Laboratoire d'Annecy de Physique des Particules

Computing, Data Analysis and Tools for HEP and Astro/Cosmo

Dominique Boutigny

June 3rd, 2026

Les Rencontres de Noirmoutier 2026

UNIVERSITÉ SAVOIE MONT BLANC

45 years of High Energy Physics

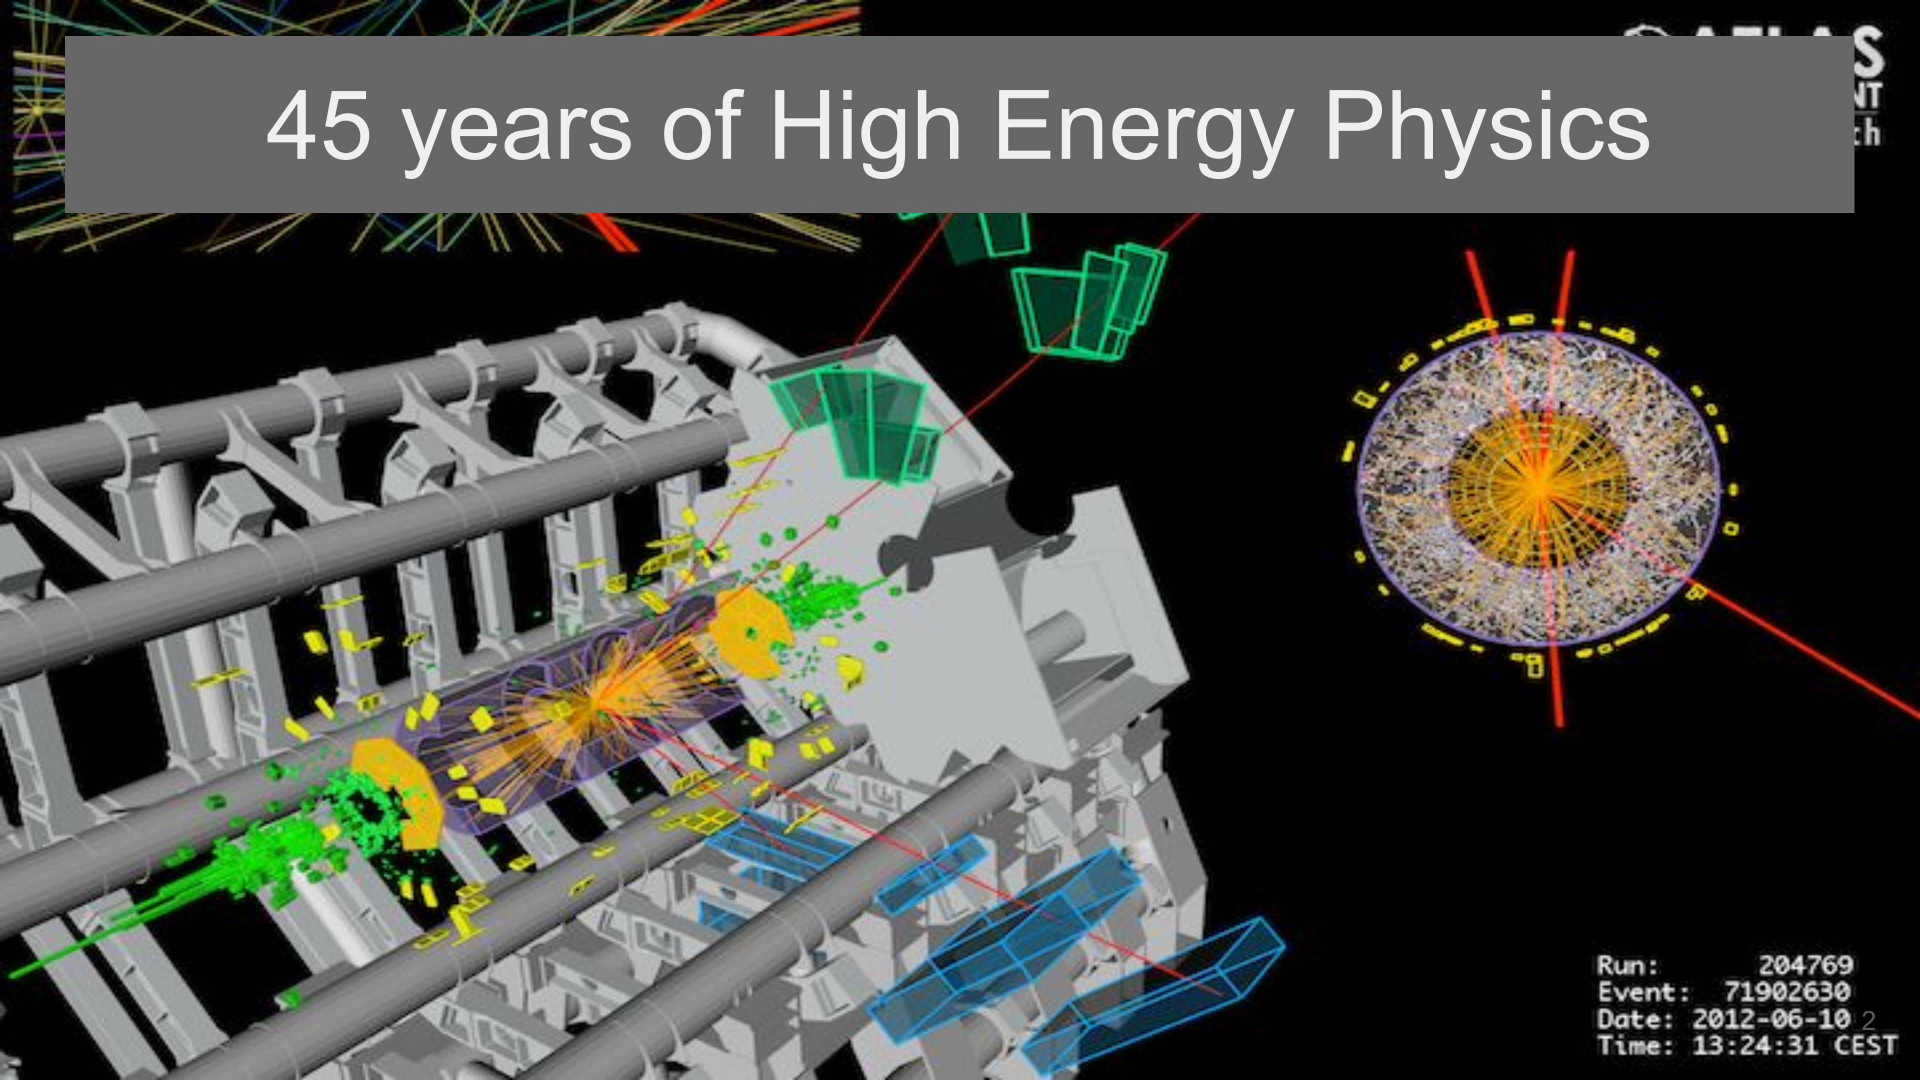

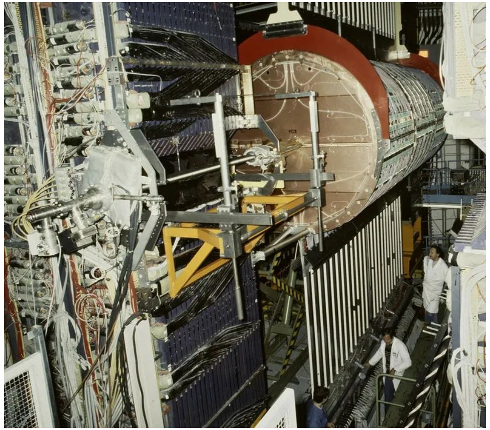

SppS - UA1 / UA2
W and Z discovery

1981 - 1990

LEP - ALEPH, DELPHI L3, OPAL - Z and W factory

1989 - 2000

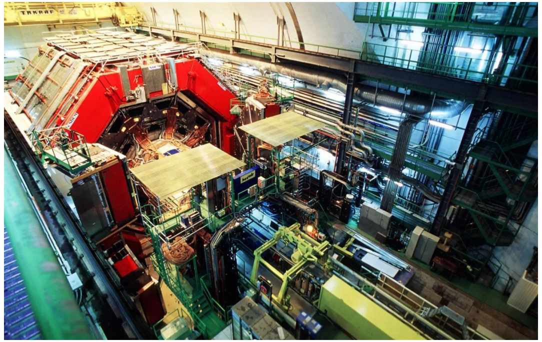

PEP2 - BaBar - CP Violation / Matter / antimatter asymmetry

1999 - 2008

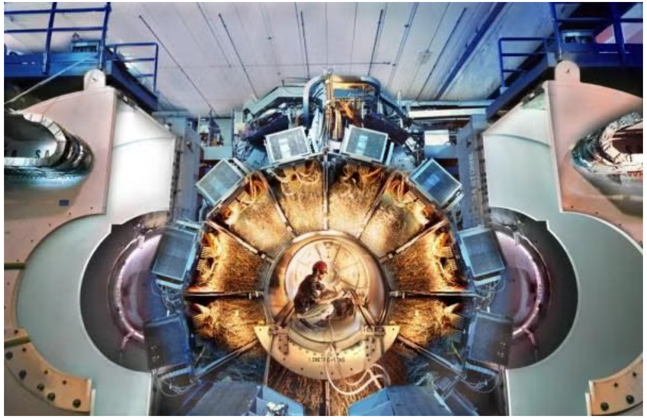

LHC - ATLAS, CMS, LHCb, ALICE - Higgs discovery

2010 up to 2041

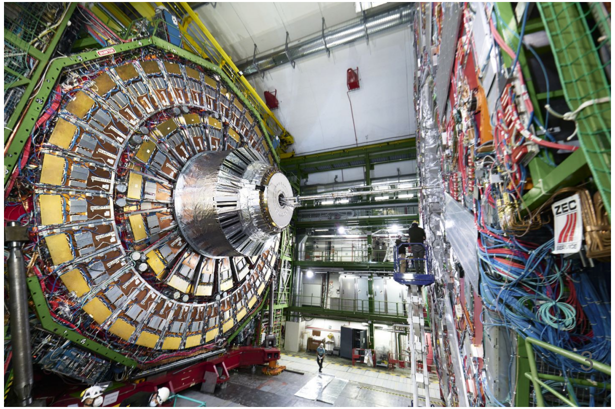

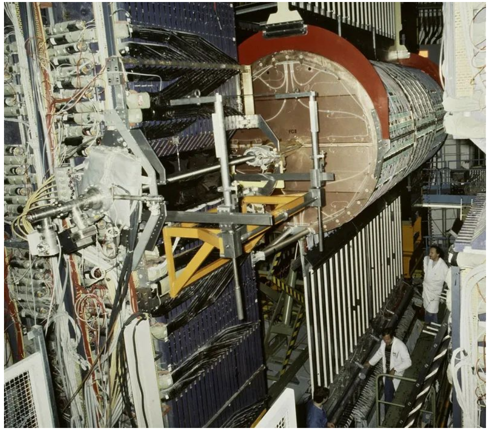

SppS - UA1 / UA2
W and Z discovery

50 10^3 channels

LEP - ALEPH, DELPHI L3,
OPAL - Z and W factory

500 10^3 channels

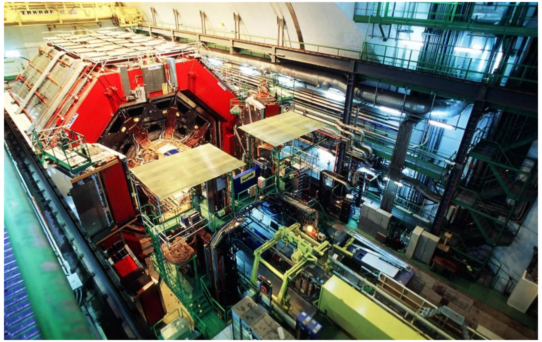

PEP2 - BaBar - CP Violation /
Matter / antimatter asymmetry

200 10^3 channels

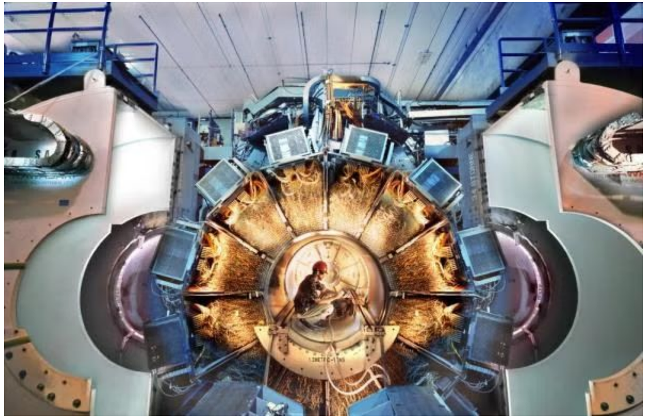

LHC - ATLAS, CMS, LHCb,
ALICE - Higgs discovery

100 10^6 channels

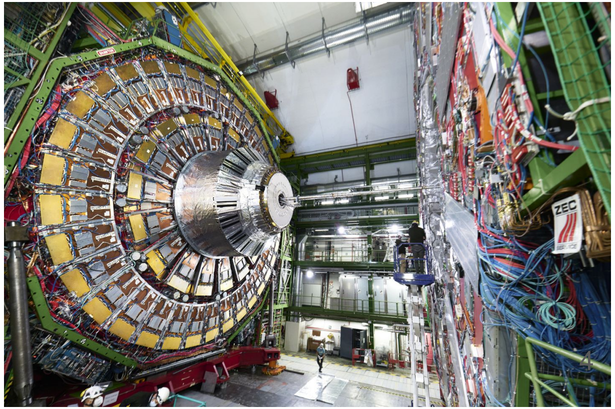

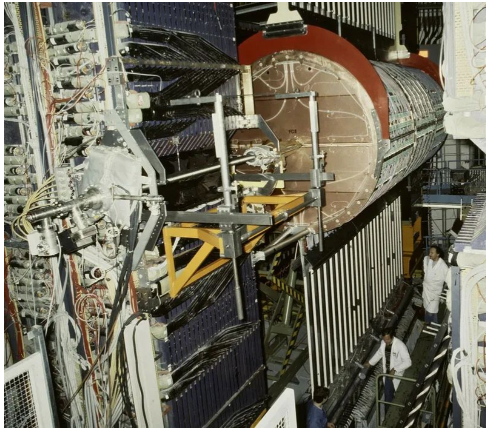

SppS - UA1 / UA2
W and Z discovery

50 10^3 channels

0.013 fb^{-1}

LEP - ALEPH, DELPHI L3, OPAL - Z and W factory

500 10^3 channels

1 fb^{-1}

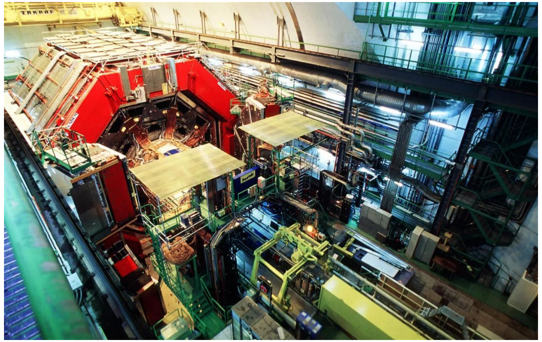

PEP2 - BaBar - CP Violation / Matter / antimatter asymmetry

200 10^3 channels **560 fb^{-1}**

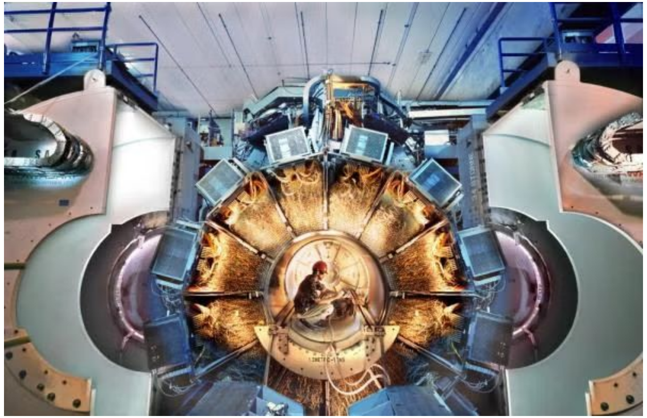

LHC - ATLAS, CMS, LHCb, ALICE - Higgs discovery

100 10^6 channels

515 fb^{-1} Run 3
4000 fb^{-1} HL-LHC

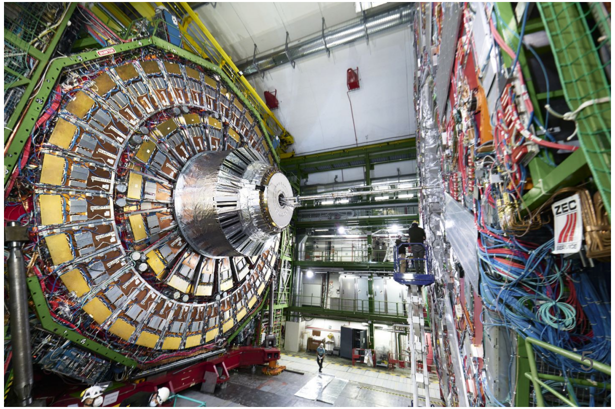

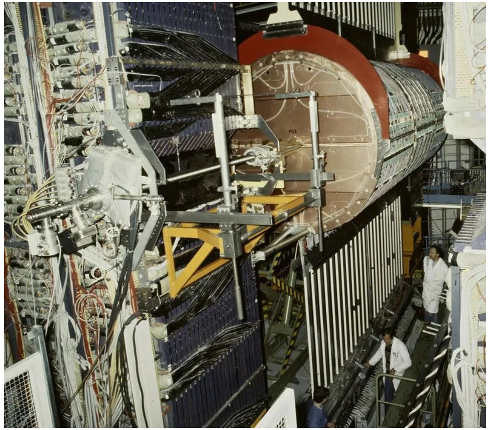

SppS - UA1 / UA2
W and Z discovery

50 10^3 channels

a few GB

LEP - ALEPH, DELPHI L3,
OPAL - Z and W factory

500 10^3 channels

a few TB

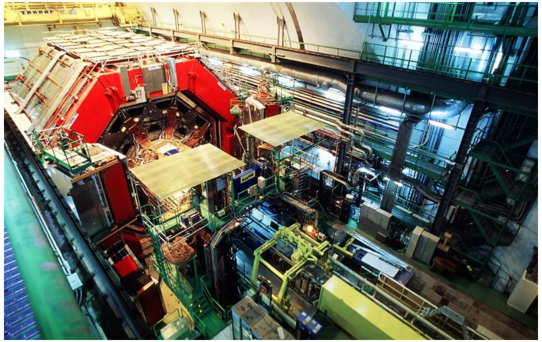

PEPII - BaBar - CP Violation /
Matter / antimatter asymmetry

