New strategies of the LHC experiments to meet the computing requirements of the HL-LHC era

The performance of the Large Hadron Collider (LHC) during the ongoing Run 2 is above expectations both concerning the delivered luminosity and the LHC live time. This resulted in a volume of data much larger than originally anticipated. Based on the status of current data production levels and the structure of the LHC experiment computing models, the estimates of the data production rates and resource needs were re-evaluated for the era leading into the High Luminosity LHC (HL-LHC), the Run 3 and Run 4 phases of LHC operation.

It turns out that the raw data volume will grow ~10 times by the HL-LHC era and the processing capacity needs will grow more than 60 times. While the growth of storage requirements might in principle be satisfied with a 20% budget increase and technology advancements, there is a gap of a factor 6 to 10 between the needed and the available computing resources.

The threat of a lack of computing and storage resources was present already in the beginning of Run 2, but could still be mitigated e.g. by improvements in the experiment computing models and data processing software or utilization of various types of external computing resources. For the years to come, however, new strategies will be necessary to meet the huge increase in the resource requirements.

In contrast with the early days of the LHC Computing Grid (WLCG), the field of High Energy Physics (HEP) is no longer among the biggest producers of data. Currently the HEP data and processing needs are ~1% the size of the largest industry problems. Also, HEP is no longer the only science with large computing requirements.

In this contribution, we will present new strategies of the LHC experiments towards the era of the HL-LHC, that aim to bring together the desired requirements of the experiments and the capacities available for delivering physics results.

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

Performance of LHC and LHC Computing grid (WLCG) during Run 2

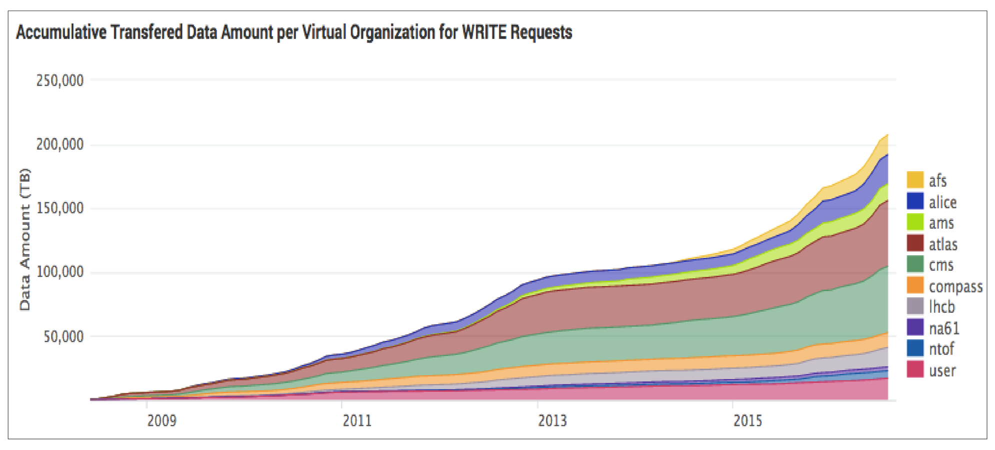

Tape archive of CERN physics data. ~160 PB of (raw) LHC data, ~ 500 million files.

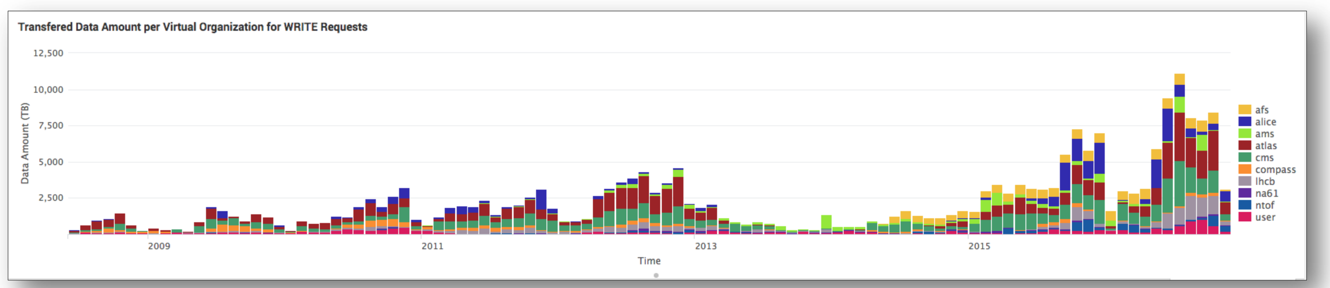

Data stored in the CERN tape archive, in TB/month. Significant rate increase in 2016. Peak >10 PB/month.

