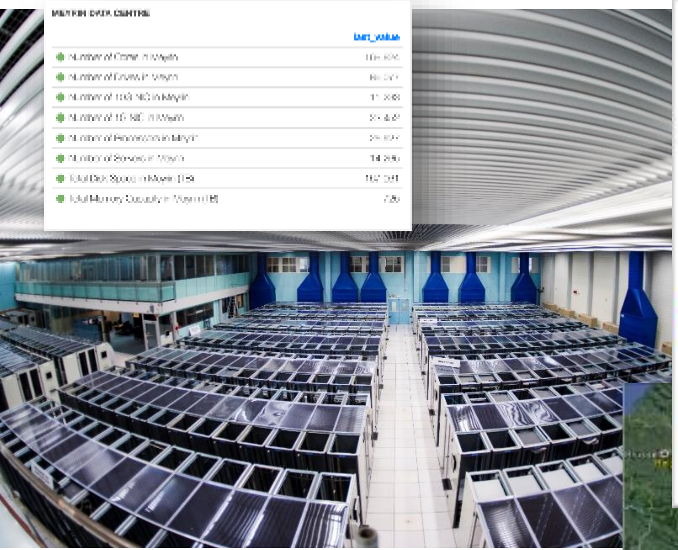

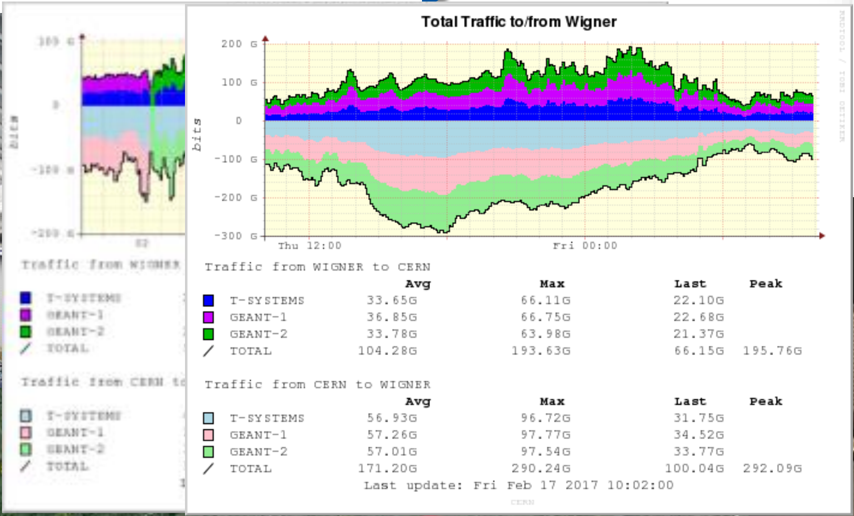

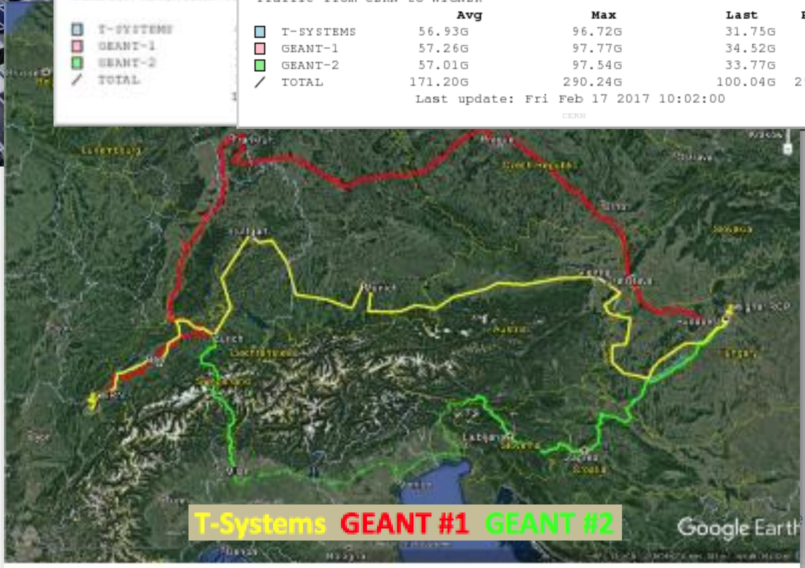

2017:
- 325 k cores
- 250 PB raw disk

2017-18/19
- Upgrade internal networking capacity
- Refresh tape infrastructure

Commercial Clouds

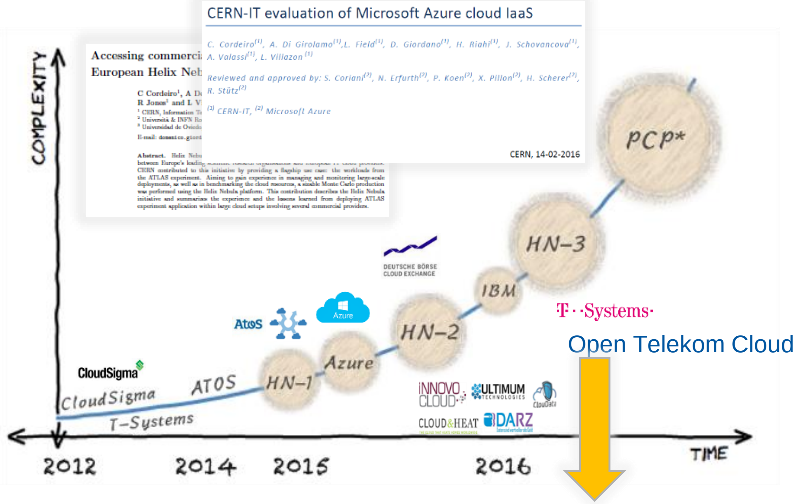