200 10^3 channels **2 PB**

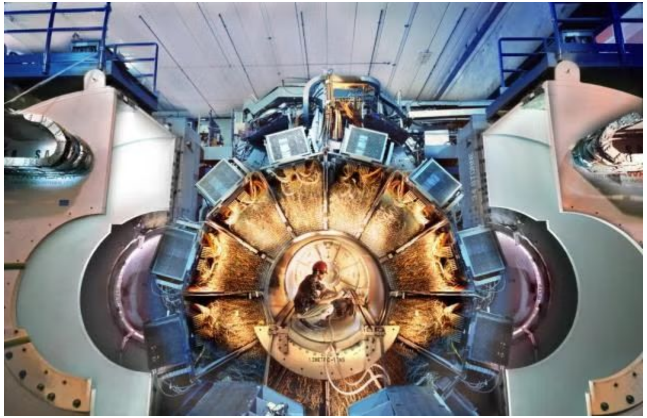

LHC - ATLAS, CMS, LHCb,
ALICE - Higgs discovery

100 10^6 channels

1.5 EB now
0.6 - 1 EB / year for HL-LHC

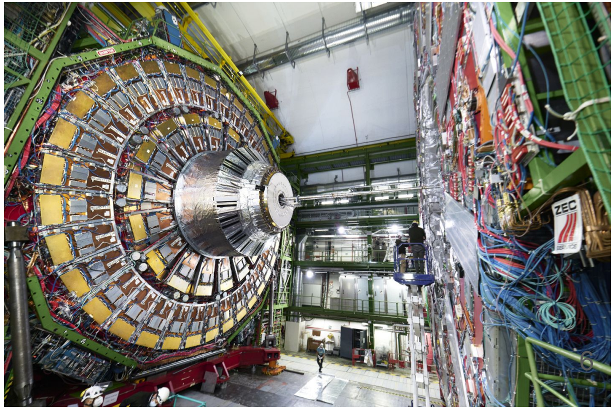

Simulation

On top of the collider data processing High Energy Physics needs a huge amount of simulations

- Physics generators
  - Simulates the physics process
- Particle tracking through the detectors
  - Detailed simulation of particle interactions with matter
- Simulation of detectors' response
- Reconstruction through the whole data processing pipelines

Simulated data may represent up to 10X the amount of real data

High Throughput Computing / High Performance Computing

Particle physics is fundamentally sequential

- 1 collision ⇒ 1 event
- Each collision is independent of the others

You can distribute the events on a computer farm

- 1 collision on 1 CPU core

In principle you can scale this to as many processors as you need

Embarrassingly parallel processing ⇒ **H**igh **T**hroughput **C**omputing

- Do not require expensive High Performance Super Computers
    - But you can use them if you have access

An incredible decrease of hardware cost over the years

Moore Law

Source: CERN - B. Panzer-Steindel

CPU cost:

- 2005: 200 CHF / HS06
- 2023: 3 CHF / HS06

Disk cost:

- 2005: 3.5 CHF / GB
- 2023: 0.03 CHF / HS06
- …
- **In 1988 the disk cost was 18 000€ / GB !!!**

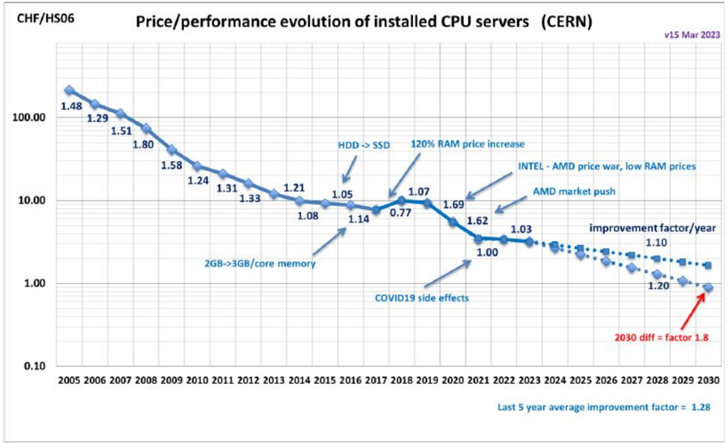

But the cost tends to stabilize now (+AI bubble…)

- No one knows how this will evolve in the coming years

A (huge) resource scaling problem (1)

LEP

- All the raw data processing was done centrally at CERN
- Part of the simulation and Physics analysis in national institutes
- No high speed network
  - Data transfer through tape copies and truck transportation !

PEPII - BaBar

- Emergence of high speed transatlantic network just in time to be able to move the data around the world
