Data Management in Particle Physics

UoB Open Research Forum
16 June 2021

Paul Newman
(Birmingham Particle Physics Group Leader)

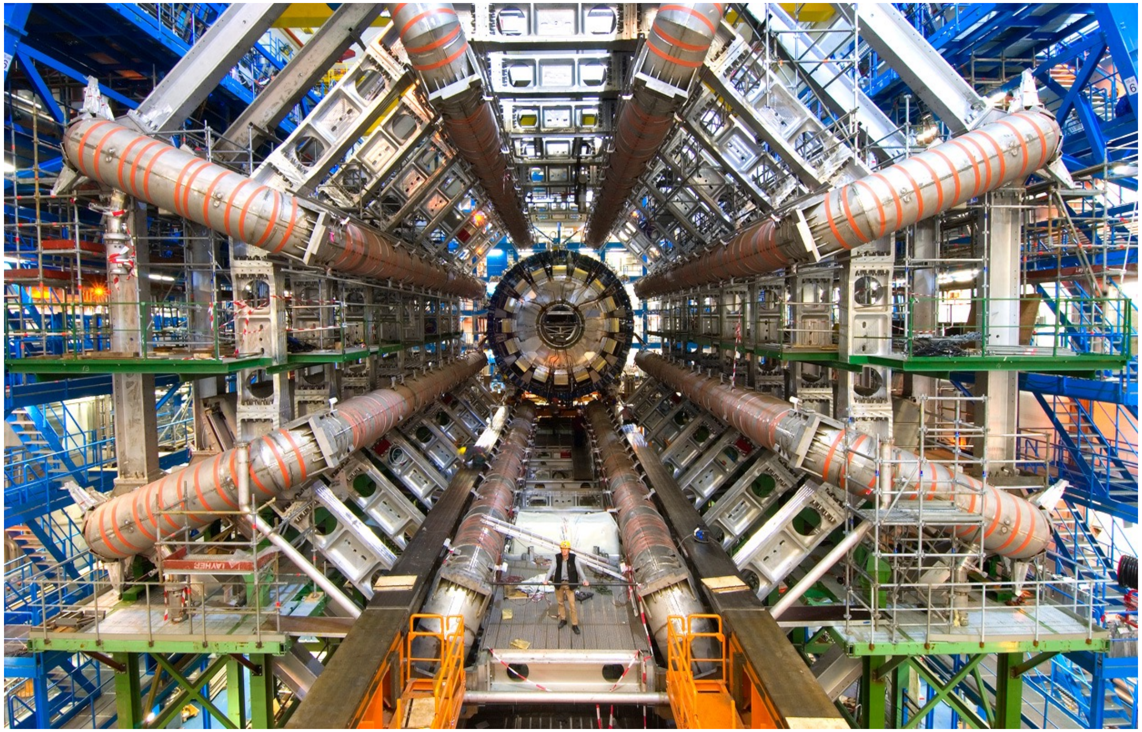

(with thanks to Mark Slater and others for borrowed material)

THE Large Hadron Collider at CERN, Geneva

- 100m underground … in a 27km long tunnel…
- Colliding protons at 0.99999996 of the speed of light

