The latest developments in preparations of the LHC community for the computing challenges of the High Luminosity LHC

After the LHC community successfully completed Run 1, the capacity of the LHC global computing infrastructure, the Worldwide LHC Computing Grid (WLCG) became tight for processing of ever growing volumes of data produced in the LHC collisions. During the last five years the LHC community launched a number of activities to increase computing performance and optimize usage of available resources because technology evolution may fall short by up to a factor 10 with respect to WLCG needs. These endeavours are in constant development and concentrate e.g. on re-structuring of computing models of the LHC experiments, to improve the efficiency of the data processing chain, adaptation of software to fast and/or cheap CPU architectures and attempts to use diverse resources like private and public clouds or high performance computing facilities.
WLCG was built in times when there were no examples of such infrastructures in industry or elsewhere. This situation changed during Run 1 when the global internet and computing industry began to provide on-demand services and developed tools and solutions which are of interest also for WLCG.
The latest strategies to increase the WLCG performance are also oriented to adoption and utilization of tools provided from outside the LHC community. Use of commercial cloud services for LHC data simulation and processing is adopted across the experiments, open source tools are employed for managing computing resources, and the formation of *data lakes* is envisaged. The process of software transformation to higher efficiency and adaptation to new technologies is orchestrated by the HEP Software Foundation. The need for external expertise and essential changes was recognized. These include e.g. adopting Python as a programming language for high-level analysis with notebooks, or further integration of Machine Learning in reconstruction. There is an effort to find and support commonalities across experiments, increase collaboration with other big-data projects like the Square Kilometer Array (SKA) and expand collaboration with IT industry through CERN openlab.
In this contribution we will present the latest strategies of the LHC community to get prepared for the computing challenges of the next ten years and especially for the demands of the High Luminosity LHC.

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

1. Performance of LHC and LHC Computing grid (WLCG) in 2017

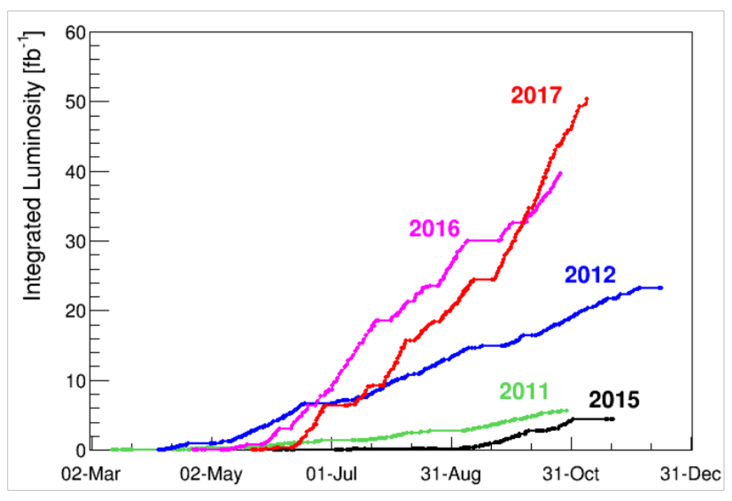

Integrated luminosity delivered by the LHC over the years of Run 1 and Run 2, in inverse femtobarns (fb^{-1}). 1 fb^{-1} corresponds to around 80 million million collisions.

2017 was the best year so far in terms of integrated luminosity
-- Exceeded 2017 goal of 45 fb^{-1} : Record: 50 fb^{-1} in one year
-- Good performance allowed to shorten 2017 run by one week
-- To the end of Run 2 LHC could produce ~5 fb^{-1} per week

95 fb^{-1} produced in Run 2 for ATLAS/CMS
123 fb^{-1} produced since start of LHC
Peak luminosity: Leveling at 1.5×10^{34} cm^{-2} s^{-1} to keep the pileup at the experiments at a manageable level (~57)

The amount of raw LHC data taken in 2017 was ~ 40 PB. It is less than in 2016 due to the later start of data taking and due to the absence of a several weeks heavy ions run.