- First distributed computing model on 3 (very asymmetric) sites
  - A few thousands CPU cores

- *Transition from Fortran77 to Object Oriented C+*

A (huge) resource scaling problem (2)

LHC

- The projected computing resource needs was growing much faster than the Moore Law
  - CERN couldn't afford to make the whole data processing alone

A new paradigm had to be invented to make the LHC data processing affordable

1998 - Here comes the Grid

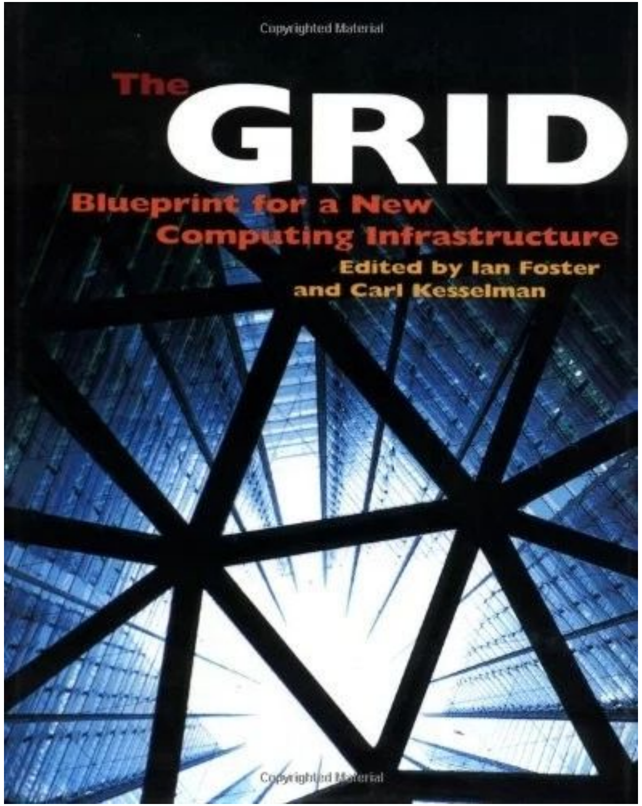

Getting computing resources should be as simple as getting electricity from the grid

As CERN couldn't afford the required infrastructure alone, the international partners have set up large, medium and small national computing centers that are interconnected and centrally orchestrated

Ingredients to make a functional Computing Grid

Some Computing resources distributed on several sites

A middleware to submit jobs on sites that have the resources to execute them

A way of handling jobs failures and retries

A global monitoring tool to understand what is going on

A powerful and reliable network to connect the sites

A way of replicating data between processing sites

- Catalog to keep track of what is where

A strong bookkeeping system

Once you have everything set up, then you need to scale this up to handle million of tasks on multi-petabyte data

Ingredients to make a functional Computing Grid

Some Comp etwork to
distributed o

A middleware between processing
sites that ha
execute the k of what is where

A way of har ystem
and retries

A global mo **ng set up, then you**
understand **handle million of**
data

© Edvard Munch

… but we did it !


1/24/2015 18:20:32
18:20
18:30
Running jobs: 277639
Transfer rate: 3.55 GiB/sec
WLCG
ALICE
ATLAS
CMS
LHCb
dashboard
US Dept of State Geographer
© 2015 Google
Data SIO, NOAA, U.S. Navy, NGA, GEBCO
Image Landsat
Google earth

… but we did it !