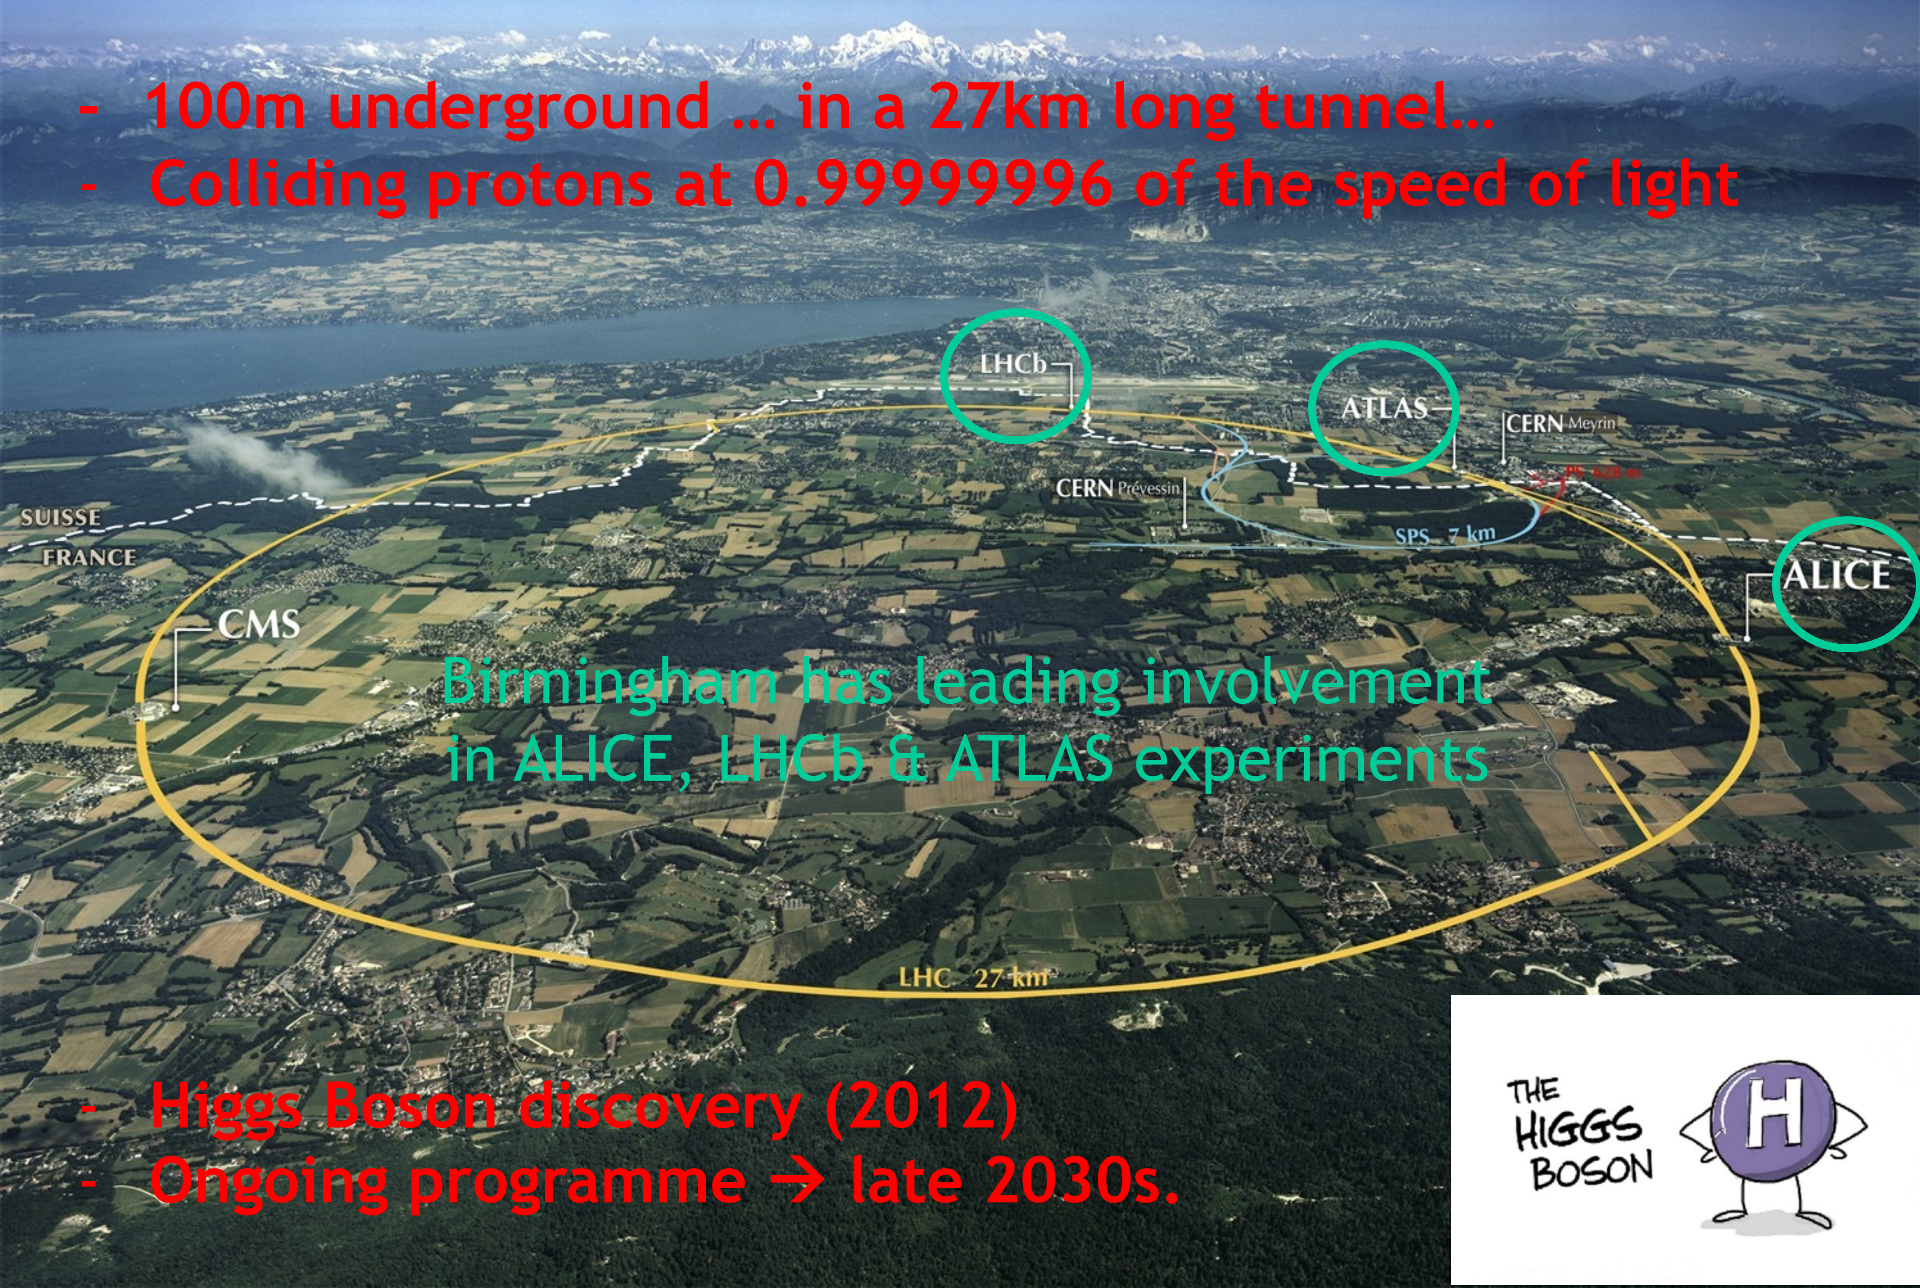

Birmingham has leading involvement in ALICE, LHCb & ATLAS experiments

- Higgs Boson discovery (2012)
- Ongoing programme → late 2030s.

