Big Data
and
Big Science

Tossing a trick coin

Gabriele Gaetano Fronzé

UNIVERSITA' DEGLI STUDI DI TORINO
INFN Istituto Nazionale di Fisica Nucleare
IMT Atlantique Bretagne-Pays de la Loire École Mines-Télécom
Subatech

Outline

- Evolving requirements: the LHC evolution, the LHC revolution
- Are we (still) alone?
- ALICE's wonderland
- Final thoughts

LHC evolution, LHC revolution

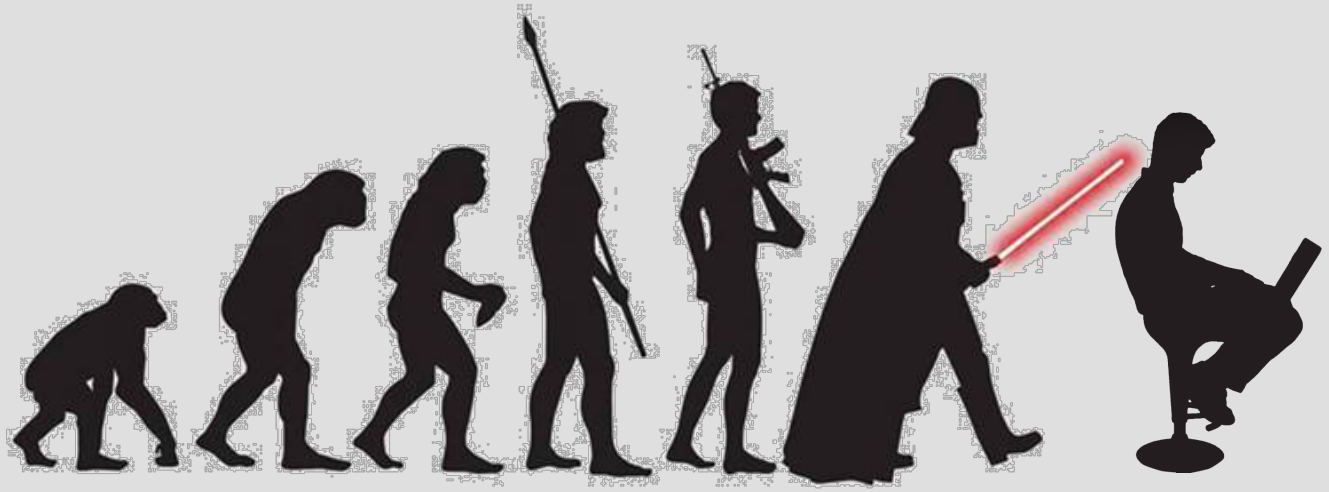

LHC ICT challenges

Historically LHC presented two ICT challenges:

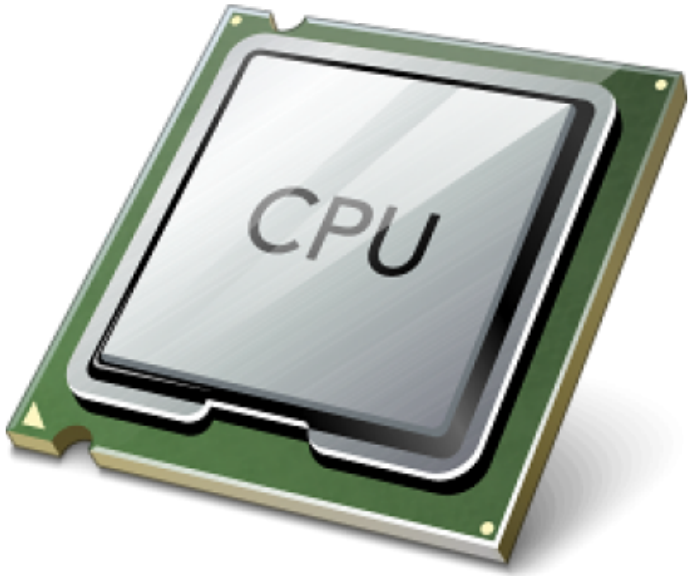

LHC ICT challenges

Historically LHC presented two ICT challenges:

Storage

Storage requirements have largely increased since the first LHC run.

Each and every luminosity upgrade triggers a data rate bump.

Data compression and selection methods are already almost ideal.

Foreseen requirements are far from being affordable

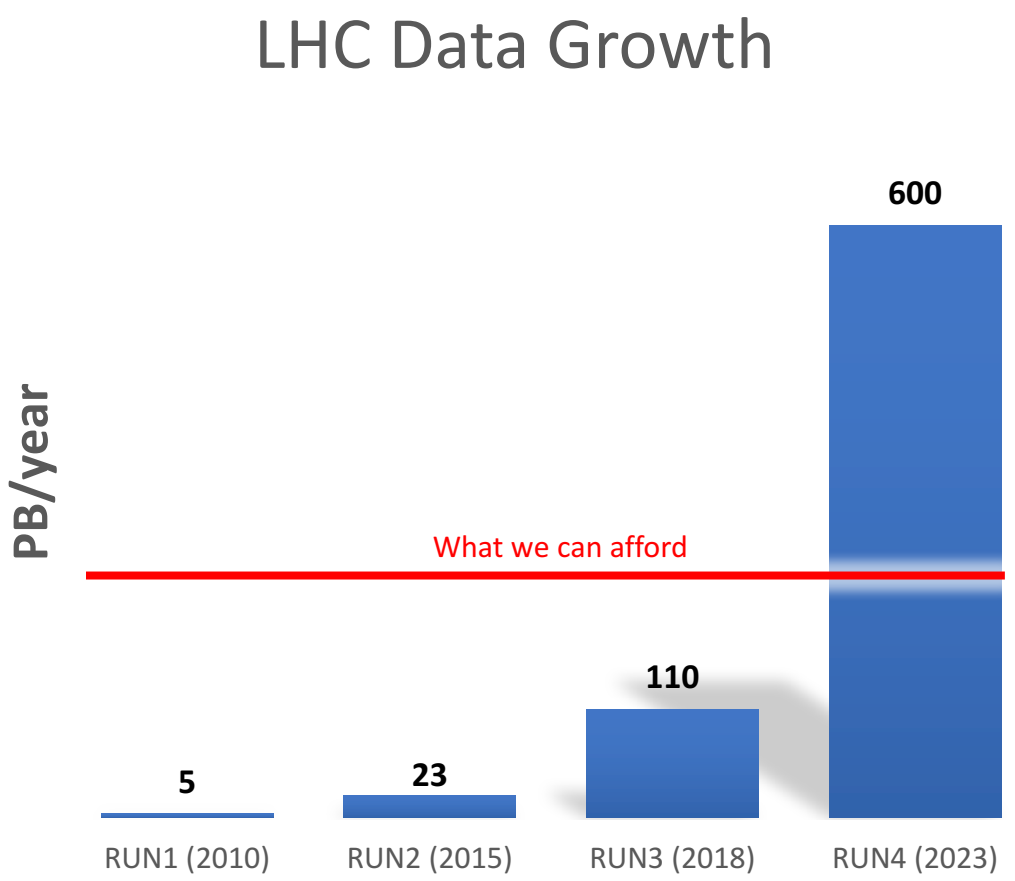

From Tim Bell (OSCON 2014) and Arne Wiebalck (German OpenStack Days 2016) talks

Storage

First idea:

Let’s wait for technology to evolve (and get cheaper)

Storage

First idea:

Let's wait for technology to evolve (and get cheaper)