2016, the second year of Run 2, was a memorable year for the LHC and CERN. LHC performance was above expectations:

- Reached design luminosity and exceeded it by 40% ($L_{peak} = 1.4 \times 10^{34}\ cm^{-2}s^{-2}$).
- Stable Beams efficiency reached 50 - 60%.
- Quite impressive progression of the integrated luminosity: ~40 fb^{-1} for ATLAS and CMS (2016 target was 25 fb^{-1}).
- More collisions in 2016 than in all previous years together.

Data production in 2016:

- Raw: ~50 PB
- Derived (1 copy): ~80 PB

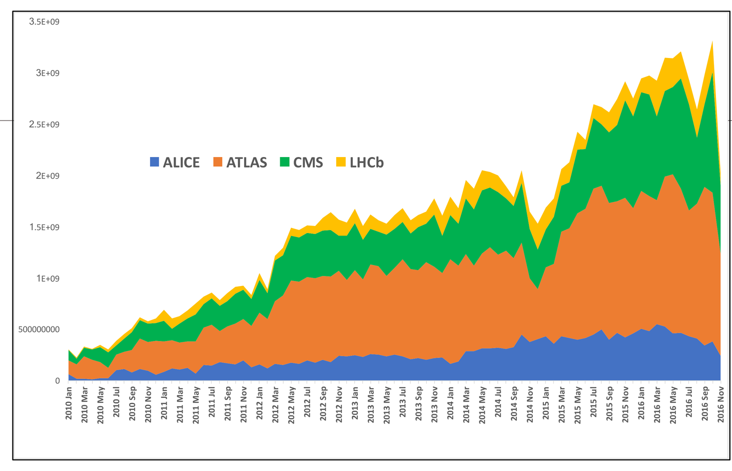

CPU usage within WLCG since 2010, in units HepSpec06. Significant increase in 2016: up to 630000 active cores.

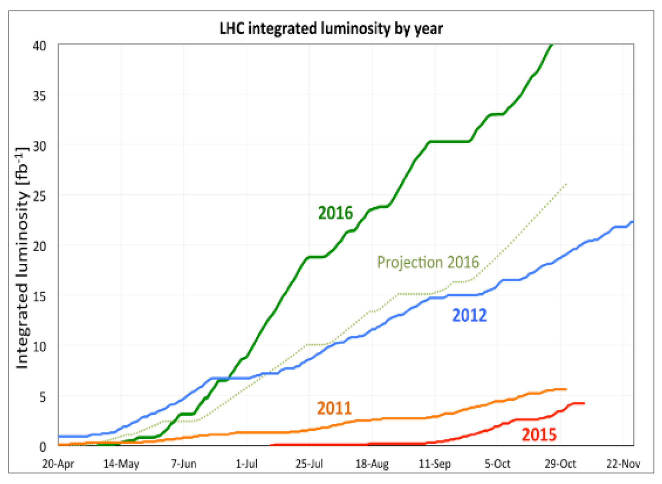

The graph shows the integrated luminosity delivered by the LHC to the ATLAS and CMS experiments in 2011, 2012, 2015 and 2016. The integrated luminosity indicates the amount of data delivered to the experiments and is measured in inverse femtobarns. One inverse femtobarn corresponds to around 80 million million collisions.