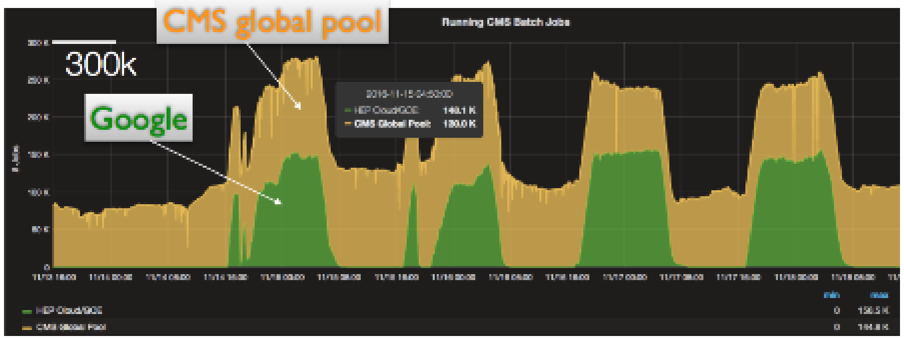

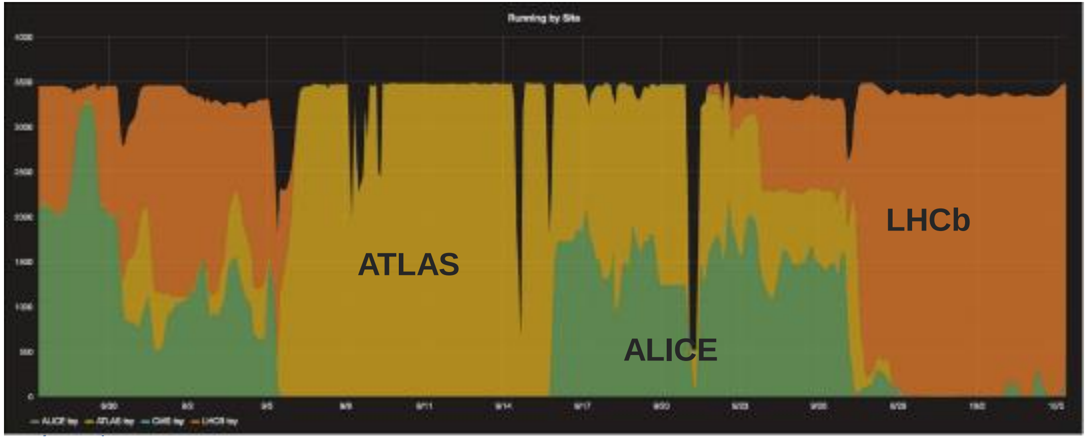

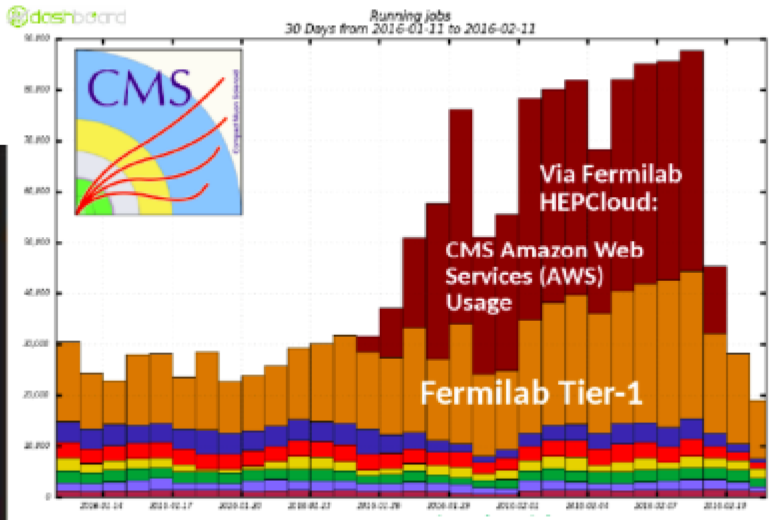

LHC Schedule

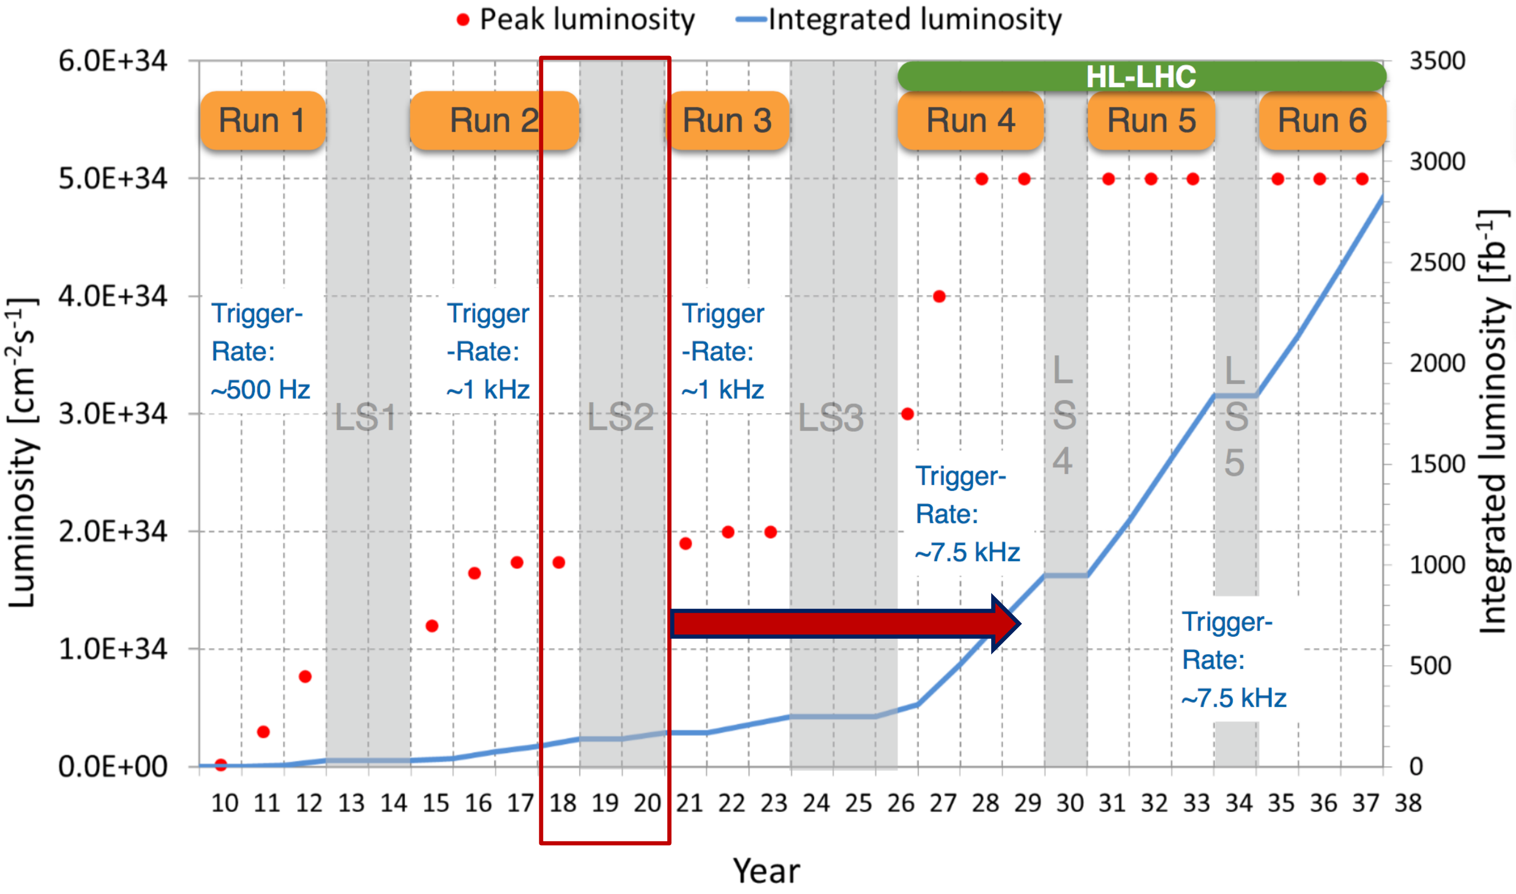

Run 3 Alice, LHCb upgrades

Run 4 ATLAS, CMS upgrades

Future Challenges

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

