Trends and strategies in the evolution of High Energy Physics computing

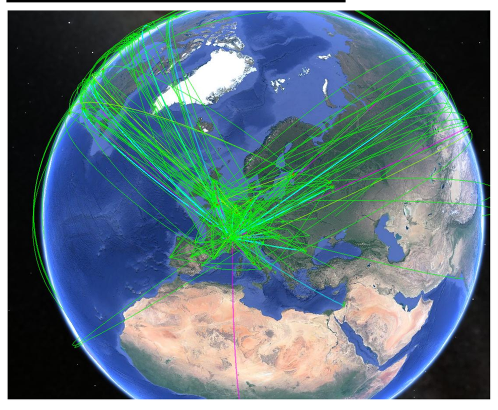

Run 1 and Run 2 of the Large Hadron Collider (LHC) at CERN produced huge volumes of data whose analysis will continue to deliver large numbers of physics results. To make this possible, the LHC experiments have been relying on services of the Worldwide LHC Computing Grid (WLCG).
At present, the LHC experiments are preparing for Run 3 of the LHC which will bring significantly higher luminosity and therefore yet higher volumes of data, which will even be exceeded in the era of the High Luminosity LHC (HL-LHC), starting with Run 4.
While during Run 1 and Run 2 the LHC experiments were the only ones to produce and analyze scientific data at a scale of hundreds of PetaBytes, the situation is gradually changing. Projects like DUNE in the USA, Belle II in Japan and SKA in Australia and South Africa will also be producing huge volumes of data and plan to make significant use of WLCG services to store and process their data. The LHC experiments will not necessarily remain the biggest scientific data producers in the future.
In this contribution we will present an overview of trends and strategies used by the LHC experiments to adapt their data processing models to future compute and storage resources available within WLCG and to the use of commercial clouds and High Performance Computing (HPC) facilities, while at the same time building collaborations with related big-data projects in order to share and evolve WLCG services together.

Dagmar Adamova (NPI AS CR Prague/Rez) and Maarten Litmaath (CERN)

Worldwide LHC Computing Grid (WLCG)