Data collection at the LHC

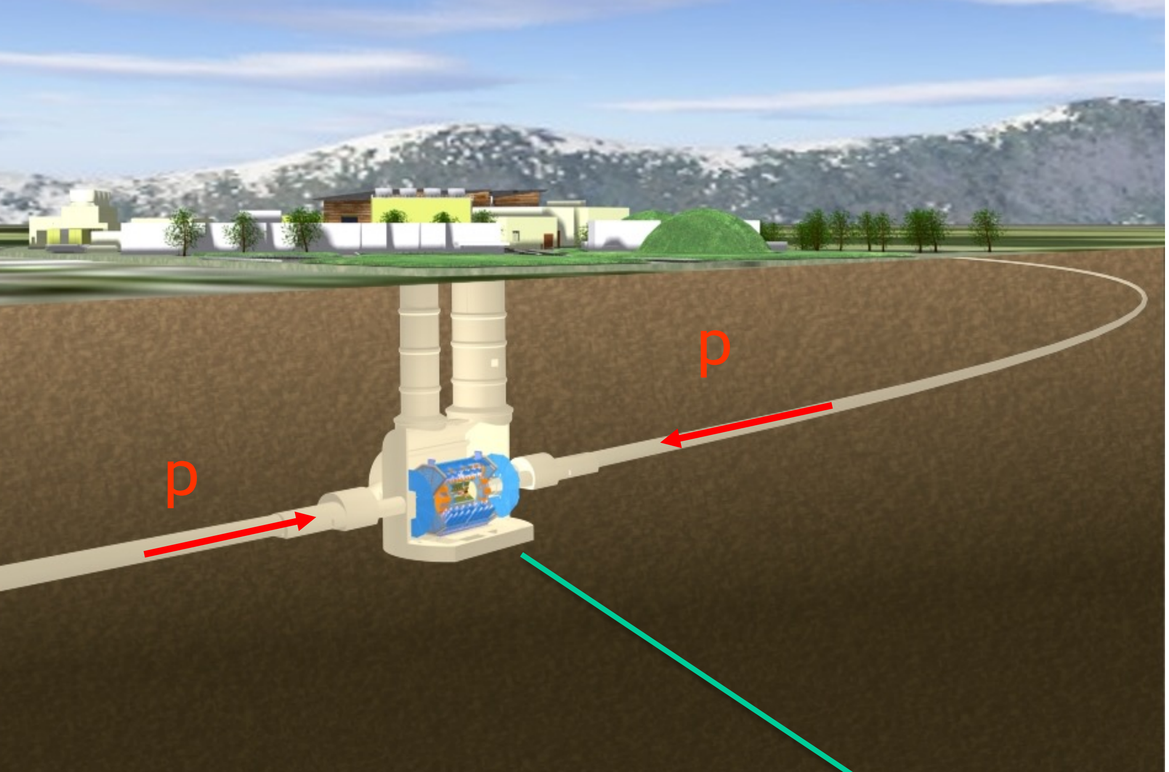

- ~10^{11} protons per bunch
- 40 million bunch crossings per second
- About 20 collisions in every bunch crossing.

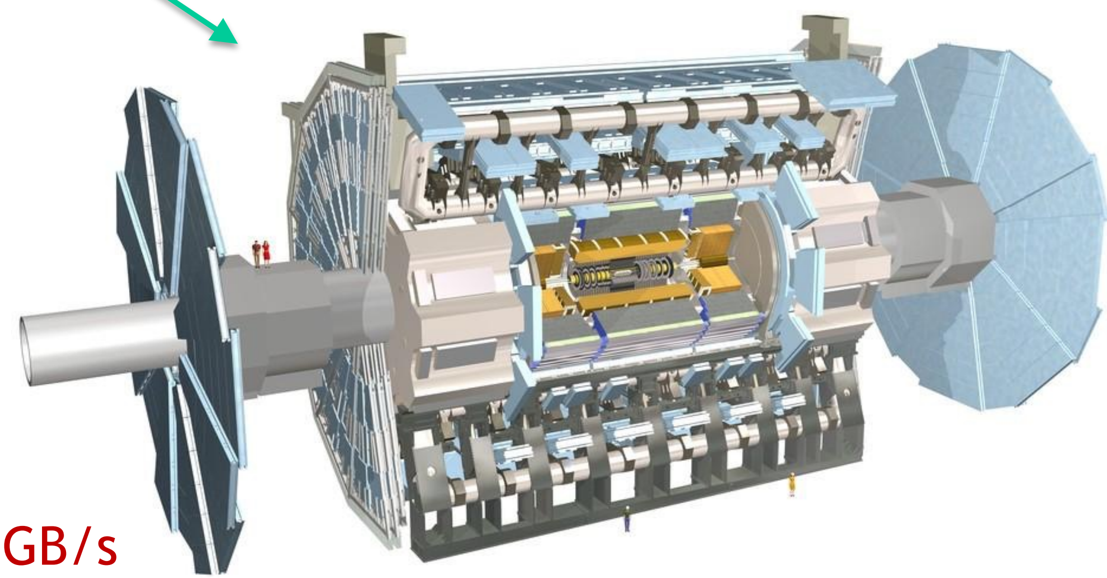

- Fast, high throughput electronics at experiments filter out most collisions
- ~1000 events per second fully read out and permanently recorded.
- Permanent storage rate ~1.5GB/s