First run | LS1 | Second run | LS2 | Third run | LS3 | HL-LHC | FCC?

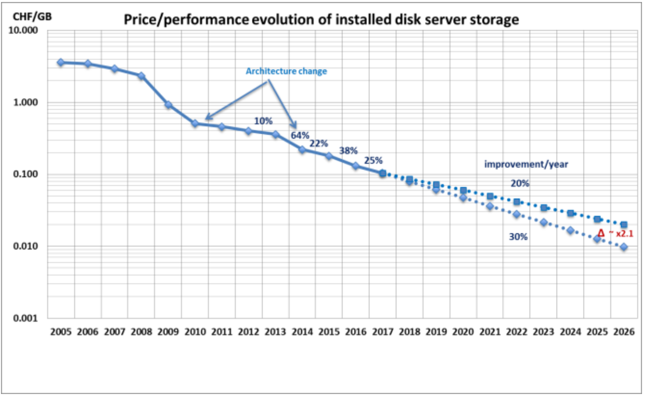

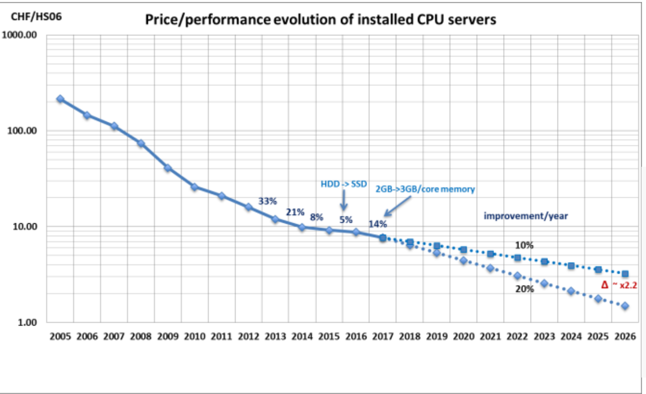

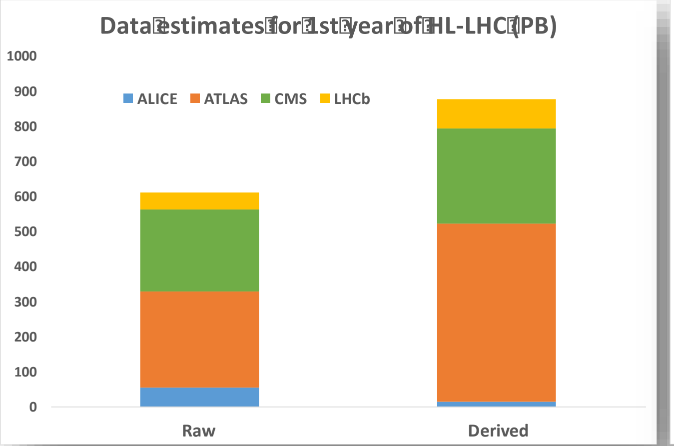

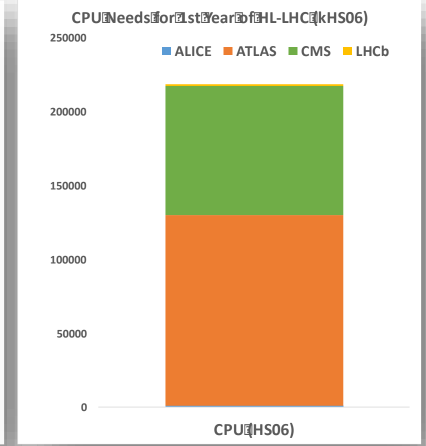

Data:
- Raw 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB

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

Software

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

HSF
HEP Software Foundation

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 infomation 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 »