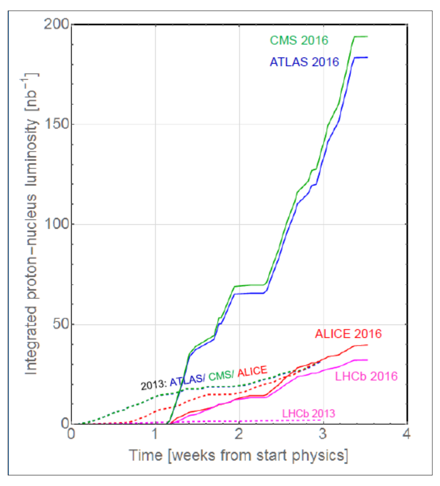

Integrated luminosity delivered to the LHC experiments during 4 weeks of proton – ion run:

- Performance of LHC and ion injectors was well beyond expectations.
- All physics goals were met and in some cases largely exceeded.

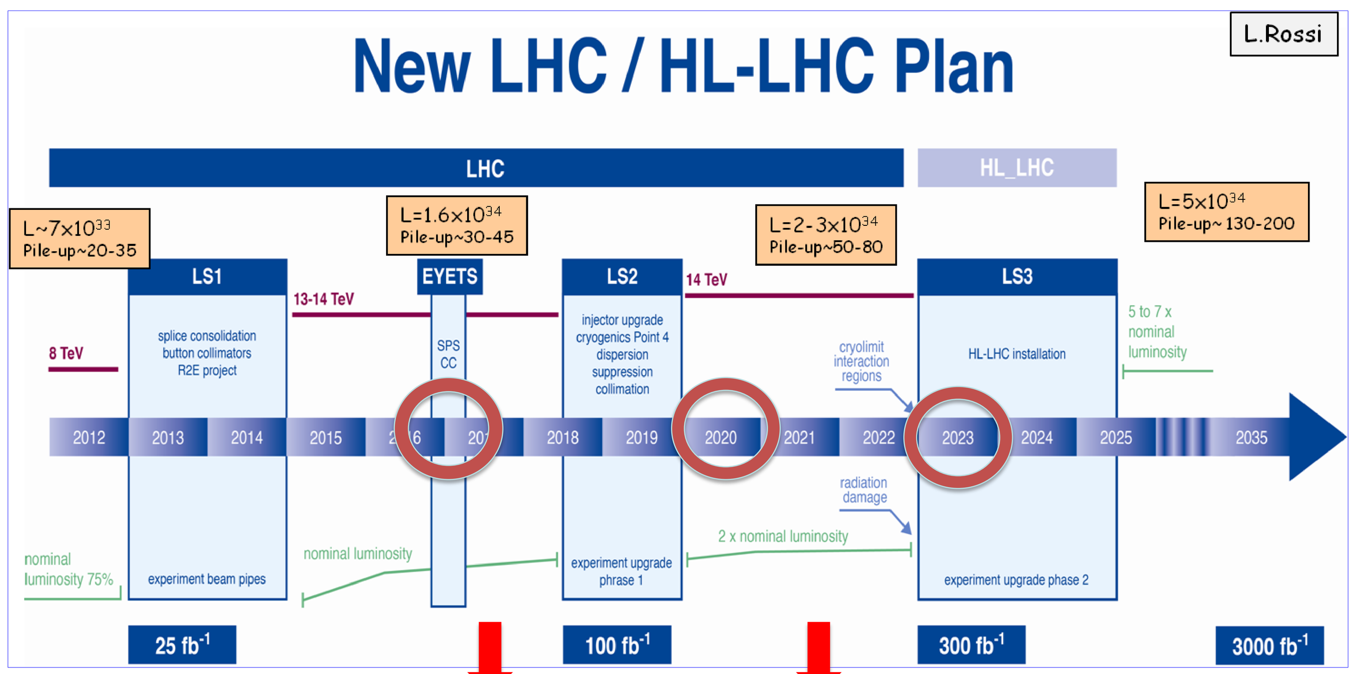

6.5 to 7 TeV
Peak luminosity 1.6-1.7x10^{34}cm^{-2}s^{-1}

7 TeV
Peak luminosity 2-3x10^{34}cm^{-2}s^{-1}

Upgrade ALICE + LHCb for Run 3
→75 Gbyte/s for 6 weeks/year

Upgrade CMS + ATLAS for Run 4
→10-40 GByte/s data rates in p-p run

Outlook: HL-LHC might/will create a computational problem

LHC performance is excellent and expected to improve. LHC and detector upgrades will bring larger luminosity, higher pile-up, more channels, higher event complexity, more complex simulations and all this translates into a massive growth of data to be processed and stored. The estimates of necessary growth of resources vary: 65 – 200 times larger needs than for LHC Run2. But experiments budgets are expected to stay constant (at least). Without significant efforts in optimisation, computing could risk to become a bottleneck for LHC science. Both experiments and CERN IT are increasingly held accountable for science throughput per investment by funding agencies.

