Analysis of LHC data using globally distributed computing resources

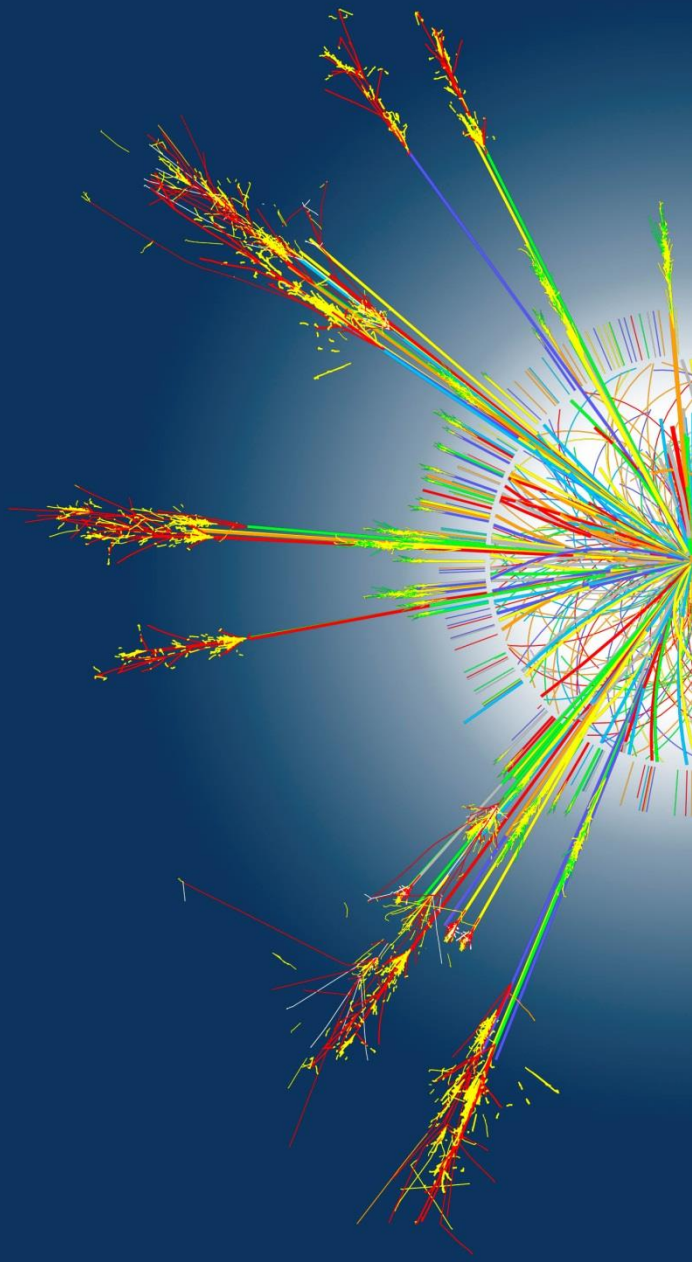

Oxana Smirnova
Div. of Particle Physics
Lund University

3 May 2017

CERN physicists don't know everything

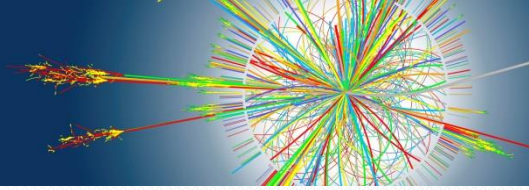

- What are **dark matter** and **dark energy** that pervade our Galaxy?
- What is **gravity**, really?
- Why is the electrical charge of the proton **equal** and opposite to that on the electron?
- Why are protons **stable**, or are they, really?
- How many space-time **dimensions** do we live in?
- Are elementary particles really elementary, or do they have **structure**?
- Why are there **3** generations of elementary particles? Why not 4 or 5?
- Where did the **anti-matter** go?
- Why are the **neutrinos** so light, and which is the heaviest?
- Are there **more states of matter** beyond solid, liquid, gas and plasma?

LHC holds answers

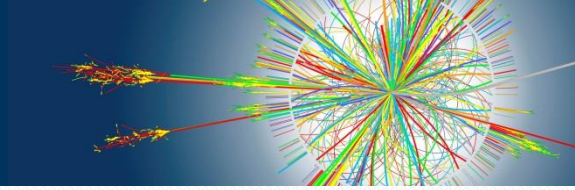

LHC is the biggest ever machine and data generator

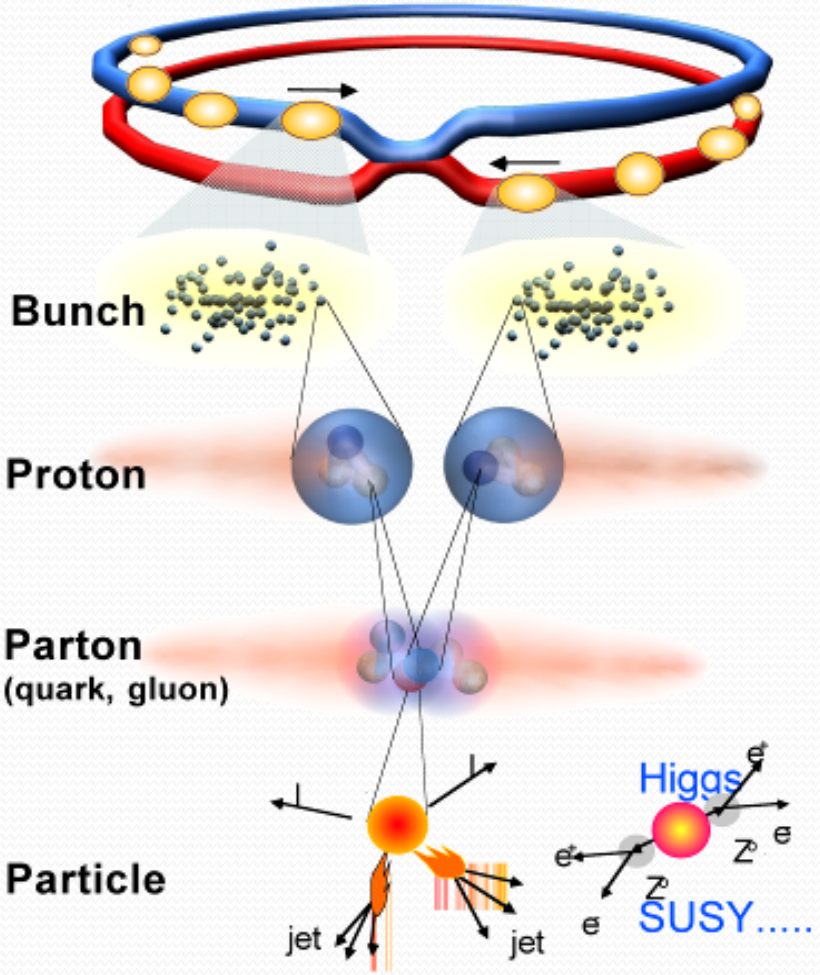

Proton-Proton	2835 bunch/beam
Protons/bunch	10^{11}
Beam energy	7 TeV (7×10^{12} eV)
Luminosity	10^{34} cm^{-2} s^{-1}
Crossing rate	40 MHz
Collisions rate ≈	10^{7} - 10^{9} Hz