Community Whitepaper

- A Community White Paper (CWP) has been written outlining a global vision for software and computing for the HL-LHC era and HEP in the 2020s
- The CWP identifies and prioritizes the software research and development investments required
  - to achieve improvements in software efficiency, scalability and performance and to make use of the advances in CPU, storage and network technologies
  - to enable new approaches to computing and software that could radically extend the physics reach of the detectors
  - to ensure the long term sustainability of the software through the lifetime of the HL-LHC
- The HSF is engaging the HEP community to produce the CWP via an inclusive process
  - Initiated as an HL-LHC planning process
  - Aiming for a broader participation (LHC, neutrino program, Belle II, linear collider, ...)
  - The resulting roadmap will be used for the HL-LHC computing TDR and other strategic plans

Main CWP themes

- Allow/help countries or regions to flexibly manage compute and storage resources internally,
  - Supporting national/regional consolidation, provisioning resources in a way that makes sense in the local situation
  - Use of federation of resources, integration of public, private, commercial, HPC, etc. as necessary
  - Foresee some Tier1/Tier2 boundaries blurred and regions with common funding can federate their facilities, in order to optimize and consolidate the resources they provide, in a way that is flexible, and not held to a history that is decades old at this point.
- Investigate the “data-lakes” concept – keep bulk data (down to derived AODs) in a cloud-like realm (data-lake). Plug in processing via traffic-managed networks, bulk processing close to the data, and:
  - Reduce the amount of data replication and distribution. Evolve from data placement to data serving.
  - Data-lakes as cloud analysis facilities to enable new analysis models (big-data tools, ML, web-based remote analysis with scalable resource back-end, etc.)
  - Review data distribution and delivery technologies, including event streaming, event serving, “FTS”, protocols, etc.