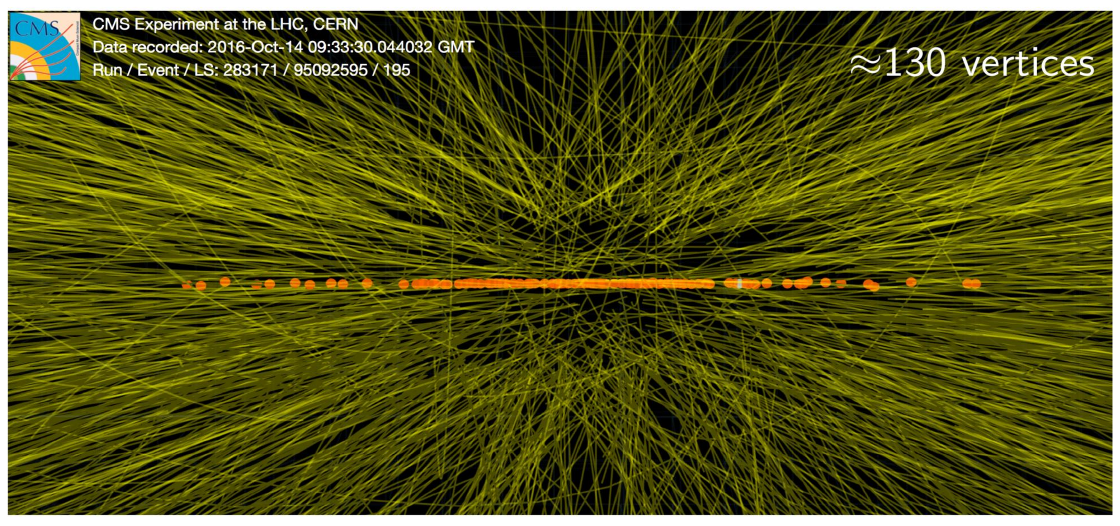

Real-life event with HL-LHC-like pileup from one special run in 2016, with individual high intensity bunches. Pileup μ ~ 130, in contrast to μ ~ 50 – 60 reached in 2017.

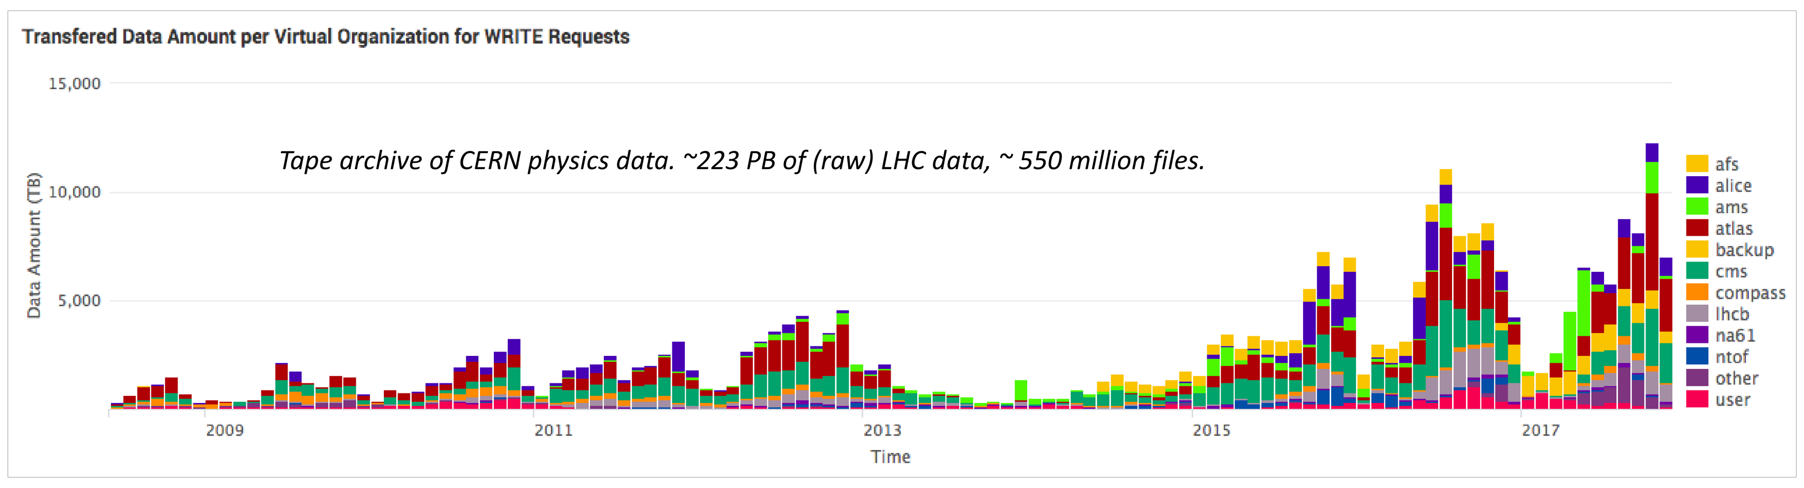

Data stored at the CERN tape archive by individual experiments, in TB/month. Distinctive increase in 2017 with a peak ~12 PB/month in October.