There are also Pb-Pb and p-Pb runs

New physics rate ≈ .00001 Hz

Event selection:
1 in 10,000,000,000,000

Graphics by CERN

A data sample: collision event at LHC

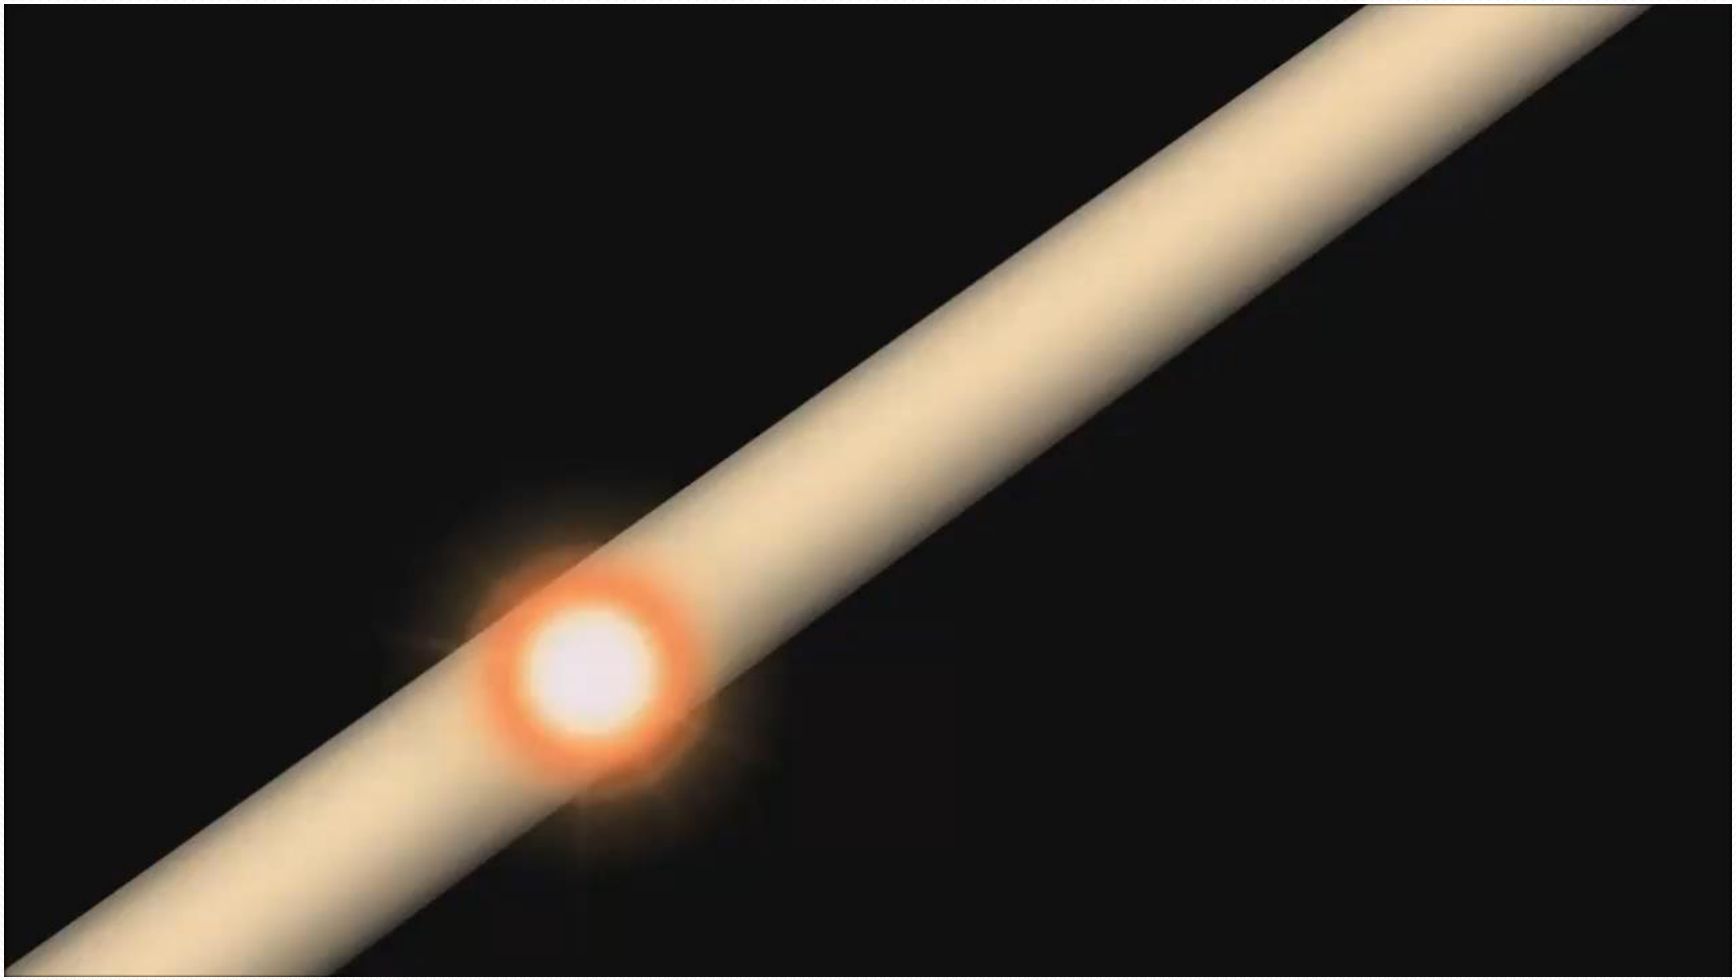

ATLAS video

An instrument at LHC: ATLAS

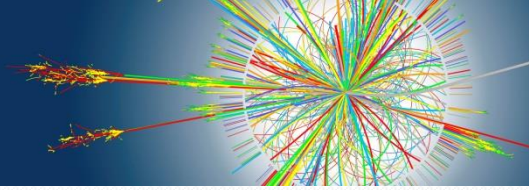


Kryogenisk fabrik för tillverkning av Helium II till dipolen


Tryckkärl i den kryo-geniska fabriken
Foto: CERN


Provlinje för dipolrör
Foto: CERN


Två dipolrör kopplas ihop


Dipolrörets ände. Vakuum-kärlet ytterst, sedan det heliumfyllda kärlet. Den massiva biten i mitten är magneten, med de två strålrören sida vid sida.