The Data Handling Challenge(s)

Total raw data volume increasing with Time (>100PB in 2018)

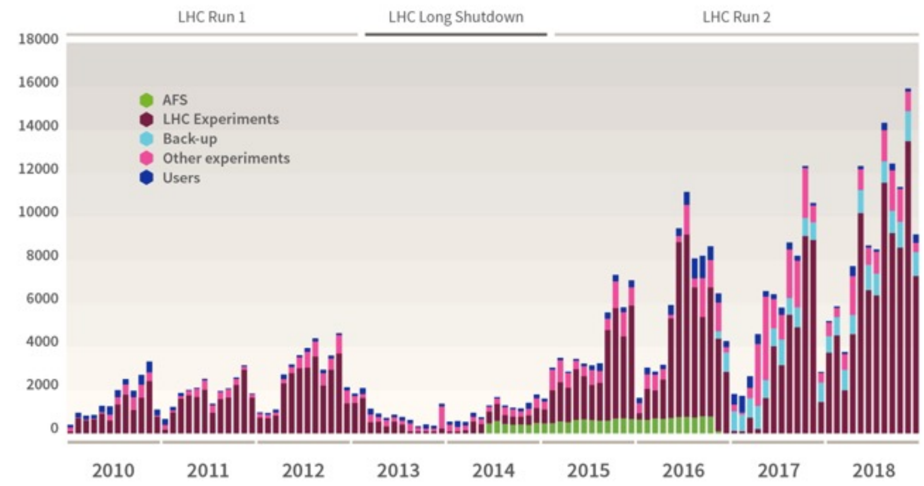

Data from experiments needs to be:

- Permanently stored
- Easily retrieved
- Processed into form for analysis
- Made available to world-wide community (~10,000 physicists) in useable format for daily use in analysis to produce physics results

In addition:

- Produce simulated data of similar scope to real data

- (At some point) make data openly available

The Solution:

WLCG
Worldwide LHC Computing Grid

- Joint project between experiments, hosted and developed at CERN

- Basic principle is to move the job to the data

- Jobs can be submitted from anywhere and may run anywhere, depending on resource availability at any given time

- All resource negotiations are centrally handled job-by-job and results sent back when completed

- [Not without glitches ☺]

