Ian Bird
CERN

EPS-ECFA Meeting
Venice, 8^{th} July 2017

The Challenges of Big (Science) Data

Evolution of the Universe

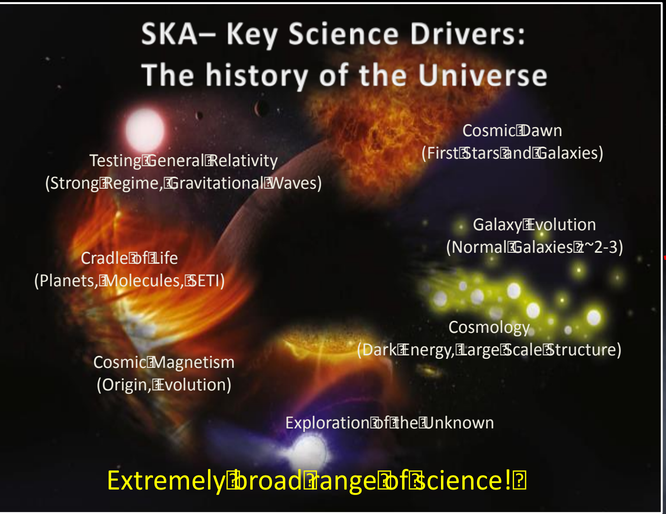

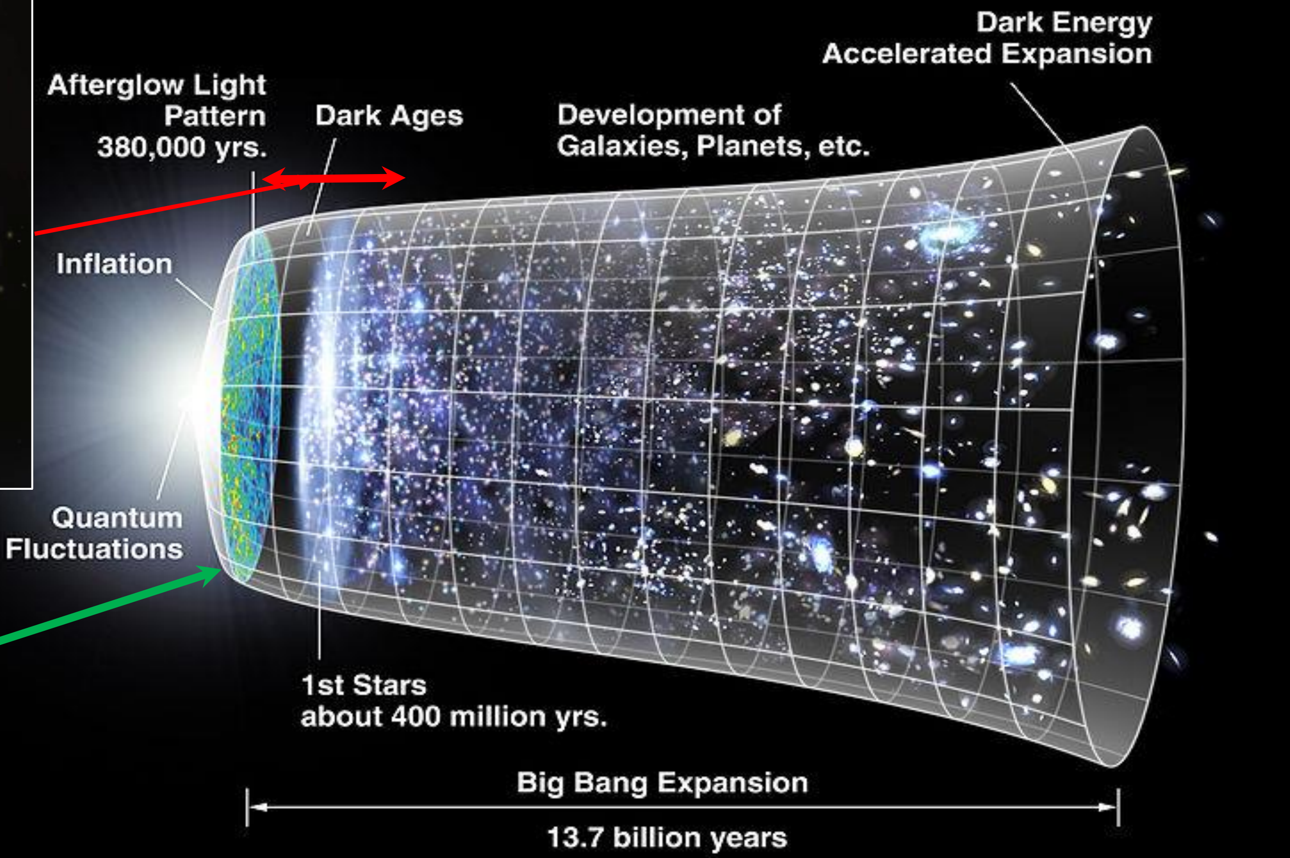

LHC drivers
Test the Standard Model?
Dark Matter?
Dark Energy?
Anti-matter?
(Gravity?)

The Large Hadron Collider (LHC)

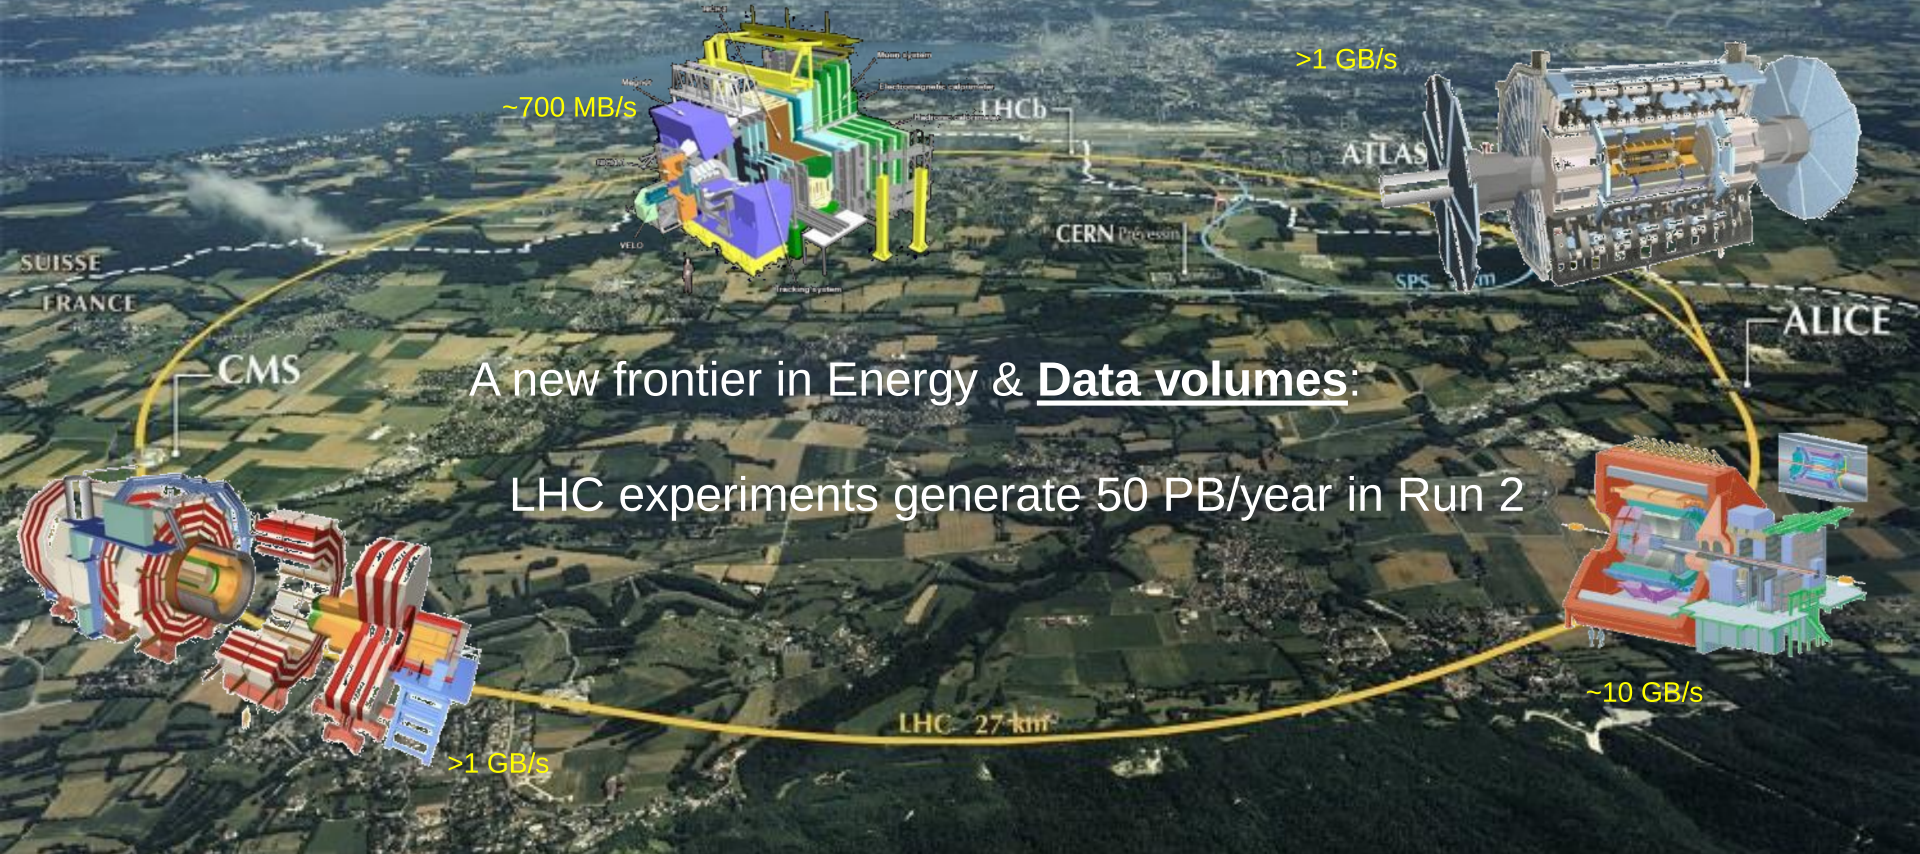

A new frontier in Energy & **Data volumes**:

LHC experiments generate 50 PB/year in Run 2

SKA - HQ in UK; telescopes in AU & SA

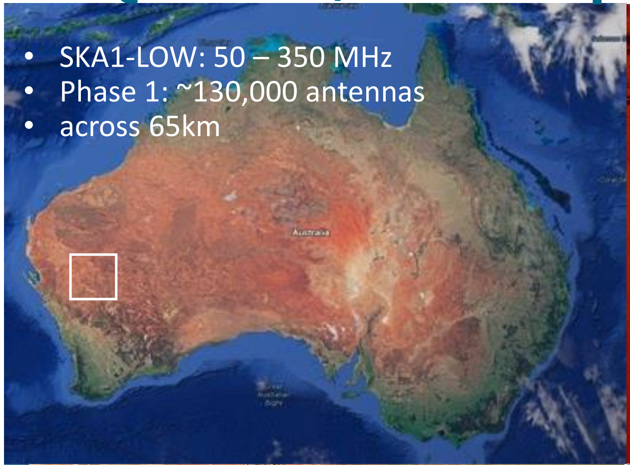

- SKA1-LOW: 50 – 350 MHz
- Phase 1: ~130,000 antennas
- across 65km

SKA1-Mid: 350 MHz – 24 GHz
Phase 1: 200 15-m dishes across 150 km

Construction Phase 1: 2018 – 2024
Phase 2: mid-2020's

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

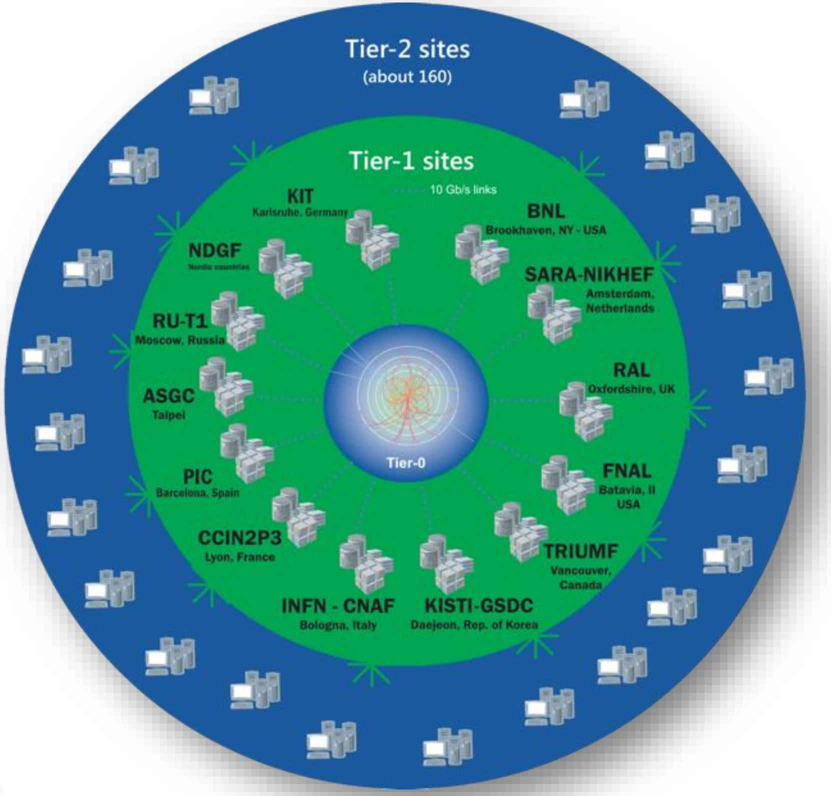

~170 sites, 42 countries

~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

Worldwide computing

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

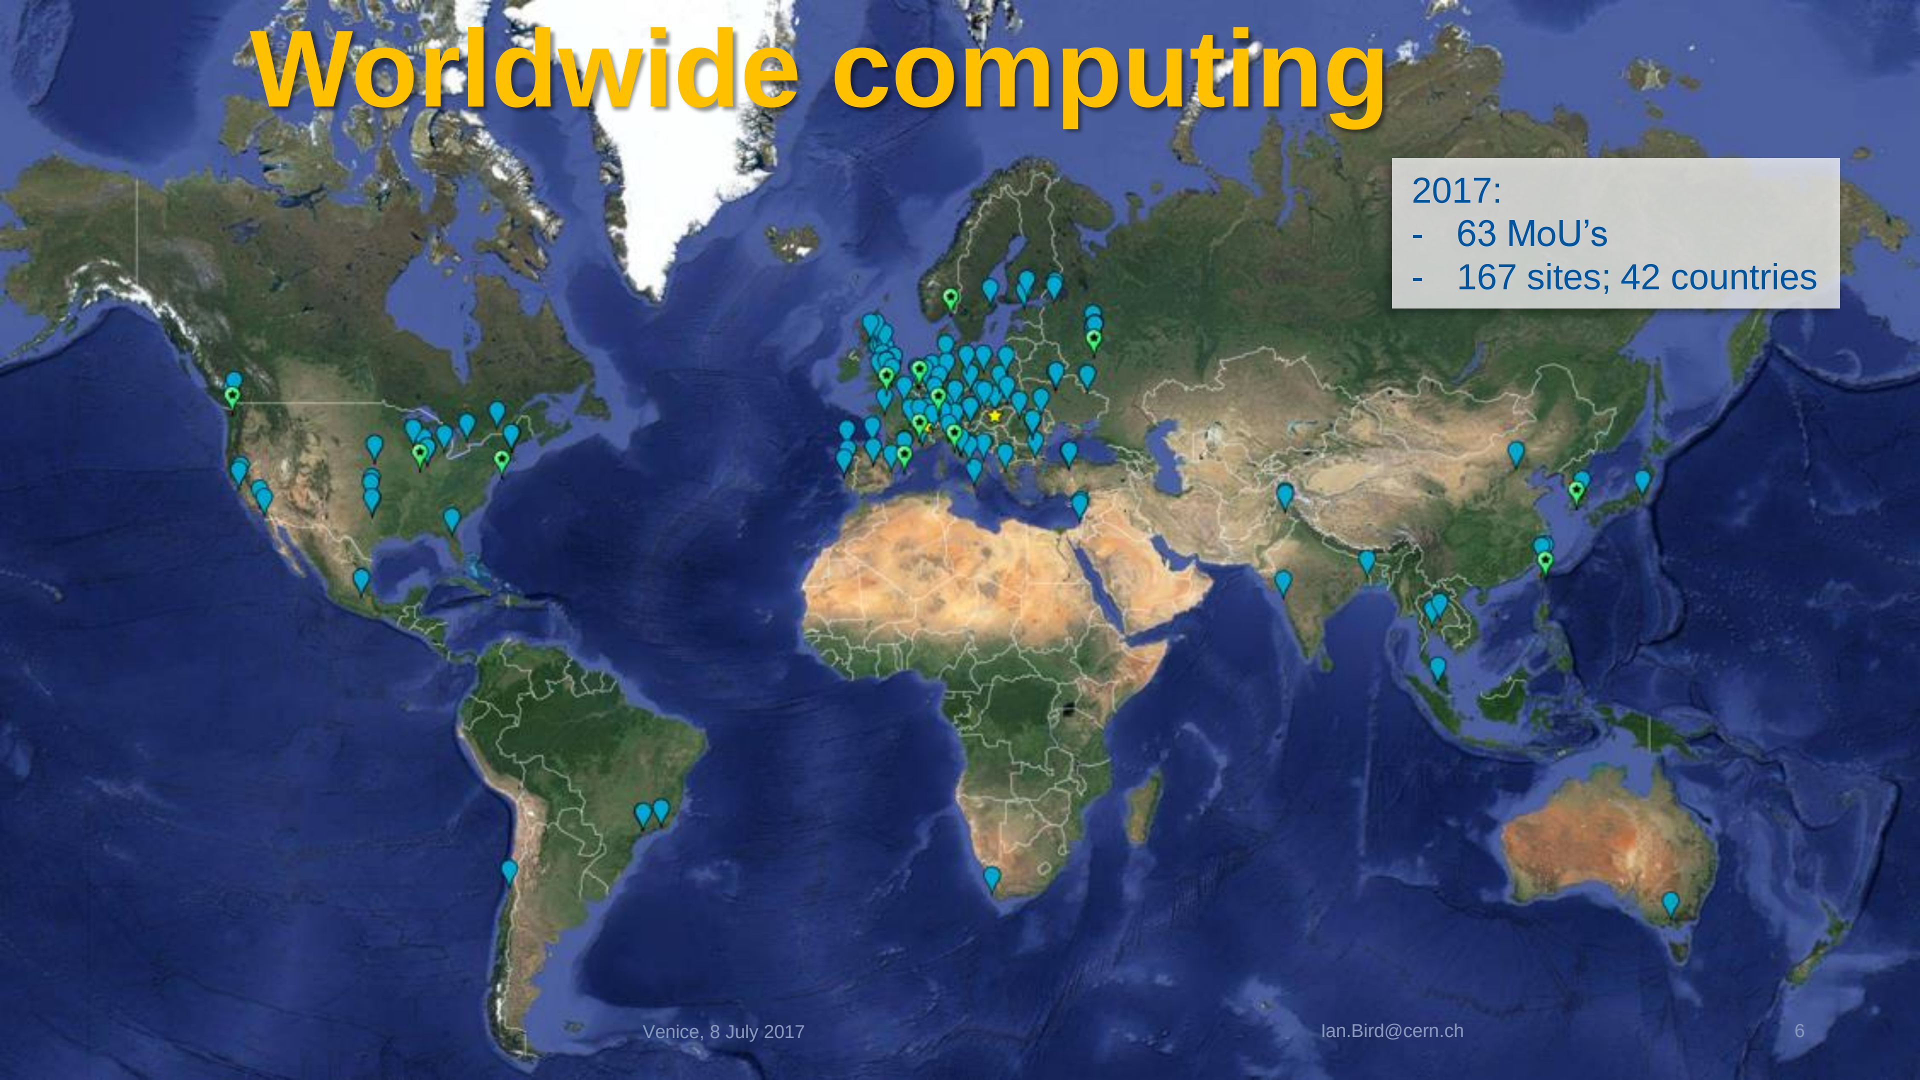

SKA Organisation: 10 countries, more to join

Australia (DoI&S)
Canada (NRC-HIA)
China (MOST)
India (DAE)
Italy (INAF)
Netherlands (NWO)
New Zealand (MED)
South Africa (DST)
Sweden (Chalmers)
UK (STFC)