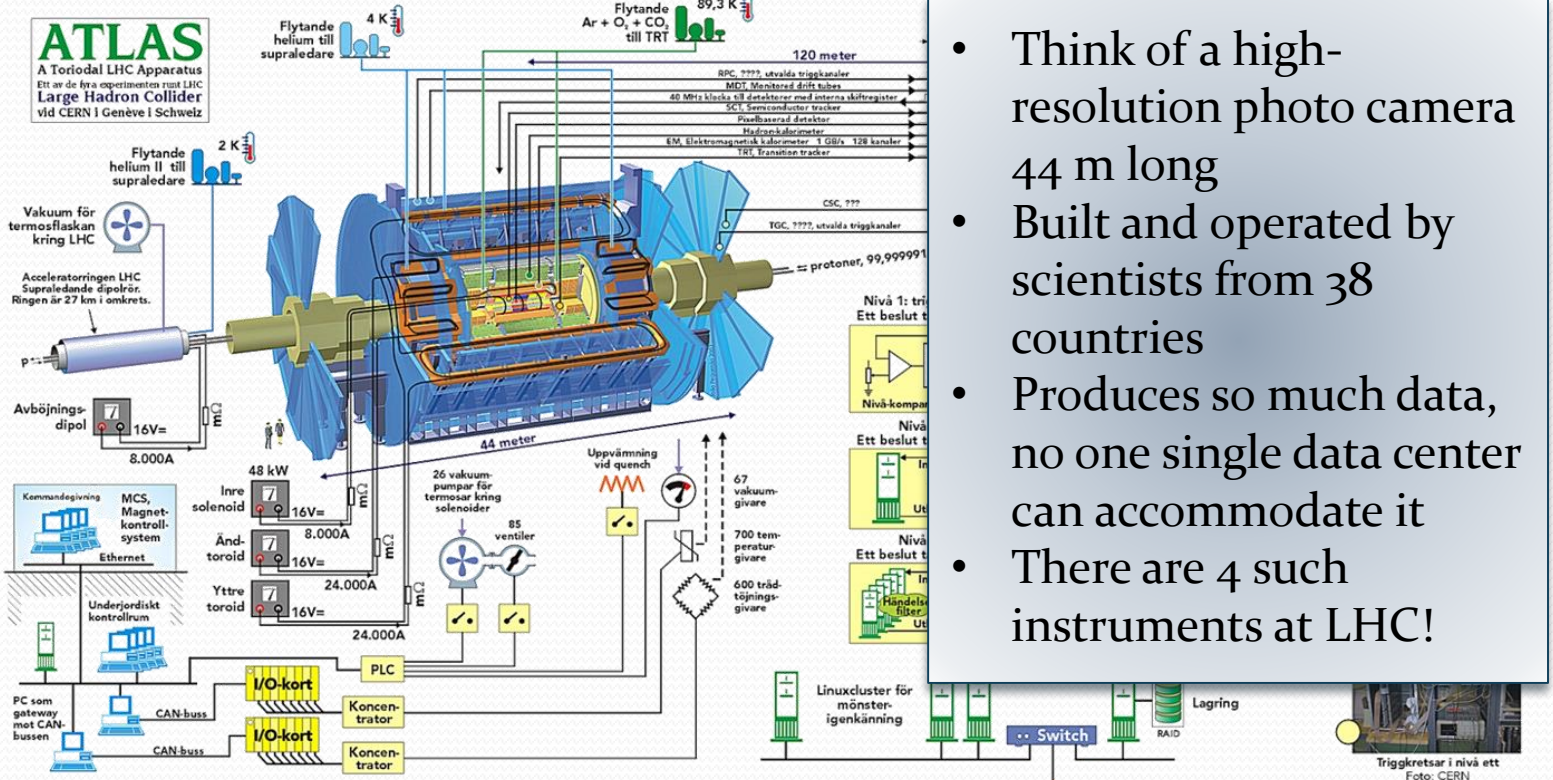

- Think of a high-resolution photo camera 44 m long
- Built and operated by scientists from 38 countries
- Produces so much data, no one single data center can accommodate it
- There are 4 such instruments at LHC!

Graphics by QED AB

LHC produces Real Big Data

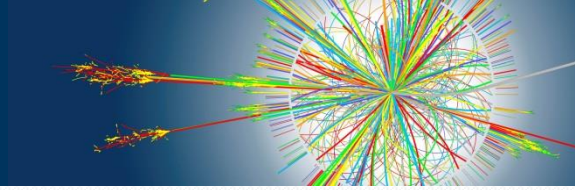

- Number of read-out channels: ~15 million
- Event rate: 40 MHz
- At 1 byte per channel per event, we'd need to write at the speed of ~**600 TB/s**
  - And process it all…
  - Luckily, not all channels get fired at once
- **But** most events are "normal" and thus <u>not interesting</u>
  - Much like mowing a field hoping to get an orchid in a haystack of daisies
  - E.g., the Higgs production rate is only ~0.1Hz, detection rate is even less
- Most of our data is "background noise"!
  - This is true for other sciences, and is characteristic of **Big Data**

Do we need to store all this data?

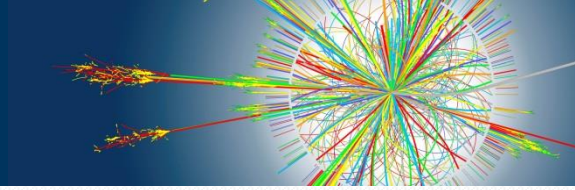

- Luckily, "interesting" events have some obvious signatures
  - For example, particles produced with high energy at large angles w.r.t. the beam direction indicate that something unusual might have happened
- We don't need high granularity to detect such signatures, thus we can quickly distinguish them
  - So we **trigger** data taking only if such signatures occur
  - Trigger involves software making a yes/no decision
  - To do it fast (real fast), serious hardware is needed
    - Mind it, this is needed before even recording data!

Trigger is a powerful system itself

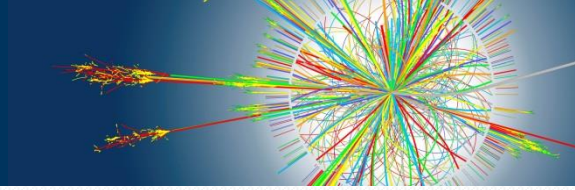

	Typical number of pieces
Detector	1
Custom links from the detector	1000
"Readout Units" for protocol adaptation	100 to 1000
Powerful Core routers	2 to 8
Edge switches	50 to 100
Servers for event filtering	> 1000