WLCG – Tiered Approach to Data Storage

Basic concept is to fan data (and computing) out in successive tiers 0 → 1 → 2

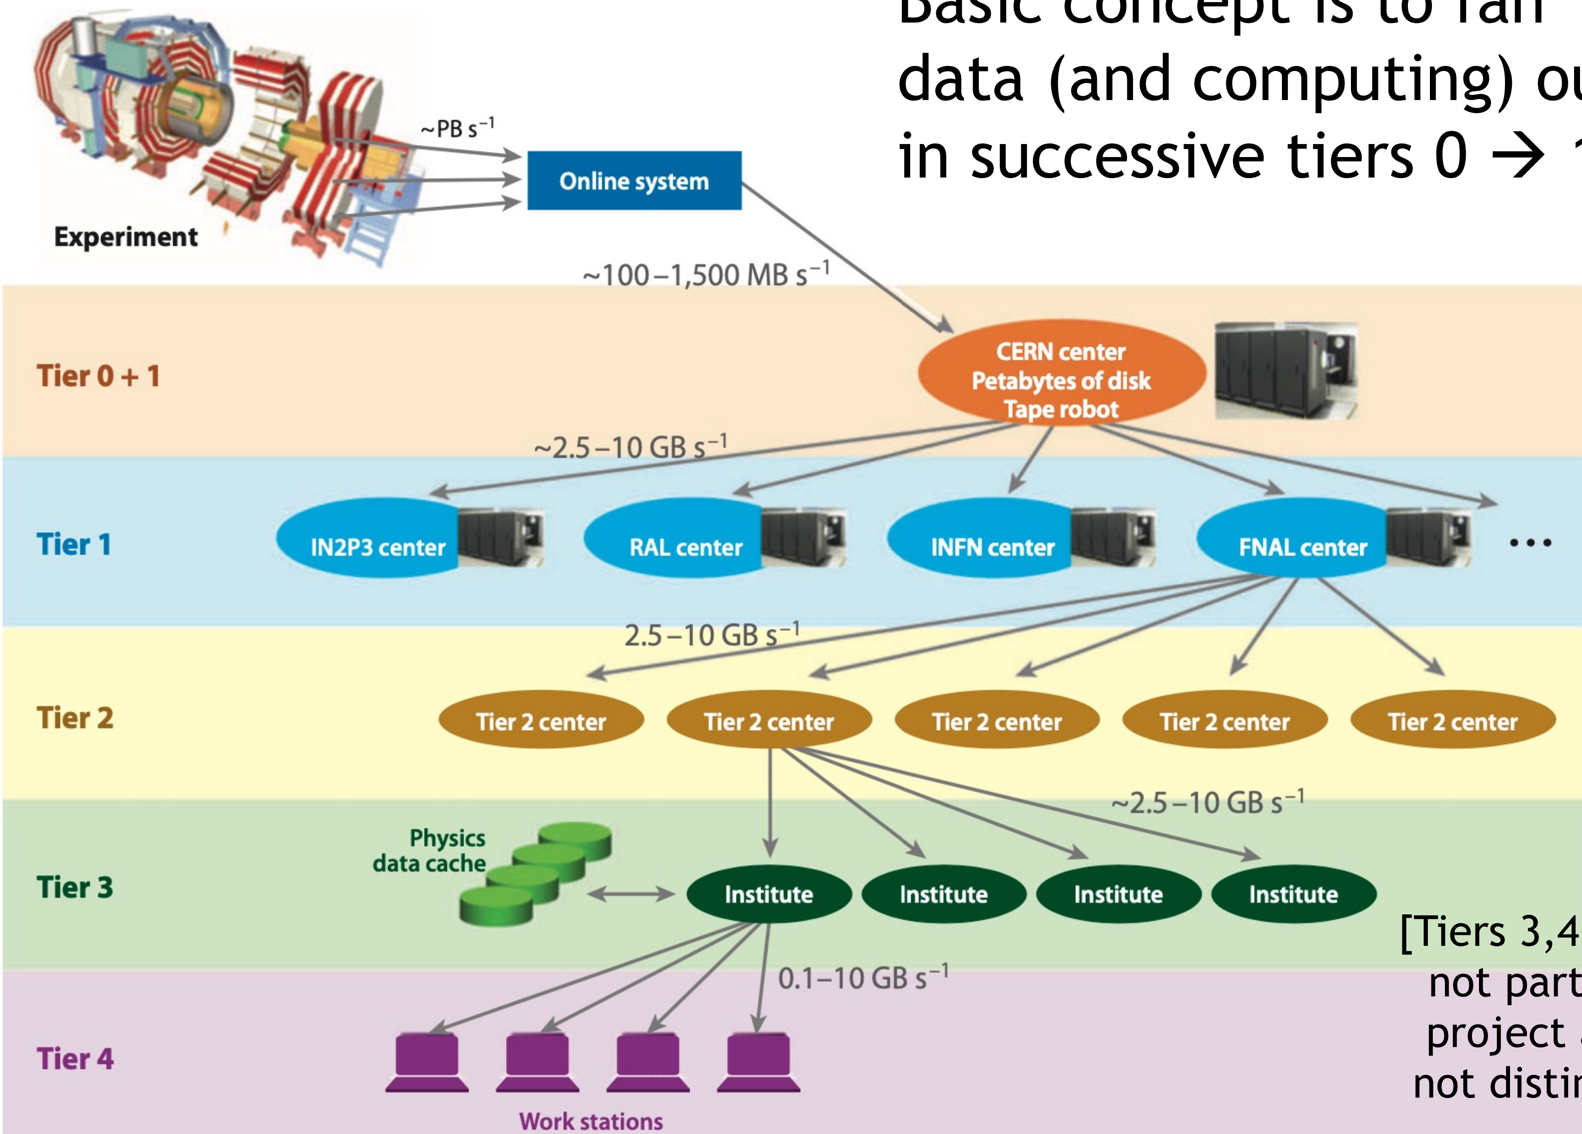

[Tiers 3,4 (end users) not part of WLCG project and often not distinguishable]