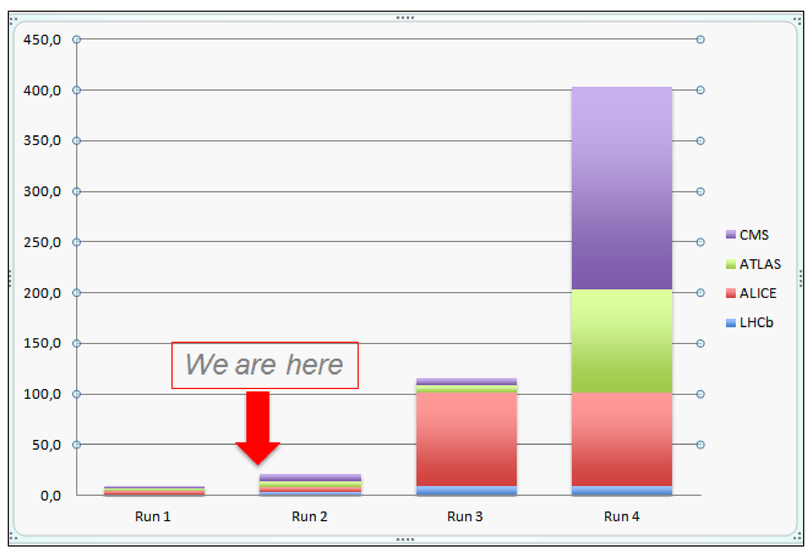

Very rough estimate of new RAW data per year of running using a simple extrapolation of current data volume scaled by the output rates (in PB, follow-up of 2015).
To be added: derived data (ESD, AOD), simulation, user data…

Estimates of resource needs for HL-LHC

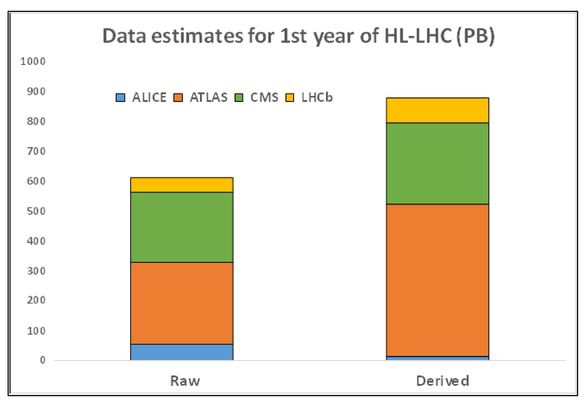

Data:
Raw 2016: ~ 50 PB → 2027: 600 PB
Derived (1 copy): 2016: ~80 PB → 2027: 900 PB

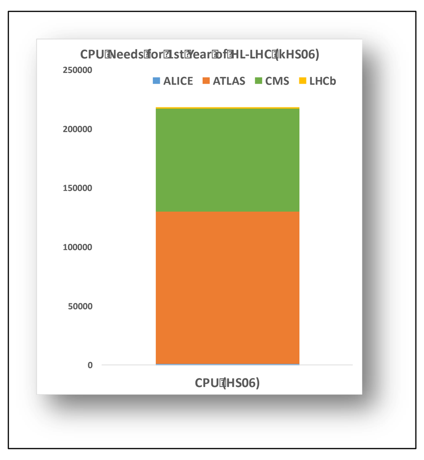

CPU: 60 times upscale from 2016 needed

These estimates are based on today's computing models, but with expected HL-LHC operating parameters. They are very preliminary and may change significantly over the coming years.

Technology development at ~20%/year will bring growth of only 6-10 in 10-11 years.