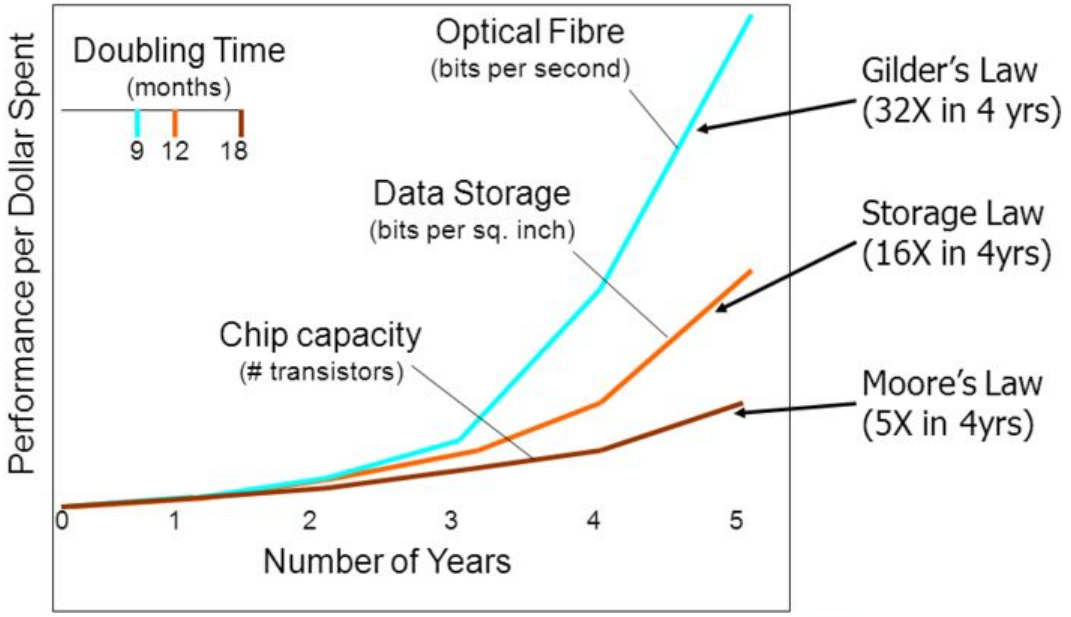

From Triumph of Light – Scientific American George Stix, January 2001

Storage

First idea:

Let's wait for technology to evolve (and get cheaper)