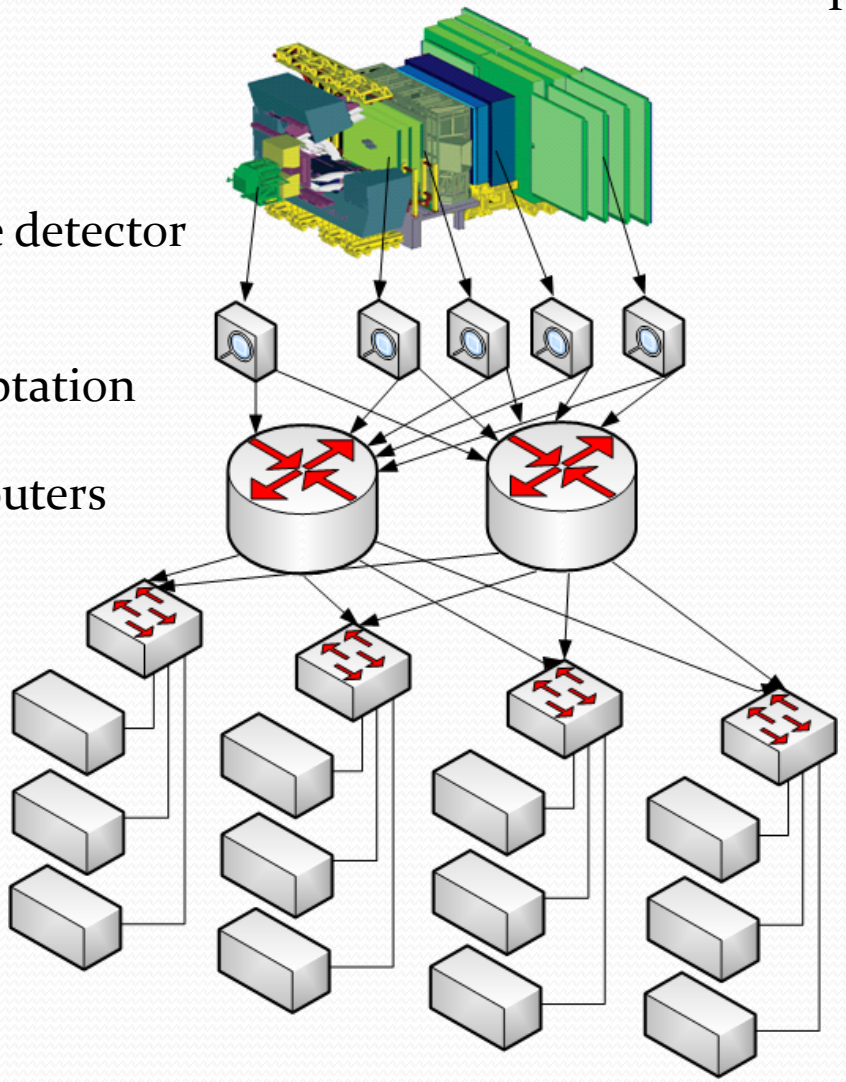

Slide from ISOTDAQ16 TDAQ for LHC - Niko Neufeld, CERN

LHC experiments and their triggers

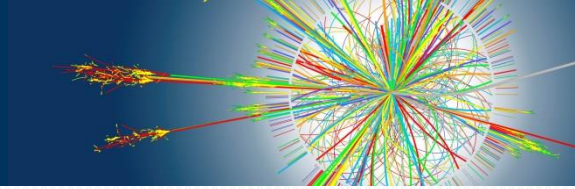

		# Trigger	Level-0,1,2 Rate (kHz)	Event Size (Byte)	Network Bandw.(GB/s)	Storage MB/s(Event/s)
ALICE	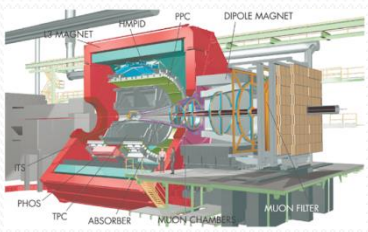	**4**	**Pb-Pb 5** **p-p 1**	**$5x10^7$** **$2x10^6$**	**25**	**4000** (10^2) **200** (10^2)
ATLAS	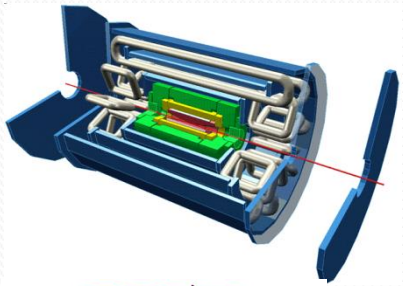	**3**	**LV-1 75** **LV-2 6**	**$1.5x10^6$**	**10**	**~1000** (>400)
CMS	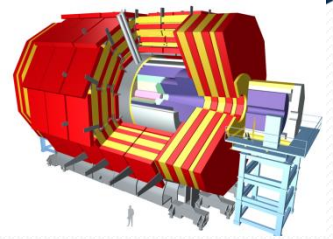	**2**	**LV-1 75**	**10^6**	**200**	**~1000** (10^3)
LHCb	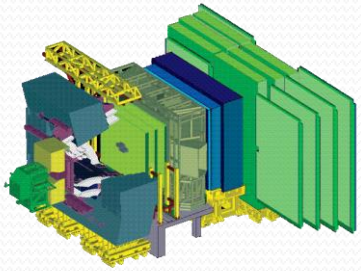	**2**	**LV-0 1000**	**$6x10^4$**	**55**	**>600** ($1.2x10^4$)

Slide from ISOTDAQ16 TDAQ for LHC - Niko Neufeld, CERN

We still have a lot of data to analyse

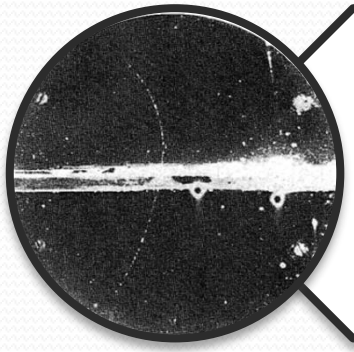

A discovery in 1930-ies

- ~2 scientists in 1country
- pen-and-paper

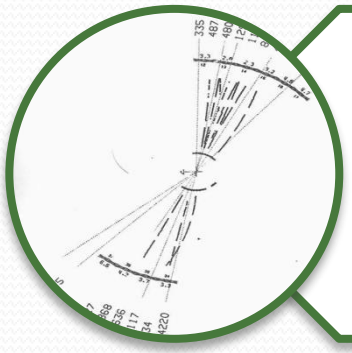

A discovery in 1970-ies

- ~200 scientists in ~10 countries
- mainframes

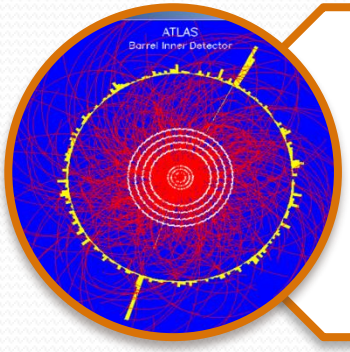

A discovery today

- ~2000 scientists in ~100 countries
- **Grids**

- Relatively simple algorithms
- Very large samples
- Analysis is easily split in very many computing jobs
- Easy to distribute data and jobs
- **Distributed computing: Grid of data centres** (as in power grid)