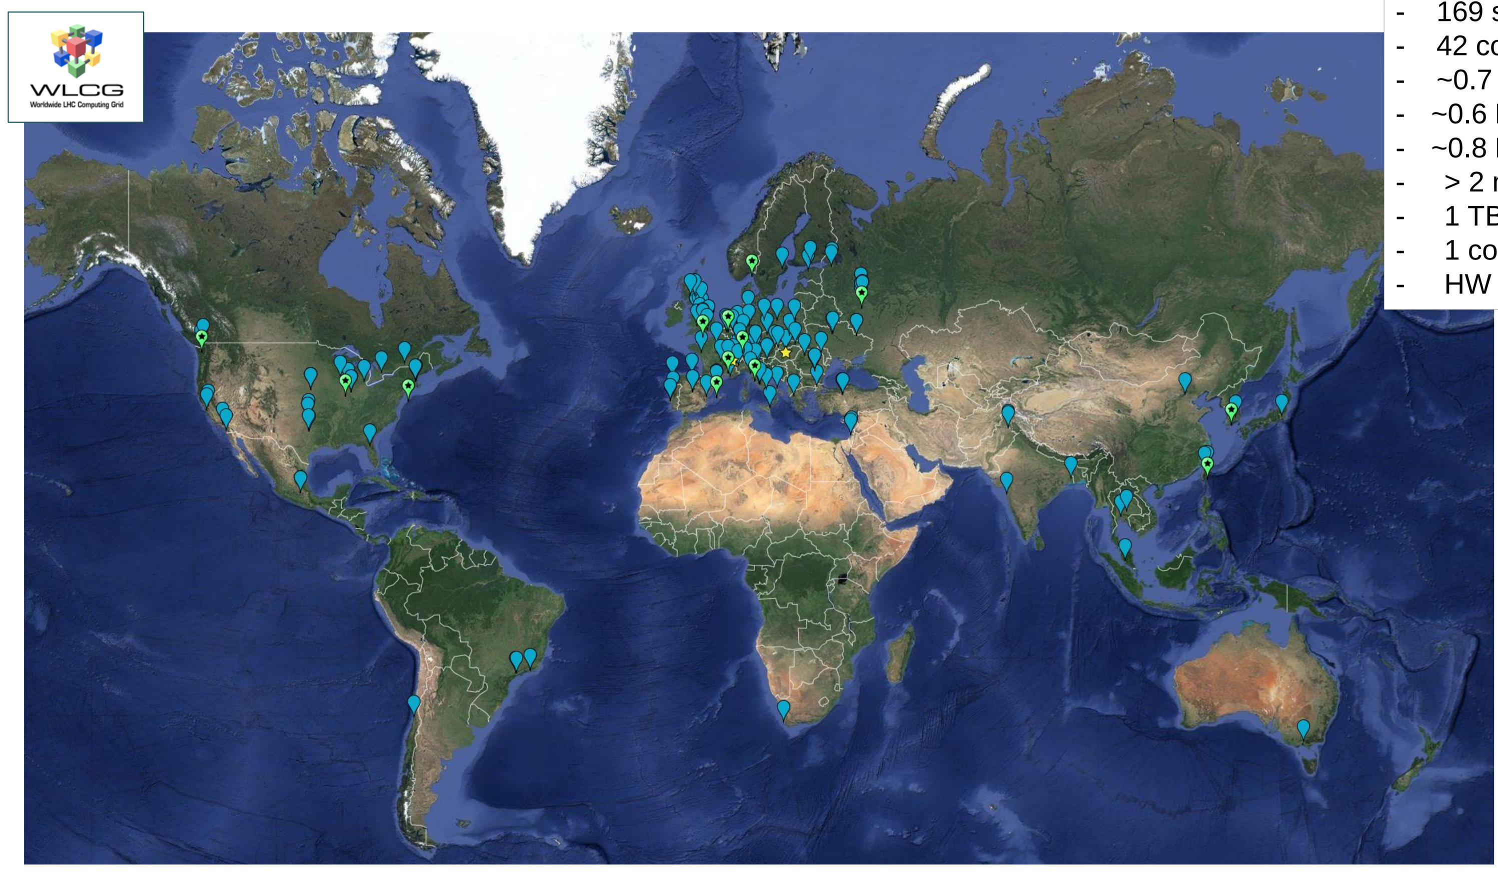

Facts about WLCG (Jan 2020):

- 169 sites
- 42 countries
- ~0.7 million CPU cores
- ~0.6 ExaByte of disk storage
- ~0.8 ExaByte of tape storage
- > 2 million jobs per day
- 1 TB ~ 10-100 CHF
- 1 core ~ 100 CHF
- HW lifetime: 3-5 years