Tasks performed at Tiers

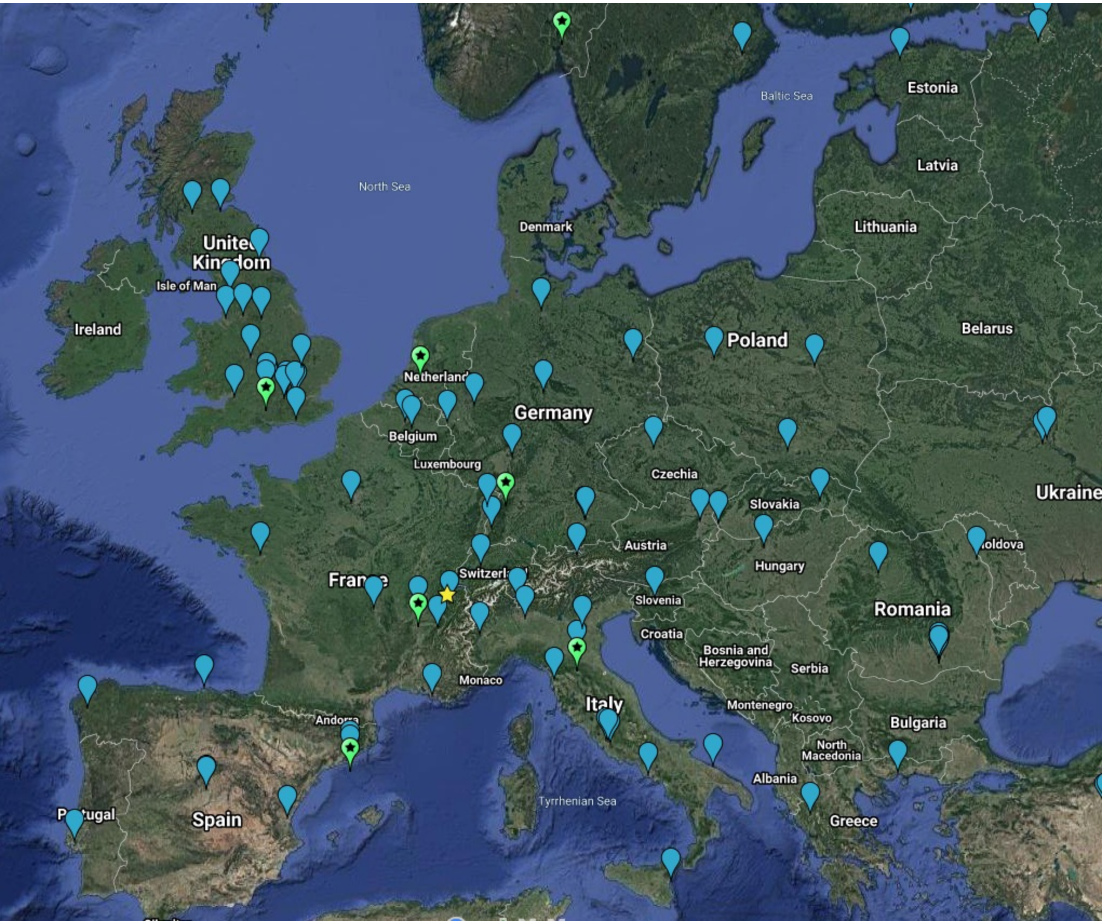

Tier 3 – Local clusters in university departments connected to tiers 0-2 and providing direct access to users.

Tier 0 (1 off - CERN)
- Gather data streams from experiments
- Long-term archiving of complete set of raw data (tape).
- First-pass processing

Tier 1 (12 off - typically 1 per country - UK tier 1 at RAL on Harwell Campus)
- Long-term archiving of a proportional share of data, amounting to 2nd complete copy
- Large-scale processing

Tier 2 (~80 off - individual institutes … eg Birmingham)
- Storing some data and providing CPU for specific analysis / simulation tasks
- No long-term archiving

WLCG Resource Snapshot

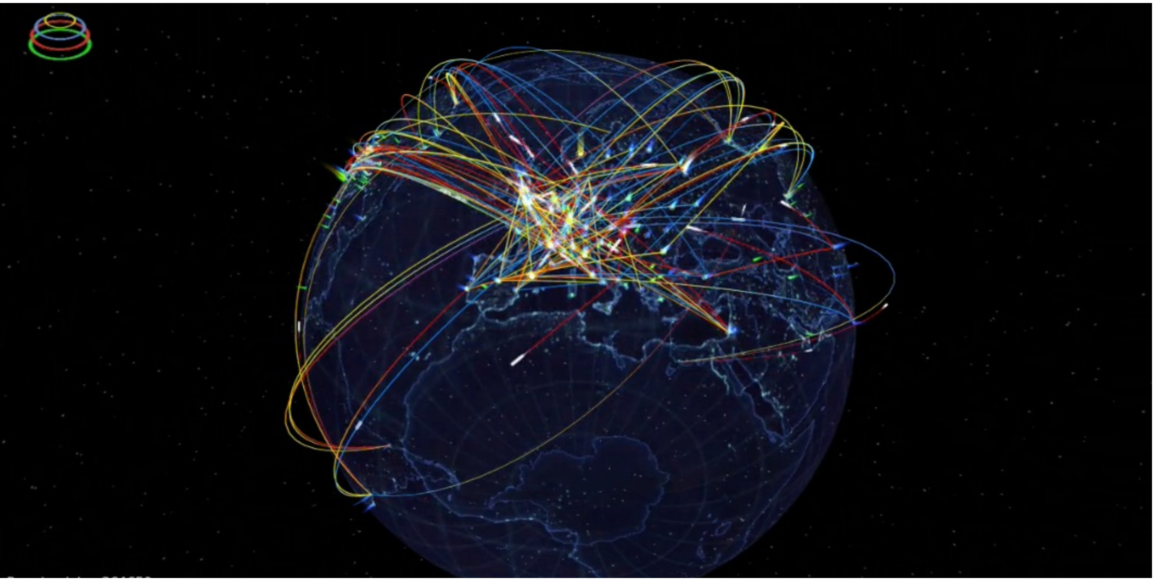

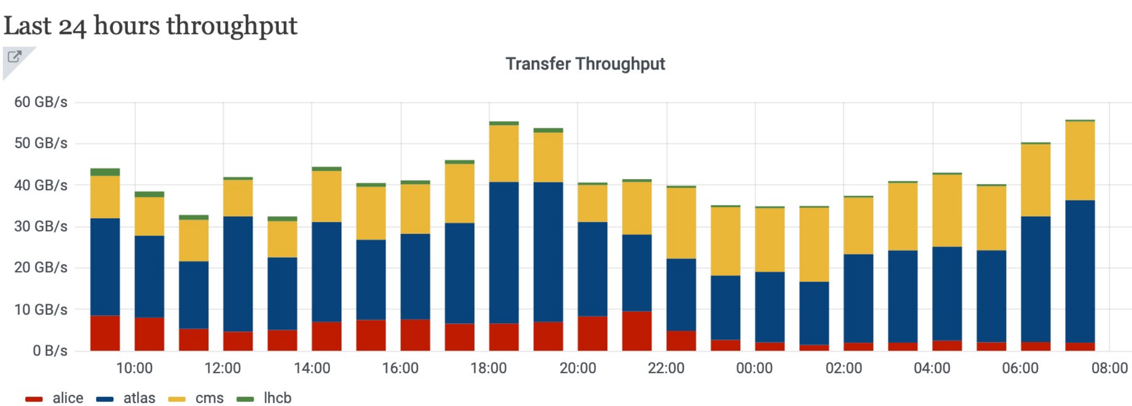

Resource requirements grow with time …

At end of 2020, tiers 0-2 contained:

- 1M logical CPU cores
- 1 Exabyte storage

CPU provision roughly equally split between Tier 1 and Tier 2 sites

Roughly speaking …

- 350k jobs running simulataneously

- 2 million jobs completed per day

- Current throughput ~50 GB/s (“idling”)

Project Management and Resource Allocation

- Producing 'middleware' and user interface software was a major task

- Ongoing need for support and evolution of provision

- Centrally agreed comprehensive policies ensuring privacy and security (CERN is under constant cyber-attack)

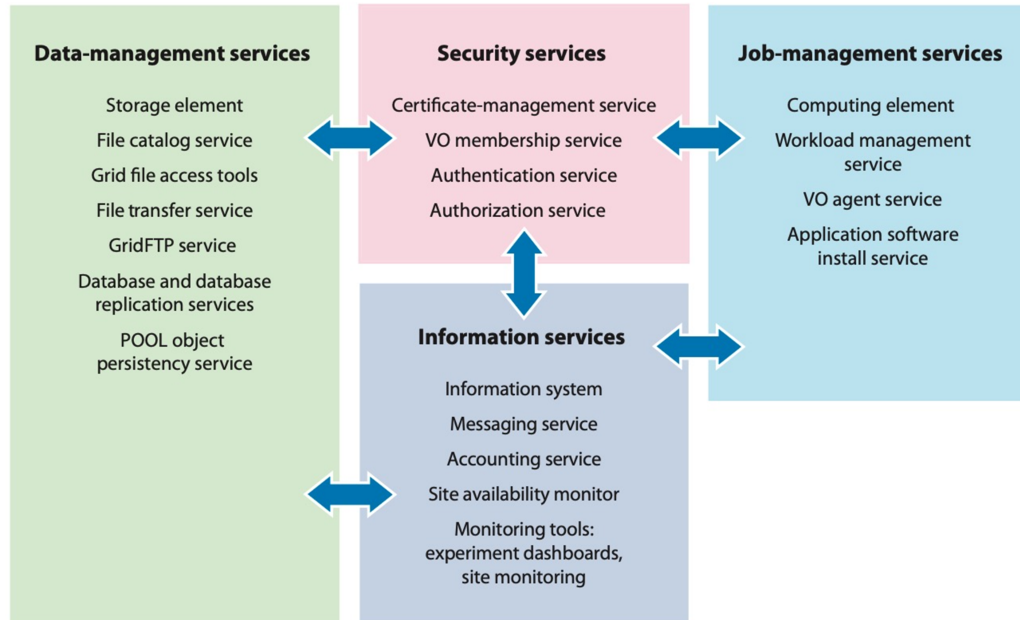

- Networks of trust for users: authentication through issuing of virtual certificates (based on employers).

- High-level agreements in place between participating countries on allocations of responsibility (fractions of CPU, storage provision)

UK Contribution: `GridPP'

- UK contributes ~10% to WLCG

- At end of 2020, this amounted to 96k cores, 65 PB disk, 67 PB tape storage

- RAL Tier 1 linked to CERN Tier 0 by 100Gb/s optical private network

- Tier 1 linked to JANET by 100Gb/s links, Tier 2 connect via 10-40Gb/s links

- Initially funded through dedicated STFC funding stream.