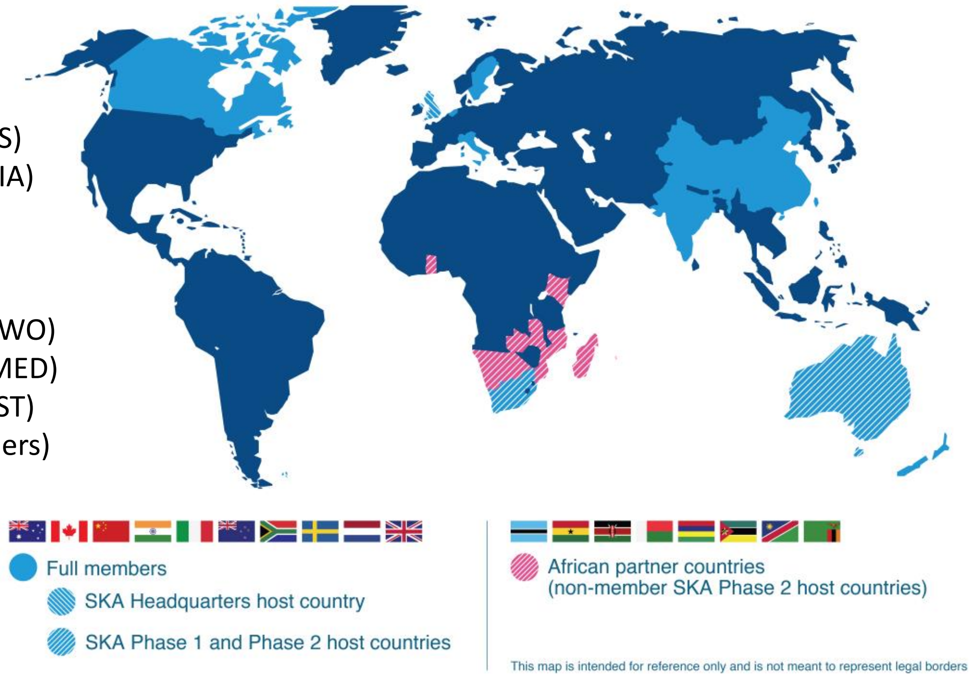

Interested Countries:

- France
- Germany
- Japan
- Korea
- Malta
- Portugal
- Spain
- Switzerland
- USA

Contacts:

- Mexico
- Brazil
- Ireland
- Russia

SKA Data Acquisition System Overview

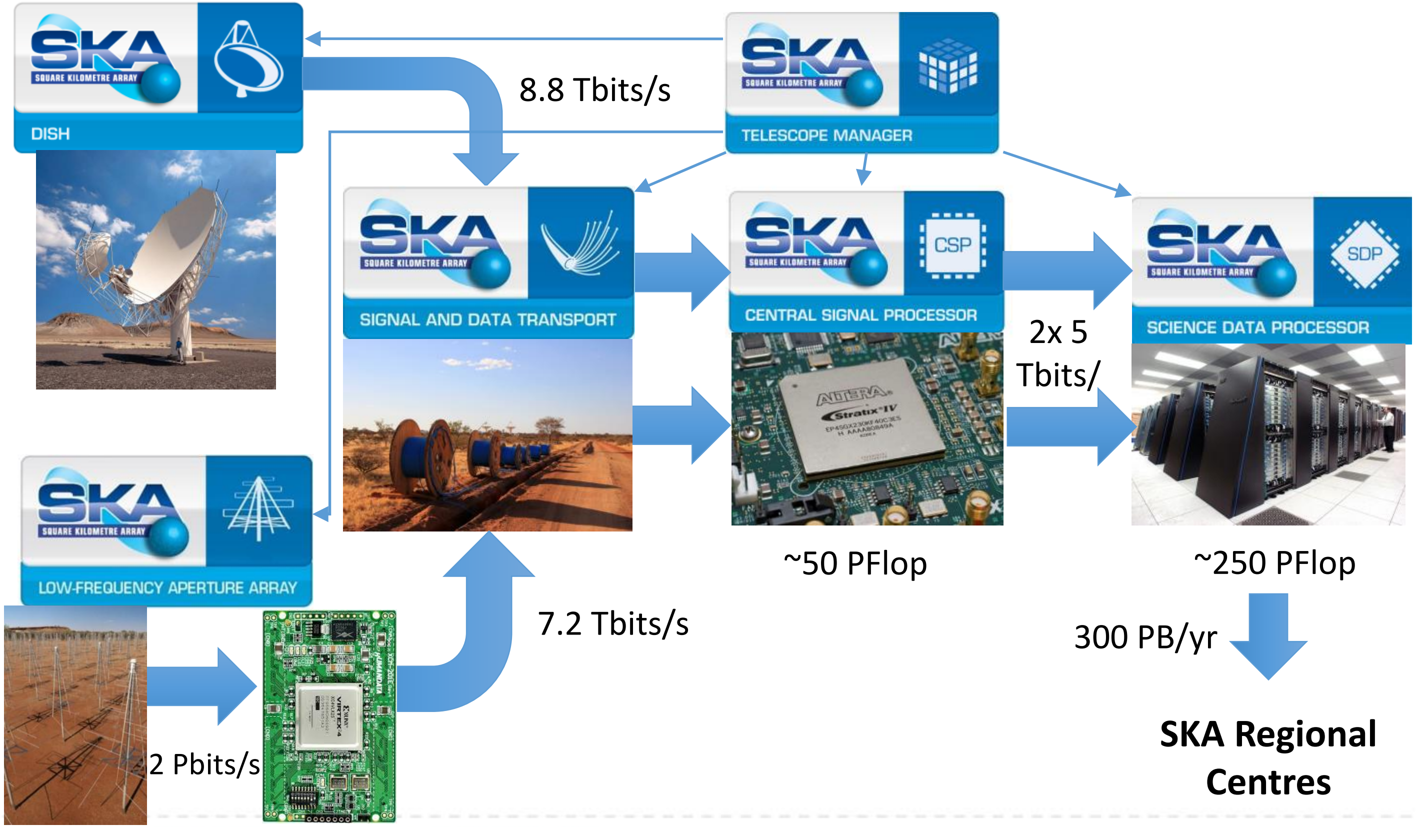

SKA Observatory

- SKA is a software telescope
  - Very flexible and potentially easy to reconfigure
  - Major software and computing challenge
- Computing challenges are significant
  - Science Data Processor (SDP) needs 25 PetaFLOPS/sec of delivered processing
    - Current estimate is that SDP needs 250 PFLOP/sec peak.
    - Tianhe-2 – 50 PetaFLOPS/sec peak.
    - Memory bandwidth is ~200 PetaBytes/sec
  - Pulsar Search is an additional 50 PFLOP/s of peak processing
  - Power efficiency required is ~40x better than Tianhe-2
- Limitations
  - Because of the bandwidth requirements, have to buy 10 x more computing than a pure HPC system would require.

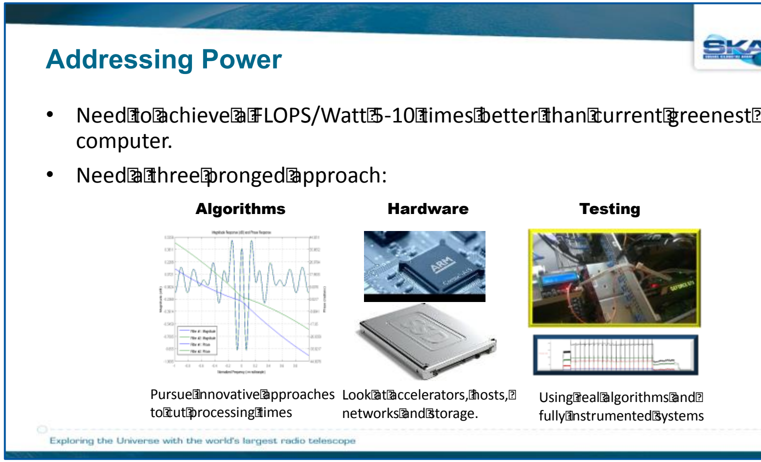

Data Products from SDP

- Image cubes of sky maps (Continuum) 18PB n*/year
  Max ~30Gbit/s Mean ~10 Gbit/s
  - SDP Fourier transform of CSP time series
  - Moderate external compute and model fitting
- Epoch of Re-ionisation 1.6PB/ 6 Hr ~600 Gbit/s
  Only 11% duty cycle so Mean rate ~ 22 Gbit/s
  - Uses calibrated aperture plane (visibility data)
  - Enormous compute of power spectra
- Relativity Gravitational lensing 70PB/2500Hr   60Gbit/s
  - Uses further processed aperture plane data
  - Considerable compute of galaxy elipsicity
- Pulsars 4-5 PBytes/y ≤10Gbits
  - Discovery and in depth study; timing 10 ns in 10 years
  - Large physics compute
- VLBI 4 beams of 4 Gbit/s (not in the 300PB)
  - Data direct from the correlator, UDP all the way
- PI led projects ? Some fraction of observing time
- Discovery science Whatever is available