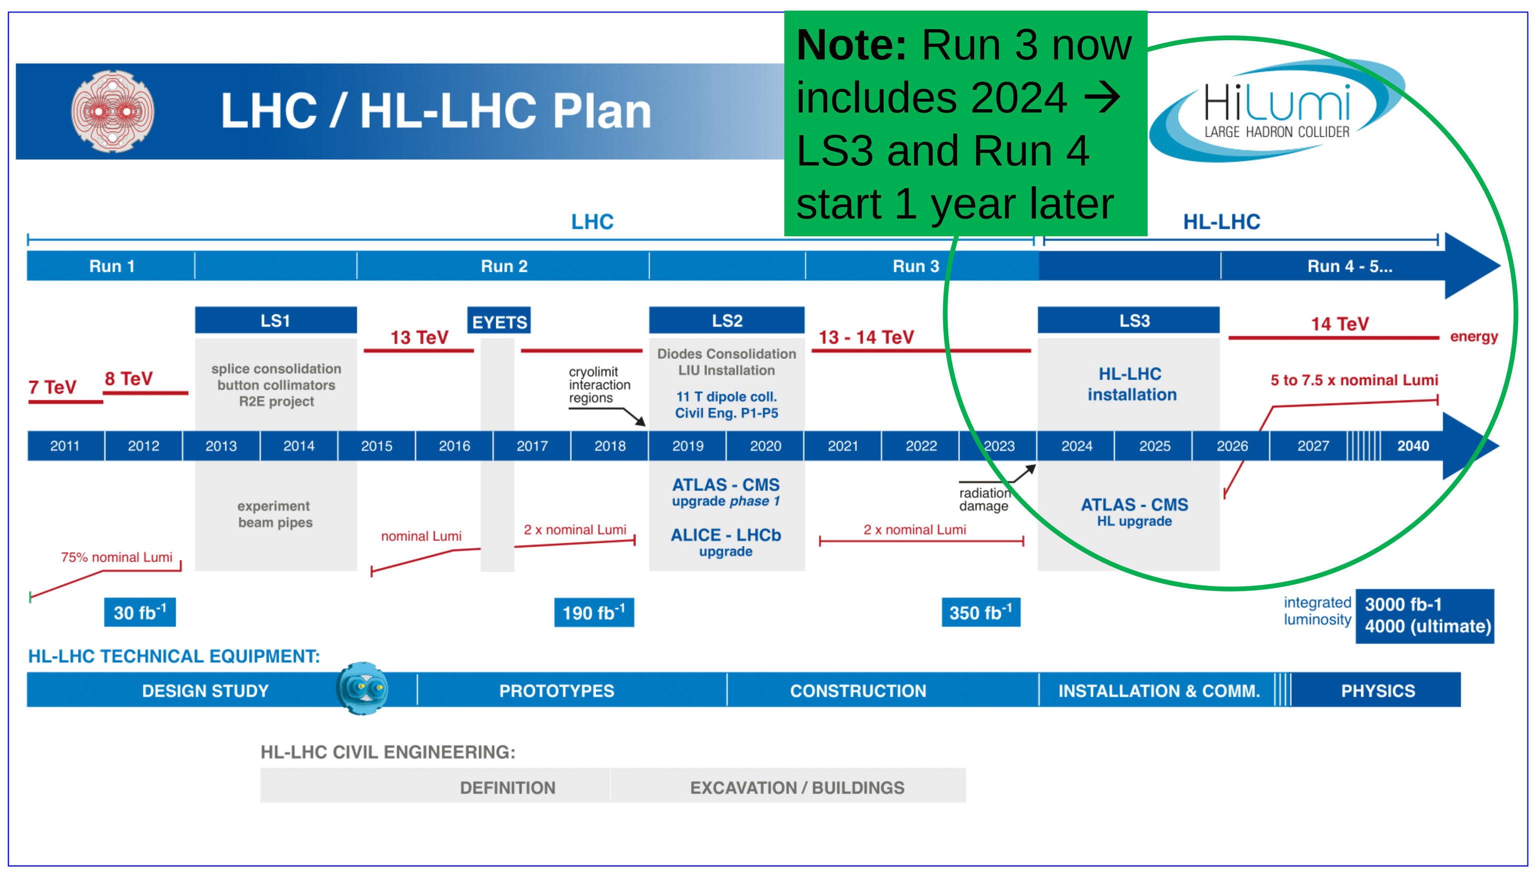

	LHC	HL-LHC
Pileup	~60	~140-200
Dataset	300/fb	3000-4000/fb
Instantaneous Lumi	~2x10^{34}	5-7.5x10^{34}

The scale of the computing challenge for LHC necessitated an organized and formal structure. A federated distributed system, the Worldwide LHC Computing Grid, was built to integrate pledged resources and make them easily usable.
The past 15 years of WLCG operation, from initial prototyping through to the significant requirements of Run 2, show that the community is very capable of building an adaptable and performant service. The WLCG and its stakeholders have continually delivered to the needs of the LHC during that time, so as to maximize the physics possibilities within a realistic and affordable budget envelope.

However, in the HL-LHC era the situation could be very different, unless there are some significant changes .