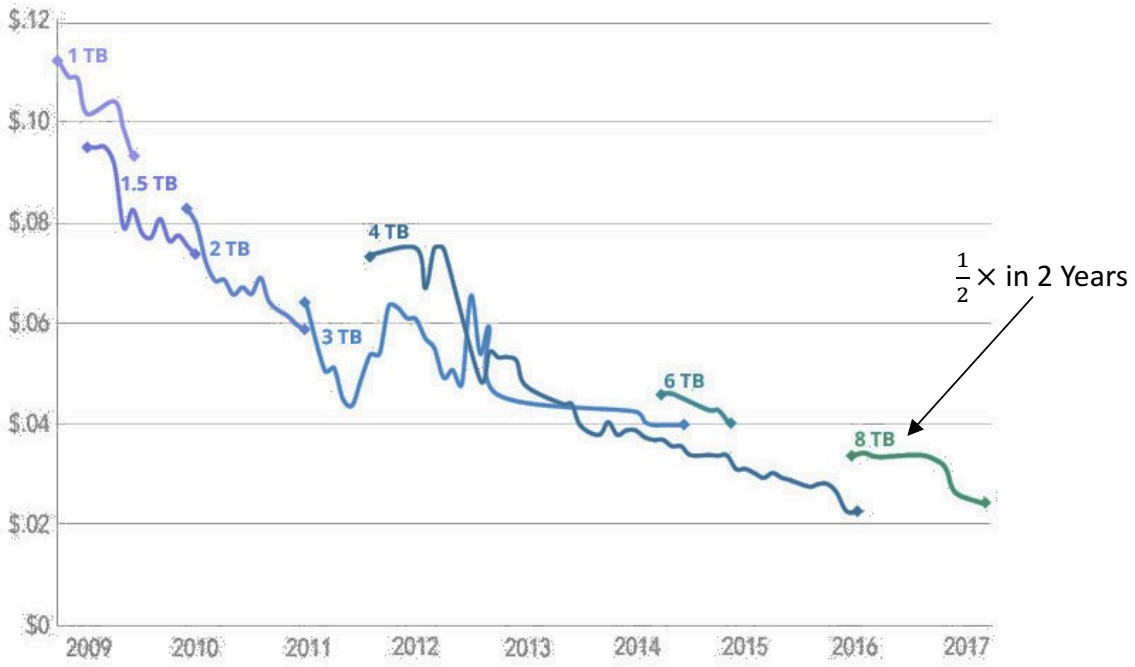

From Backblaze.com

Storage

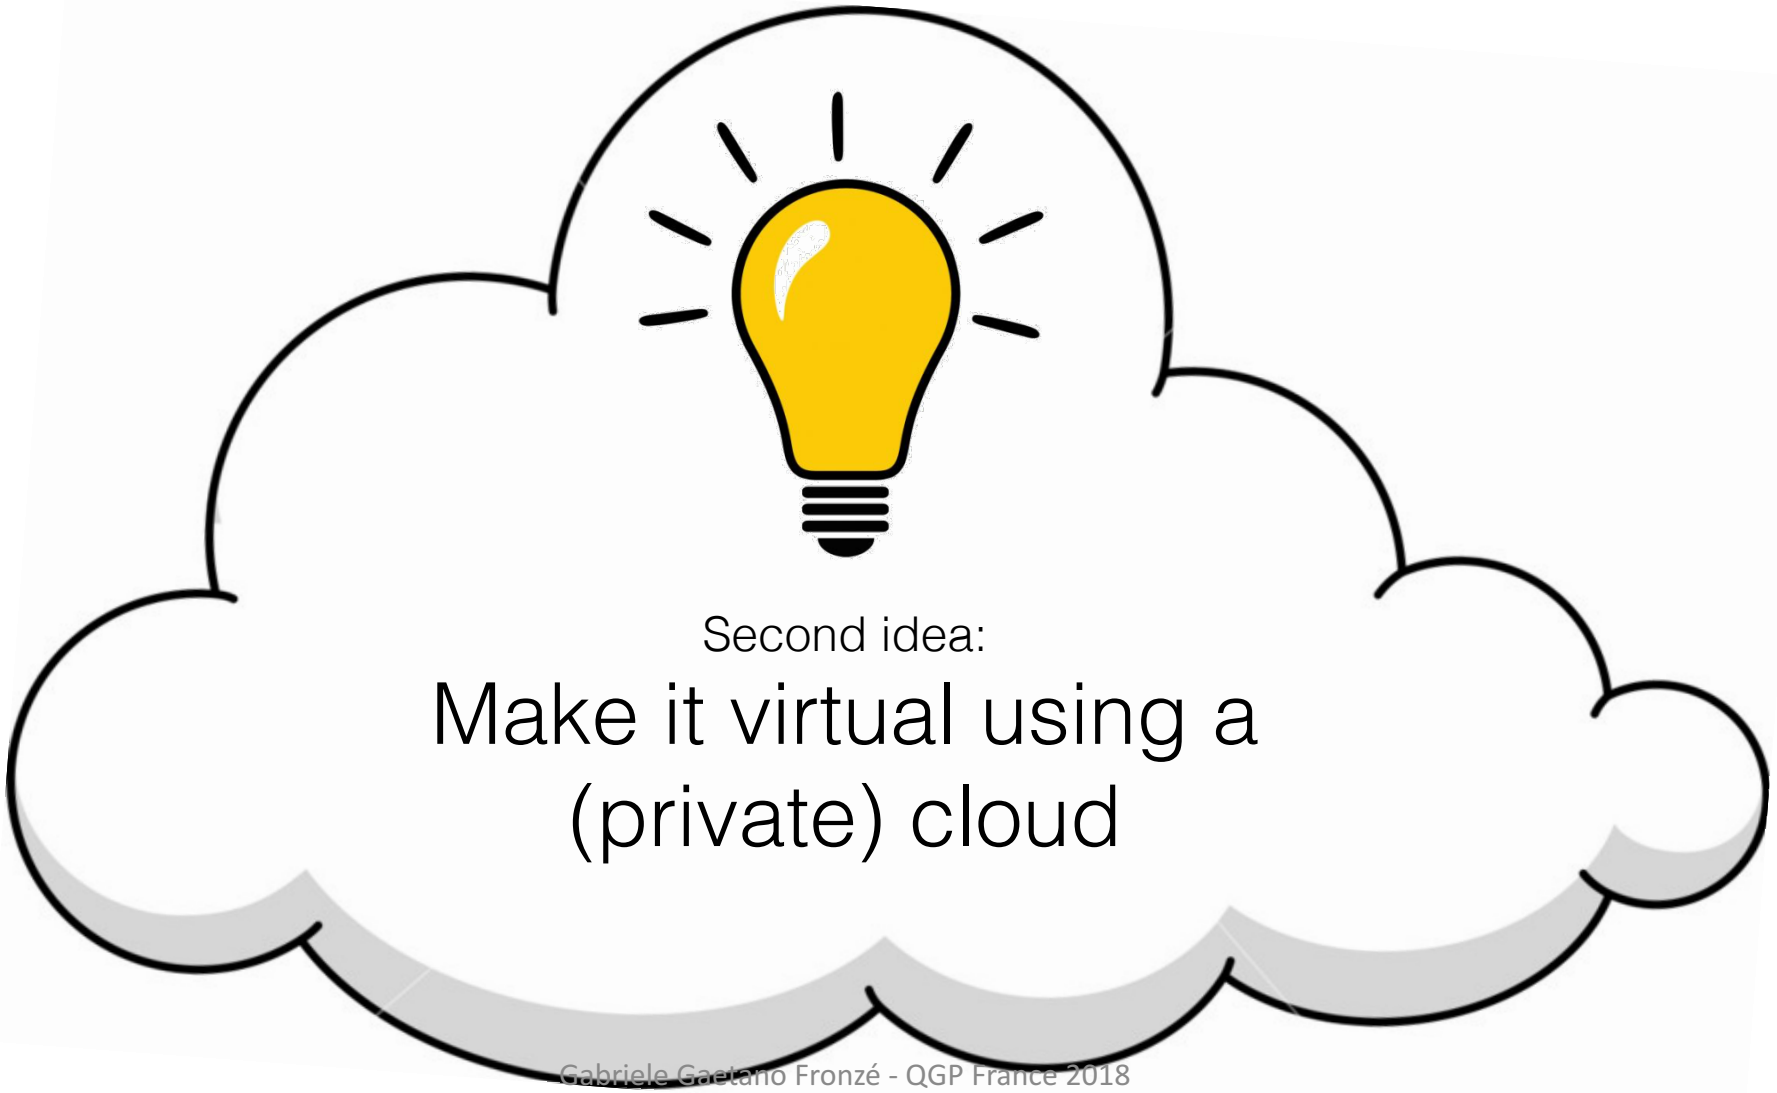

Second idea:

Make it virtual using a (private) cloud

Storage

Second idea:

Make it virtual using a (private) cloud

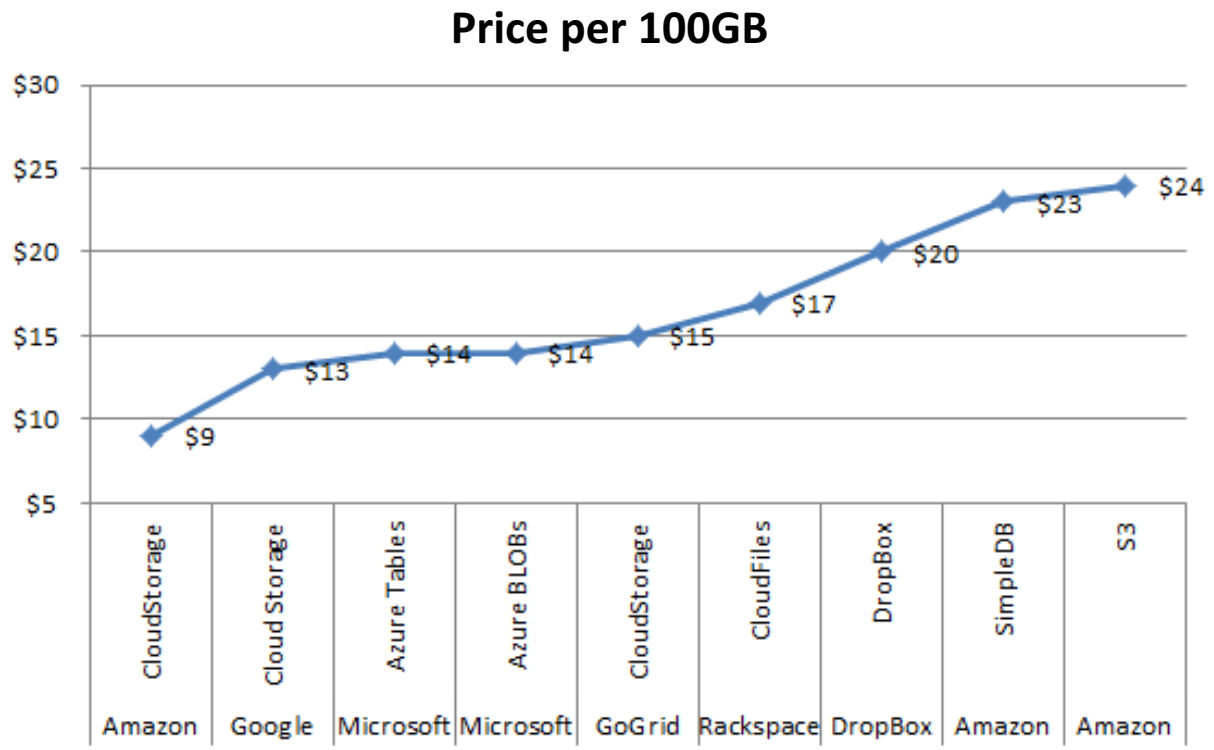

From lynnlangit.com

Storage

Second idea:

Make it virtual using a (private) cloud

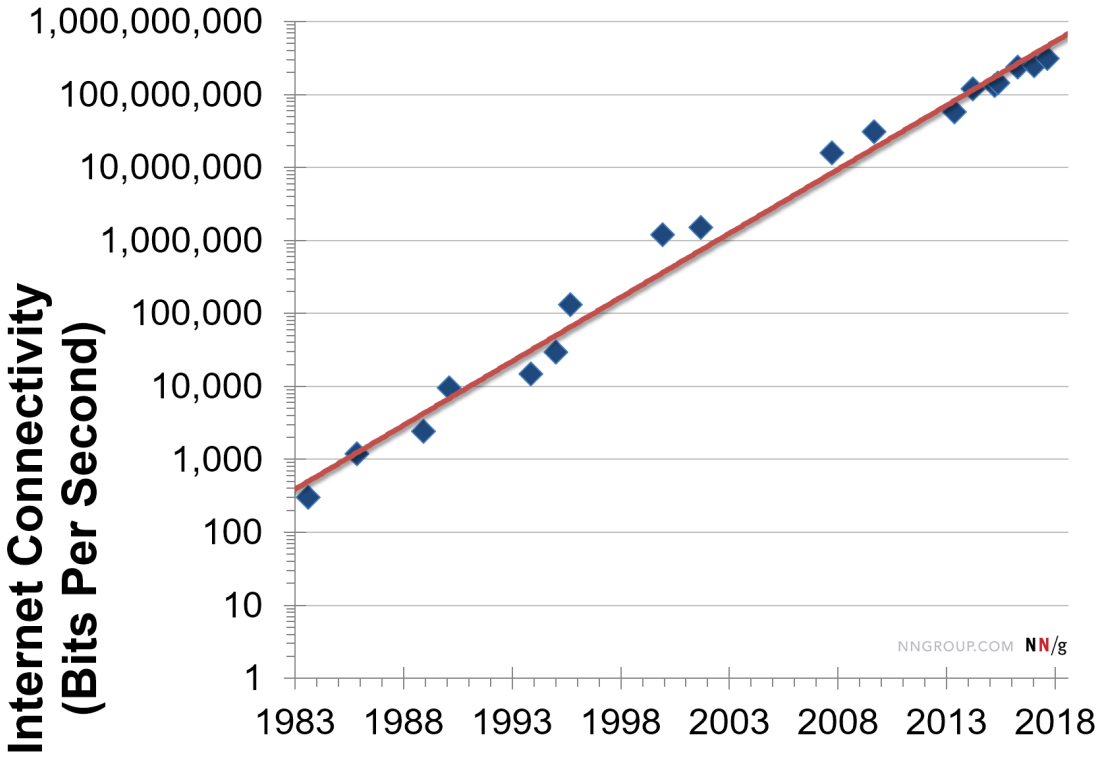

From nngroup.com

Computing power

Historically LHC presented two ICT challenges:

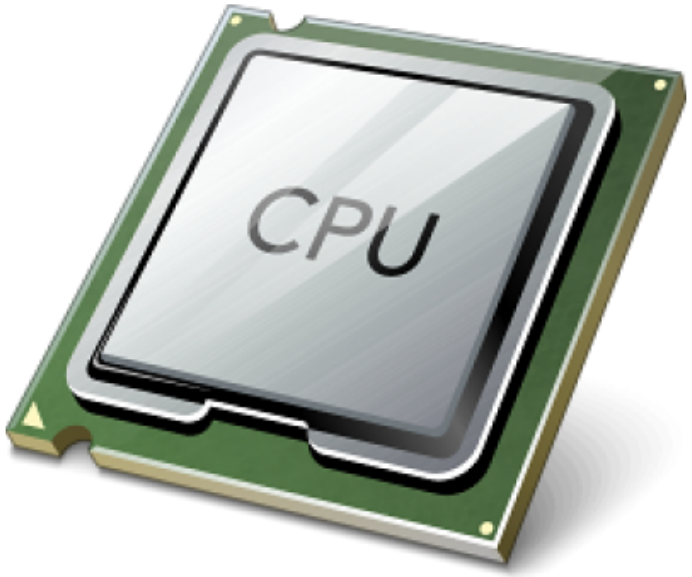

Computing power

CERN WLCG has been designed when most computers were single core.