Commercial providers of computing services run infrastructures (computing clouds) much larger than HEP

There are various data driven projects outside HEP

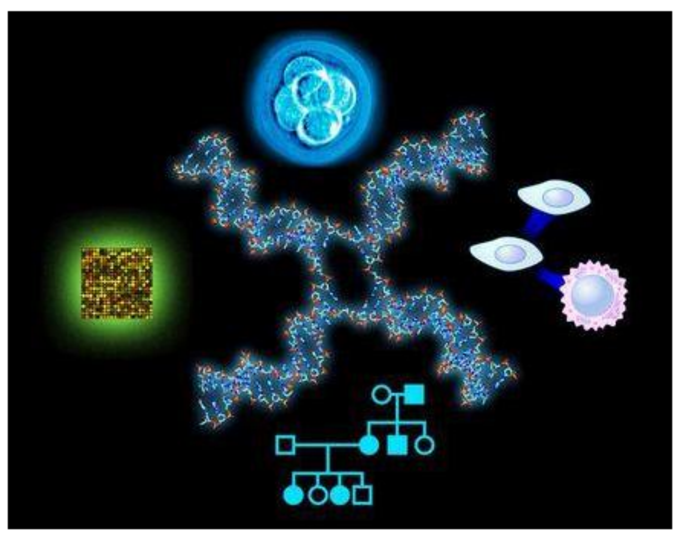

Example of non-HEP data driven scientific discovery:
human genome = 1 PetaByte. Finding patterns in clinical and genome data at scale can help cure cancer and other diseases

Nowadays scales are significantly different than in the first years of building the computing grid for LHC (WLCG). The enormous outburst of modern technologies caused a hectic development of data driven activities. Both industry and commerce enormously increased their demands for IT services which were met by the providers of cloud computing services. The providers dominating the market manage infrastructures up to orders of magnitude larger than the resources within WLCG. By charging for services, unlike the WLCG, the cloud providers reached remarkable revenues. For the HEP community, the standard commercial cloud services are still too expensive to be used on regular basis.

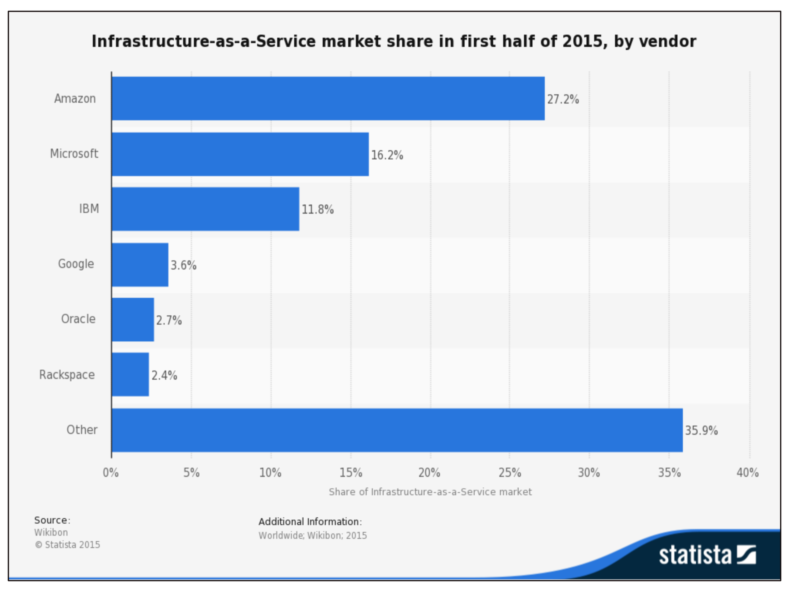

Comparison of the largest cloud providers by the delivery of Infrastructure as a Service (IaaS). Amazon, Microsoft, IBM and Google dominate the market.