Some real data volumes

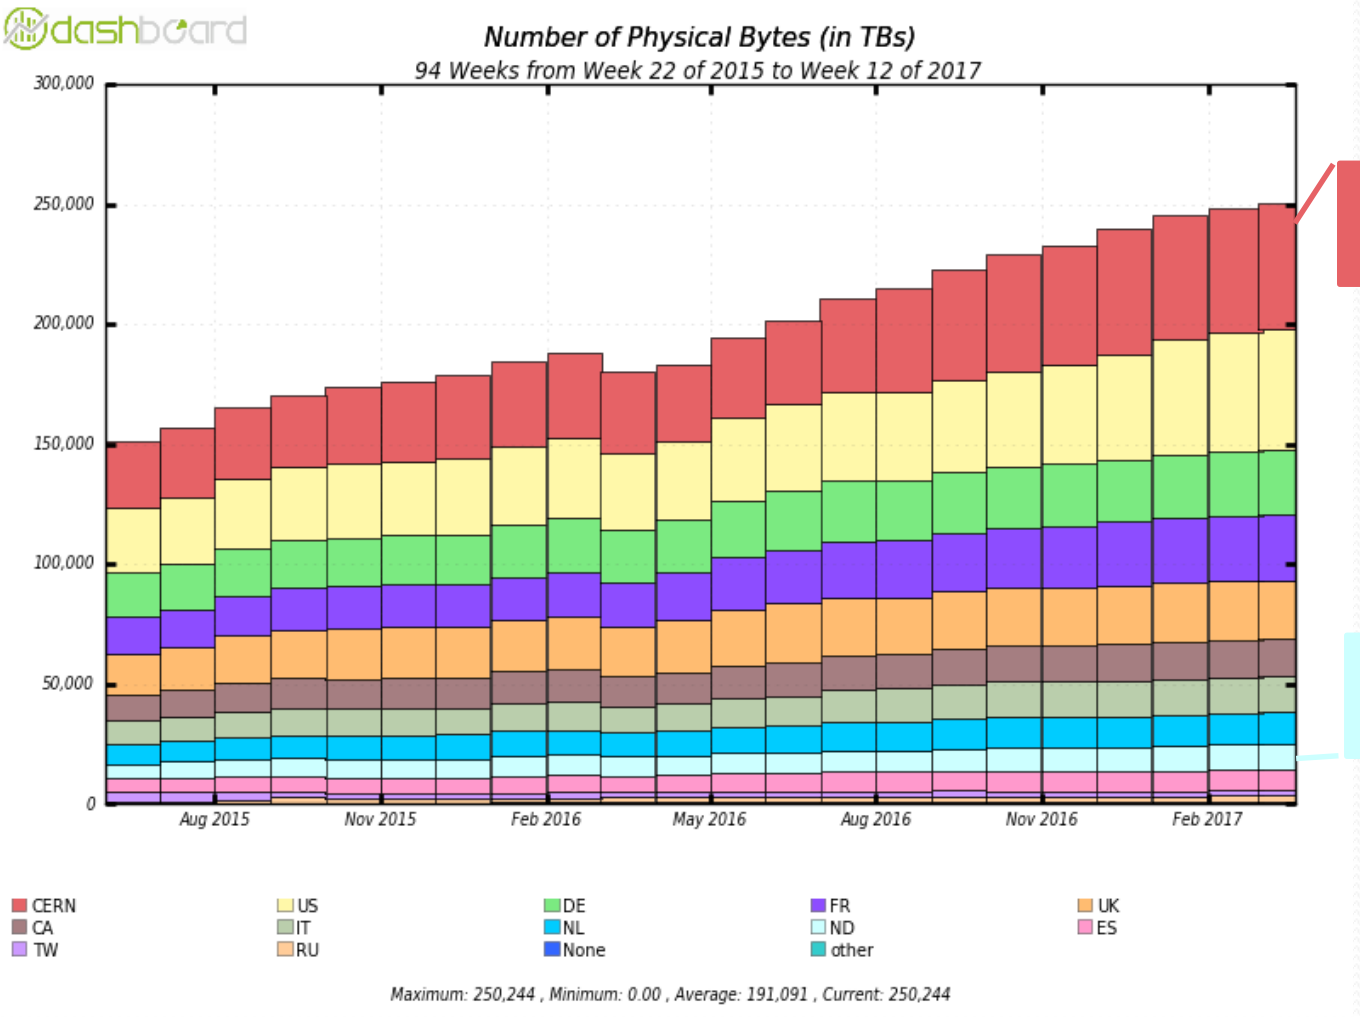

CERN

Nordic share
(3 countries)

Data stored by the ATLAS experiment alone so far, by region

(comparable with annual Facebook uploads)