Single core performance plateau has been reached.

Core count increase is now a leading trend.

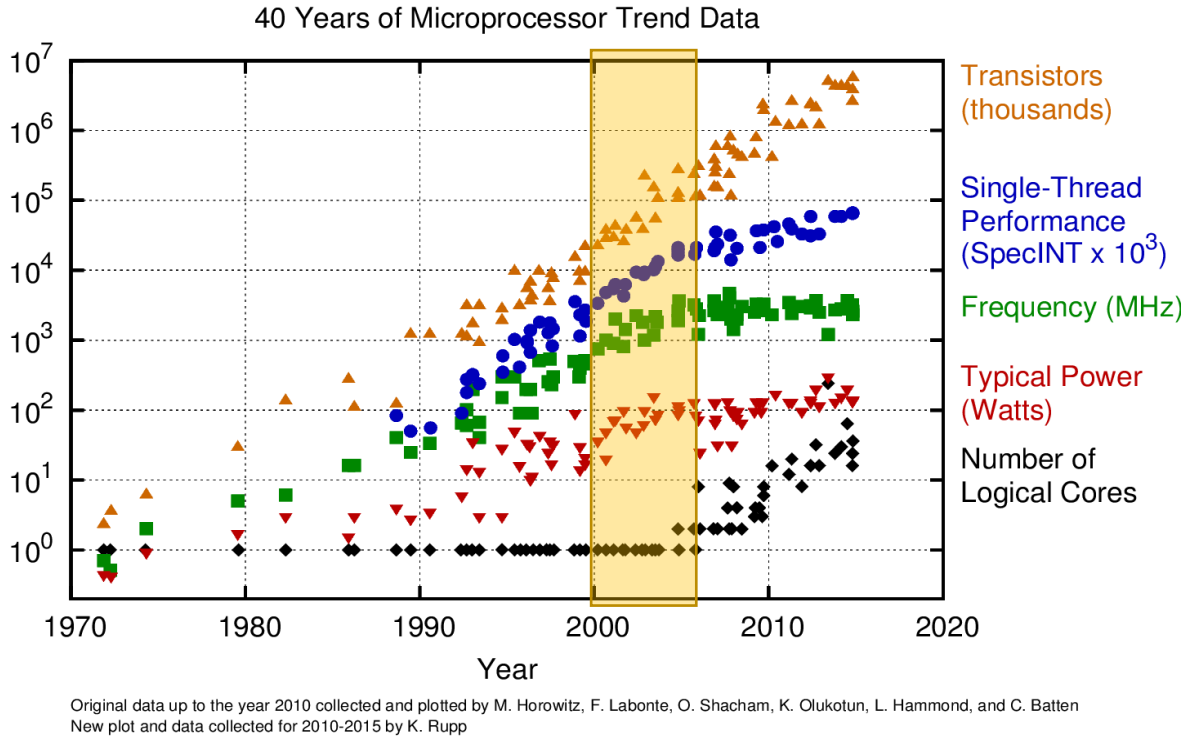

Computing power

Ground-breaking technologies in computing accelerators are being tested by CERN OpenLab.

In the future FPGAs and GPUs will be the main computation tools for the research.

X86 and x64 instructions are still diffused but PowerPC and ARM architectures are being evaluated.

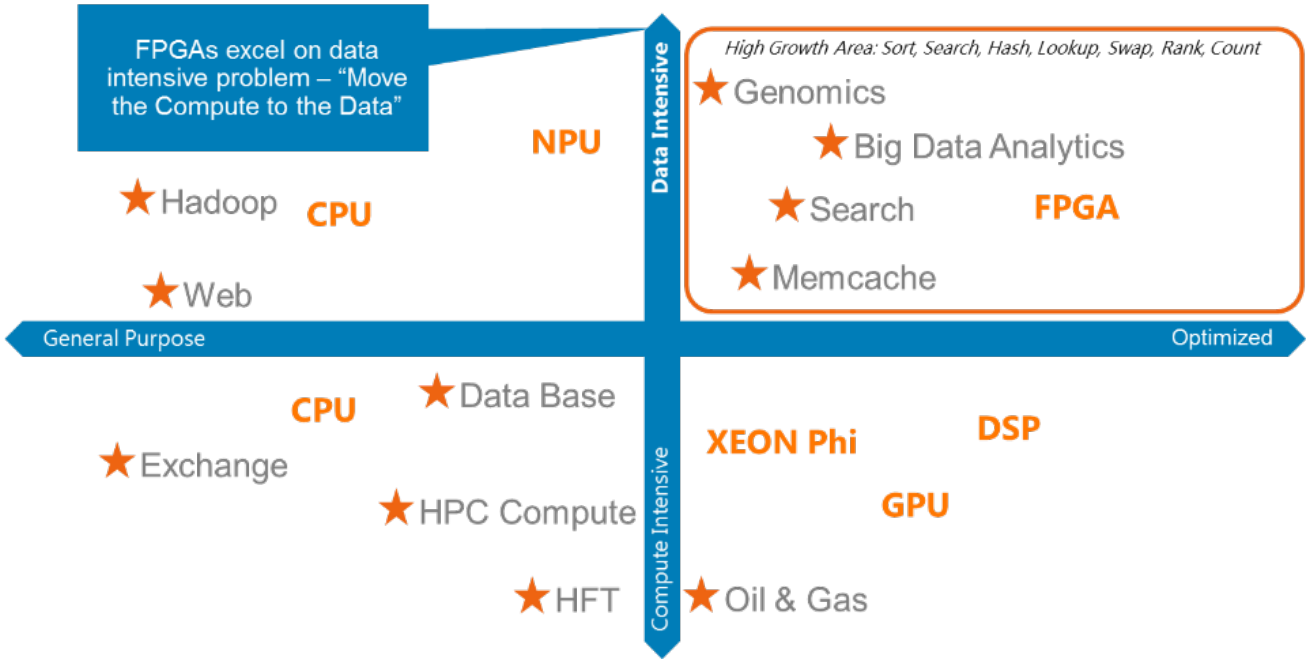

From Dell (blog.dellemc.com/en-us/fpgas-use-cases-in-the-data-center)

Computing power

CPUs with around 20 cores per chip are currently available.

The performance density is higher than ever in dual and quad CPU systems.

BUT…

Computing power

Are we able to use the full power of these engineering prodigies?

Computing power – doing the right choice

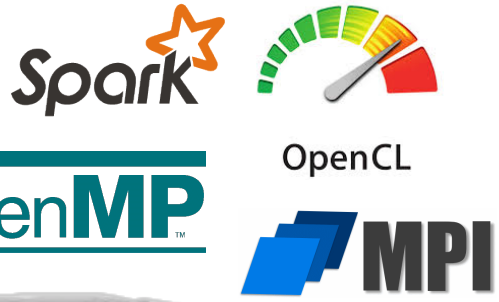

We have to choose between the low level and the high level roads

Difficult to use but extremely powerful and manageable.
Ideal results with maximum effort.

Less customization but open to common interfaces.
Can hide hardware complexity.

Computing power – doing the right choice

Cloud computing means:

- Ubiquitous access to computing resources
- Pool of (heterogeneous) resources
- Configurable (and fast reconfigurable)
- Over the internet (?)