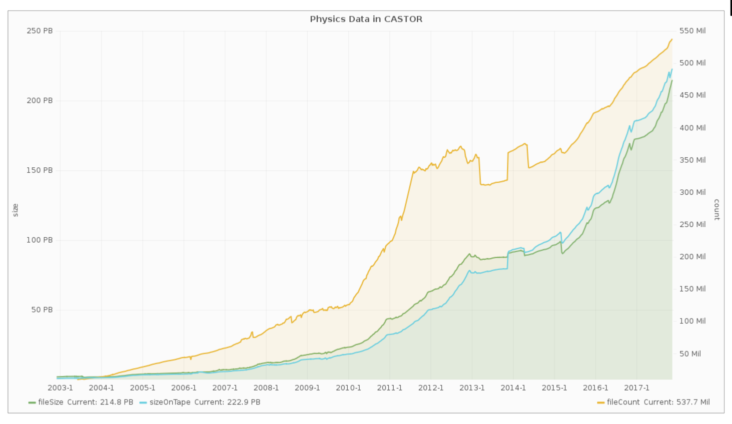

Data archived at the tape storage @ CERN since 2003. ~223 PB of data (turquoise line) and ~550 million of archived files (pink line) stored by the end of November 2017 .

The plot shows significant increase of time needed for reconstruction of one event depending on the pileup (ATLAS experiment).