Initially: 140 computing centers in 35 countries

July 4th, 2012

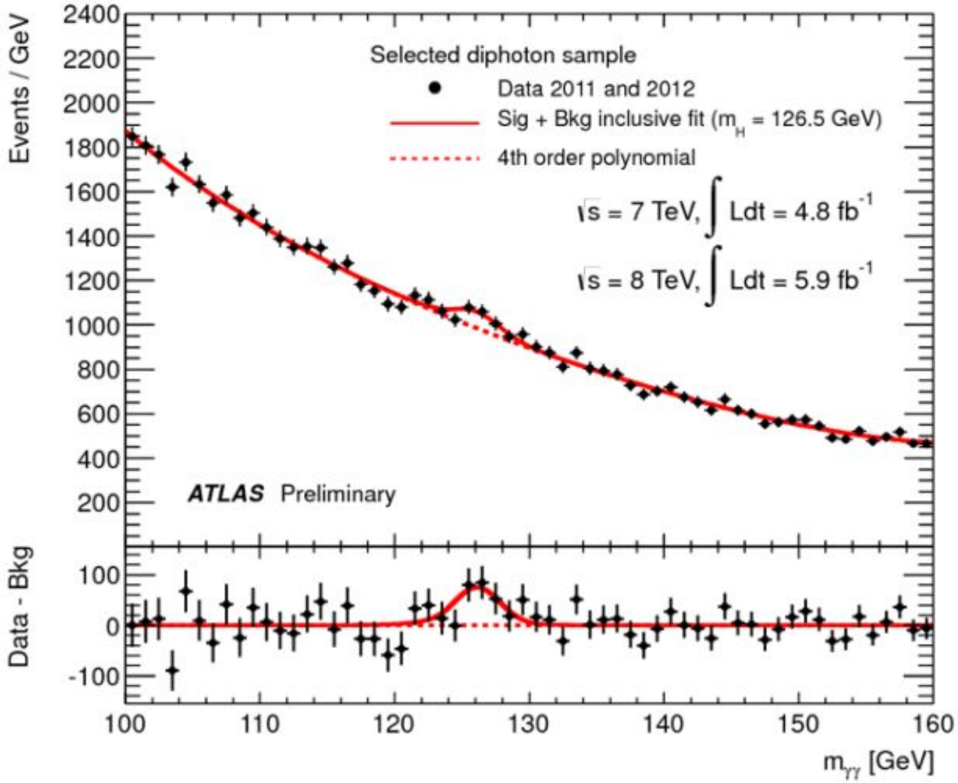

The LHC Grid was fully operational and allowed to process the LHC experiments data as soon as they were acquired

Nowaday physicists run their analyses by submitting jobs on the grid exactly as it it was a big computer with a simple user interface

Network

The network is an essential piece of the computing edifice

- The functionalities that are offered within a data processing system are most of the time limited by the network capacity
  - Network speed / bandwidth
  - Latency

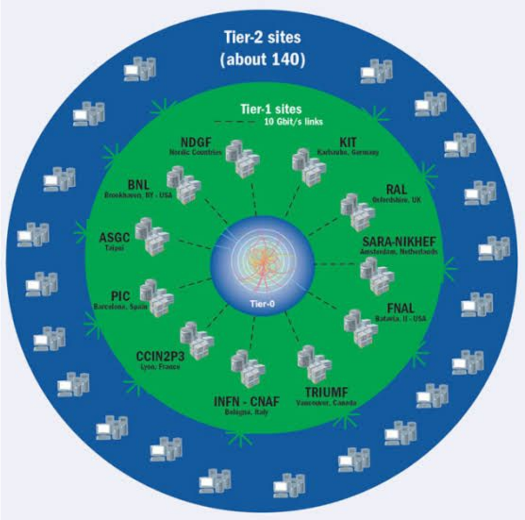

From the MONARC model to Full Mesh through the LHCONE network architecture

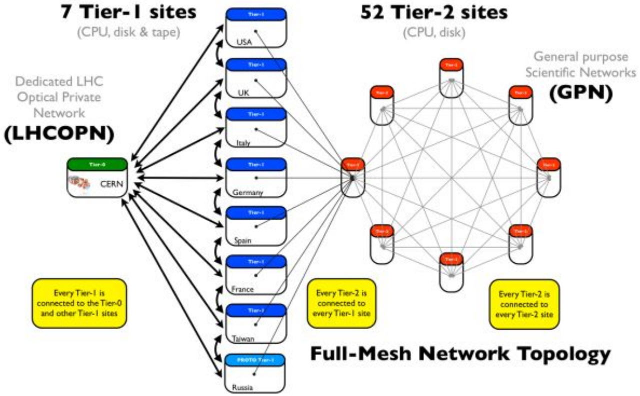

HEP as a leading developer of data processing technologies

Up to ~2010, HEP was the highest data producer / processor in the world far beyond any other scientific domains or private companies

HEP as a leading developer of data processing technologies

Up to ~2010 HEP was the highest data producer / processor in the world far beyond any other scientific domains or pri

Intelligent Data Access

HEP as a leading developer of data processing technologies

Up to ~2010 HEP was the highest data producer / processor in the world far beyond any other scientific domains or private companies

Data Management & Replication

HEP as a leading developer of data processing technologies

Up to ~2010 HEP was the highest data producer / processor i
beyond any other scientific domains or private companies

Job Workload Management

dCache.org
distributed storage for scientific data

XRoo

SCIENTIFIC DATA MANAGEMENT
RUCIO

FTS

DIRAC
THE INTERWARE

PanDA

HTCondor
Software Suite

Around 2010 the technological landscape shifted toward the major Internet companies

- Virtualization ⇒ Cloud computing ⇒ Google Cloud, AWS, Azure,...
- Large scale data mining ⇒ MapReduce / Hadoop…
- Extremely large databases ⇒ Google BigQuery…
- Scalability and Agility ⇒ Kubernetes
- … and many more…

And of course:

- Machine Learning
- AI
- GPUs

The HEP ecosystem had to integrate and coexist with these new tools

… and on the cosmology side…

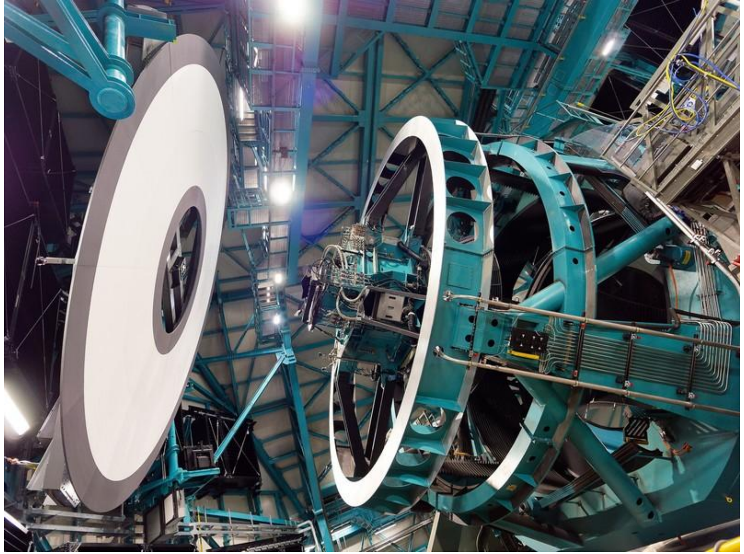

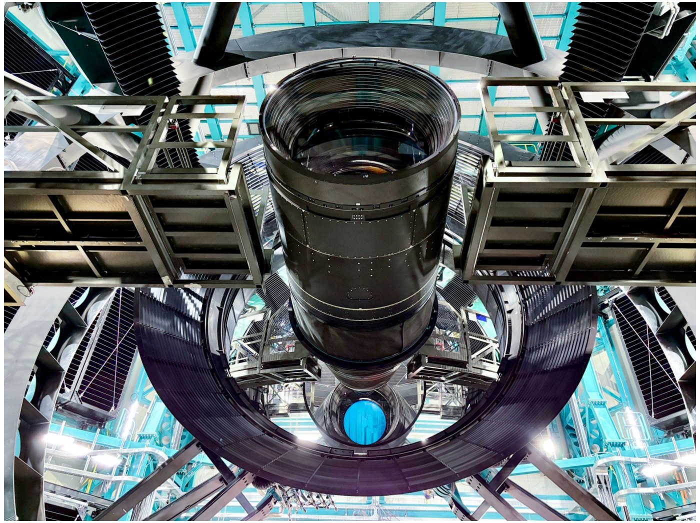

The Vera C. Rubin Observatory
and the Legacy Survey of Space and Time (LSST)