Computing power – doing the right choice

A cloud infrastructure is network-bandwidth bound

BUT…

Computing power – doing the right choice

Network evolution is one of the fastest growing ICT departments

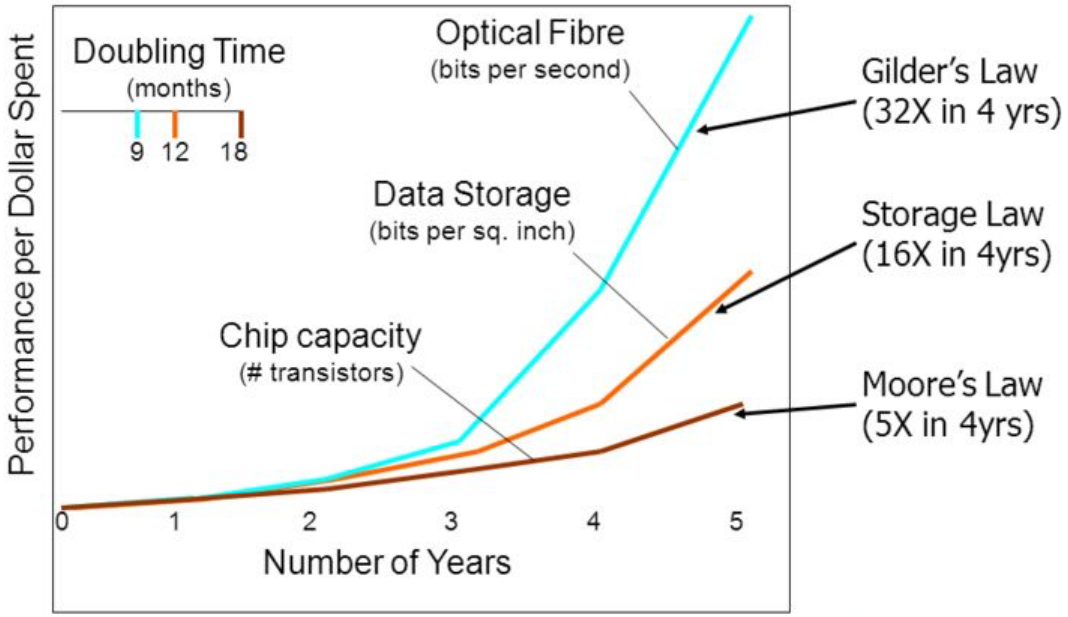

From Triumph of Light – Scientific American George Stix, January 2001

Are we (still) alone?

Are we (still) alone?

Once CERN was one hotspot of a kind for the use of ICT infrastructures and internet/storage traffic.

Today the Googles, Amazons, etc. are data analysis and storage powerhouses.

They make profit out of their proprietary tools, but they (mostly) make these tools commonly available.

Are we (still) alone?

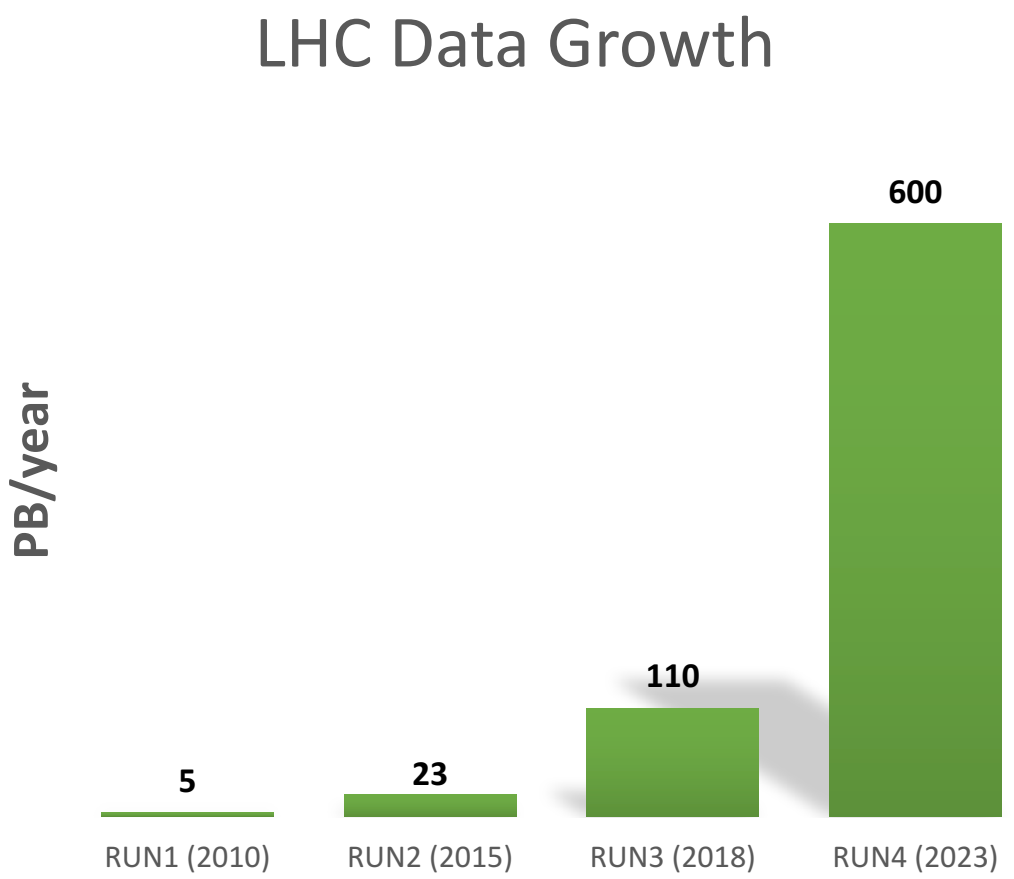

Are we (still) alone?

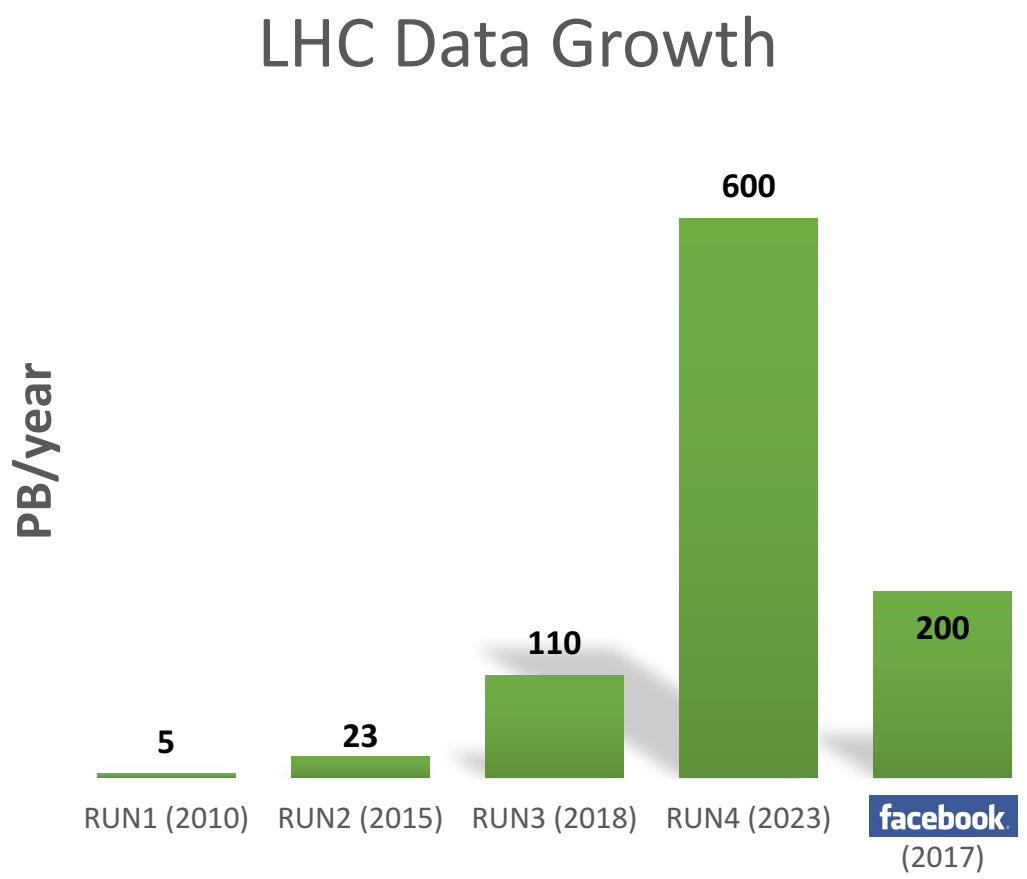

Are we (still) alone?