Stored and analysed around the World

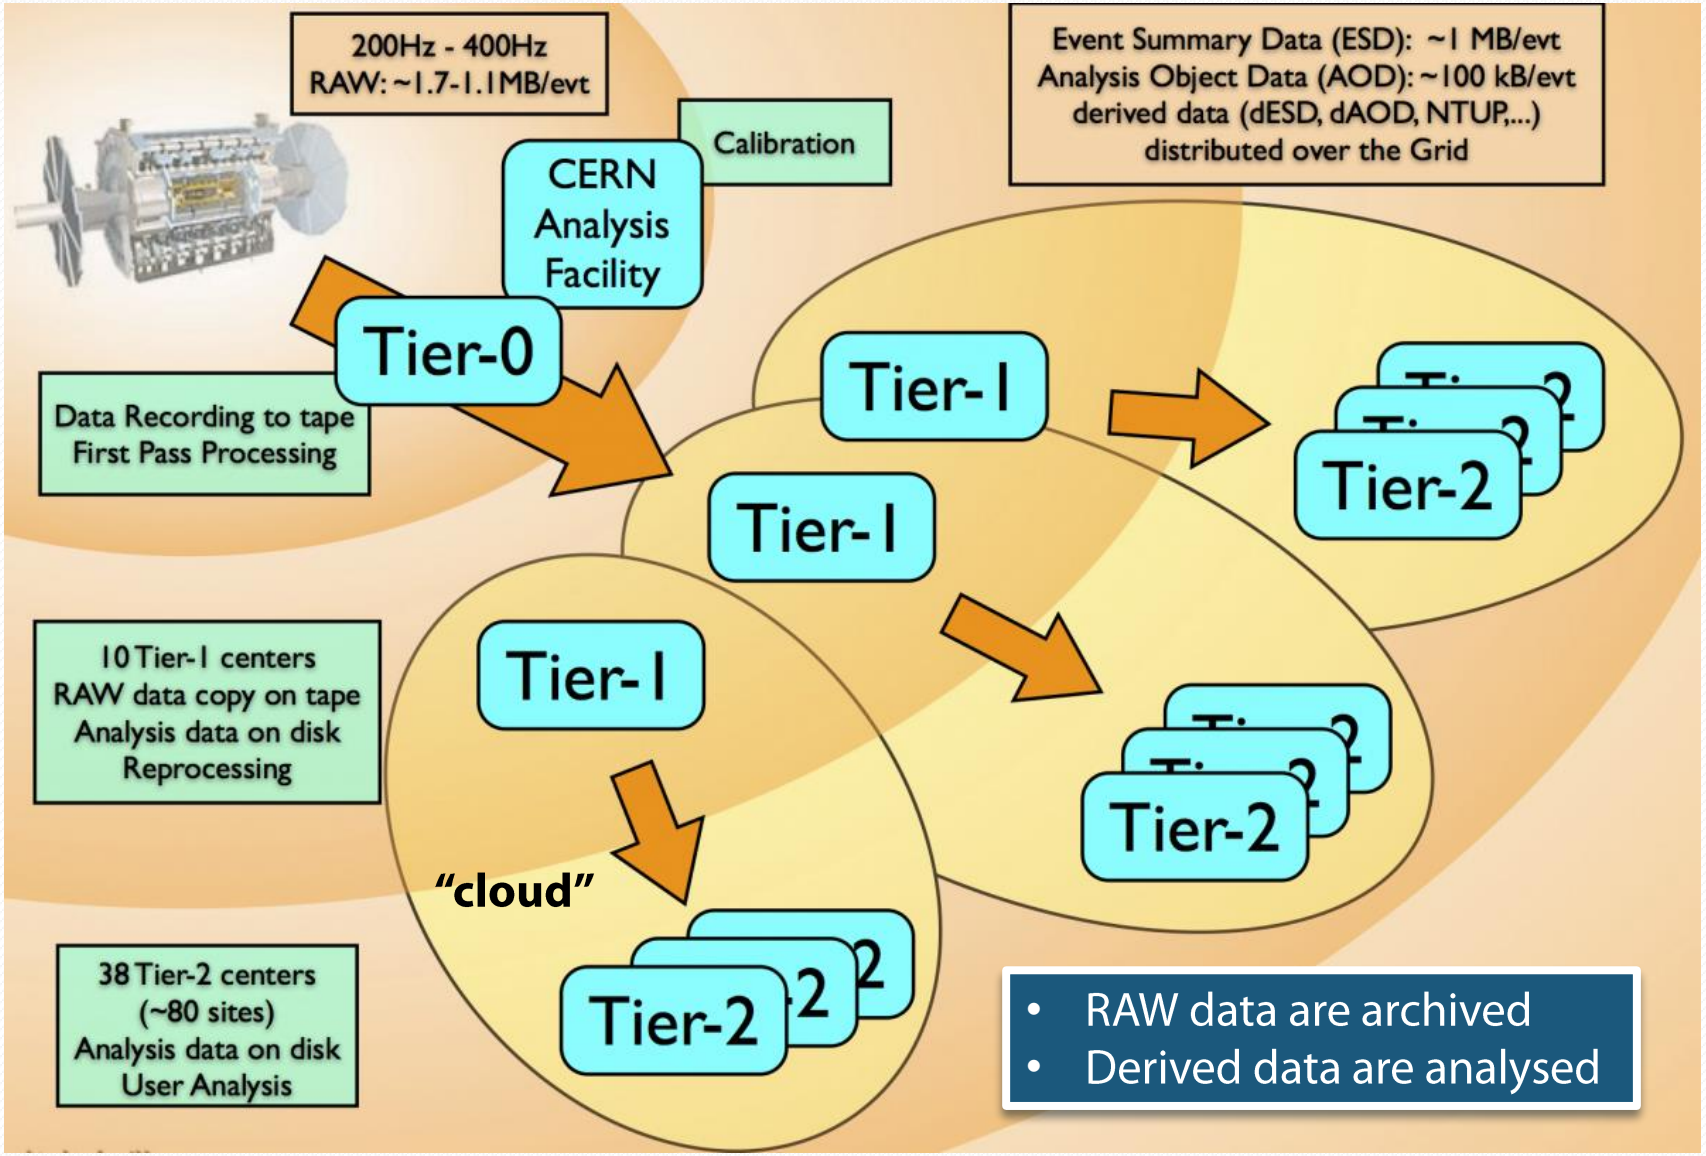

Graphics by iSGTW

LHC Computing Grid sites

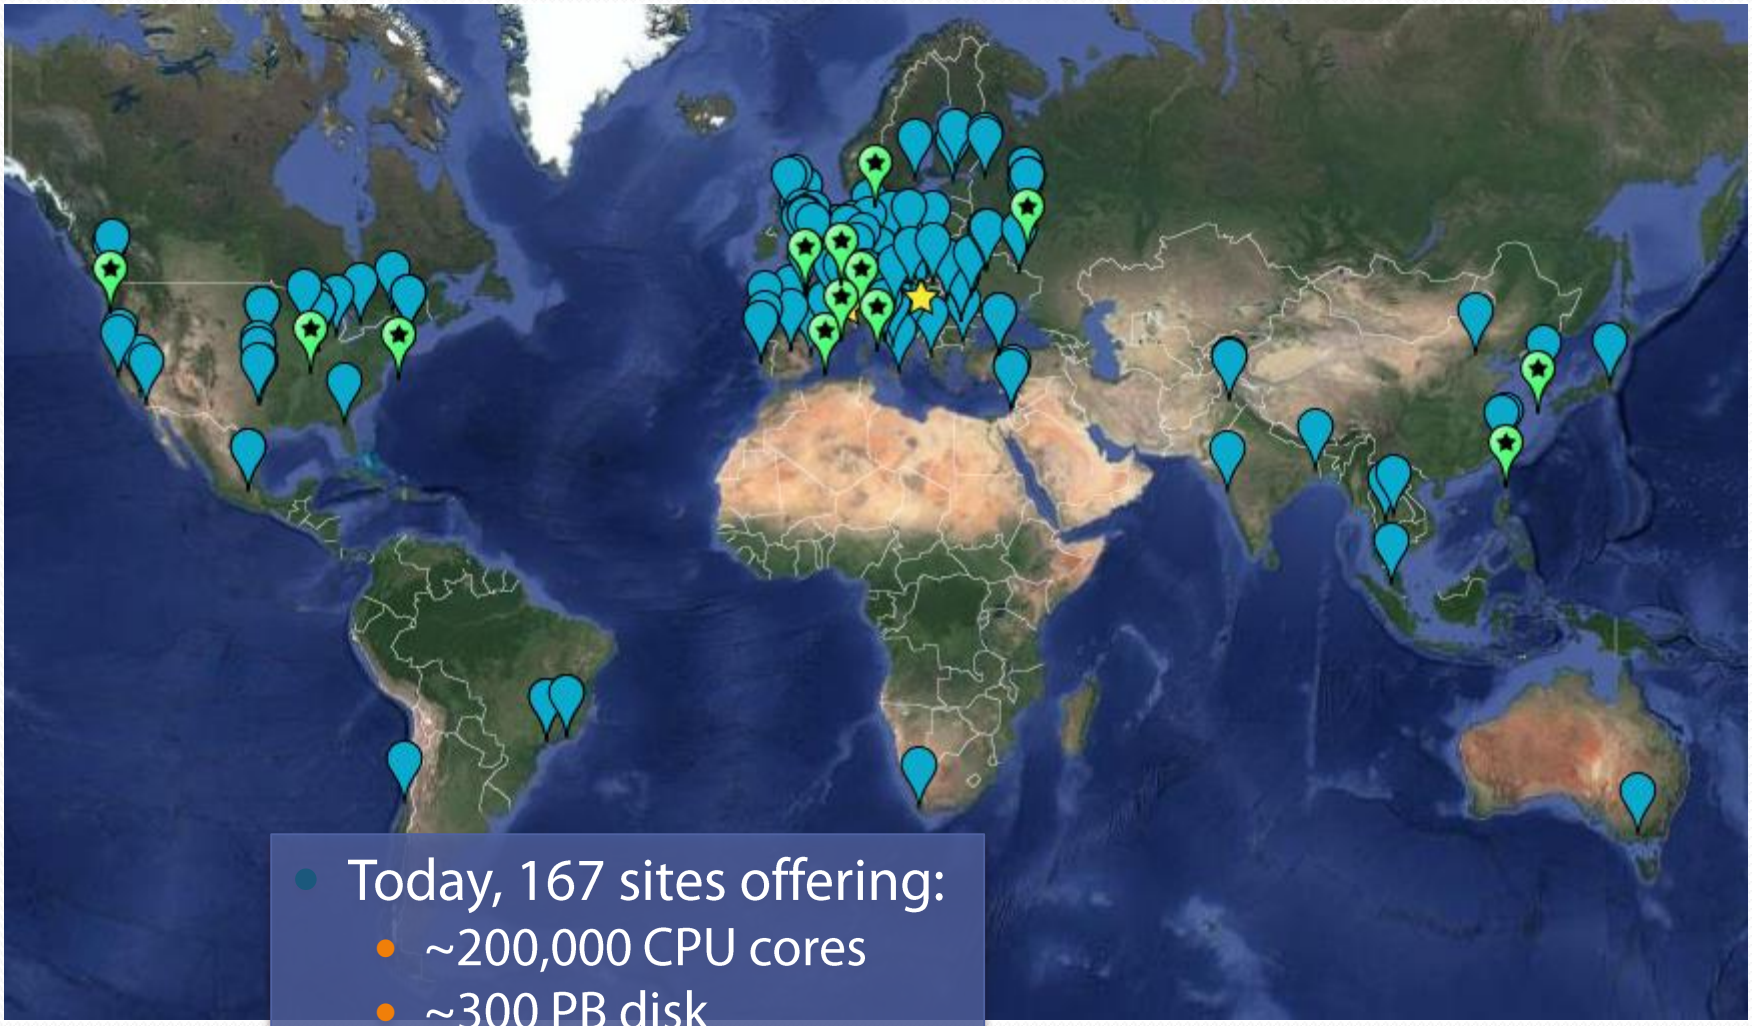

- Today, 167 sites offering:
  - ~200,000 CPU cores
  - ~300 PB disk
  - ~380 PB tape