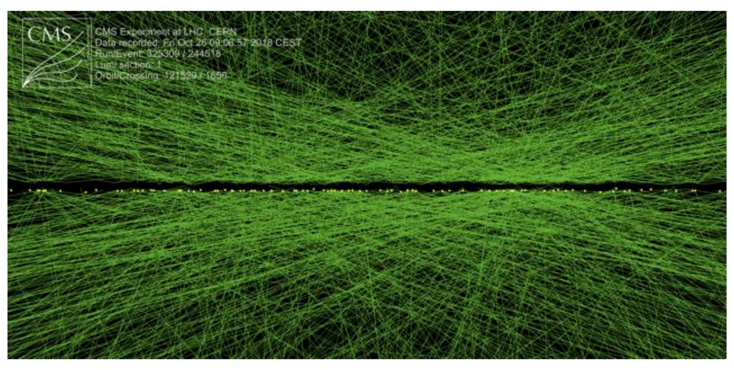

14 TeV center of mass energy - 6000 primary tracks per event
Simultaneous events (Pileup) increases from ~60 to 140-200

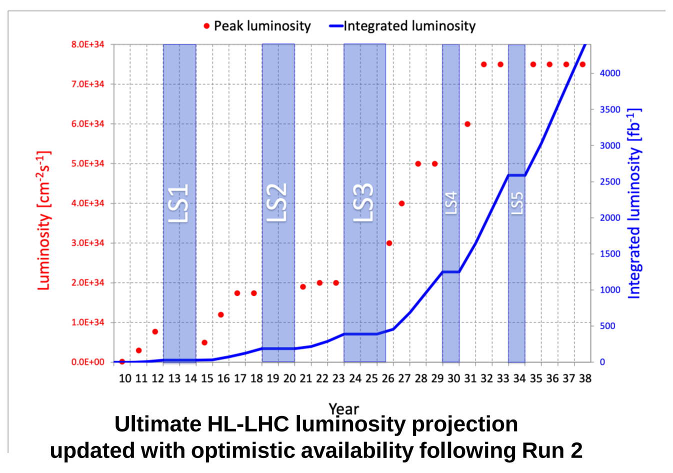

Ultimate HL-LHC luminosity projection
updated with optimistic availability following Run 2

Big challenges exist to meet the computing requirements for Run4.
The current resource provisioning model is not sufficient.
Technology and architecture changes are essential: $/resource unit is not a continuous and monotonic function, there may well be surprises.

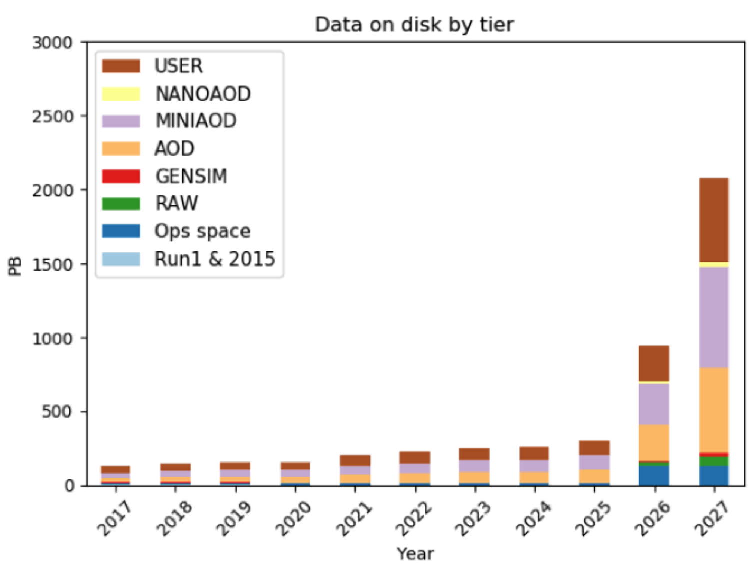

HL-LHC - Run 4: ATLAS and CMS resource challenges.
With respect to 2019 pledges, they need:
22x more CPU power, 13x more disk capacity and 15x more tapes.