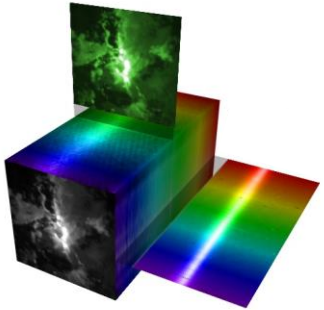

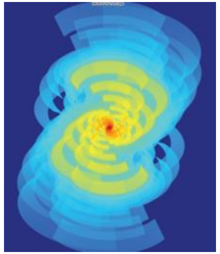

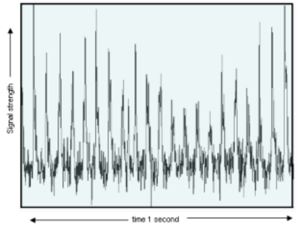

LHC Schedule

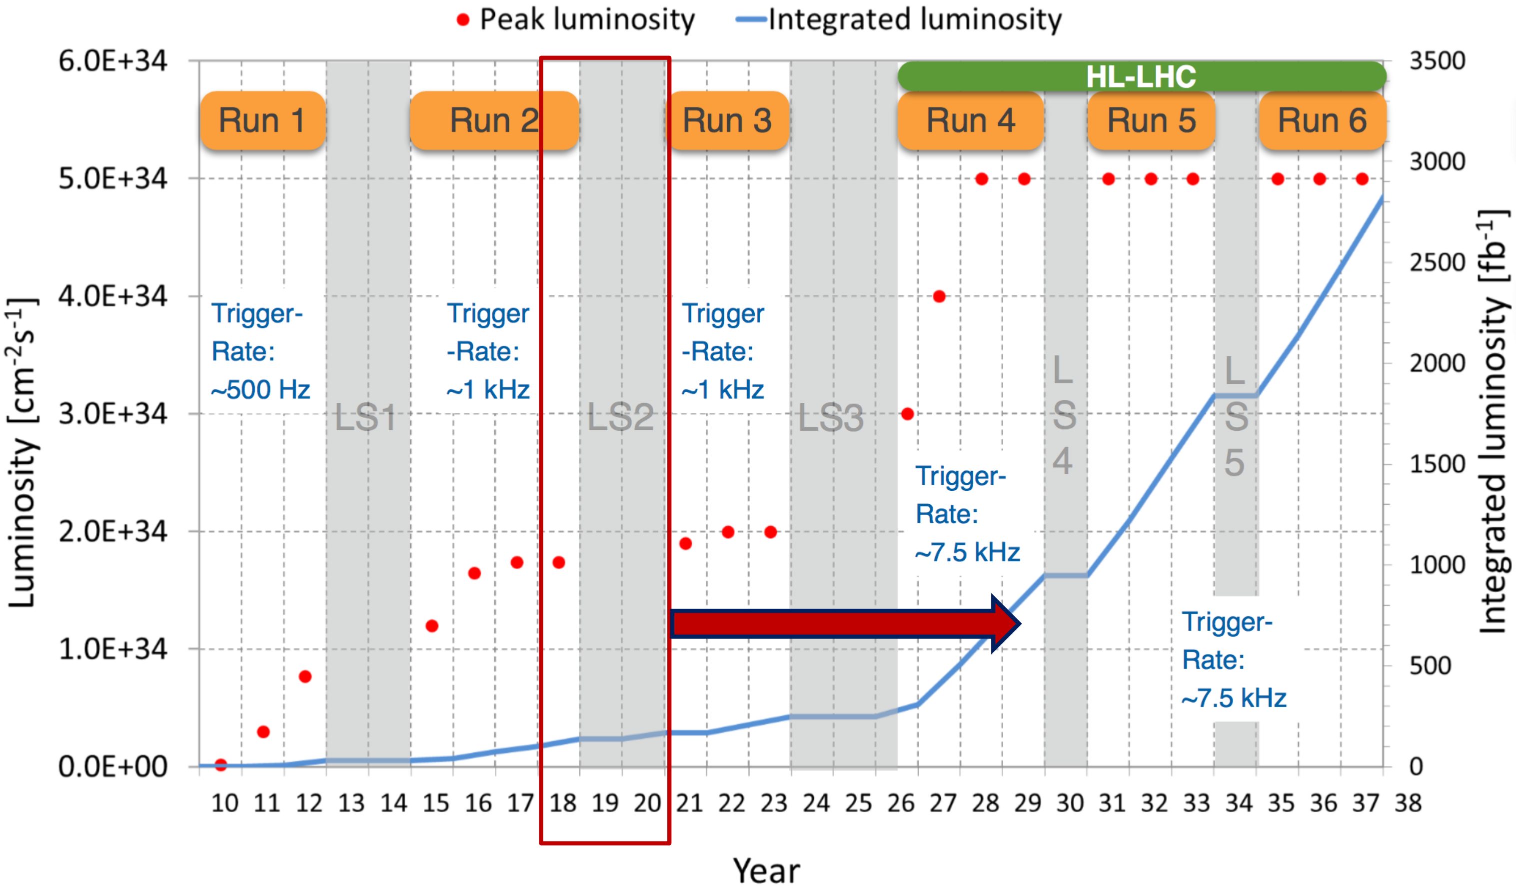

Run 3 — Alice, LHCb upgrades

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?

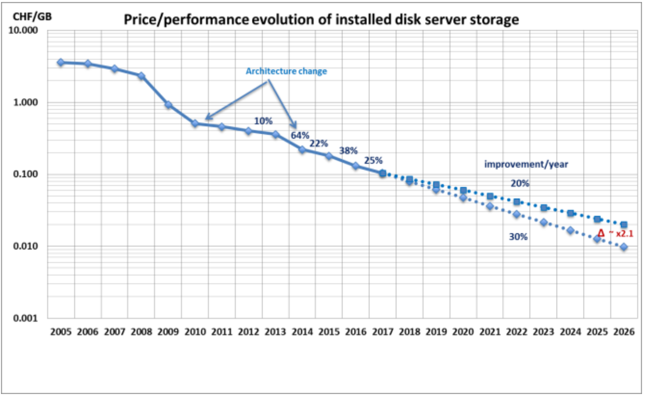

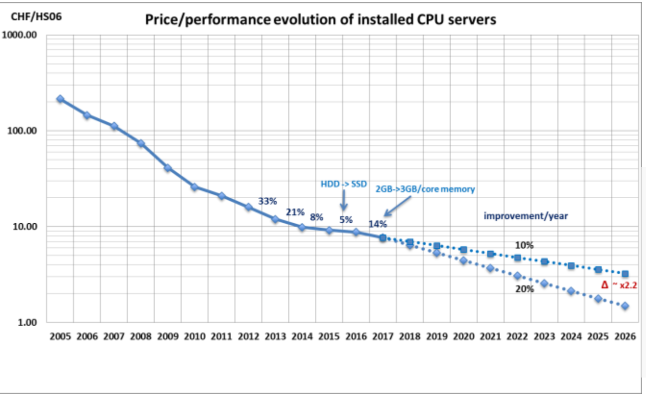

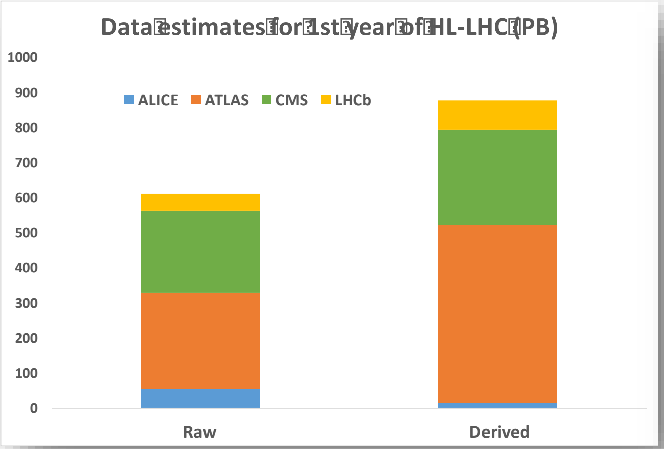

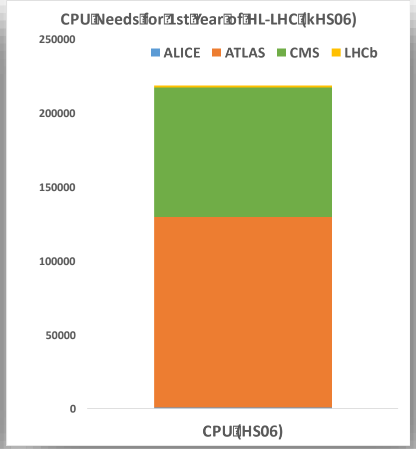

Data:
- Raw 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB

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

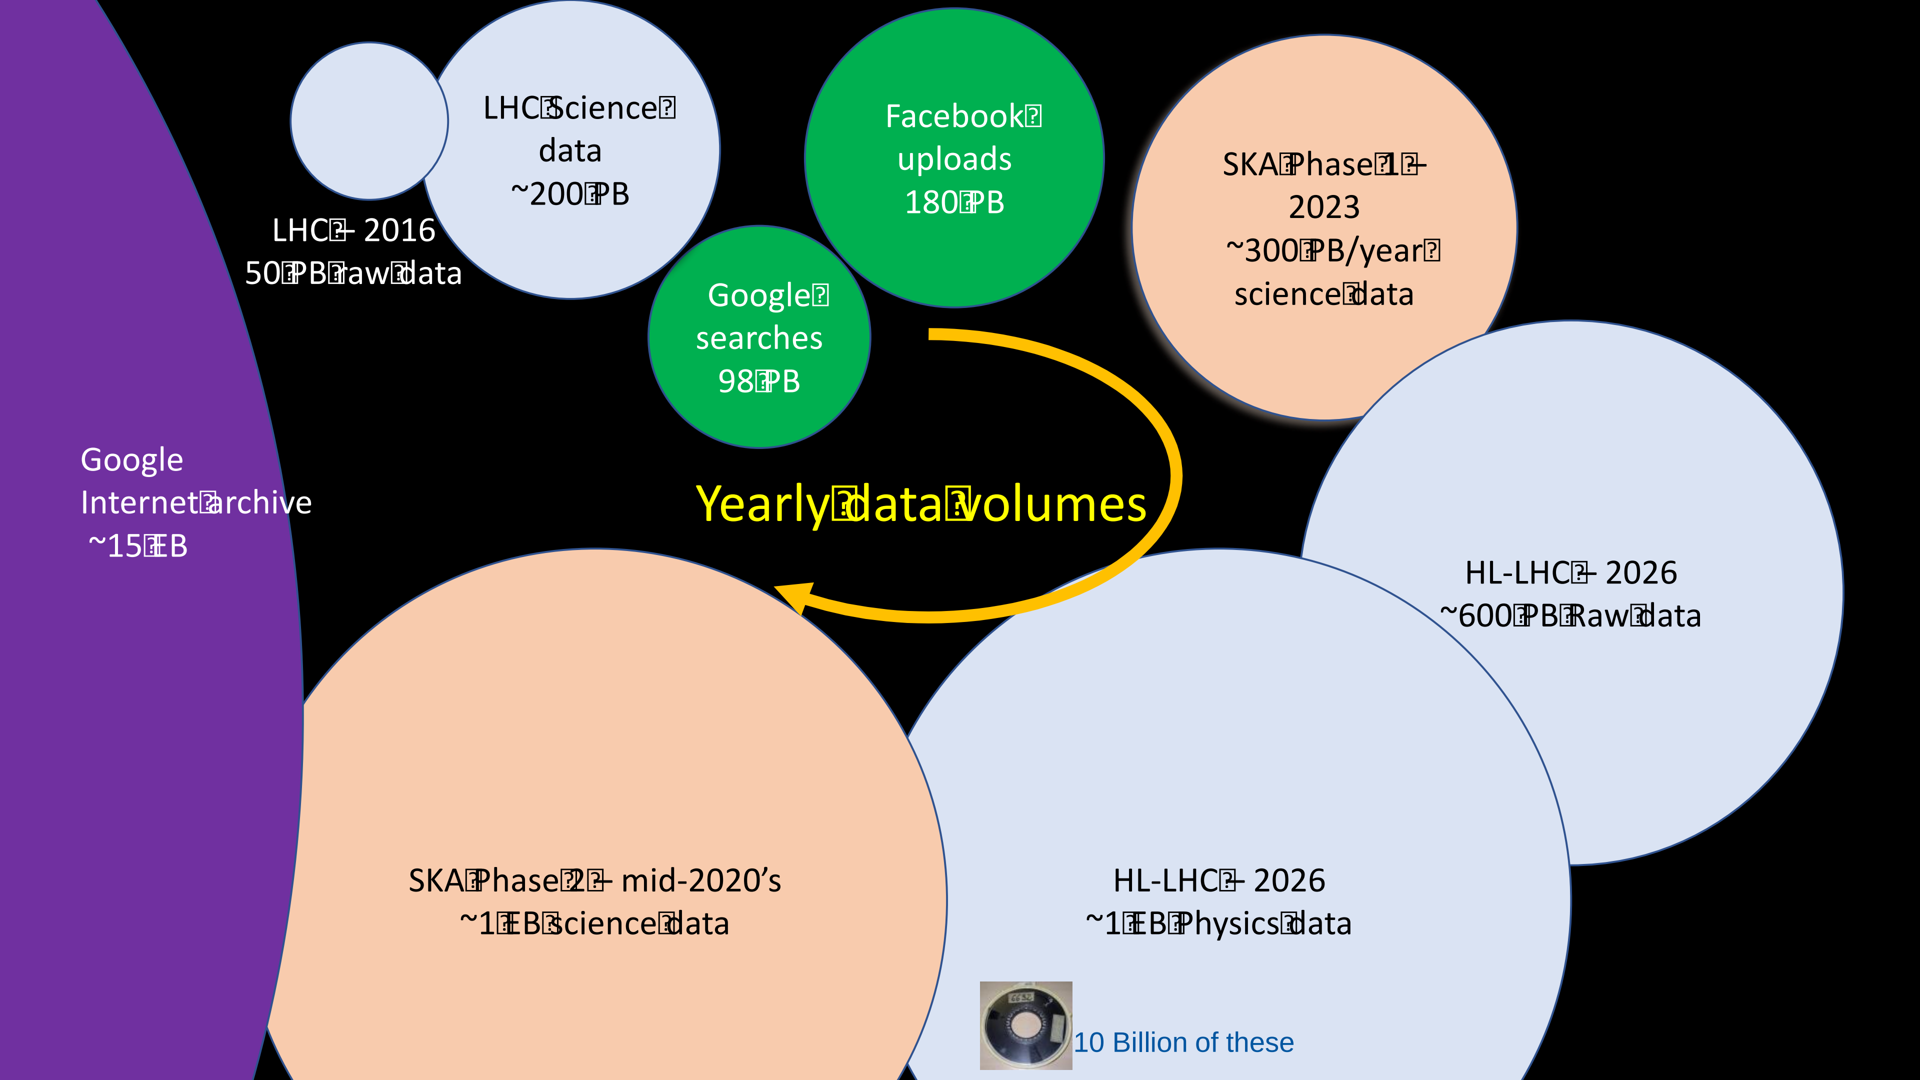
LHC Science data ~200 PB
LHC – 2016 50 PB raw data
Facebook uploads 180 PB
Google searches 98 PB
SKA Phase 1 – 2023 ~300 PB/year science data
Google Internet archive ~15 EB
Yearly data volumes
HL-LHC – 2026 ~600 PB Raw data
SKA Phase 2 – mid-2020's ~1 EB science data
HL-LHC – 2026 ~1 EB Physics data
10 Billion of these