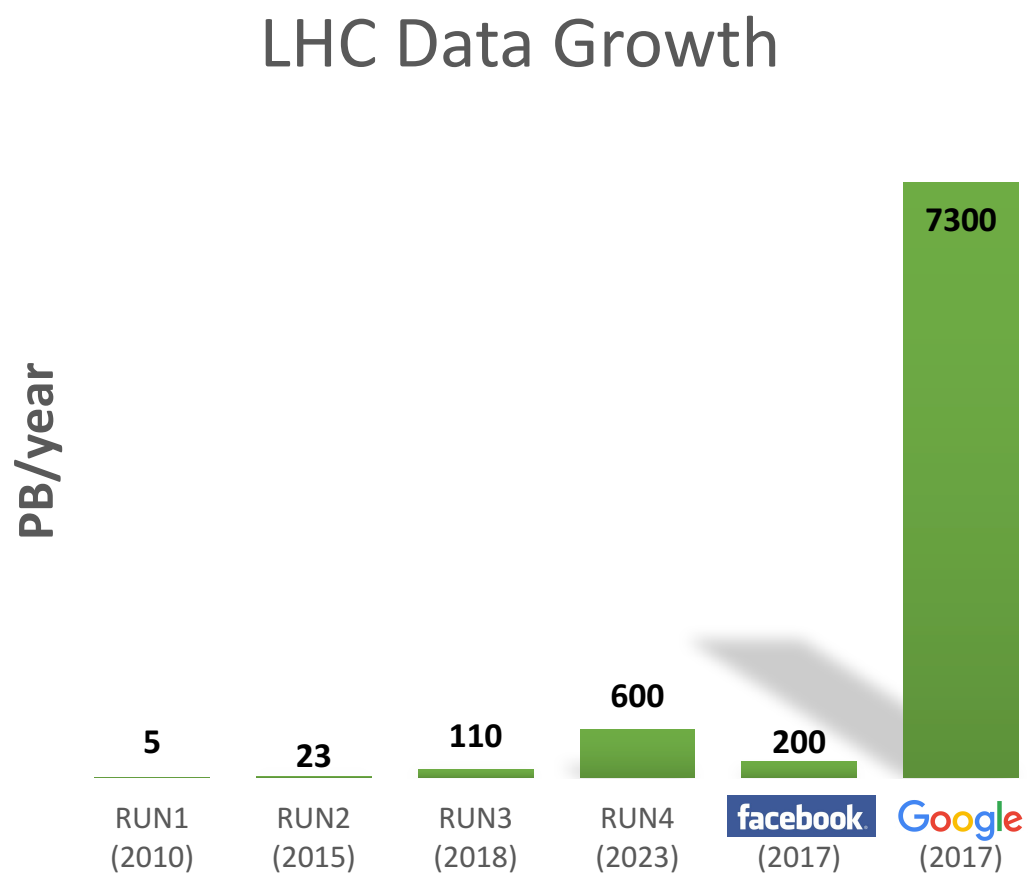

Are we (still) alone?

Data mining and machine learning are leading the development of tools for high performance computing over enormous datasets.

HPC of big datasets is the frontier CERN has always addressed.

The CERN requests and market offerings (both in terms of services and tools) are more aligned than ever.

Are we (still) alone?

Collaboration between Big Data companies and CERN is taking place with direct collaboration and adoption of tools such as:

openstack®

Cross platform data-center grade resources manager (IaaS).

OS-level virtualization. Containerisation (hence shipping) of single applications.

docker

MESOS

Cross-framework orchestrator and resources manager. Can run MPI, spark, hadoop (and many others) clusters.

Cloud storage with file and block level access. Used by EOS (hence CERNbox)

ceph

Are we (still) alone?

CERN is taking part in the development of tools with external Big Data companies.

A bigger user base enhances the testing and debugging capabilities.

Like never before CERN is accepting to avoid the wheel re-invention.
Adopting leading edge tools is easier and can lead to great results.

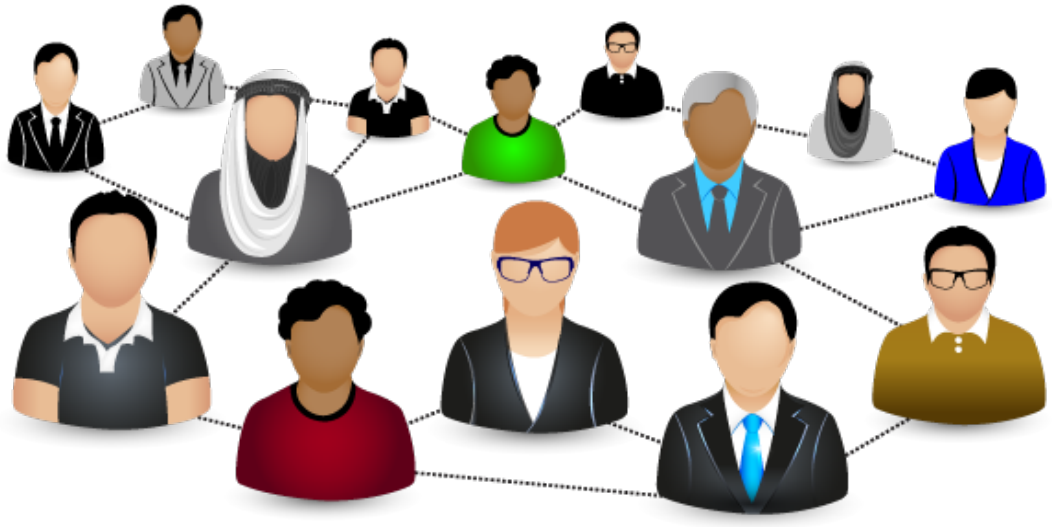

Are we (still) alone?

Not anymore!

ALICE's wonderland

Acquisition challenge

DAQ will still rely on local and private resources (at least in the near future).

A local cloud, isolated from the internet for security purposes, can serve well to dynamically allocate resources for the detectors and run parallel instances of the reconstruction routines.

ALICE O^2 project follows this pattern, introducing modern routines for DAQ and data reconstruction.

ALICE acquisition paradigm

ALICE

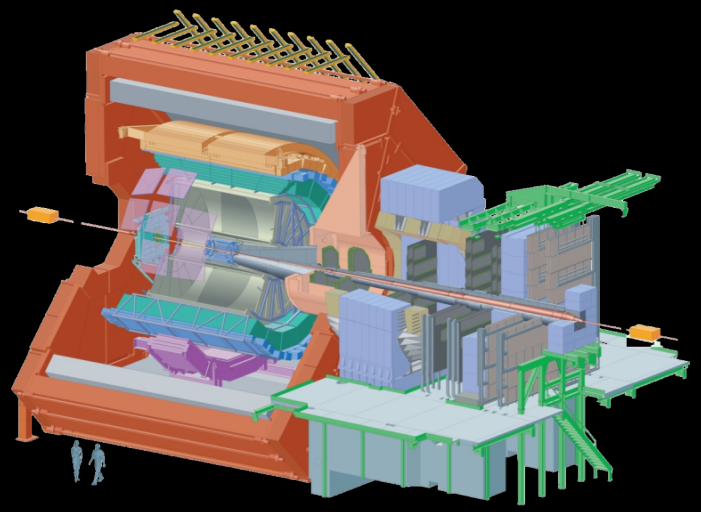

3 TB/s

Common
Readout Unit

90 GB/s

Storage

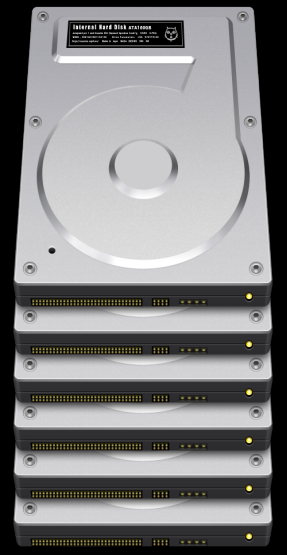

ALICE's present

RUN2

(2015-2018, present)

8 kHz PbPb @ 1 kHz w/ triggers

ALICE's future

RUN3

(2021-2024)