- Increasingly tensioned against other aspects of STFC core programme

- Streamlining resulted in restricting number of Tier 2 sites

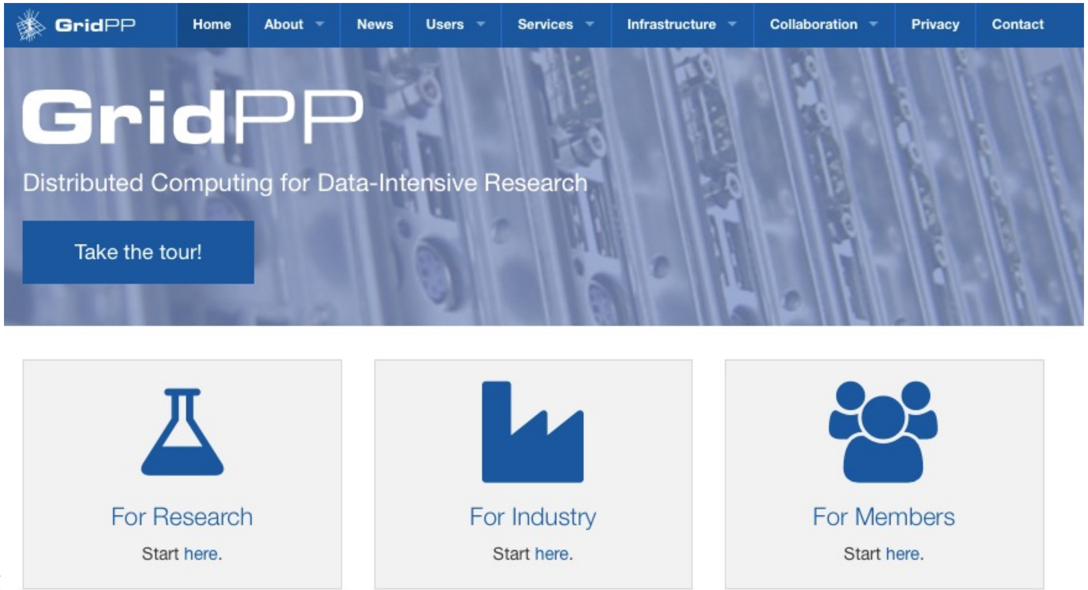

- Generating support through collaborations elsewhere in academia and with industry.

GridPP Beyond Particle Physics

User Community	Sector	Compute	Storage	CernVM	CVMFS	DIRAC	Ganga
Galactic Dynamics (GalDyn)	Astrophysics	✔	✖	✔	✔	✔	✖
Large Synoptic Survey Telescope	Astrophysics	✔	✔	✖	✖	✔	✔
LUCID	Space	✔	✔	✔	✔	✔	✖
PRaVDA	Healthcare	✔	✔	✖	✔	✔	✔
HTC for Biology	Biology	✔	✔	✖	✖	✖	✖
MoEDAL	Physics	✔	✔	✔	✔	✔	✔

LUCID – Surrey Satellite Technology Limited (SSTL)

The Langton Ultimate Cosmic ray Intensity Detector (LUCID) was designed by students at the Langton Star Centre and built by engineers at SSTL. Launched on Tuesday the 8th of July 2014 aboard SSTL's TechDemoSat-1, LUCID uses five Timepix detectors to measure properties of the space radiation environment in Low Earth Orbit (LEO). GridPP is working with the Langton Star Centre, CERN@school, and SSTL to process, store, and share the data with schools and scientists across the country. Read more about how LUCID uses the Grid in the case study here.

Econophysica

Econophysica is a small company working with researchers at Queen Mary,University of London looking at mathematical models for commodity trading.

IMENSE (formerly Cambridge Ontology)

Formed in 2004 by researchers at Cambridge University, Cambridge Ontology developed a text-driven, content-based image retrieval technology. The technology can automatically categorise all elements of an image without human tagging or other metadata.

Dell

Oxford and Dell-UK worked with Dell more closely in leading Grid applications and to help build up the Dell-UK research programme.

HP Labs

The SouthGrid consortium had a joint development project in collaboration with HP Labs, Bristol. HP is a major player in Grid development and standardisation, and the first industrial partner in the CERN wLCG project. A researcher at the University of Bristol, was funded jointly by HP and GridPP, worked to attach SouthGrid hardware resources, including those of HP Labs, to the UK particle physics Grid. HP is also a leading vendor of 64-bit computing platforms, and will provide valuable assistance to GridPP and LCG in the porting of both physics applications and Grid middleware in order to take full advantage of increasingly cost-effective 64-bit processors.

Projects with Commercial Partners

The LHC, Data Preservation and Open Data

Level 1 Data (final published analysis results)

- Published with open-access under CERN SCOAP agreement with journals
- Accessible through recognised repositories (e.g. Durham HEPDATA)

Level 2 Data (For outreach / education / non-LHC science)

- Subsets of data e.g. for `Masterclass' outreach to schools
- No preservation requirements.
- Open by construction

Level 3 Data (Reconstructed data typically used in analysis)

- Selected data preserved for re-analysis, re-use and verification
- Made public eventually (individual experiment policies vary)

Level 4 Data (Un-reconstructed raw data produced by experiments)

- All other data in principle can be re-derived from raw data (CPU-heavy)
- 2 copies preserved indefinitely
- Not practical (or useful) to make publicly accessible

Monte Carlo Data (Simulations matching real data at all 4 levels)

- Some long-term preservation alongside real data

Data Preservation and Public Access

- Level 3 data availability: agreed policy

  "The LHC experiments will release calibrated reconstructed data with the level of detail useful for algorithmic, performance and physics studies. The release of these data will be accompanied by provenance metadata, and by a concurrent release of appropriate simulated data samples, software, reproducible example analysis workflows, and documentation. Virtual computing environments that are compatible with the data and software will be made available. The information provided will be sufficient to allow high-quality analysis of the data including, where practical, application of the main correction factors and corresponding systematic uncertainties related to calibrations, detector reconstruction and identification. A limited level of support for users of the Level 3 Open Data will be provided on a best- effort basis by the collaborations."

- Available through CERN Open Data Portal: http://opendata.cern.ch

- Proprietary period varies between experiments (50% by 5 years and 100% by 10 years or at the end of the lifetime of the collaboration)

Open Data Portal

- Primarily used in education

- Some take-up by interested amateurs hoping to solve the mysteries of the universe

Summary

- Particle Physics coped with explosion in data volumes and CPU needs for shared use by large numbers of geographically diverse collaborators through world-wide distributed computing network

- Works (usually) well and scalable

- Based on ideas / technology of >10 years ago ... could certainly be improved if started again from scratch.

- Open data policies include open-access for all publications of LHC results and (eventual) public release of all data / associated analysis frameworks in format ready for analysis