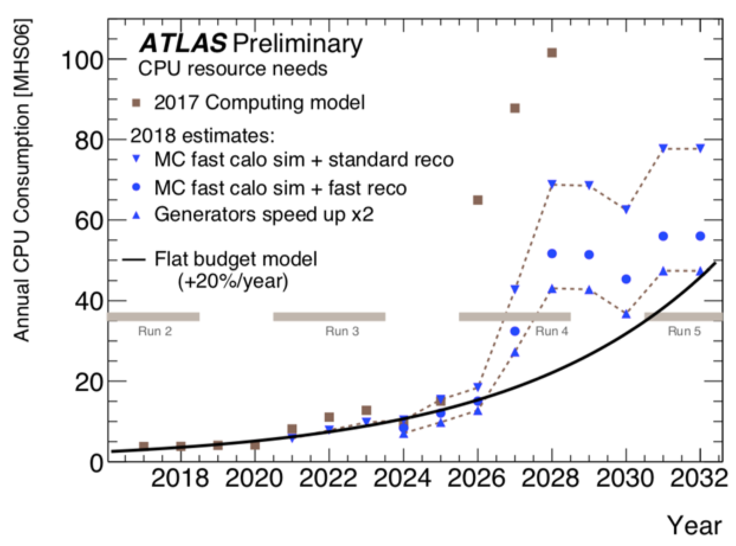

ATLAS: CPU projections for Run 4.
"flat budget" assumption is identified risk

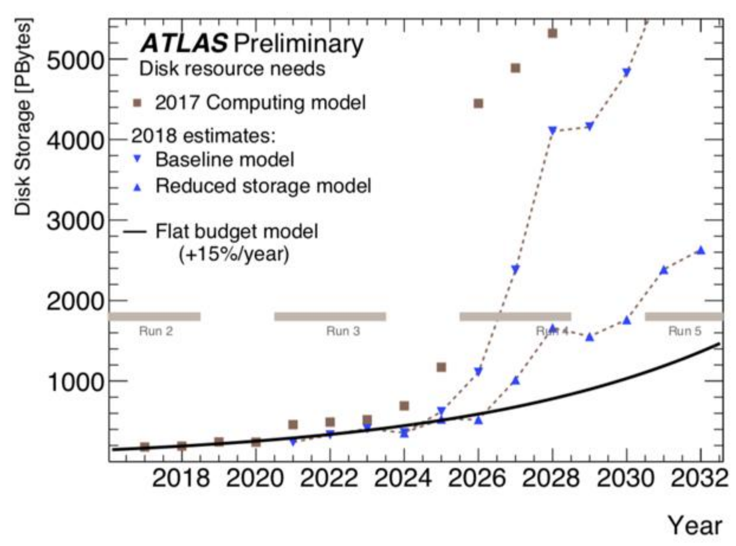

ATLAS: Disk storage projections for Run 4.
"flat budget" assumption is identified risk