50 kHz PbPb @ 50 kHz triggerless

ALICE's old acquisition paradigm

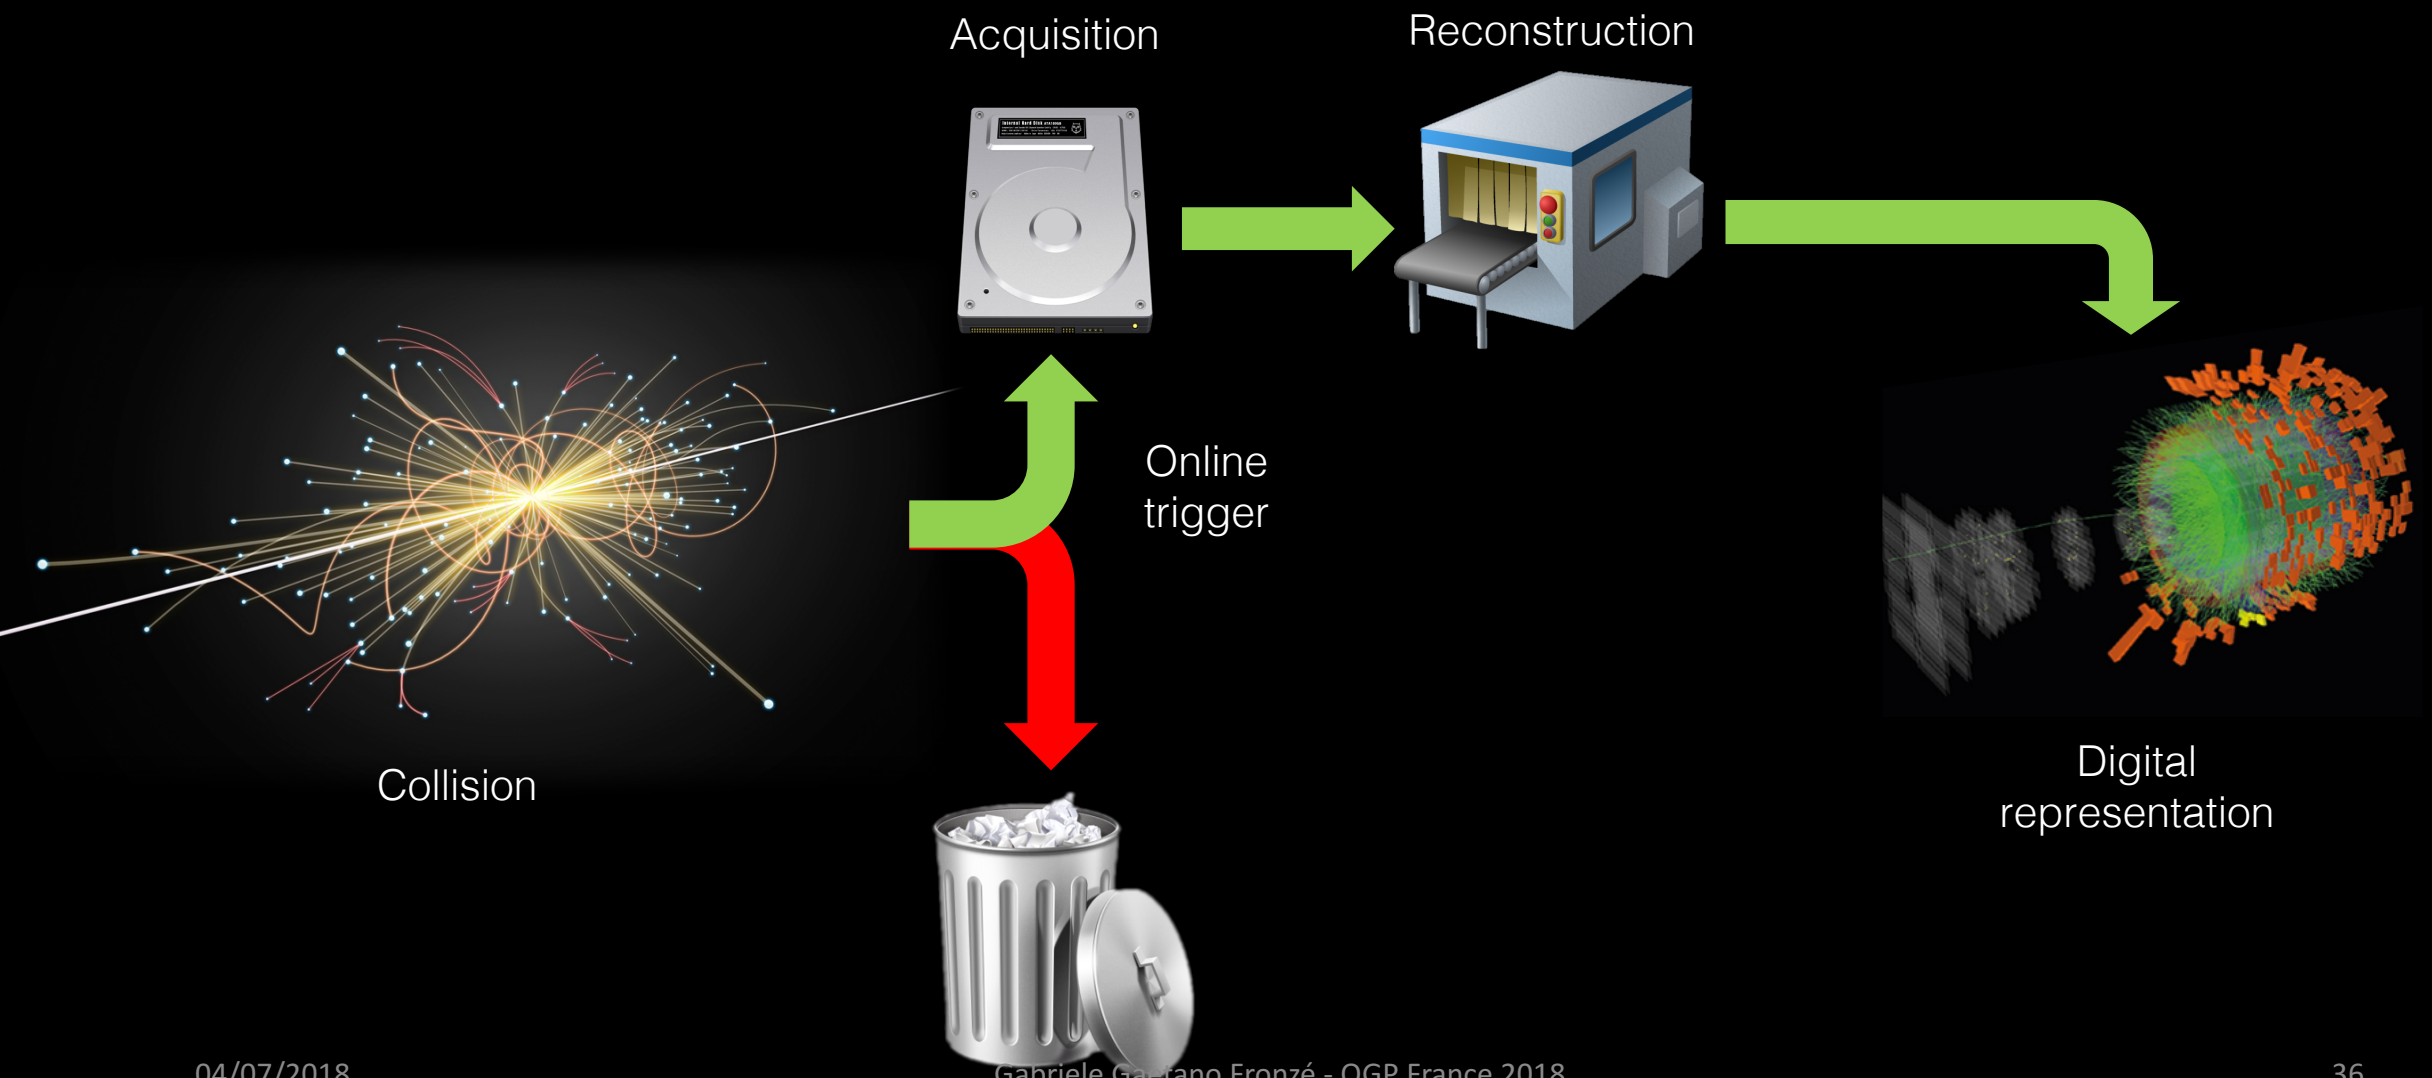

ALICE's new acquisition paradigm

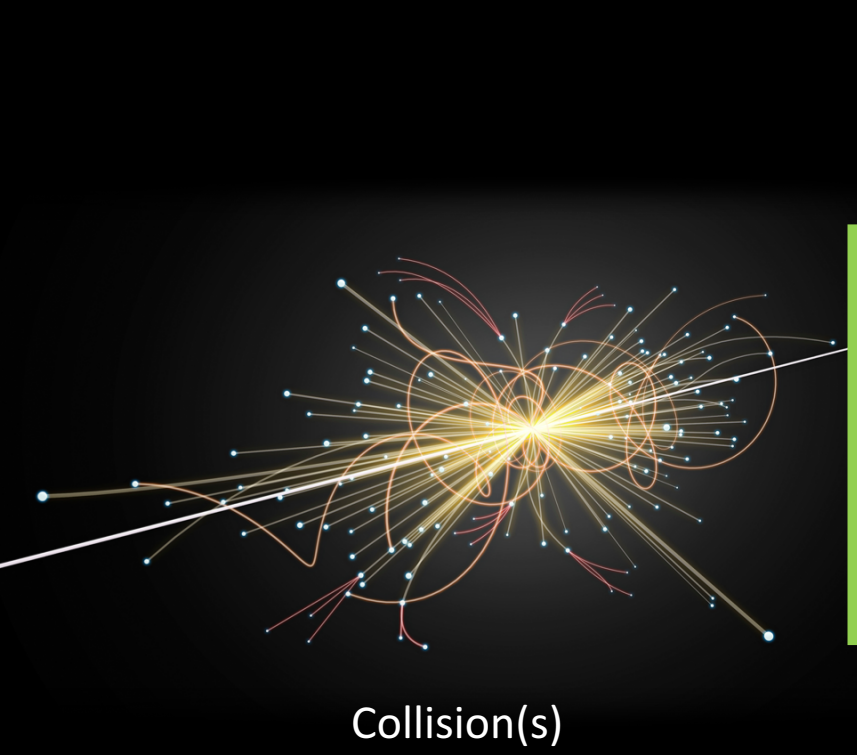

Collision(s)