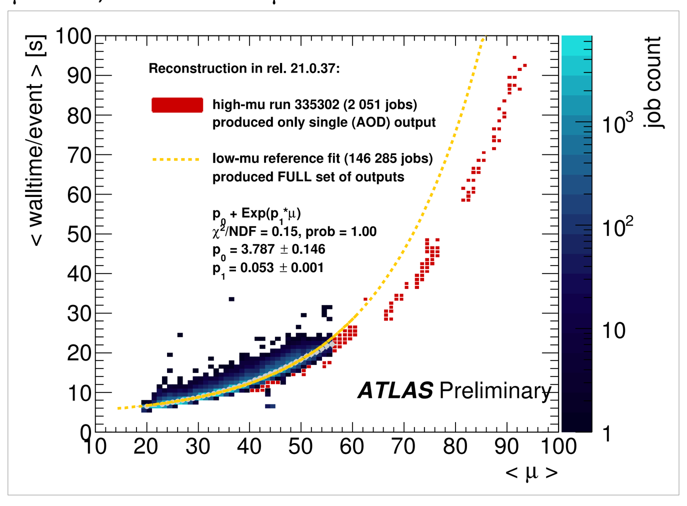

2. Anticipated development and resource needs for HL-LHC

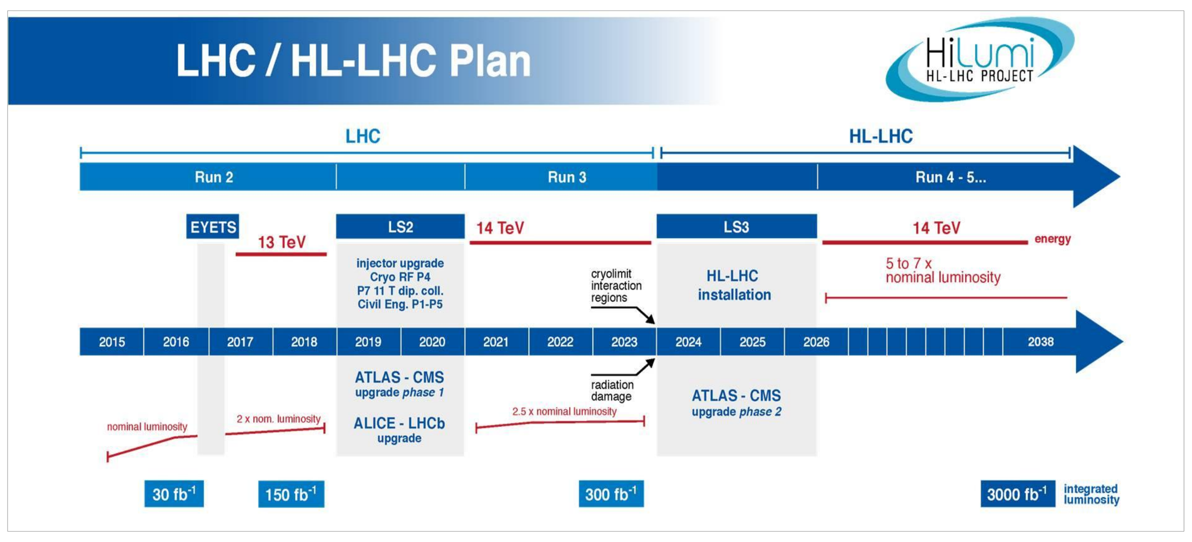

The ultimate target of the High Luminosity LHC:
A peak luminosity of $L_{peak} = 7.5 \times 10^{34}$ $cm^{-2}s^{-1}$ with levelling, allowing for an integrated luminosity of > 300 fb^{-1} per year, enabling the goal of L_{int} = 4000 fb^{-1} twelve years after the upgrade.
This luminosity might be twenty times the luminosity reach of the first 10 years of the LHC lifetime.

The LHC performance beats any expectations year by year. The anticipated spectacular increase in the luminosity will bring much higher pileup and massively larger amounts of data produced by the detectors. The higher pileup will translate into correspondingly longer times needed for event reconstruction, and higher volumes of raw data will require significantly more simulation campaigns. In addition, the more complex event structures will require development of new physics generators for the Monte Carlo simulations with Next-to-Next-to-Leading-Order corrections. All this will result in a need for so many resources that computing may become a limiting factor for physics.

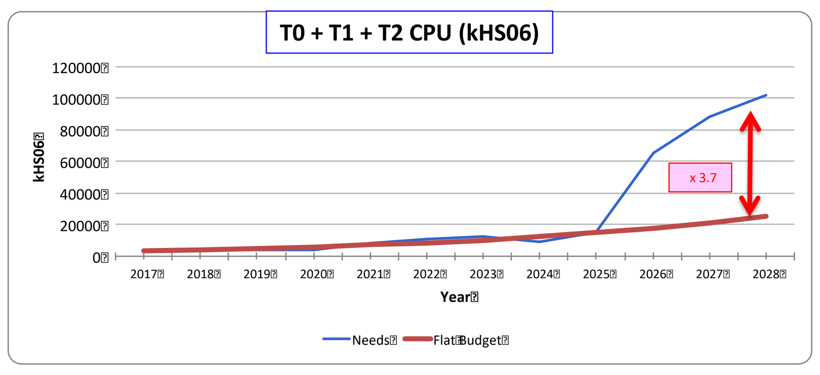

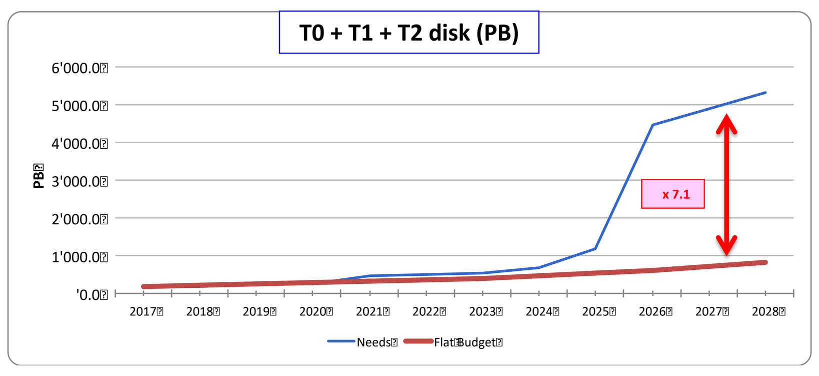

Anticipated development of available resources (red line: expected flat budget growth of ~ 20% for CPU/year and ~ 15% for disk/year) and resources needed to leverage the amount of data produced in Run 3 and especially in Run 4 (blue line). There will be a sharp rise in the amount of recorded data in Run 4 after the upgrade of the ATLAS and CMS detectors. The estimates elaborated in 2017 predict a shortfall in disk needs by ~ 7 times and CPU needs by ~ 4 times.