Processing/CPU:
Amazon has more than 40 million processor cores in its Elastic Compute cloud, EC2
Google has ~1million servers and ~20 million cores
WLCG typically has ~500 thousand physical cores in use, which is ~1% of Amazon EC2 capacity

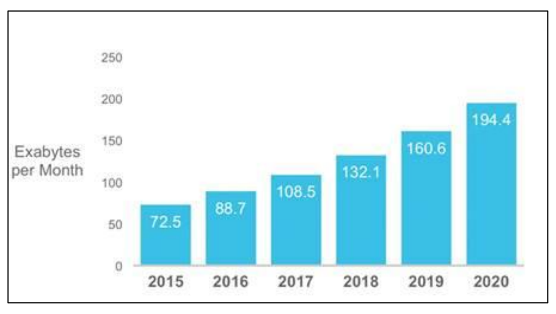

Forecast of CISCO Systems for worldwide data traffic rates, in ExaBytes (= 1000 PB) per month. Expect 194 EB per month by 2020.
WLCG is good at managing transfers of large volumes of data, reached a peak of ~100 PB/month in 2016.

Storage
Amazon supports millions of queries per second
Google has 10-15 thousand PB under management
Facebook manages 300 PB which roughly corresponds to what is in WLCG
eBay collected and accessed the same amount of data as LHC Run 1

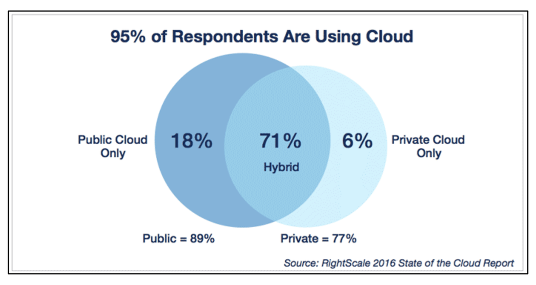

Survey of the latest cloud computing trends in 2016. Respondents: 1060 IT professionals, ~half of them represented enterprises with more than 1000 employees.

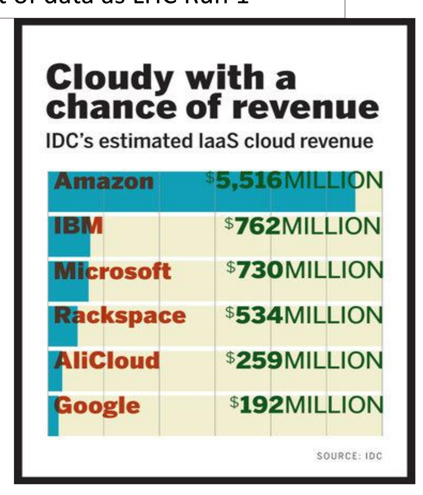

In addition to the revenues from IaaS, providing also other types of cloud services makes the revenues of Amazon ~$7.8 billion, Microsoft ~ $11.6 billion and IBM ~$4.5 billion.

Current status and strategies of WLCG

Run 2 has set new records in all areas of performance and capabilities.
The WLCG and experiments have managed the challenges successfully.
Continuous efforts in various areas like improvements of experiment software performance, changes of computing models, adoption of new data management strategies, upgrades of network capacity or utilization of external resources – high performance computing centers (HPC) and computational clouds – make it possible to adapt the system to the constantly evolving demands on performance.

CERN organizes ongoing procurements to use services of different cloud providers at least for particular LHC data processing use cases. Such a schema looks likely to be adopted for the future LHC runs, while usage of commercial cloud services on a very large scale is not yet anticipated.

New Computing TDR is under preparation for 2020

It will be a checkpoint on a roadmap – the current working system will always continue to evolve.
In addition to technological challenges, there are also sociological aspects: without recognizing computing and software activities as being equally important as analysis in physics careers, the amount of effort available for computing matters might become a problem for LHC science.