Acquisition and raw reconstruction

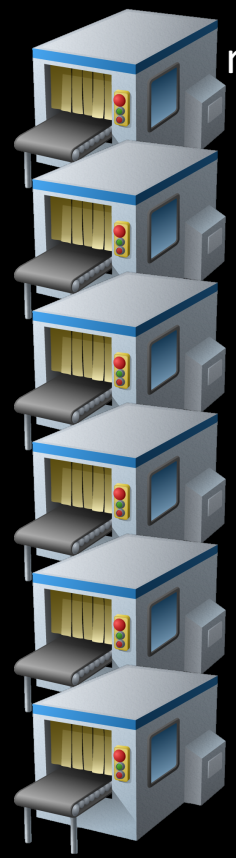

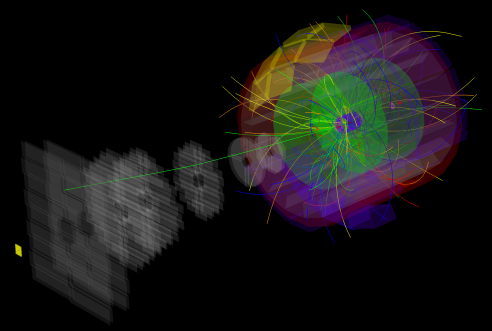

Raw reconstruction

Storage

Data based decision

ALICE's new acquisition paradigm

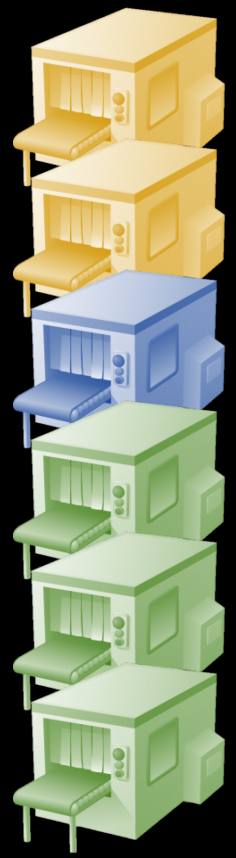

Detector A → Taking too much time

Detector B → Good performance

Detector A → Faster than needed (dead time)

ALICE's new acquisition paradigm

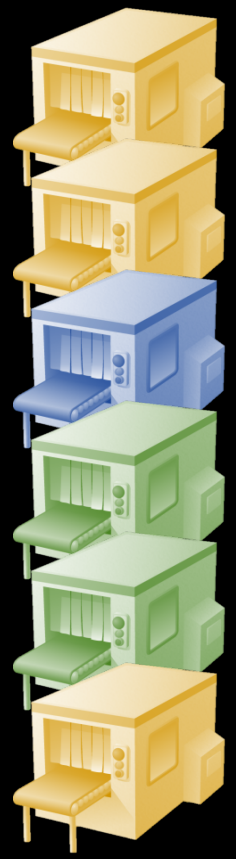

Detector A → Good performance

Detector B → Good performance

Detector A → Good performance

Detector A → Good performance

ALICE O^2

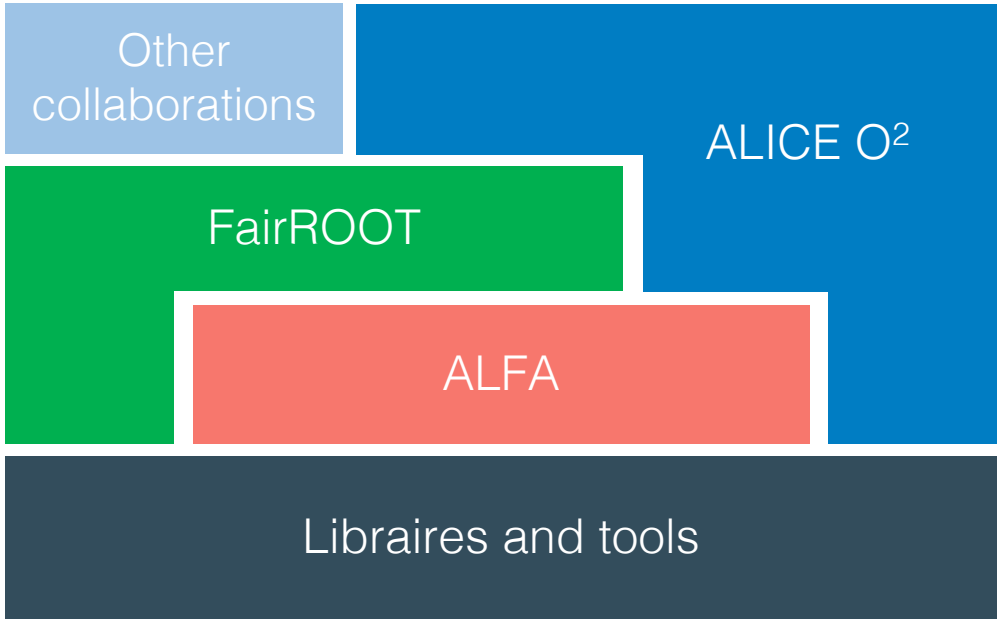

- Simulation and reconstruction developed by ALICE et Fair together
- Communication layer using ZeroMQ and nanomsg
- Common, expandable, base of libraries and tools

O^2 = libraries + optimisations + standards

ALICE O^2

Communication

- Language-agnostic
- Device-agnostic
- Simple interfaces

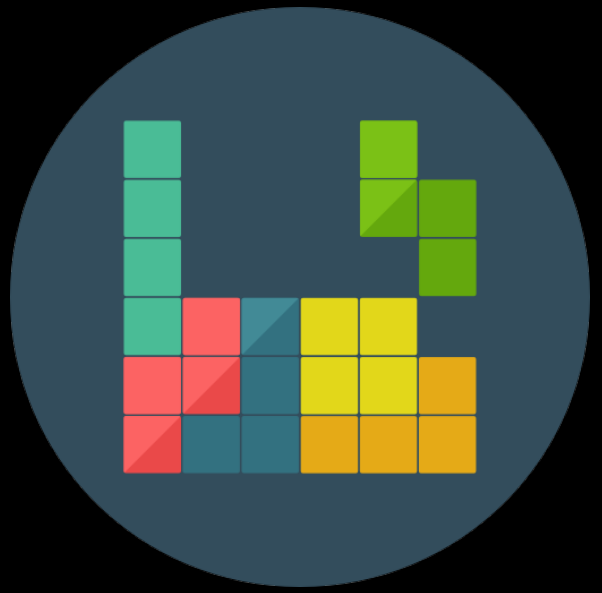

Modularity

- Expandable
- Reconfigurable
- Debug-friendly

Standard

- General approach to common problems
- Reusability of code
- Optimisation

Conclusions
and
Perspectives
On a trick coin history

Conclusions

We are undergoing a new era in scientific computation.

Large scientific institutions are not anymore the biggest producers of informatics data.

Collaboration between scientific institutions and (ICT) industry is possible and should be encouraged.

Powerful tools, able to handle (or to be repurposed for) the typical HEP data, are available to the public.

Perspectives

The availability and disposability of a cloud infrastructure is unrivalled.

The (growing) complexity of hardware solution can be hidden inside elegant and comfortable interfaces.

The future will present even harder challenges, lets try to avoid to reinvent the wheel each time!


That's all Folks!