LHCOPN

Networks

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

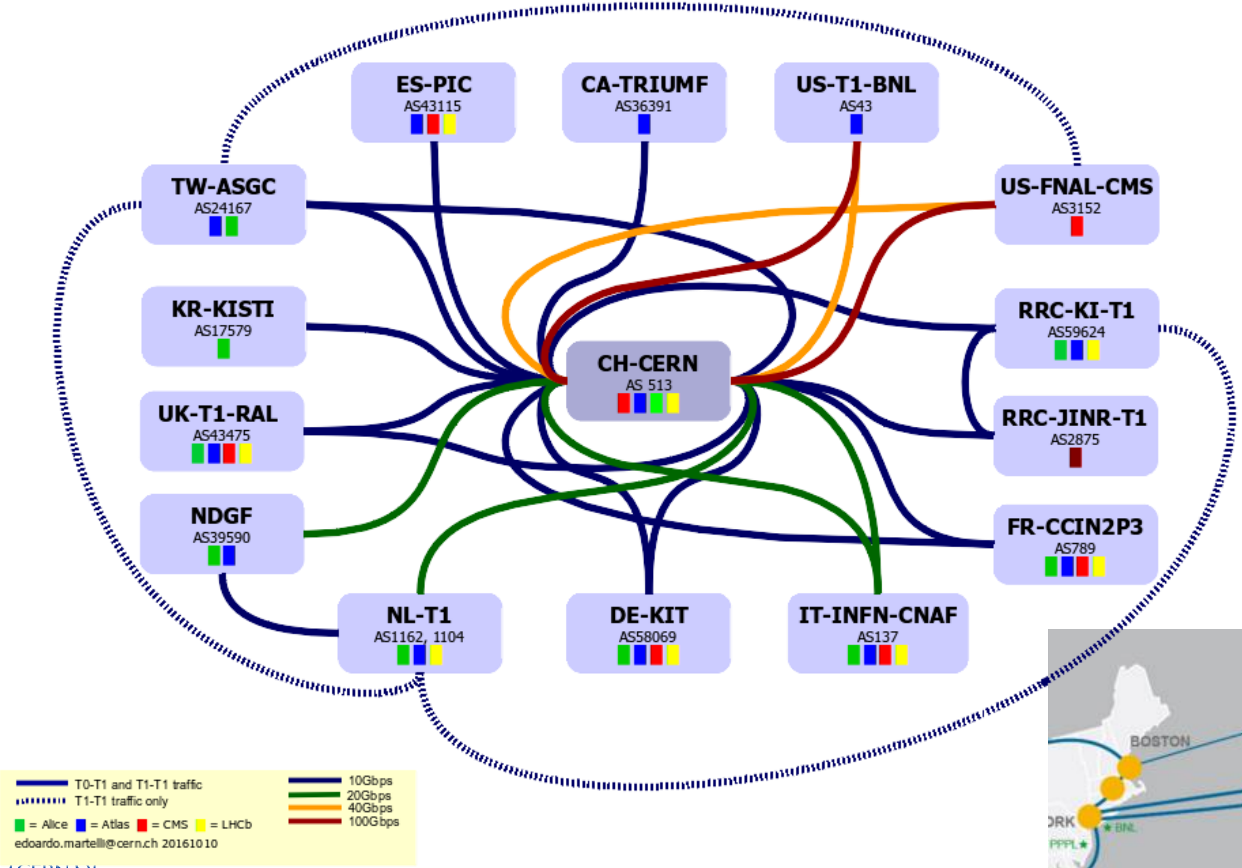

Up to 340 Gbps transatlantic

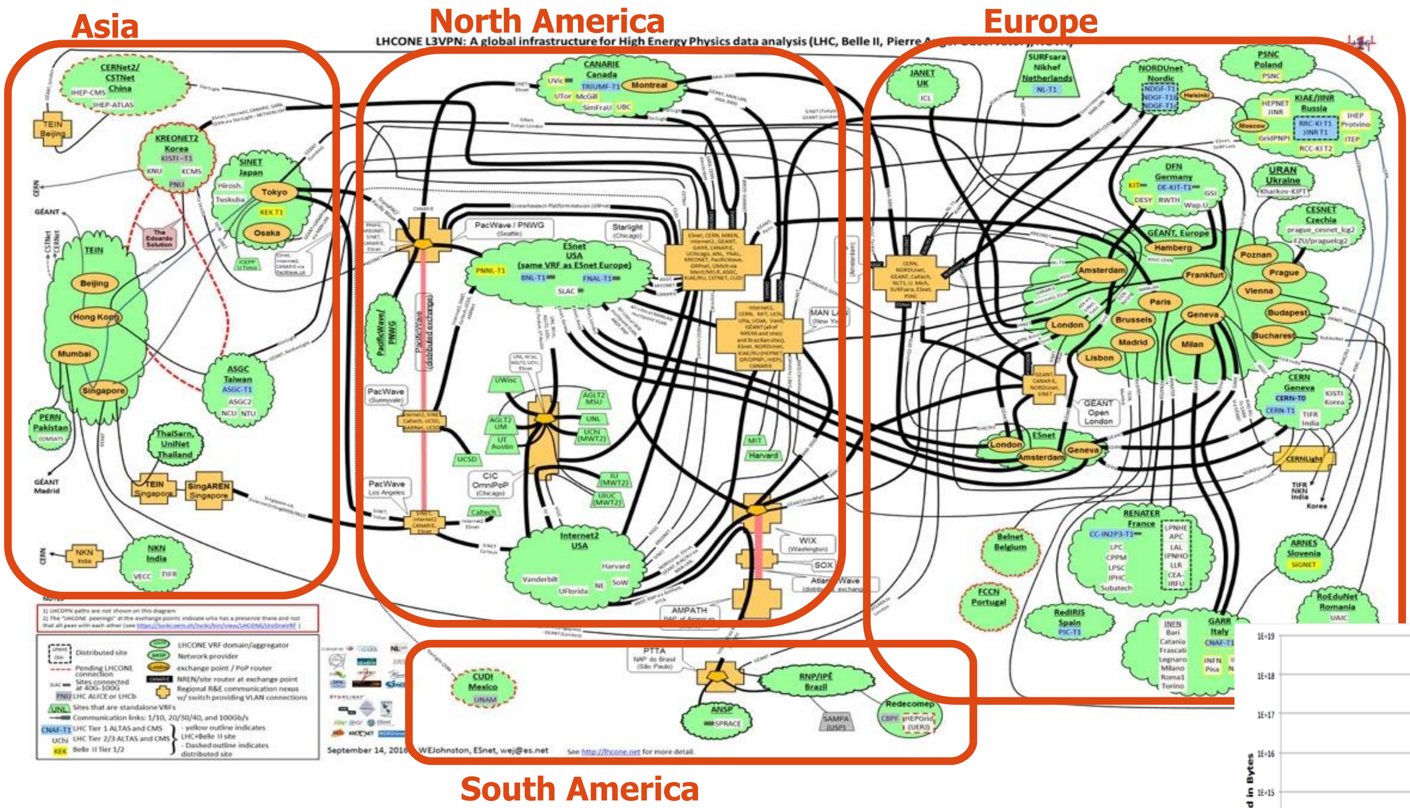

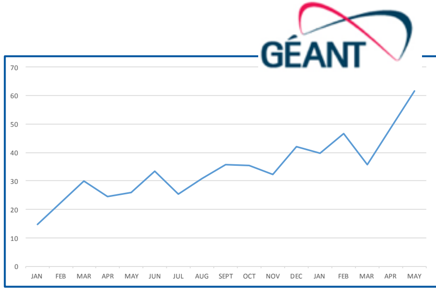

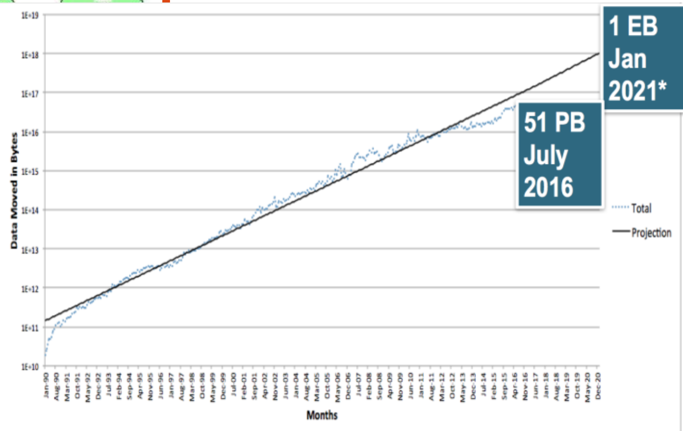

LHCOne: Overlay network

Allows NREN's to manage HEP traffic on general purpose network

Managed by NREN collaboration

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

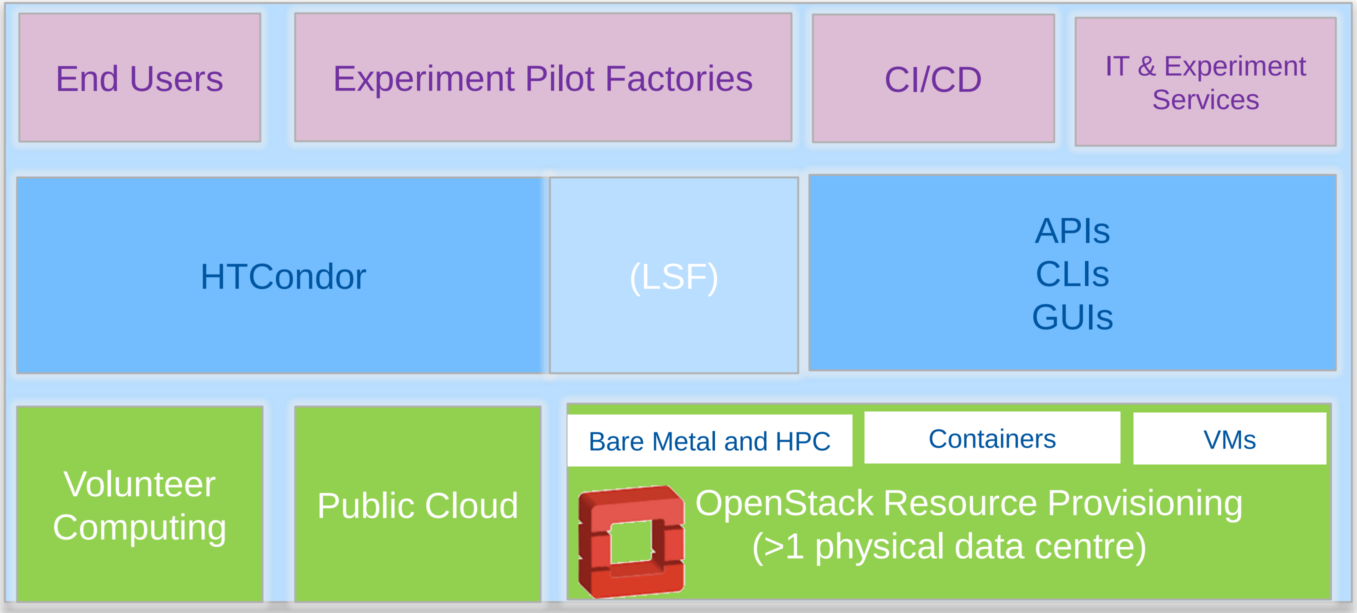

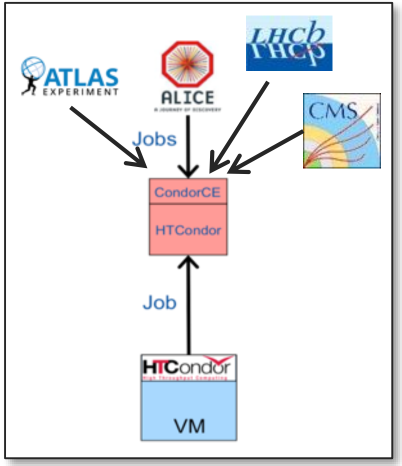

Ian.Bird@cern.ch

Commercial Clouds

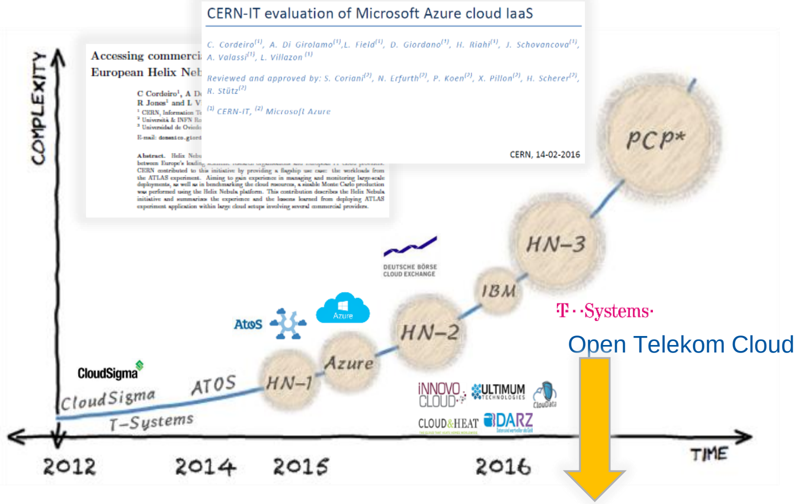