Possible Model for future HEP computing infrastructure

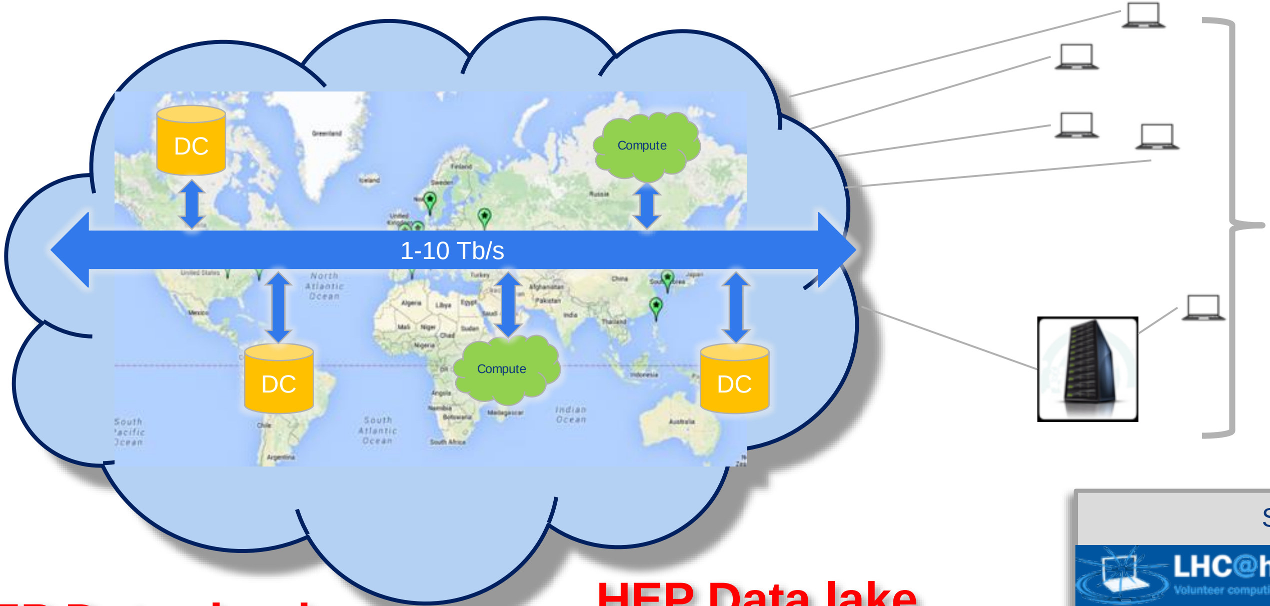

Cloud users:
Analysis

HEP Data cloud
Storage and compute

HEP Data lake
Storage and compute

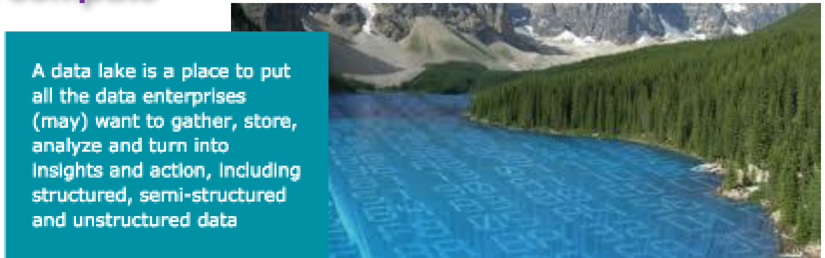

CWP themes – 2

- Enable the adaptation to use very heterogeneous resources:
  - HPC, specialized clusters, opportunistic, clouds (commercial or not) -> managing cost and quotas. Elasticity vs fixed capacity.
- Software/libraries adaptation and validation for wide variety of processor types:
  - Many/multi core; multi-threading, vector units, GPUs, all common CPU types
  - Need capability to rapidly port to and validate on new architectures, even processor generations (new instruction sets)
- Need ongoing and continual review and assessment of performance, bottlenecks, to understand where to direct next investment

CWP themes – 3

- Assess the utility of implementing commonality at various layers of the architecture
  - Commonality across well-understood functionalities,
  - Interest from experiments in working together on common data management, resource provisioning (partly integrated in facilities), workload and workflow management, frameworks, etc.
  - Room for innovation and change, within common frameworks

Evolution vs change

- The running system of today has to evolve into the system for HL-LHC
  - This does not mean there cannot be a major change in components; new facilities alongside old ones, new services phased in and old ones phased out
  - But there will never be a stop of the current system and building of a new one
  - The paradigm is one of managed change and evolution

Cost

- There is no "cost model"
  - Conditions are very different across countries, sites, funding agencies, etc., and change with time.
- What is needed is a way to continually optimize the system:
  - To understand current performance bottlenecks;
  - A process of continual measurement, review, optimization, change; requires:
    - A useful set of understandable metrics (which may evolve)
    - A good understanding of the performance
  - To guide where next investment of effort, resources will be of greatest benefit;
    - how to balance between CPU, storage, network, etc., and FTEs!

R&D

- A program of R&D and prototyping is being drafted
  - Covering main topics called out in the paper
  - Providing testbeds (building on things like Techlab, openlab)
  - Program of in-depth performance understanding, metrics, to optimize the system across CPU, Storage, networks
    - End-end performance and *overall* optimisation of the system
      - Recreate vs keeping datasets; ties together DM, Workflows, data analysis
  - Etc.
- Such an R&D program could inform the needs of the EOSC for large scale data infrastructures
  - A focused set of needs, but a broad engagement of several science communities can help bring together various potential players in the EOSC
    - EOSC, EGI, GEANT, PRACE, EUDAT, etc.