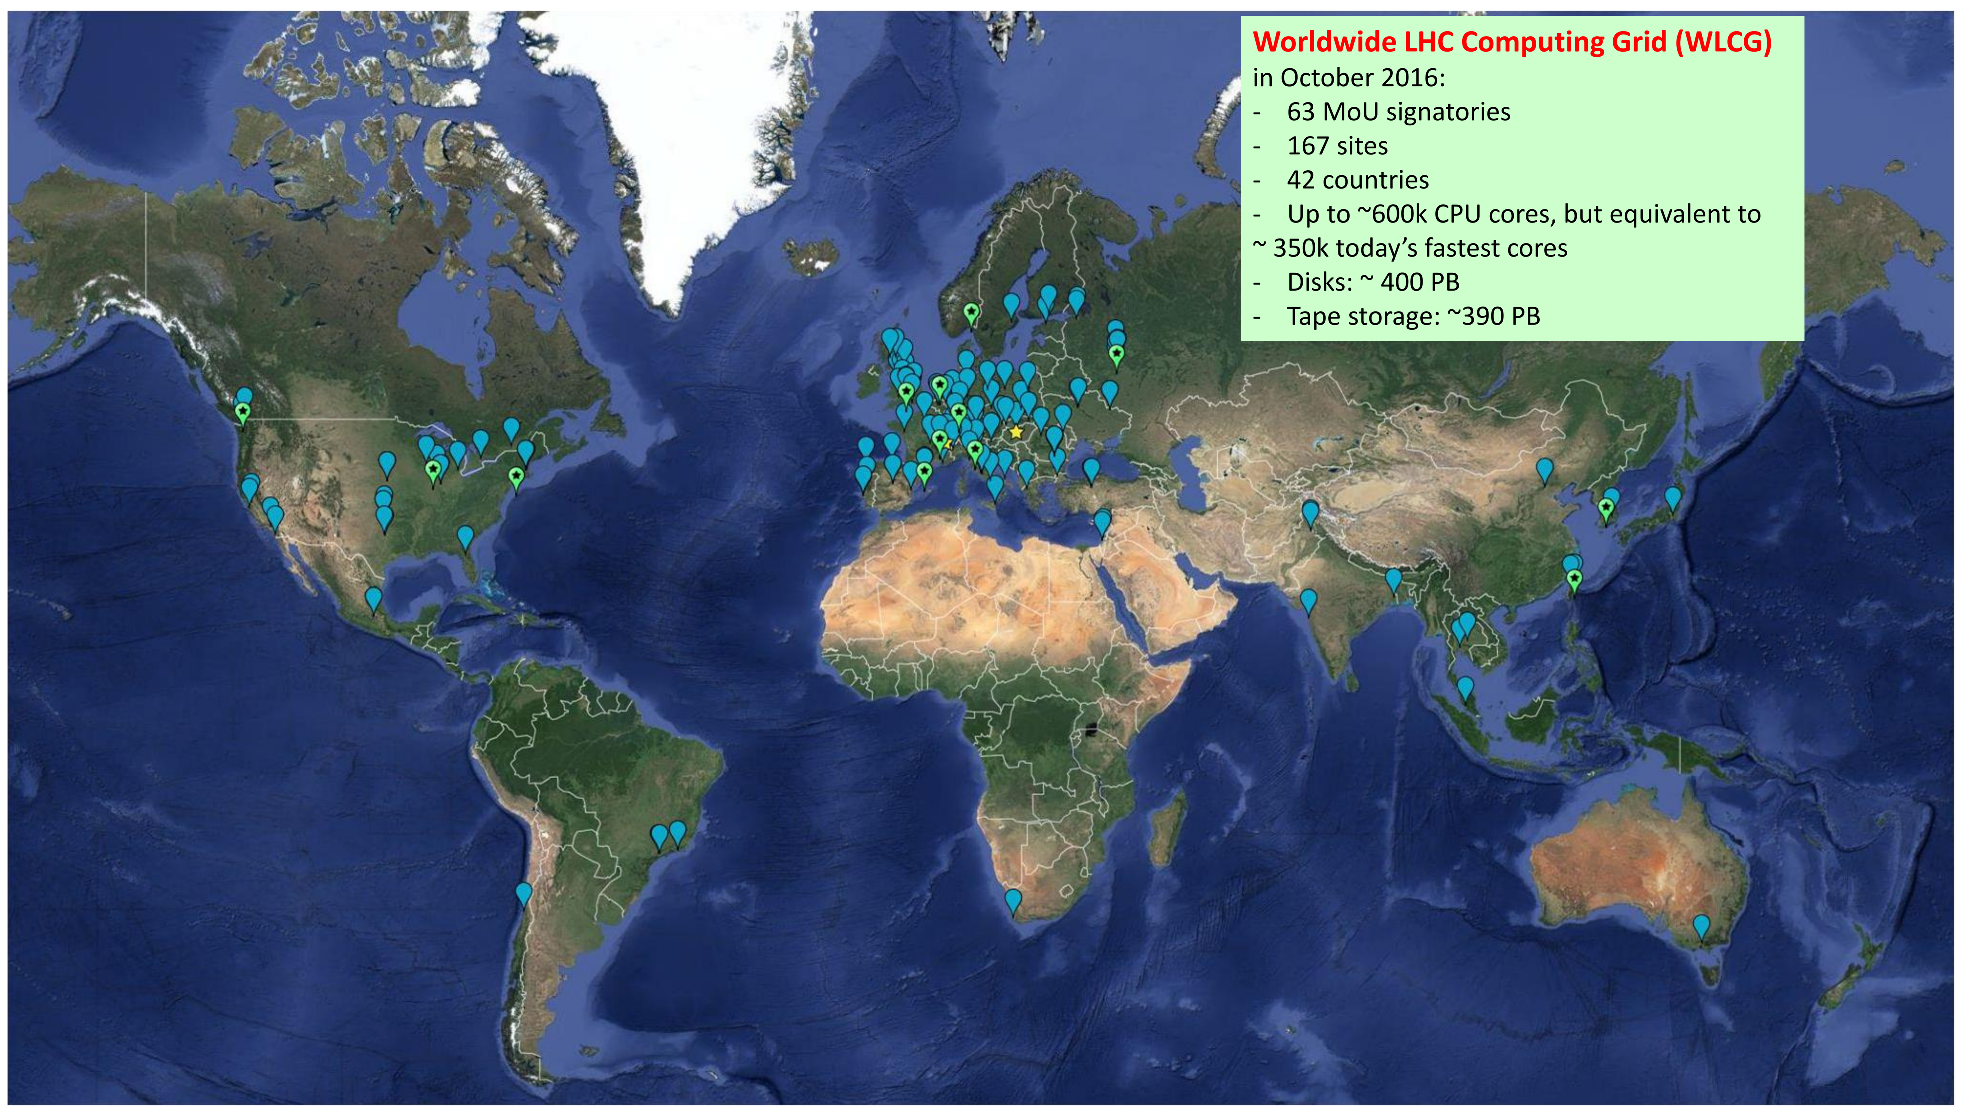

Worldwide LHC Computing Grid (WLCG)
in October 2016:

- 63 MoU signatories
- 167 sites
- 42 countries
- Up to ~600k CPU cores, but equivalent to ~ 350k today's fastest cores
- Disks: ~ 400 PB