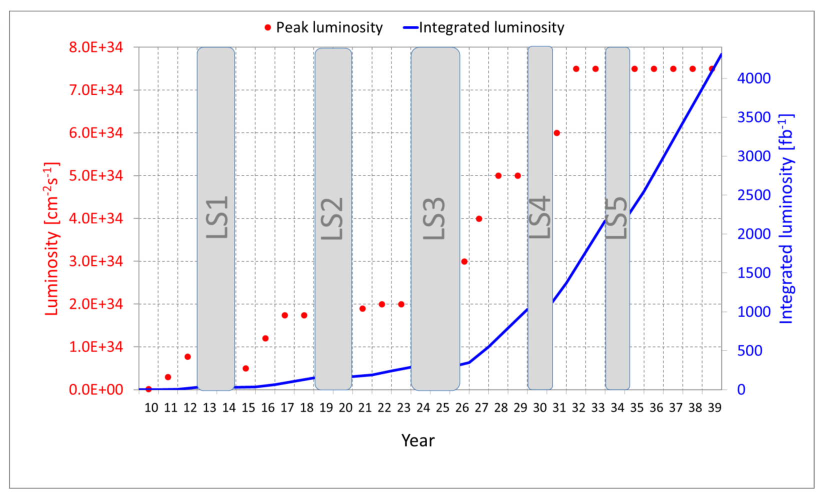

The ultimate future LHC luminosity profile. After LS4, proton physics days increase from standard 160 days to 200 and after LS5 to 220.

3. Preparations for Run 3 and HL-LHC : strategies in computing and software development

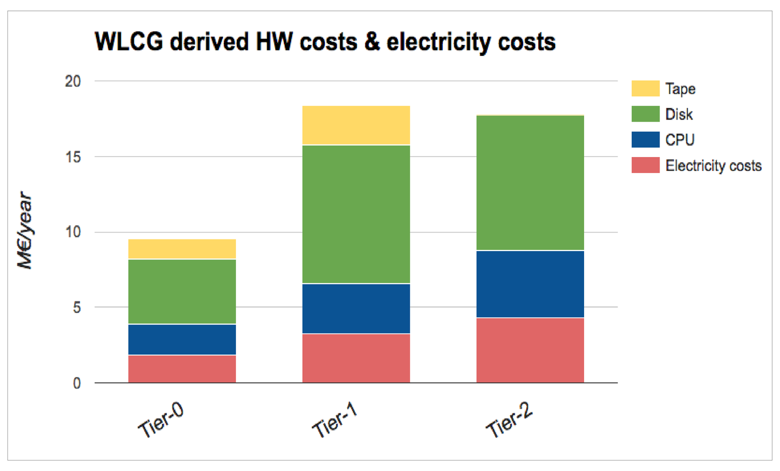

WLCG started a process to understand the resource needs for Run-3 and beyond. The ultimate goal is to provide the computing capacity needed for the LHC physics program, managing the cost. The currently elaborated estimates predict a lack of computing and storage resources which is impossible to satisfy with the standard growth of ~20% CPU/year and ~15% storage/year. External resources like private and public computational clouds, High Performance Computing centers (HPC) or volunteer computing can help but do not guarantee delivery of all needed resources. The WLCG community started preparations for the WLCG Strategy document for the HL-LHC computing challenges. This will lead to prototyped solutions which will be the foundation of the WLCG Technical Design Report (TDR) for HL-LHC, planned for 2020.
The challenges of HL-LHC lie not only in computing capacity but also in the software domain of the LHC experiments. It comprises over 12 million lines of code mostly written over the last 15 years . The software challenges motivated the startup of the HEP Software Foundation (HSF).