New analytics

- Big Data tools
- Machine/Deep learning

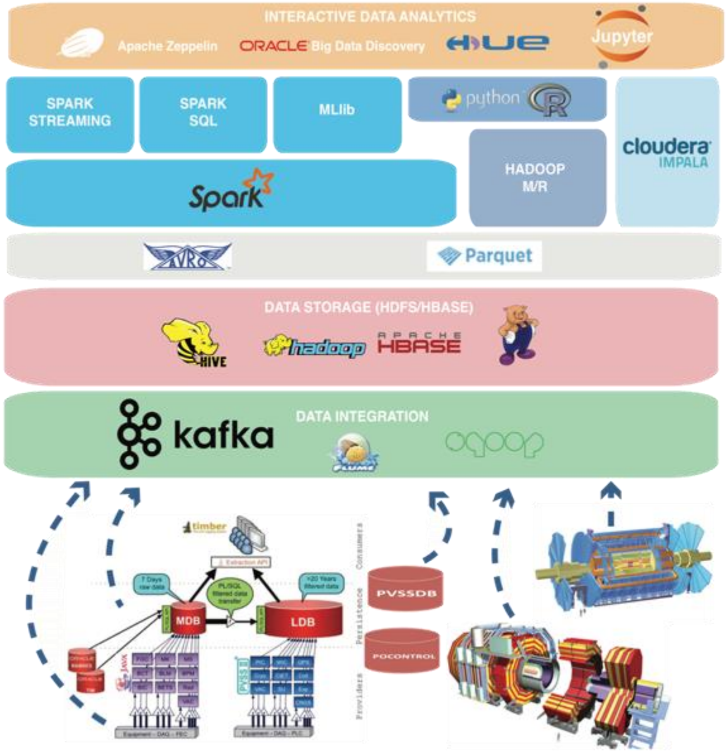

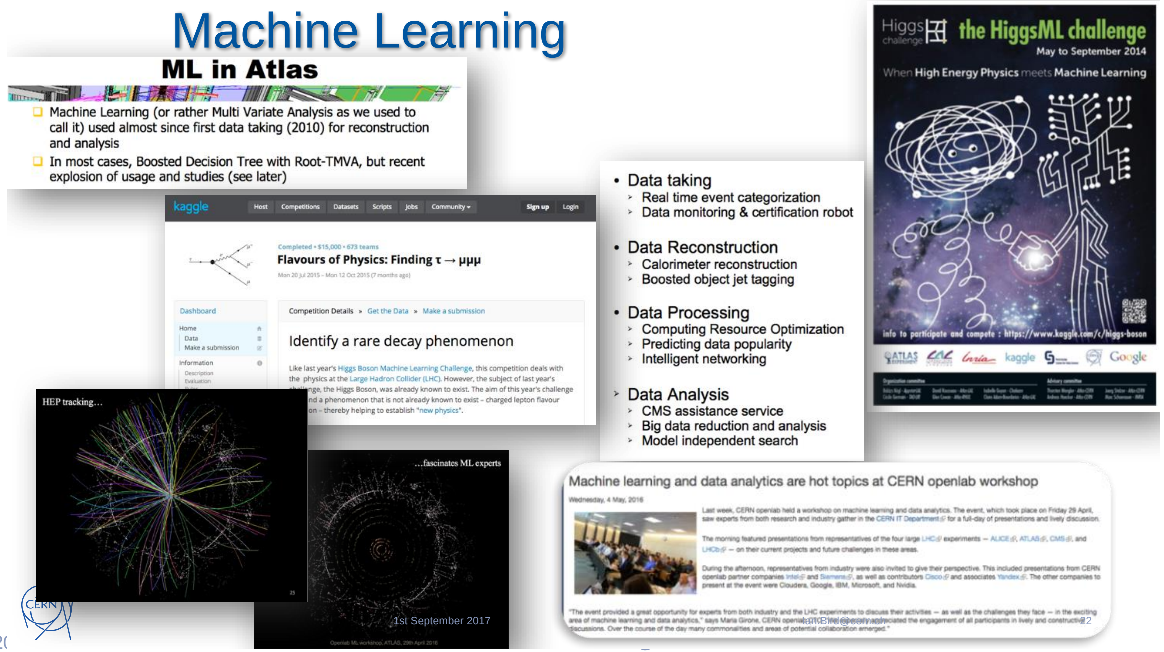

SWAN

- 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

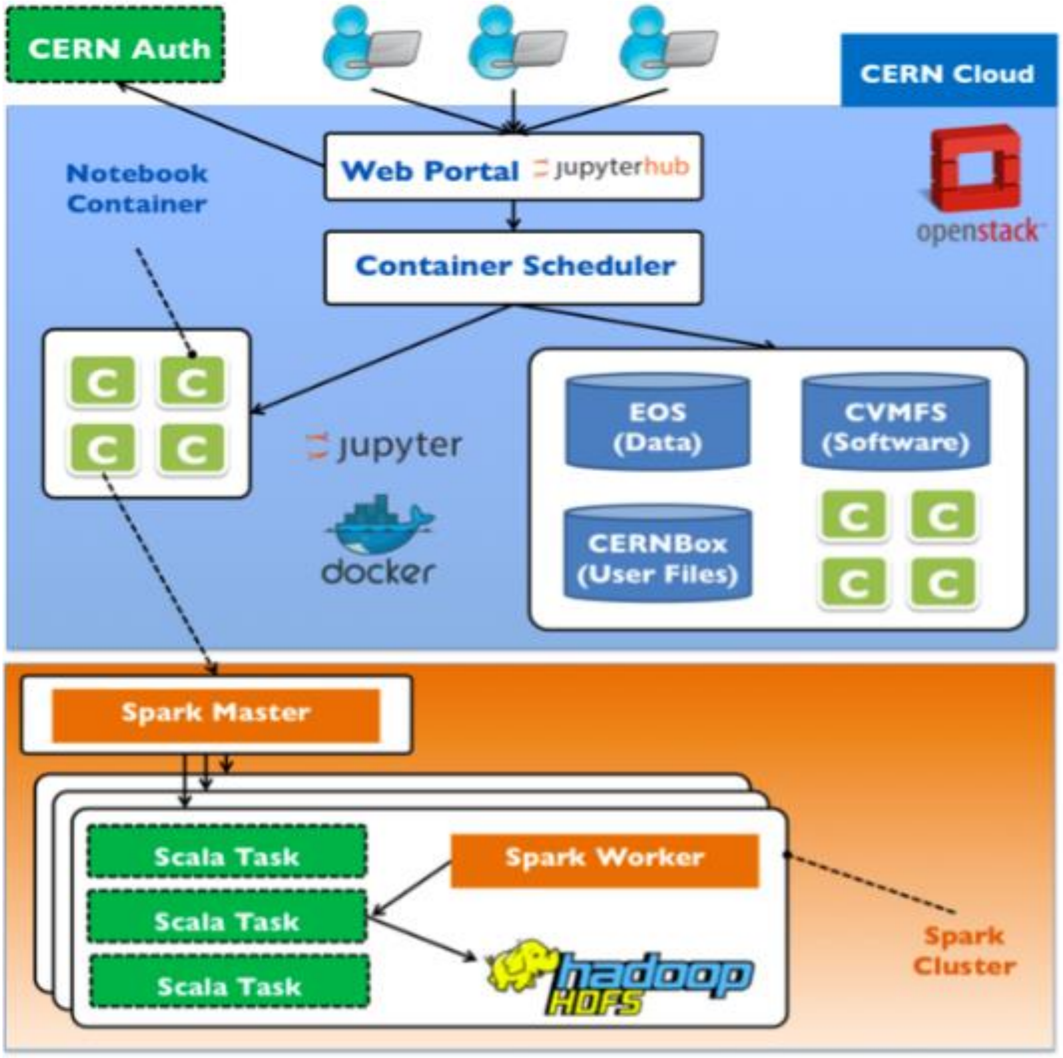

Collaboration CERN – SKA

- Recognition on both sides of potential synergies and requirements
  - Various ad-hoc interactions between communities
  - Reviews and panels etc.
  - Planning a CERN-SKA “Big data” workshop in the UK Alan Turing Inst. in Spring 2018
- On July 13 CERN and SKAO DG’s signed a collaboration agreement on computing, data management, etc.
  - Recognizing that both HL-LHC and SKA will be Exabyte-scale scientific experiments on a 10-year timescale
- ➔ More details in Miles Deegan’s talk

COLLABORATION AGREEMENT

And

CERN COURIER

Aug 11, 2017

SKA and CERN co-operate on extreme computing

Big-data co-operation agreement

Summary: European Open Science Cloud

- HL-LHC and SKA have formally recognised their potential synergies
  - Of course LSST, CTA, Virgo-Ligo, and others have very aligned requirements and constraints
    - E.g. co-incident regional centres, use of heterogeneous regional/national resources, …
    - WLCG significant experience & tools for hybrid clouds (academic, commercial, EC)
- Together this grouping would represent a very considerable influence on the direction and implementation of the EOSC
  - Understanding the commonalities *and* the differences
- It would make sense to work with EOSC as a single grouping of aligned interests
  - Requirements and needs from across the range of service providers
  - Adds weight to ensure that science drives the EOSC implementation and evolution