- Tape storage: ~390 PB

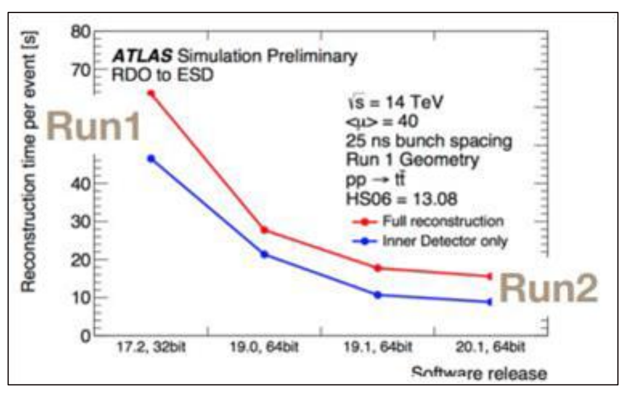

ATLAS event processing performance improved significantly since Run 1

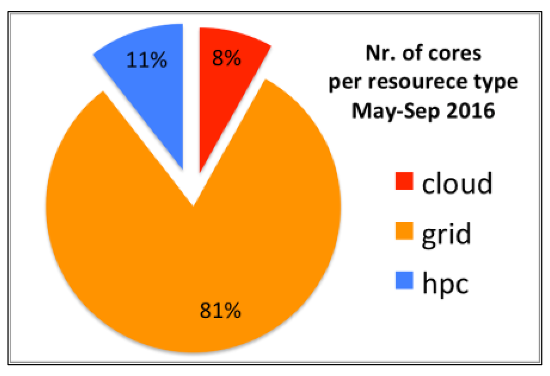

Usage of external resources for LHC data in 2016