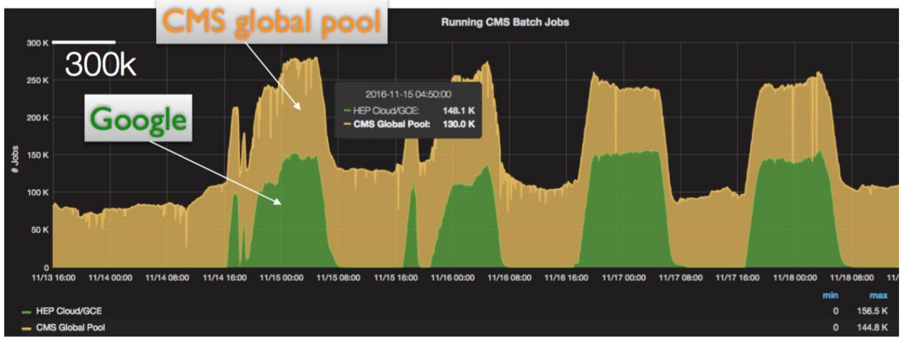

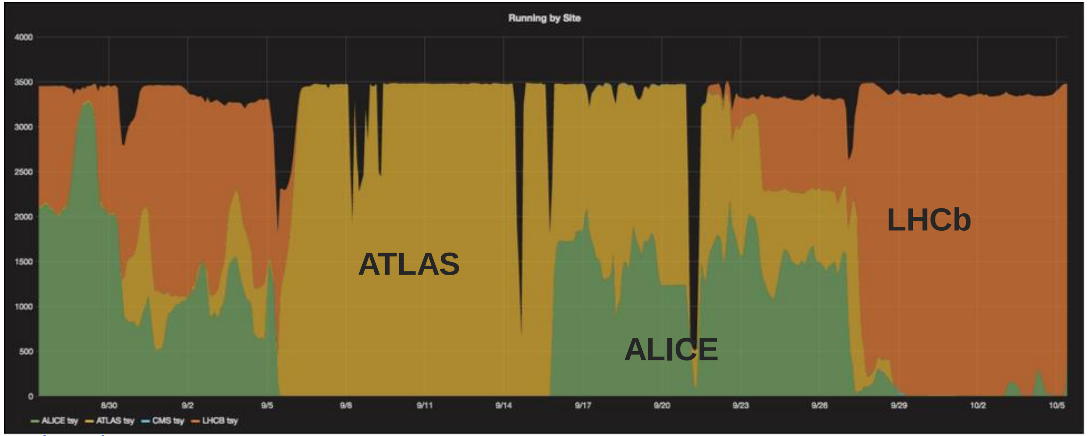

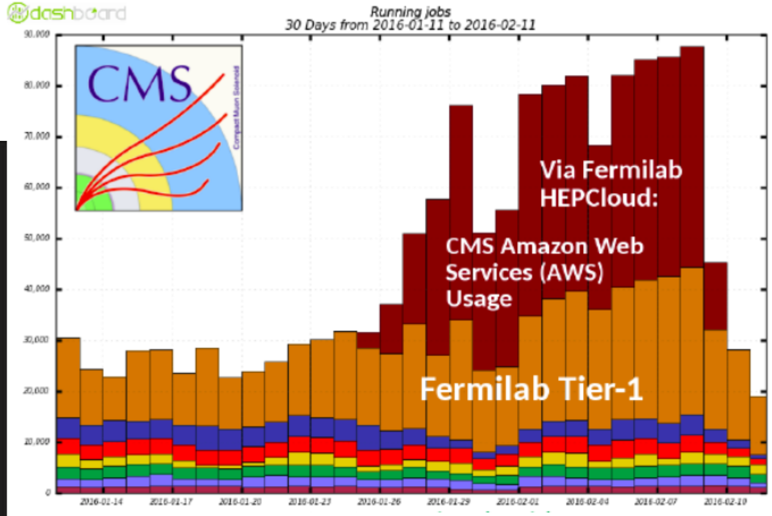

SKA Regional Centres

- SKA envisage a network of regional centres
- Modelled on the LHC Tierw
- Not part of SKA, or funded by SKA, but:
  - Essential to generate science
  - Coordinated with assistance from SKAO and accredited with SKAO
- Principle functions
  - Take data products generated by SDP and turn them into science
  - Support regional astronomers with their data processing.
  - Act as centres for domain expertise.

Regional Centre Concept:

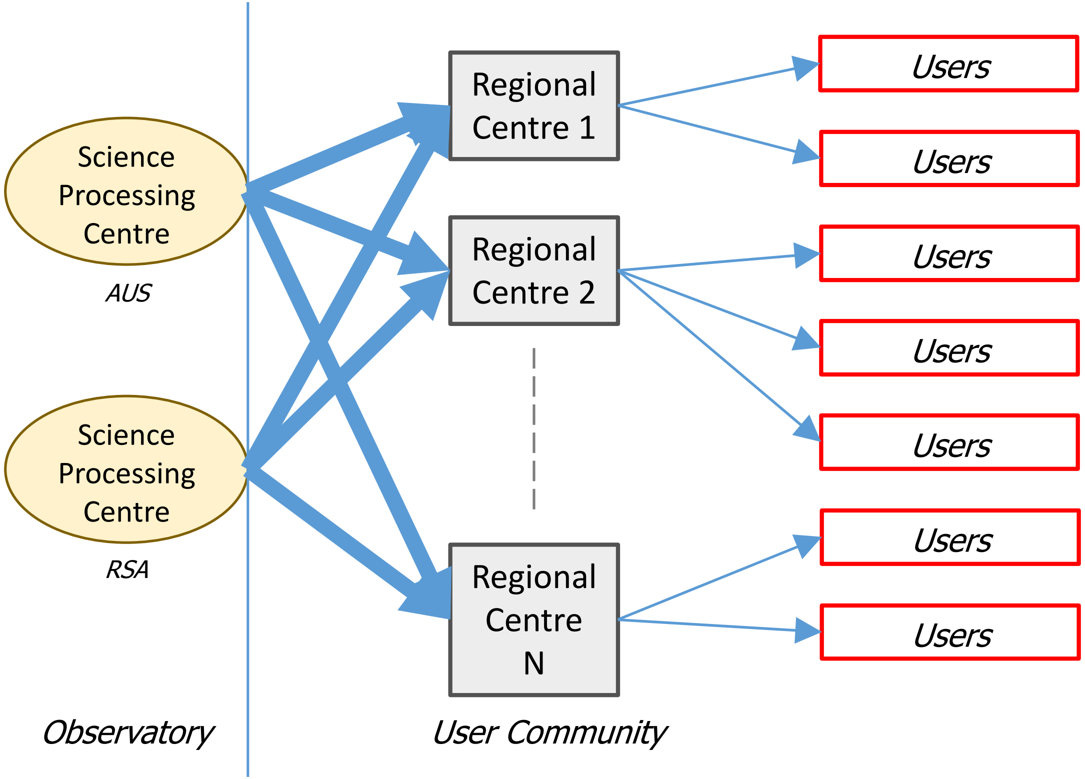

Regional Centre Network

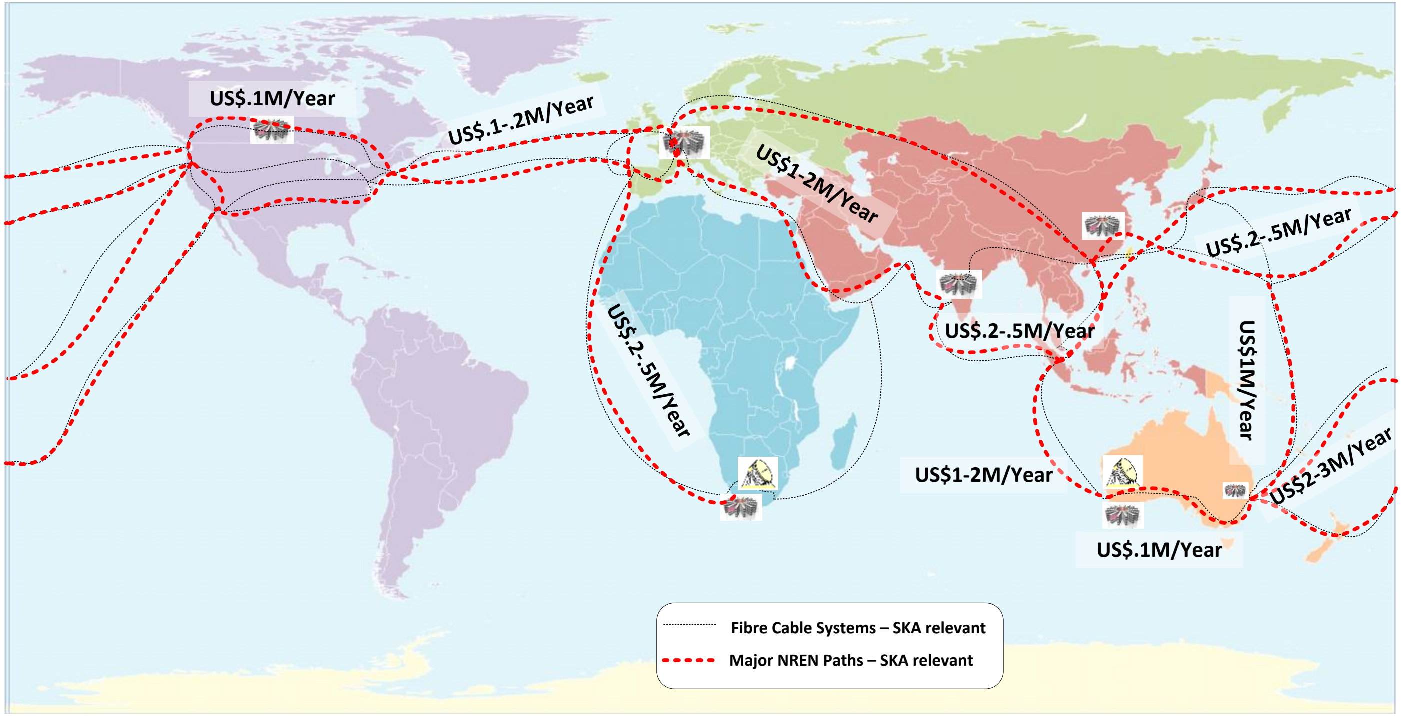

v

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
- SKA will have similar data volumes on the same time-scale
- Opportunity for synergy – in particular in large scale facilities
  - SKA and LHC likely to be co-located in major facilities
- 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

HL-LHC computing cost parameters

Business of the experiments:
amount of Raw data, thresholds;
Detector design long term
computing cost implications

Business of the experiments:
reconstruction, and
simulation algorithms

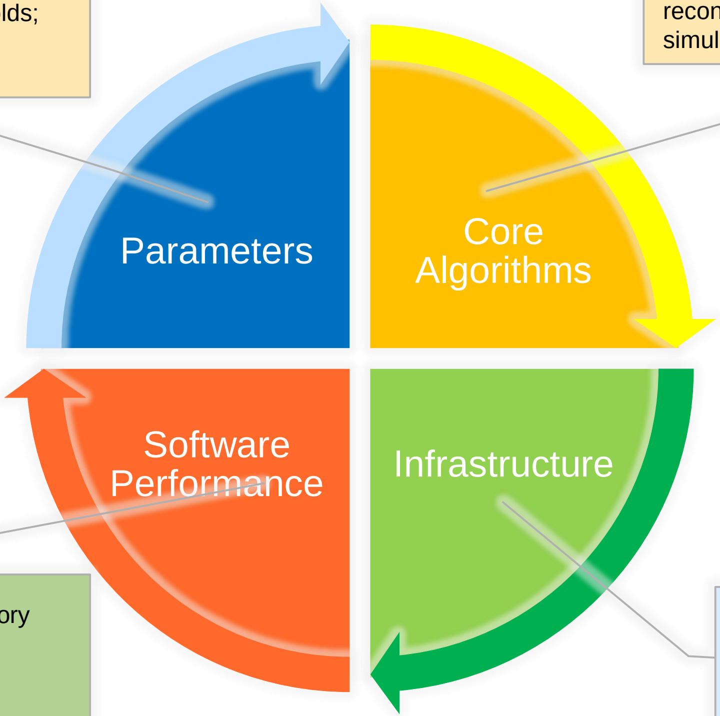

Performance/architectures/memory etc.;
Tools to support: automated build/validation
Collaboration with externals – via HSF

New grid/cloud models; optimize CPU/disk/network; economies of scale via clouds, joint procurements etc.

Software

- Not just a HEP problem
- Transistors go into many cores, co-processors, vector units, etc.
- Memory access and I/O paths also become problematic
- Getting performant software requires significant investment in skills

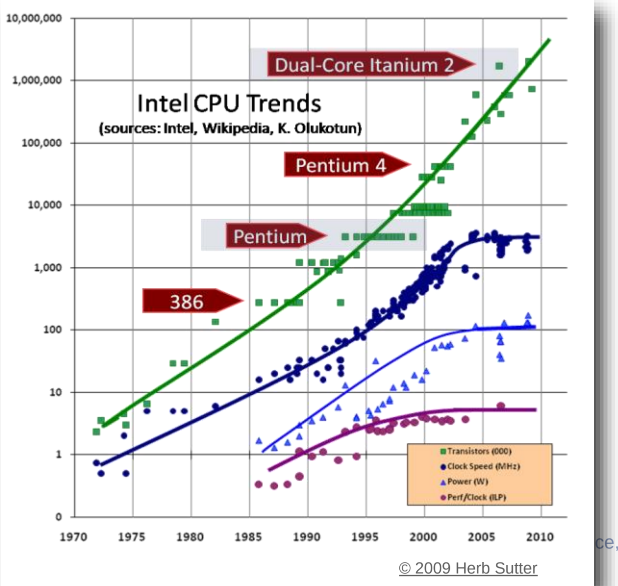

© 2009 Herb Sutter

The new realism: software runs slowly on supercomputers

No supercomputer runs real applications faster than five percent of its design speed. Robert Roe and Tom Wilkie report on recalibrating expectations of exascale and on efforts to tune software to run faster

he speed with which supercomputers

director of the HLRS supercomputer centre in Stuttgart, Germany. The HLRS was not targeting exascale per se, he said, because the centre's focus was on the compute power it could actually deliver to its users.

According to Resch, simple arithmetic meant that if an exascale machine achieved a sustained performance of only 1 to 3 percent, then this would deliver 10 to 30 Petaflops. So buying a 100 Petaflop machine that was 30 per cent efficient – which should be achievable, he claimed – would deliver the same compute power, for a much lower capital cost and about one tenth the energy cost of an exascale machine.

The Jülich Supercomputing Centre (JSC) in Germany is tackling the issues of software

encourage our users to try and reach exascale readiness.'

At the Lawrence Livermore National Laboratory, effort is going into developing APIs and tools to create applications and effectively optimise how code is run on a cluster. Supinski stated that the LLNL's plan was to use a programming tool developed at Livermore called Raja. 'The idea of Raja is to build on top of new features of the C++ standard. The main thing that we are looking at is improving application performance over what we are getting on Sierra, Sequoia, and Titan. Application performance requirements are what we really care about.'

'The US Coral programme means that applications will be running on systems with in excess of 100 ple of years, and US national labs will uirements as much as proach to the problem ns to run on the new