For comparison: Hazel Hen

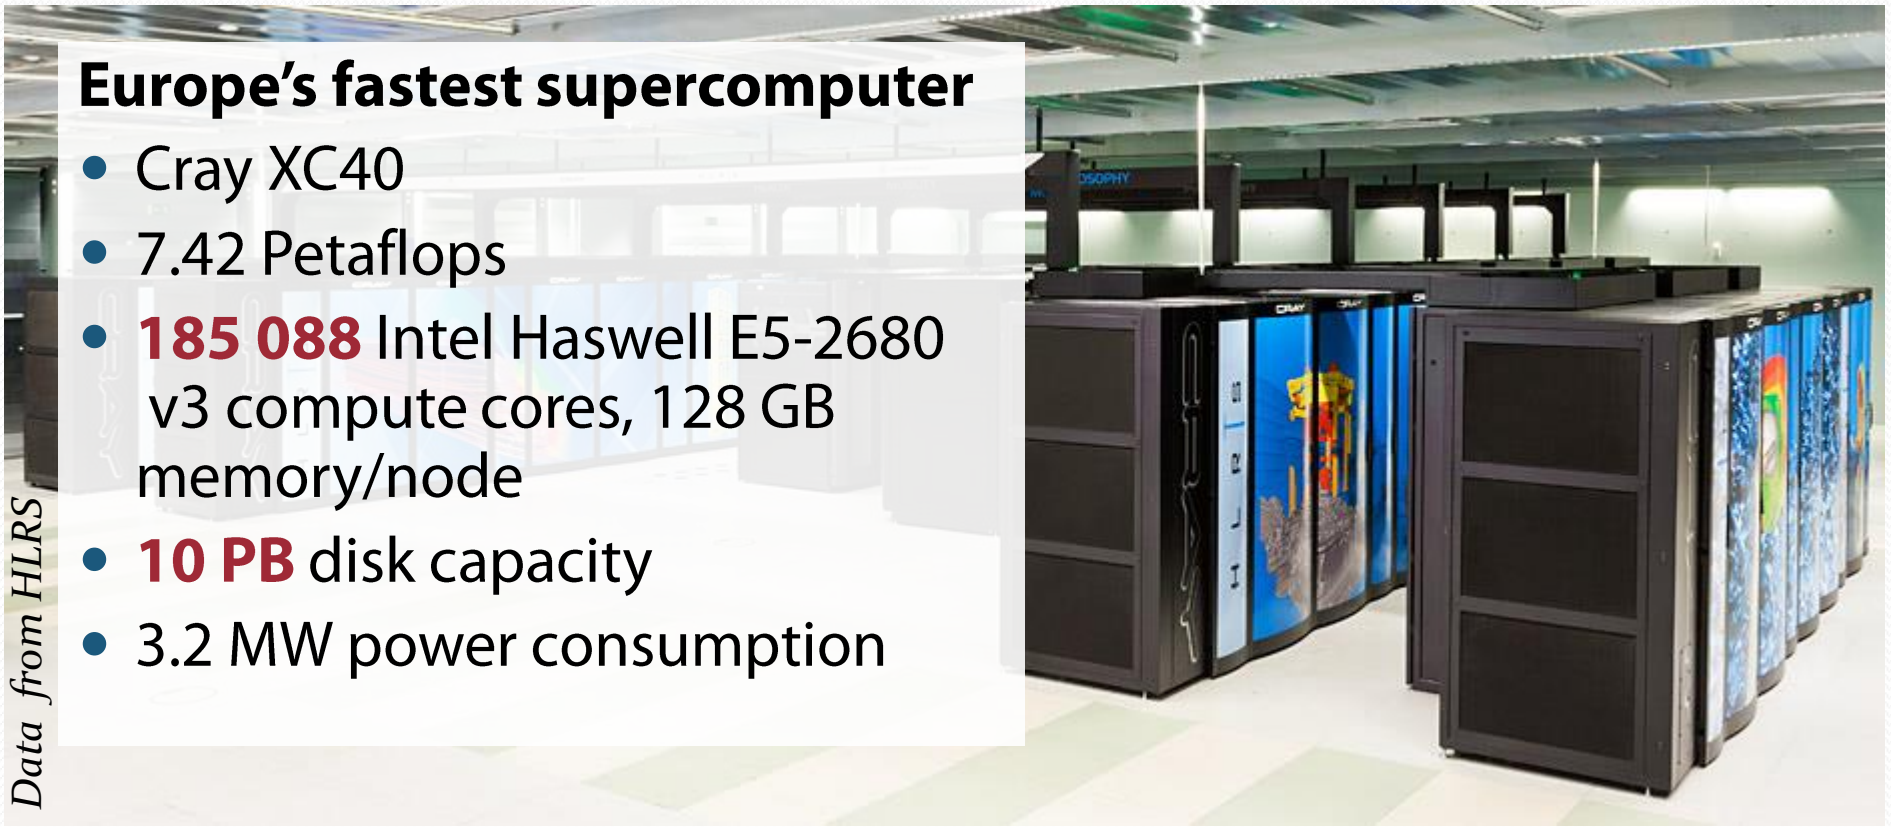

Europe's fastest supercomputer

- Cray XC40
- 7.42 Petaflops
- **185 088** Intel Haswell E5-2680 v3 compute cores, 128 GB memory/node
- **10 PB** disk capacity
- 3.2 MW power consumption

Data from HLRS

Supercomputers: focus on data generation,
less on data processing

We develop Grid solutions ourselves

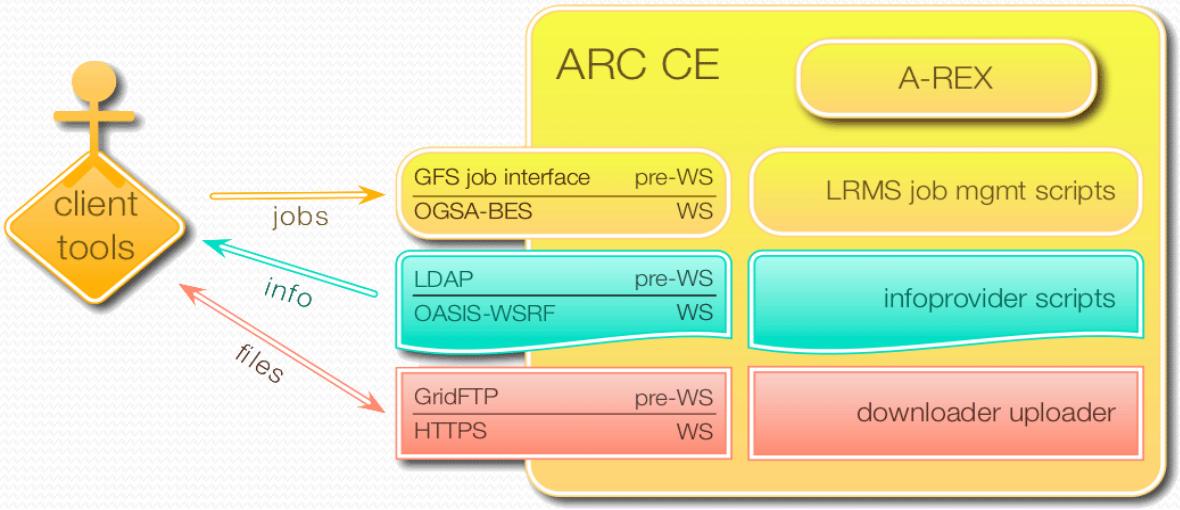

- Big businesses are all doing Clouds, so we do things ourselves
- ARC– the key component of the Nordic Grid infrastructure
  - File handling on behalf of the user
  - Universal front-end for different batch systems
  - Status information publishing
  - Standard and custom interfaces
  - Developed in Nordic countries (in Lund, too), used by ~20% of LHC Grid
  - Is used by many other scientists (not just LHC)