Vera C. Rubin Observatory

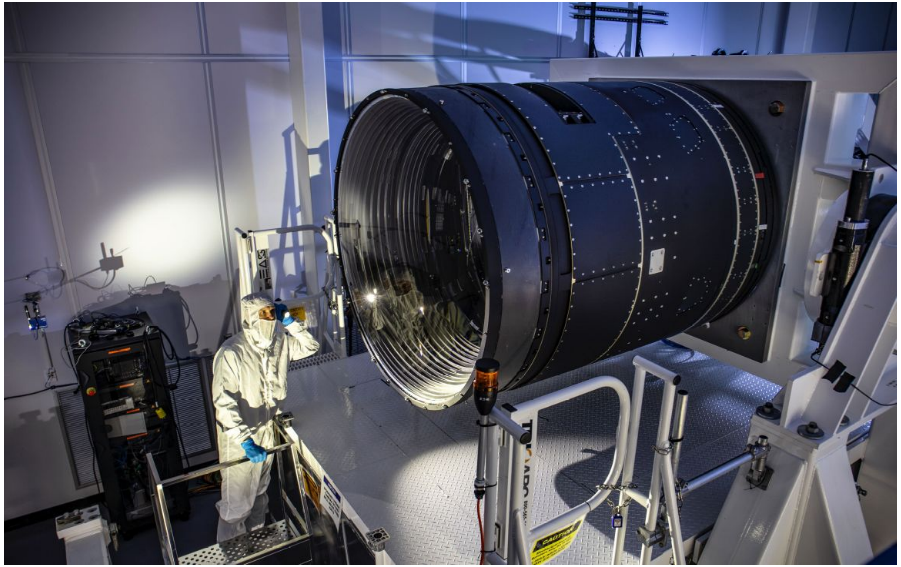

Every year the whole dataset is processed
⇒ 10 Data Releases

Giant Camera with 3.2 10^9 pixels

- 1 image every 40s each night during 10 years
- The whole sky is imaged every 3 to 4 night
- 20.10^9 galaxies
- 17.10^9 stars
- 6.10^6 orbits of solar system bodies

- 10 million alerts / night on transient objects

- 10 TB of raw images every night
- 30.10^{12} measurements recorded in a database after 10 years


UK Data Facility
SLAC/USDF
French Data Facility
HQ Site (NOIRLab)
Brazil
Cerro Pachón
CHILE
Santiago