e a peak performance

wares. As part of its logy to the Editor

entific-computing.com

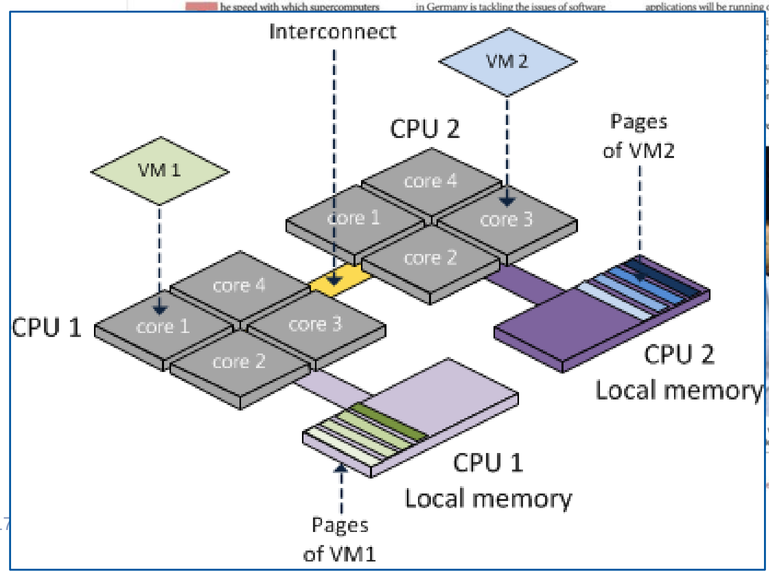

Possible Model for future HEP computing infrastructure

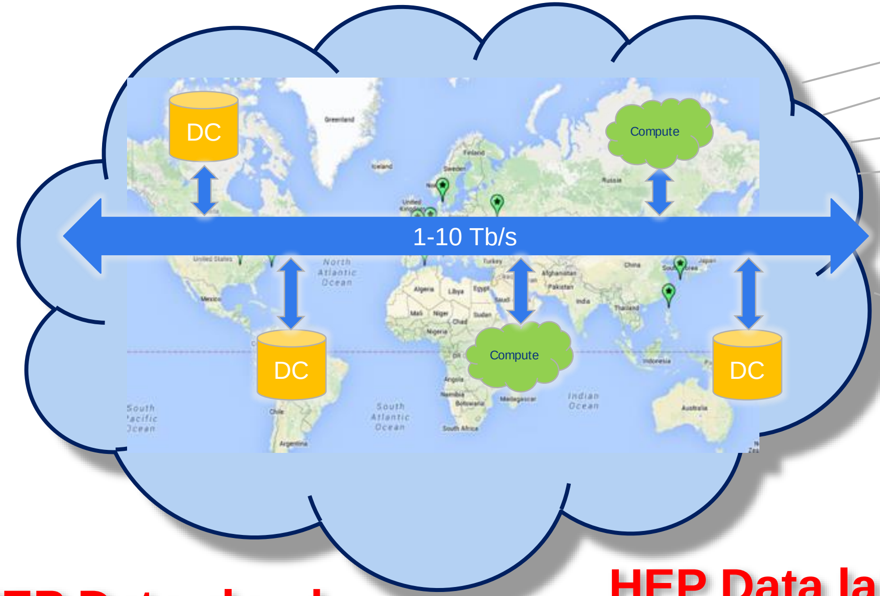

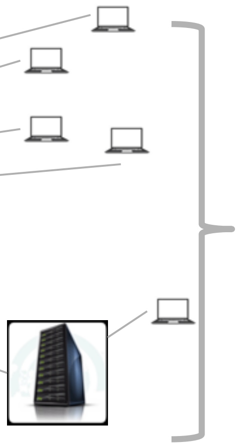

Cloud users:
Analysis

HEP Data cloud
Storage and compute

HEP Data lake
Storage and compute

Hadoop and Analytics

- **Hadoop Production Service**
- New scalable data services
  - Scalable databases
  - Hadoop ecosystem
  - Time Series databases
- Big Data Analytics
- Activities and objectives
  - Develop projects and services with/for users
  - Support of Hadoop Components
  - Further value of Analytics solutions
  - Define scalable platform evolution based on requirements

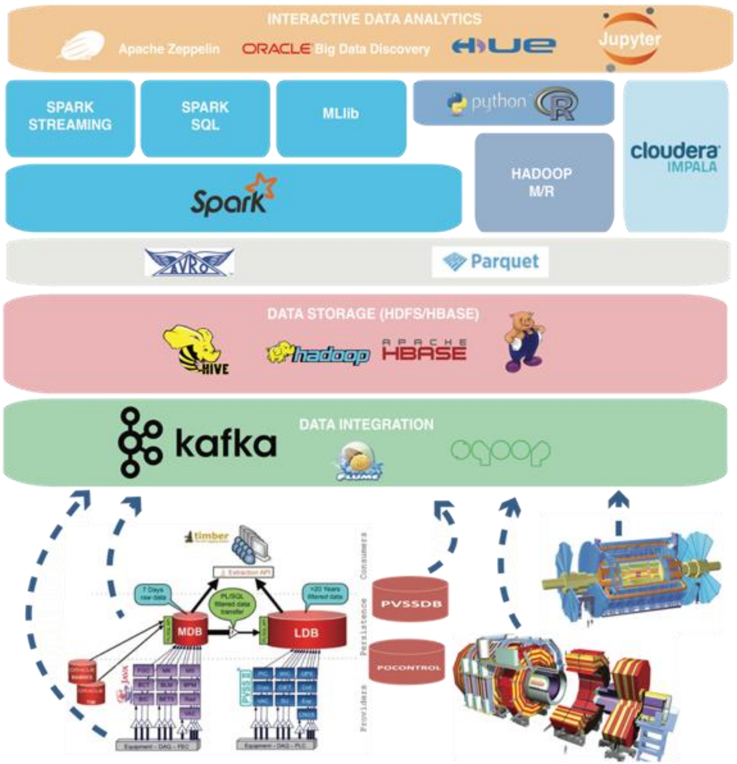

Ian.Bird@cern.ch

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)

kaggle | Host | Competitions | Datasets | Scripts | Jobs | Community ▾ | Sign up | Login

Completed • $15,000 • 673 teams

Flavours of Physics: Finding τ → μμμ

Mon 20 Jul 2015 – Mon 12 Oct 2015 (7 months ago)

Dashboard

Home
Data
Make a submission

Information
Description
Evaluation

Competition Details » Get the Data » Make a submission

Identify a rare decay phenomenon

Like last year's Higgs Boson Machine Learning Challenge, this competition deals with the physics at the Large Hadron Collider (LHC). However, the subject of last year's challenge, the Higgs Boson, was already known to exist. The aim of this year's challenge ...nd a phenomenon that is not already known to exist – charged lepton flavour ...on – thereby helping to establish "new physics".

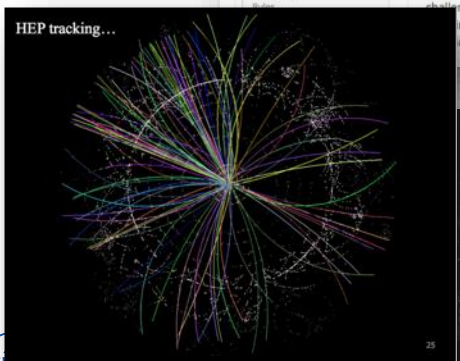

- Data taking
  - Real time event categorization
  - Data monitoring & certification robot
- Data Reconstruction
  - Calorimeter reconstruction
  - Boosted object jet tagging
- Data Processing
  - Computing Resource Optimization
  - Predicting data popularity
  - Intelligent networking
- Data Analysis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

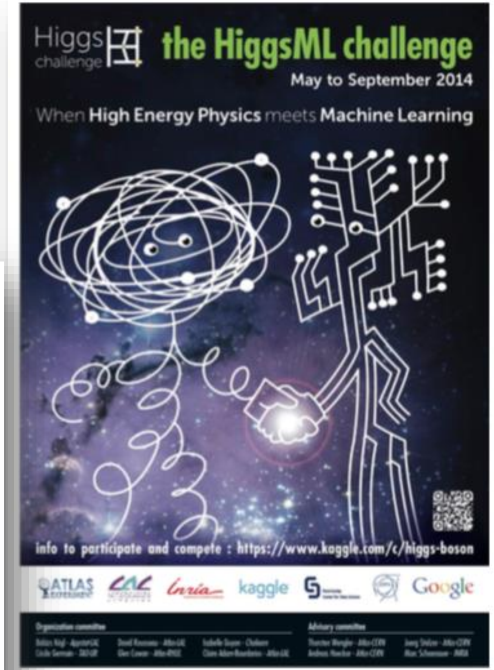

Machine learning and data analytics are hot topics at CERN openlab workshop

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 openlab partner companies Intel and Siemens, as well as contributors Cisco and associates Yandex. 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 openlab CTO. "I greatly appreciated 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

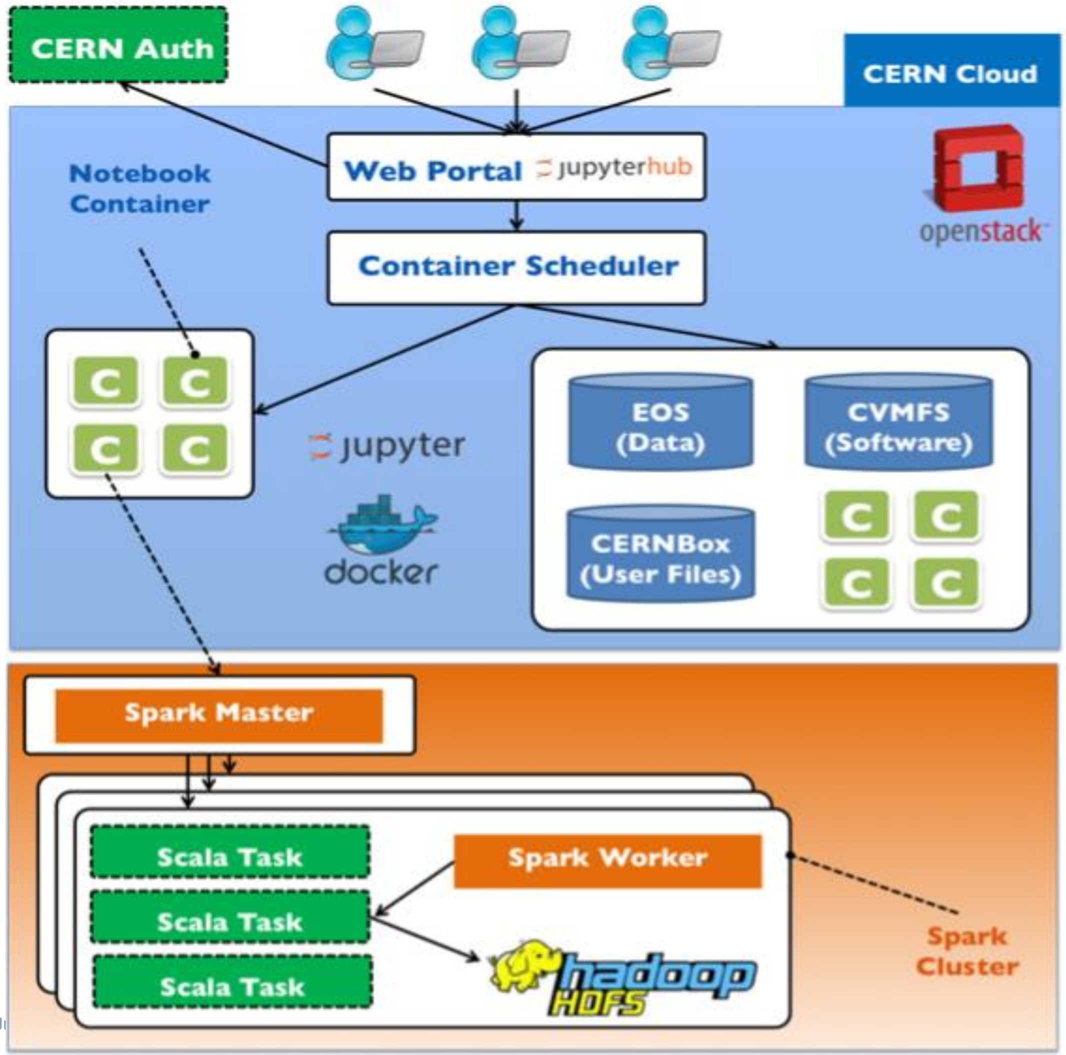

Conclusions

- In the next decade HL-LHC and SKA will be the leading Exabyte-scale scientific big data challenges
- Computing & software challenges and concerns are often similar or complementary
- Evolution of computing models must be very agile to be able to adapt to a rapidly evolving landscape of new technologies & funding opportunities
- Many opportunities for synergy, collaboration, and leadership in scientific big data

Acknowledgements

- Thanks to:
  - Nick Rees, Rosie Bolton, Richard Hughes-Jones – SKA Organisation & Project Office, for background, discussions & material on SKA

Ian.Bird@cern.ch

CERN