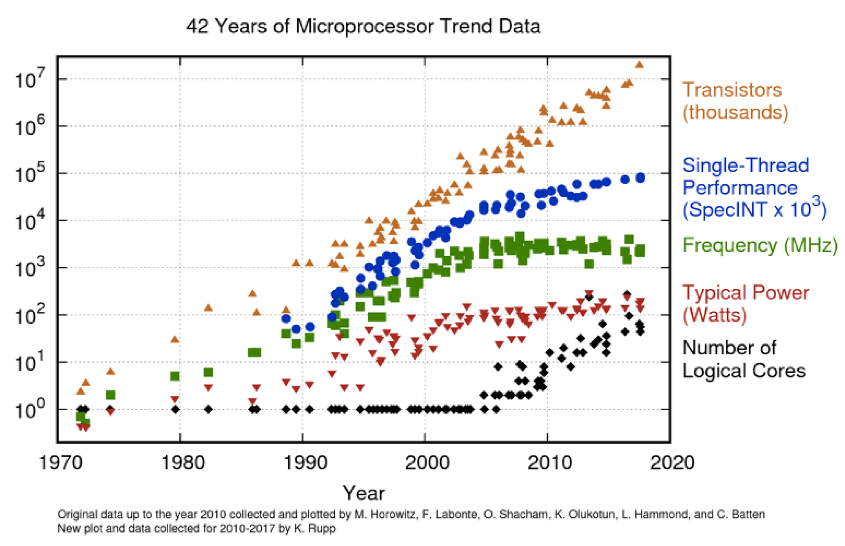

In general, it is necessary to continue modernizing our applications to take advantage of hardware evolution.

Currently recognized strategies and work packages to deliver solutions which would match the compute/storage requirements with available resources, see below.
Due to continuous, fast and unpredictable technology developments, the strategies will need to keep adapting to new circumstances. An example would be the use of commercial clouds and HPC resources, which still was unrealistic not so long ago.

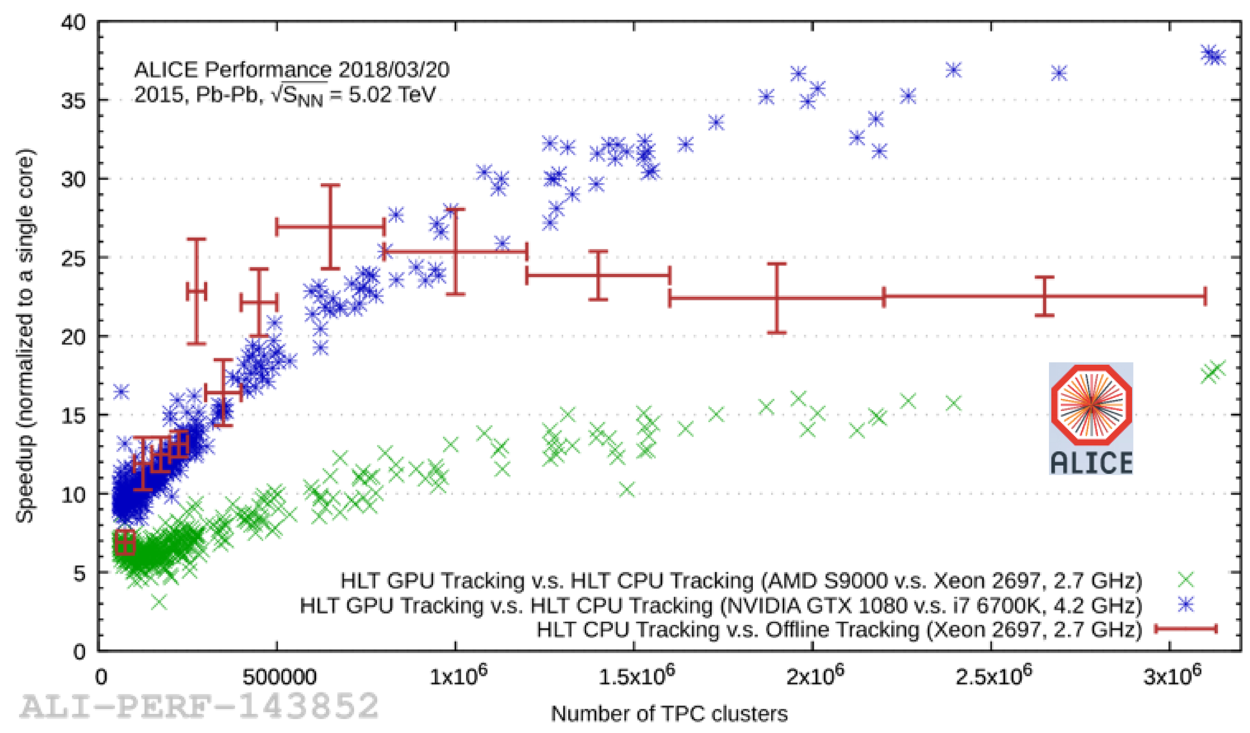

ALICE: Speed up tracking from GPU usage + from algorithmic improvements + tuning on CPU