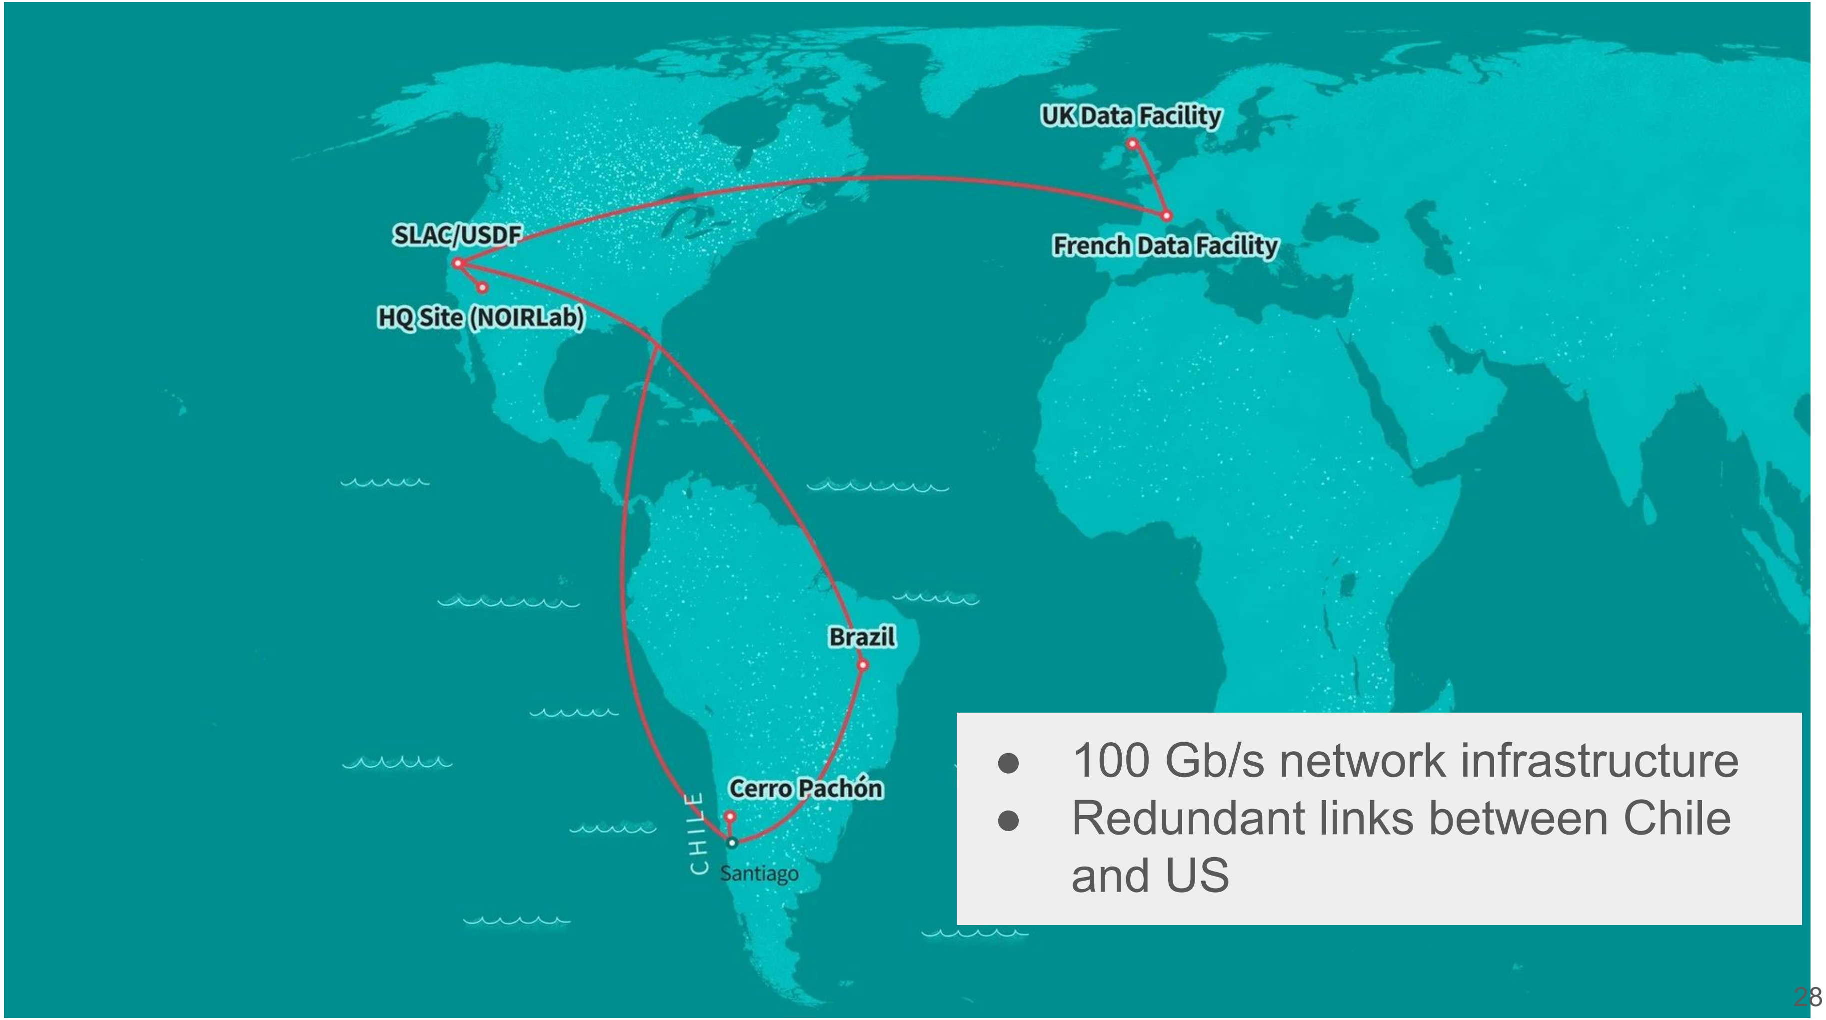

- 100 Gb/s network infrastructure
- Redundant links between Chile and US

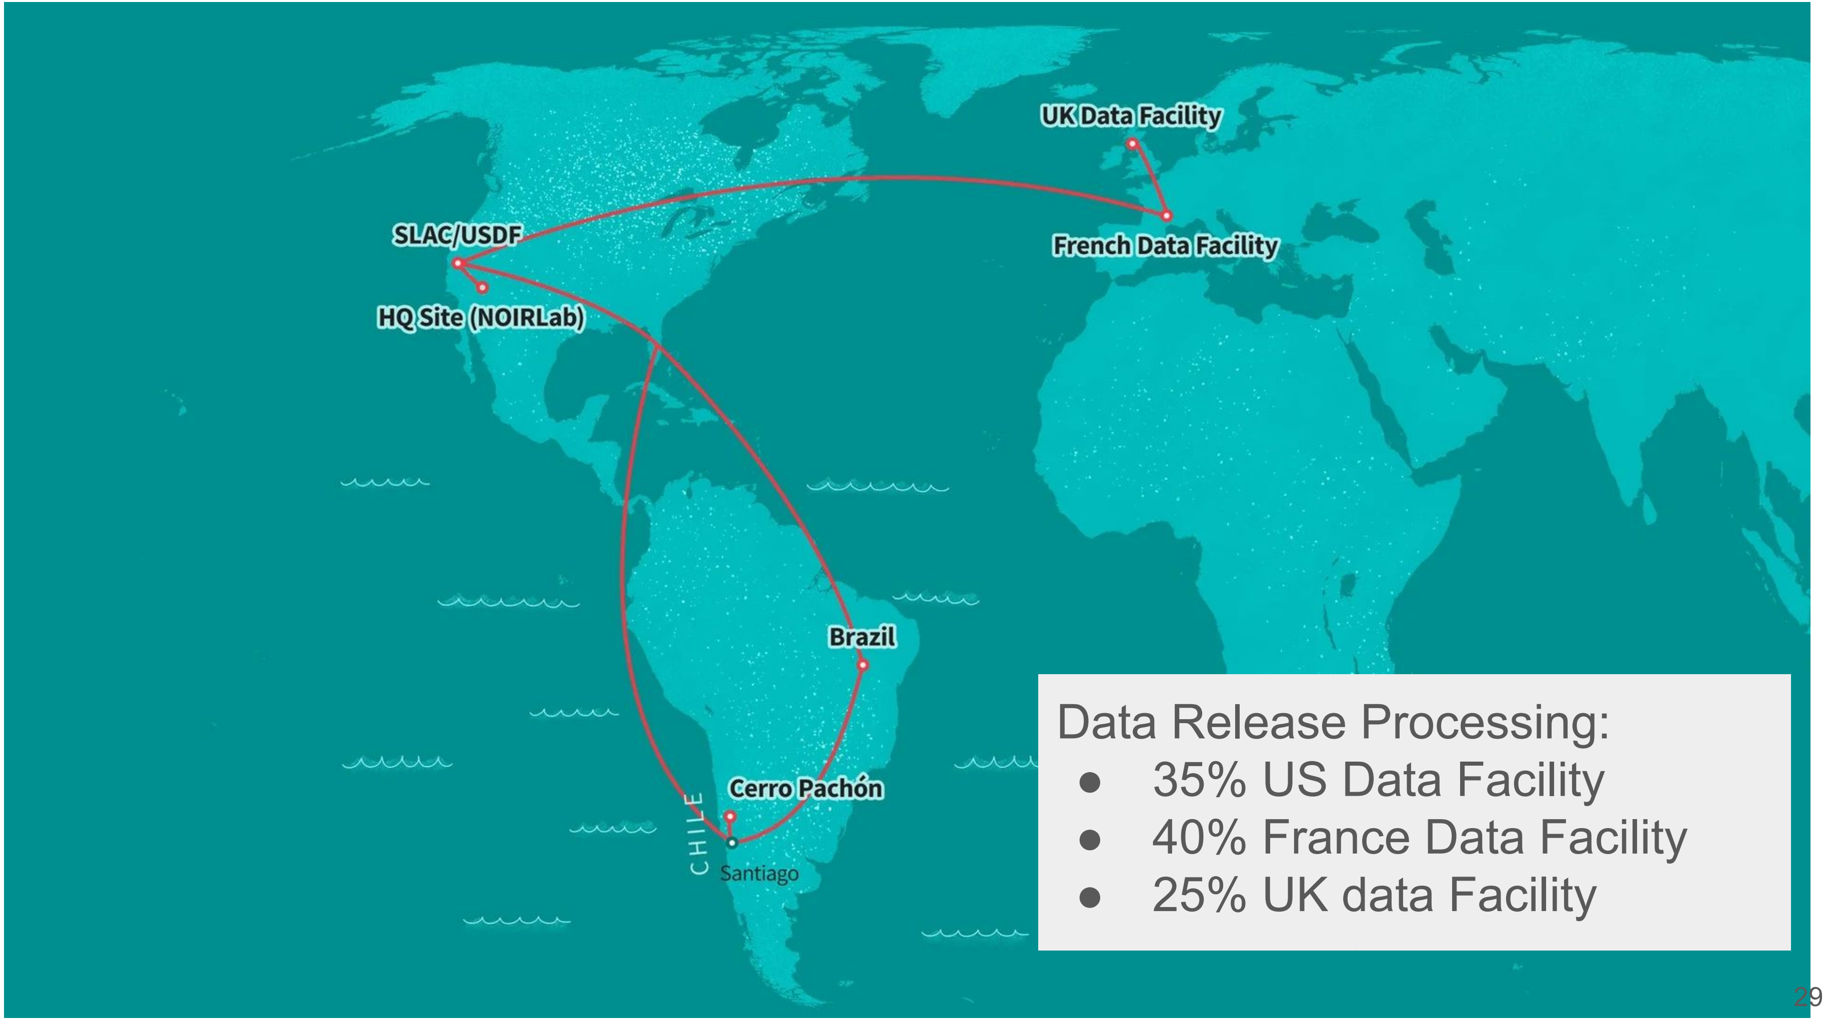

Data Release Processing:

- 35% US Data Facility
- 40% France Data Facility
- 25% UK data Facility

Prompt processing

https://lsst.fink-portal.org/170411112000913487

Each image is transferred from Chile to the US Data Facility

- Comparison to image template of the same region
- Any difference triggers an alert within 60 seconds

Stream of alerts sent toward 7 alert brokers:

- Classify
- X-Match
- Enrich
- …

The science communities can then connect to the brokers to filter and select the alerts

France: Fink https://lsst.fink-portal.org

The Rubin Software

The Rubin Software development was part of the project since the very beginning

- 70 M$ invested in software development (not infrastructure)
- Mostly done by professionals

Mixt between C++ (low level algorithms) and Python (interfaces)

Complex but justified by the amount of pixels to be processed and the need to serve multiple scientific communities

More complicated than the HEP embarrassingly parallel processing due to the need to combine multiple images to detect fainter objects (co-addition)