Nordic Grid infrastructure: NDGF-T1

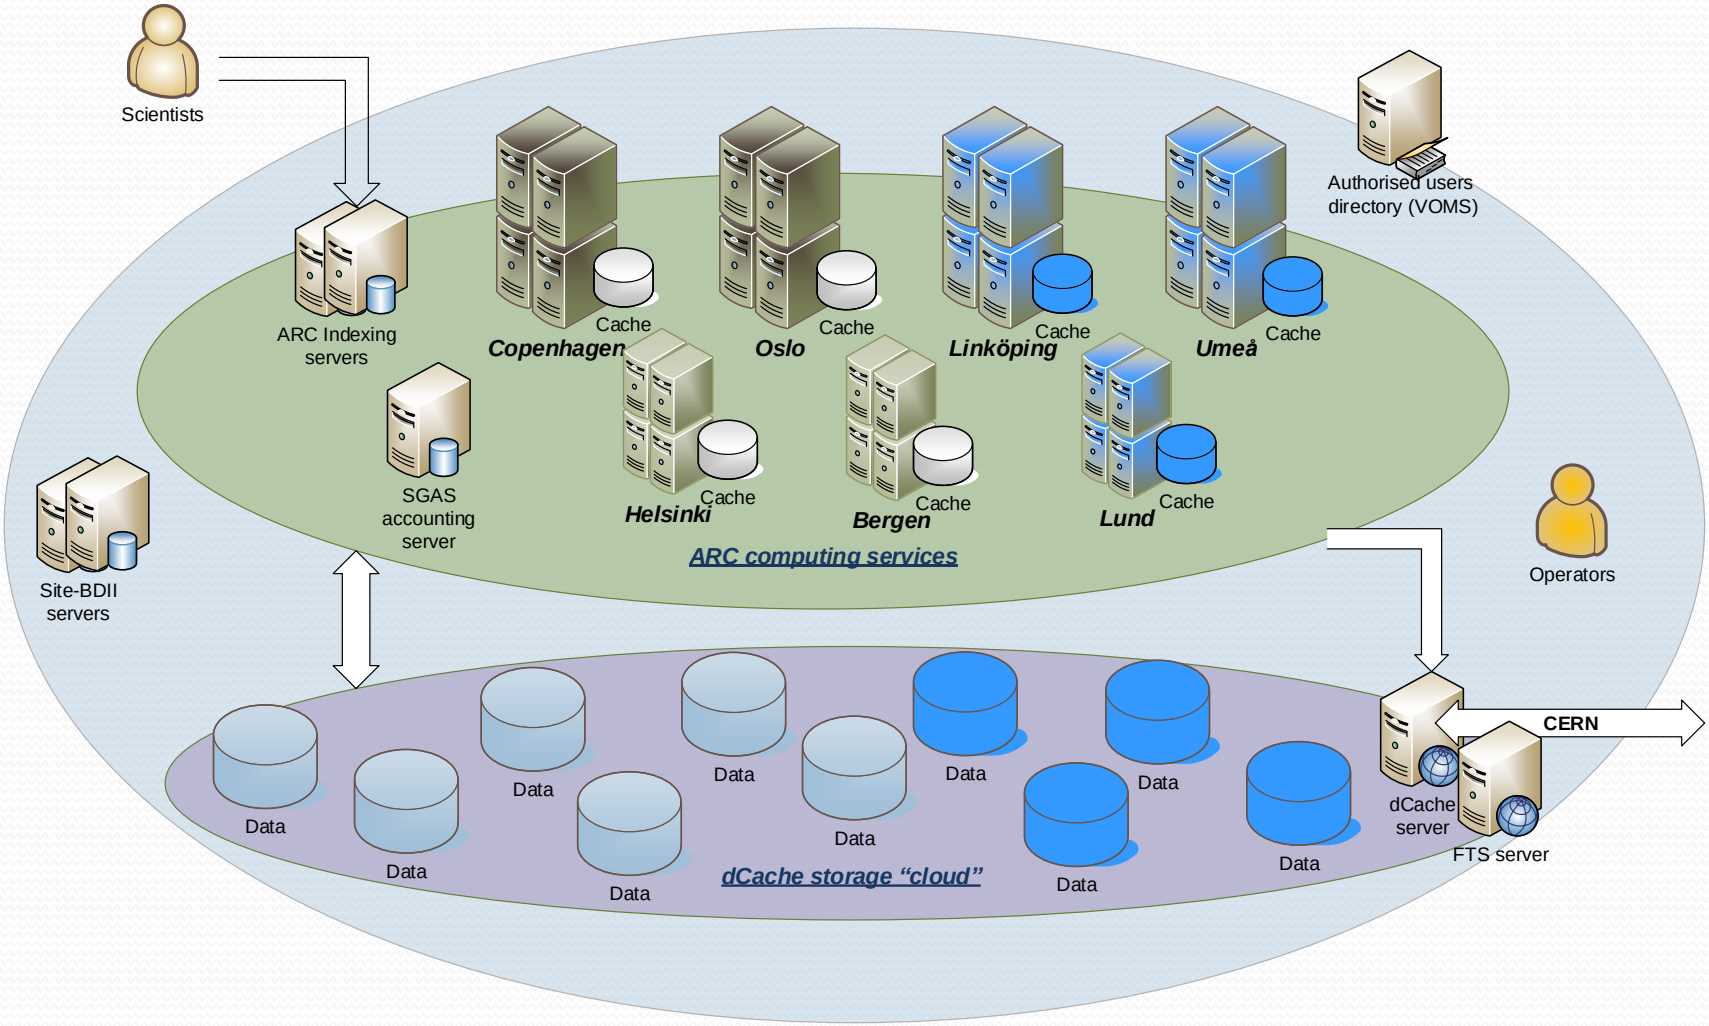

How is Grid different from Clouds?

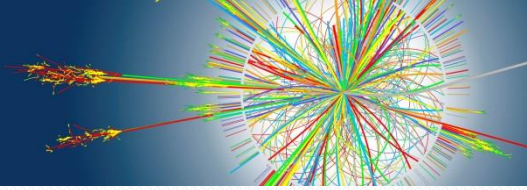

- Technical view:
  - In Clouds, environment is tailored for your tasks
  - In Grids, your tasks have to be tweaked for different environments
- The economy, stupid:
  - Clouds are capitalism
  - Grids are communism
- The reality:
  - We plug Cloud resources into our Grid

Summary

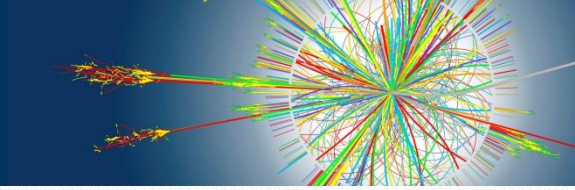

- All sciences face rapid increase in digital data volumes
- LHC developed a working solution for very large volumes of data
  - Do not try to store and process data in one place: **share and distribute** instead
  - Use own open source software
  - This allows LHC to achieve scientific results almost instantaneously, despite huge data sets
- LHC solutions can be generalized to other data-intensive sciences and even industry