- **Done: technical integration of cloud resources into WLCG systems. What remains is the TCO, viz. where and when clouds are more cost-effective than in-house resources.**
- **To be continued: integration of HPCs.**
- **Work to enable more use of GPUs.**
- **Strategy for WLCG data storage: create WLCG Data Cloud.**
- **Adopt newest network technologies to enable fast/cheap data streaming/transfers.**
- **HW trends/cost development unpredictable. Technology and architectures may change quickly and may dictate adjustments in WLCG.**
- **Work to enable use of modern CPU features.**
- **Simulation: use of fast simulation and Machine Learning algorithms and tooling.**
- **Reconstruction: improve data models/data structure.**
- **Analysis: use of techniques beyond ROOT.**
- **Generators: become more complicated with increasing pile-up, better performance needed.**
- **Collaboration with compute intensive experiments beyond LHC (Belle II, DUNE, FAIR, SKA, …) should be beneficial: different projects/sciences ←→ same infrastructure.**

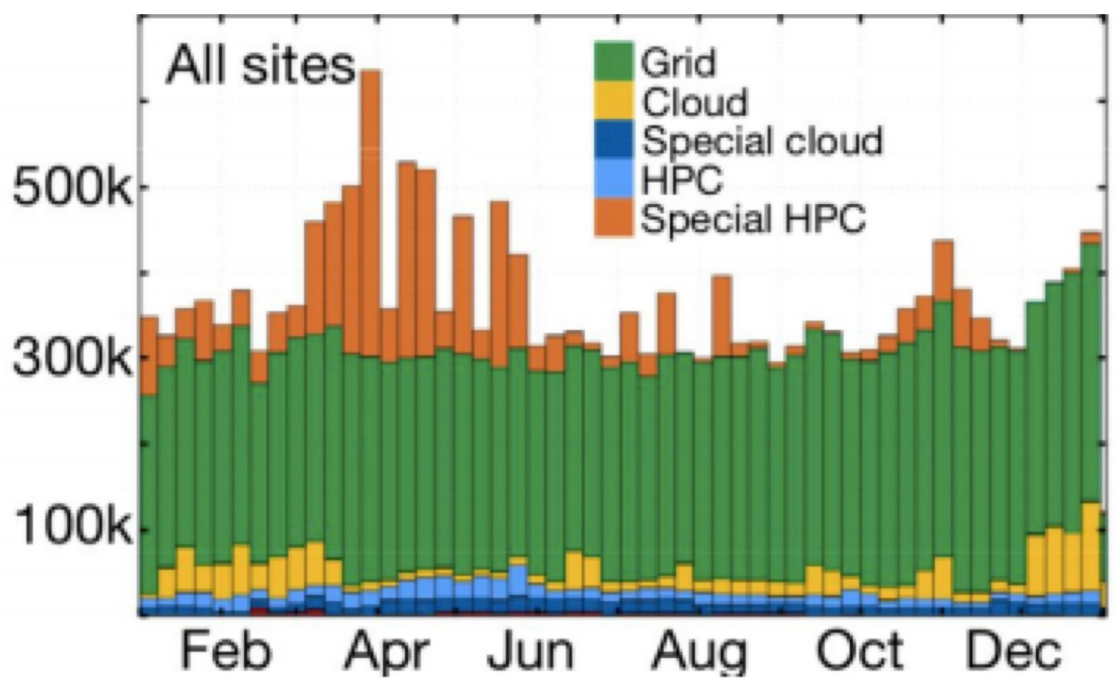

Up to 25% of resources currently used by LHC experiments comes from non-Grid facilities:
Cloud Computing, High Performance Computing, HLT farms.

ESCAPE (European Science Cluster of Astronomy and Particle physics ESFRI research infrastructures):
EU funded initiative. Prototyping new compute infrastructure for Exascale science projects.