- Every year we reprocess everything since the beginning of the project

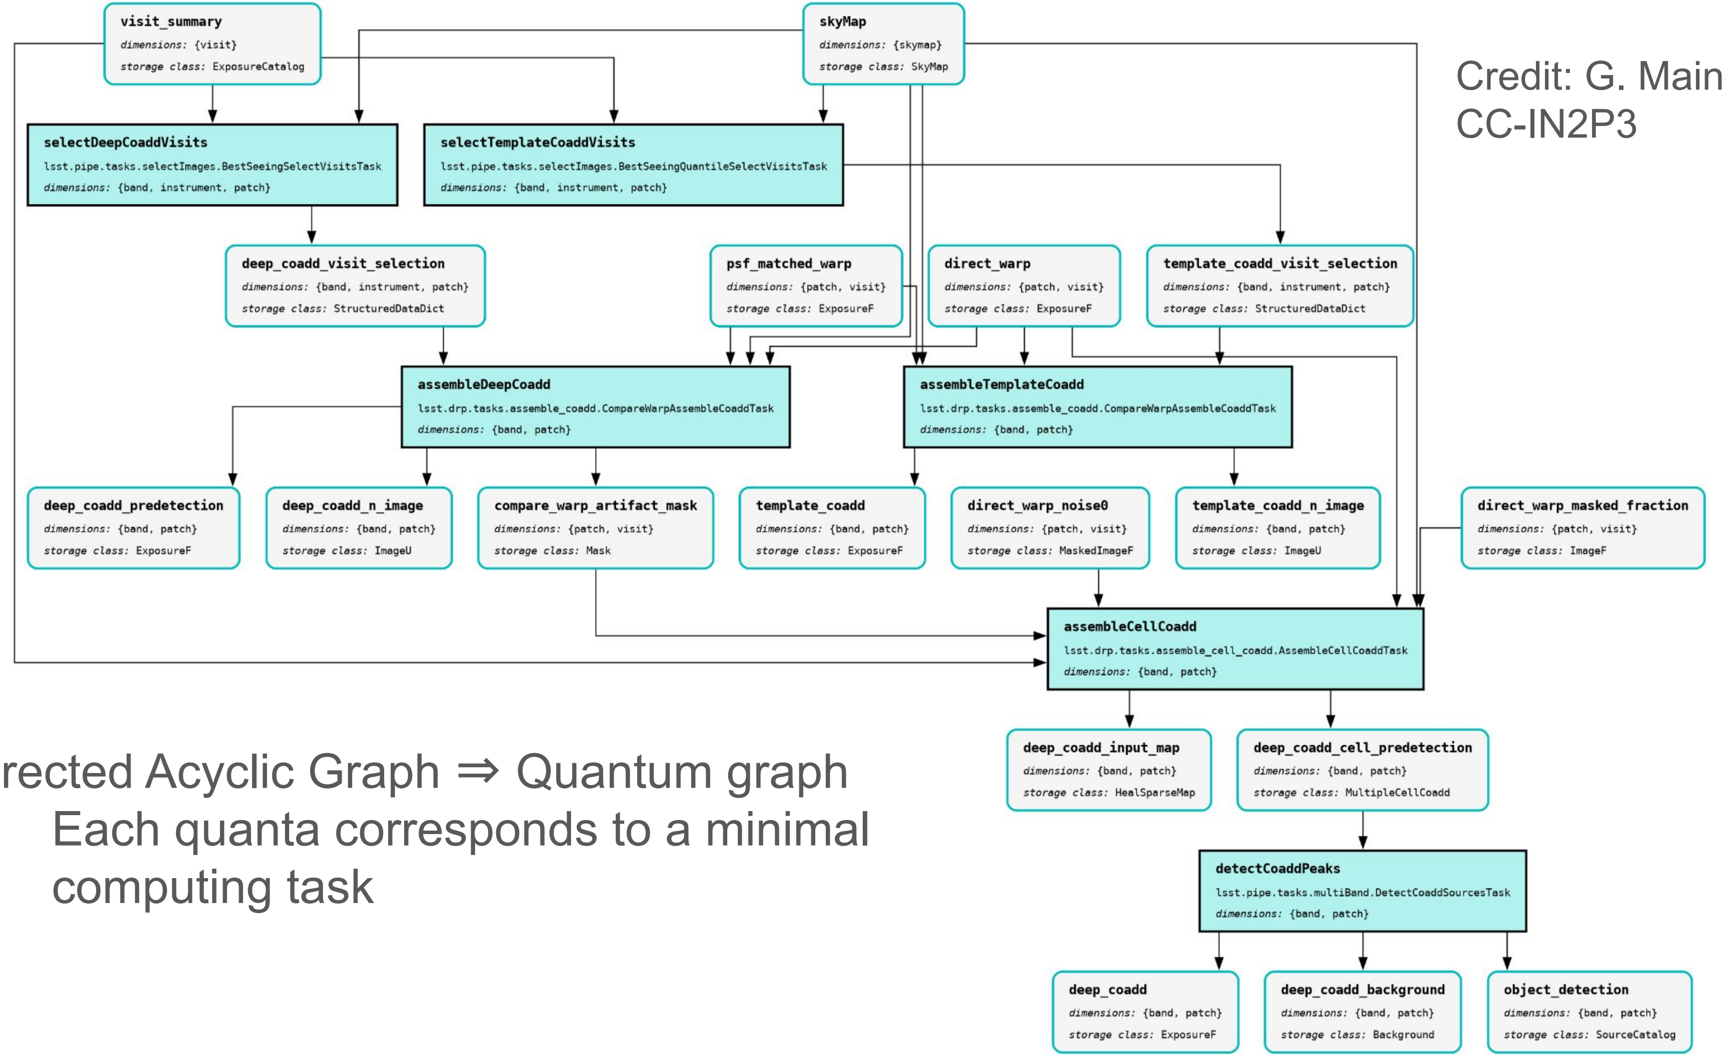

Credit: G. Mainetti
CC-IN2P3

Directed Acyclic Graph ⇒ Quantum graph

- Each quanta corresponds to a minimal computing task

An extra-large database to store the sky "objects"

Rubin decided to develop its own database system **QServ**

- Shared-nothing distributed database built on top of MariaDB and xRootD

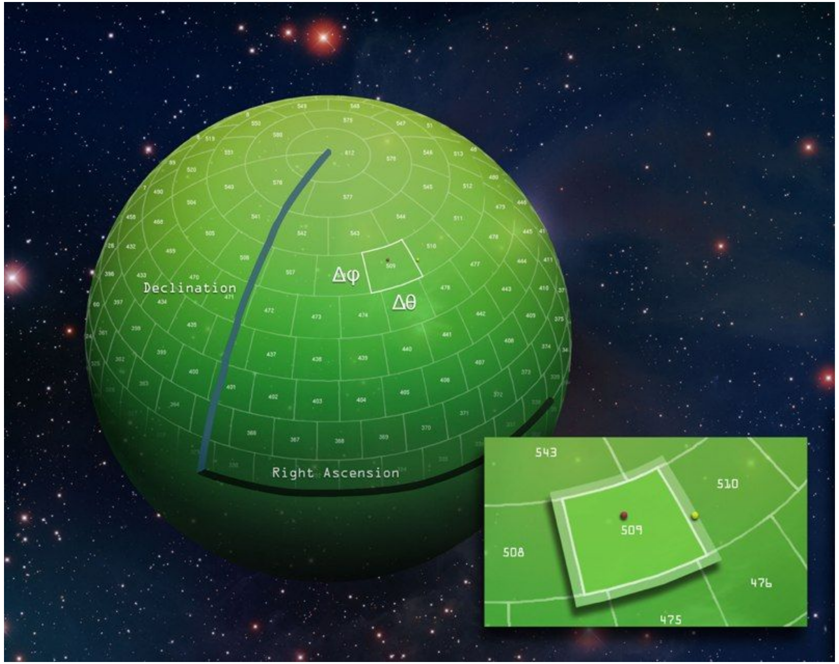

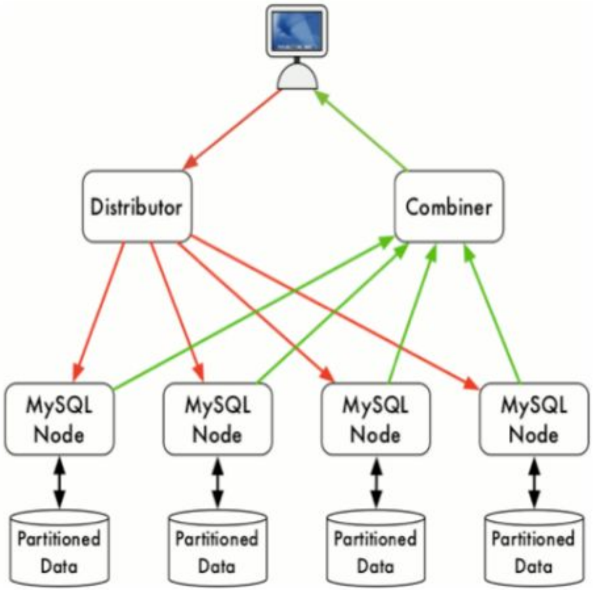

Rubin Science Platform

Rubin has developed a Science Platform able to support the analysis from ~8000 users

- Deployed on Google cloud + several other instances
- Heavily rely on Kubernetes
- Fully integrated with the Rubin software environment
- Implement Jupyter Notebooks + Portal to access QServ
- Implement the International Virtual Observatory Alliance standards

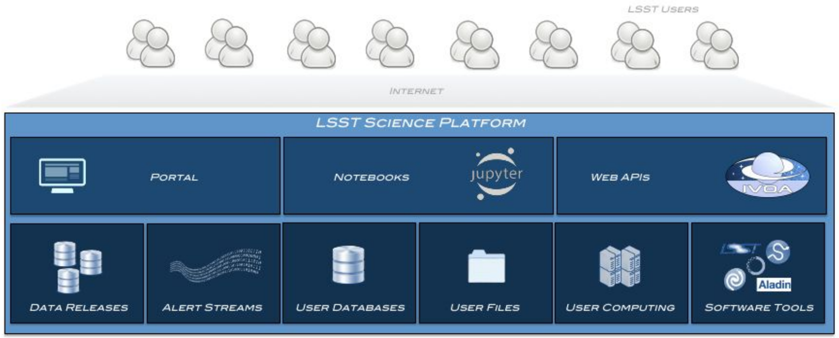

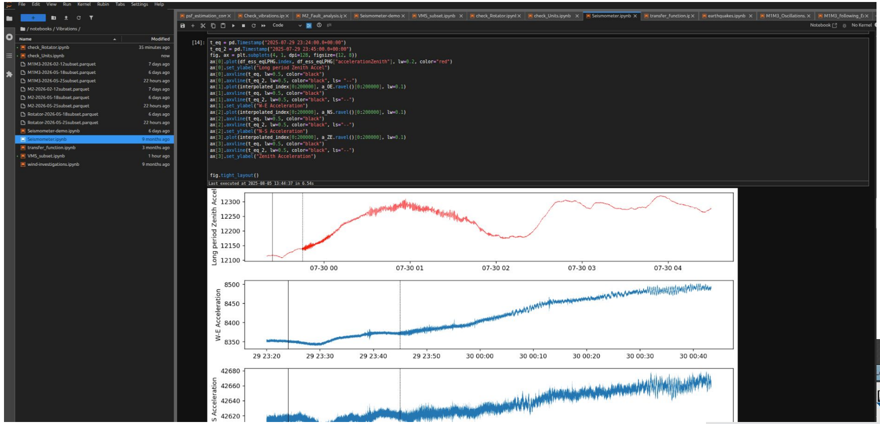
Long period Zenith Accel
W-E Acceleration
S Acceleration

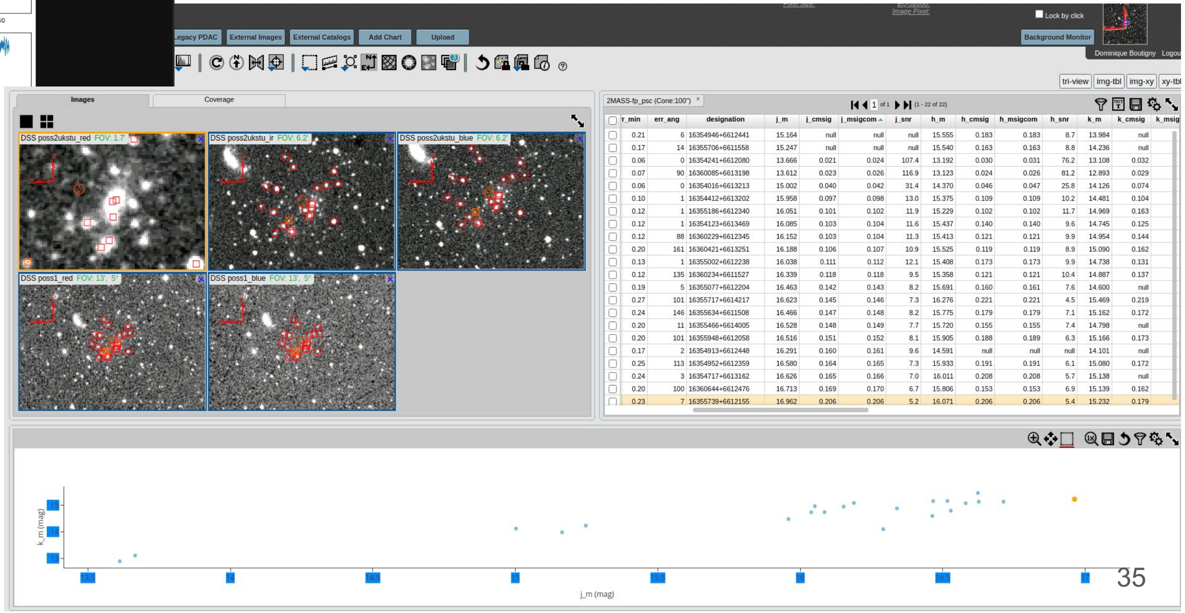
Legacy PDAC
External Images
External Catalogs
Add Chart
Upload
Background Monitor
Images
Coverage
DSS poss2ukstu_red
DSS poss2ukstu_ir
DSS poss2ukstu_blue
DSS poss1_red
DSS poss1_blue
2MASS-fp_psc (Cone:100")
k_m (mag)
j_m (mag)

The Rubin Science Platform is an implementation of the Infrastructure as Code paradigm

- The infrastructure is fully described in configuration files

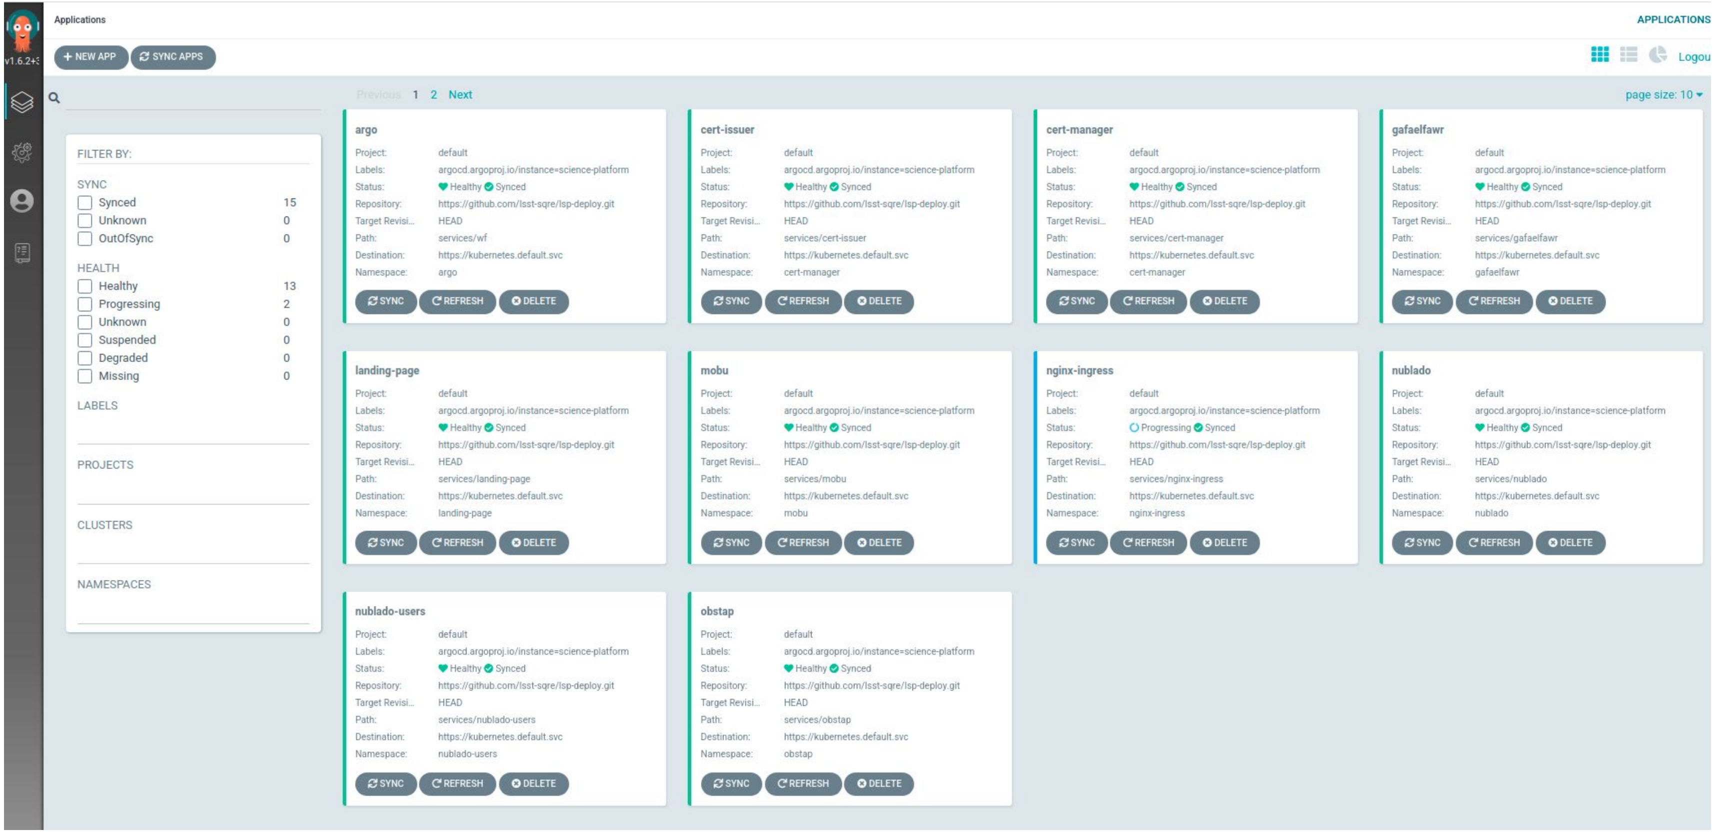

Sustainability (1)

We can't minimize the environmental footprint of HEP + Astro / Cosmo computing

- Resources for device manufacturing (CPU, memory, storage…)
- Electricity
- Cooling
- (Water consumption)
- …

But I would like to put a few numbers in perspective:

LHC Computing:

- 1 Tier-0 + ~11 Tier-1 centers ⇒ ~15 MW

AI Training

- people are speaking of a number of (?) (multi)GW sites !

Sustainability (2)

- It doesn't mean that we should consider our resources as negligible

- Enormous efforts are invested to optimize the code and the hardware architectures
  - Very complex
  - Requires highly skilled people

- A side effect of the electronics component shortage is that the costs is already increasing and will continue to do so
  - Nobody can reasonably predict where we will be in the coming years
  - But this is another reason to work on optimization !

Conclusions

- The implementation of the Grid for LHC and other experiments is a success story
  - International collaborative development effort
  - Resource sharing
- HEP is not anymore a leader in data processing (by far)
- We have to adapt to new technologies developed elsewhere and integrate what we can in an opportunistic way

- We have to face the increasing complexity of the code development and of the data processing infrastructures
  - Shift toward highly skilled professional developers
  - **We still need to understand how to make sure that the technical complexity isn't impairing the quality of Science**

- Code optimization is crucial