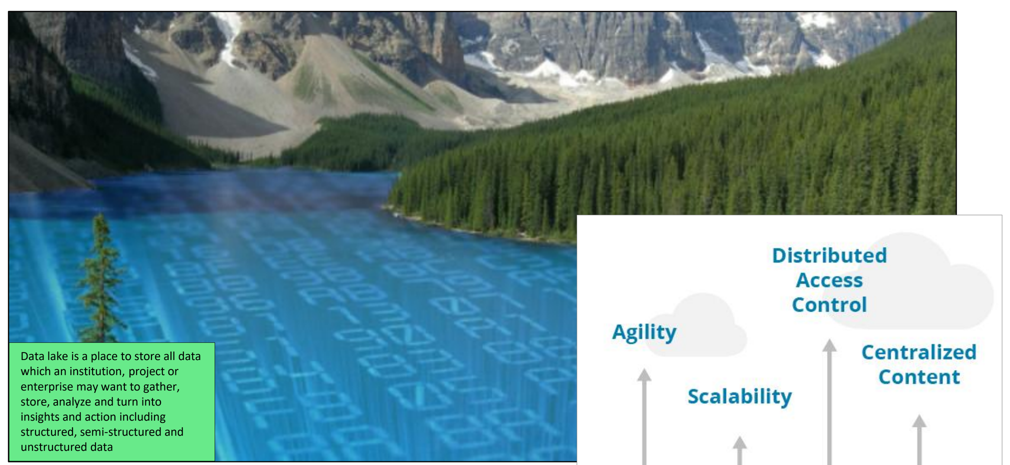

One of the strategies pursued within the area of Data Organisation, Management and Access (DOMA) is the building of data lakes within the WLCG infrastructure.

Cost ("toy") model for WLCG
Hardware : ~ 36 M€/year
Electricity: ~ 9 M€/year
Assumptions of the model:
→ 12.5 (3) FTEs to operate a Tier-1 (Tier-2) site
→ 50 k€/FTE → manpower costs = 32 M€/year
→ does not include network costs
This model yields WLCG Operations + Infrastructure costs ~100 M€/year.
Ongoing work on building a decent cost model: vital to quantify (dis)advantages of various development strategies like e.g. massive adoption of commercial cloud services (current evaluations still discourage this).

Main aspects to be elaborated by the WLCG Strategy document
- Build a cost model (hardware, operations, people)
- Improve software performance and efficiency
- Cutting across all aspects of event processing
- Reduce data volume
- Triggers and selection, data formats.
- Data replication, caching
- Reduce processing costs
- Improved algorithms, fast simulation
- Role of Machine Learning
- Reduce infrastructure and operation costs
- Storage consolidation and harmonization
- Data reproducibility, data preservation
- Ensure sustainability
- Common solutions and services at infrastructure and higher layers

The mission of the HEP Software foundation (HSF) is to foster communication and collaboration across HEP and beyond on possible strategies and a roadmap for the needed fundamental HEP software upgrade.

15 Working groups:
1. CWP (Community white paper) Roadmap
2. Careers & Training
3. Conditions Data
4. Data Organisation, Management and Access
5. Data Analysis and Interpretation
6. Data and Software Preservation
7. Detector Simulation
8. Event/Data Processing Frameworks
9. Facilities and Distributed Computing
10. Machine Learning
11. Physics Generators
12. Security
13. Software Development, Deployment and Validation
14. Software Trigger and Event Reconstruction
15. Visualisation

HSF Priorities:
- achieve improvements in software efficiency and performance
- enable new approaches in software and computing to extend physics reach of the detectors (e.g. Machine Learning (pattern recognition) in event reconstruction)
- ensure the long term sustainability of the software
An important keyword in the HSF endeavour is commonality: support of cross-project collaboration and creation of tools and solutions usable by various HEP experiments. Specific contacts in place with FNAL muon and neutrino experiments, Belle II, Linear Collider community, Square Kilometer Array (SKA) and national computing organisations.

4. Outlook

The distributed computing and storage infrastructure for the LHC experiments, WLCG, has been providing a safe framework for processing of LHC data and delivery of physics results ever since the startup of the LHC operations.
In 2017, in the middle of Run 2 and in the Exabyte era, the prospects of ever growing amounts of produced data especially anticipated for the HL-LHC, raise concerns about managing the costs of providing the desired resource levels for HL-LHC.
HSF and WLCG produced two strategic documents defining roadmaps to provide capacity and software necessary for the HL-LHC physics program.
Nevertheless, the current evaluations show, that

"The amount of data that experiments can collect and process in the future will be limited by affordable software and computing, not by physics"

(HSF Community white paper, 2017).

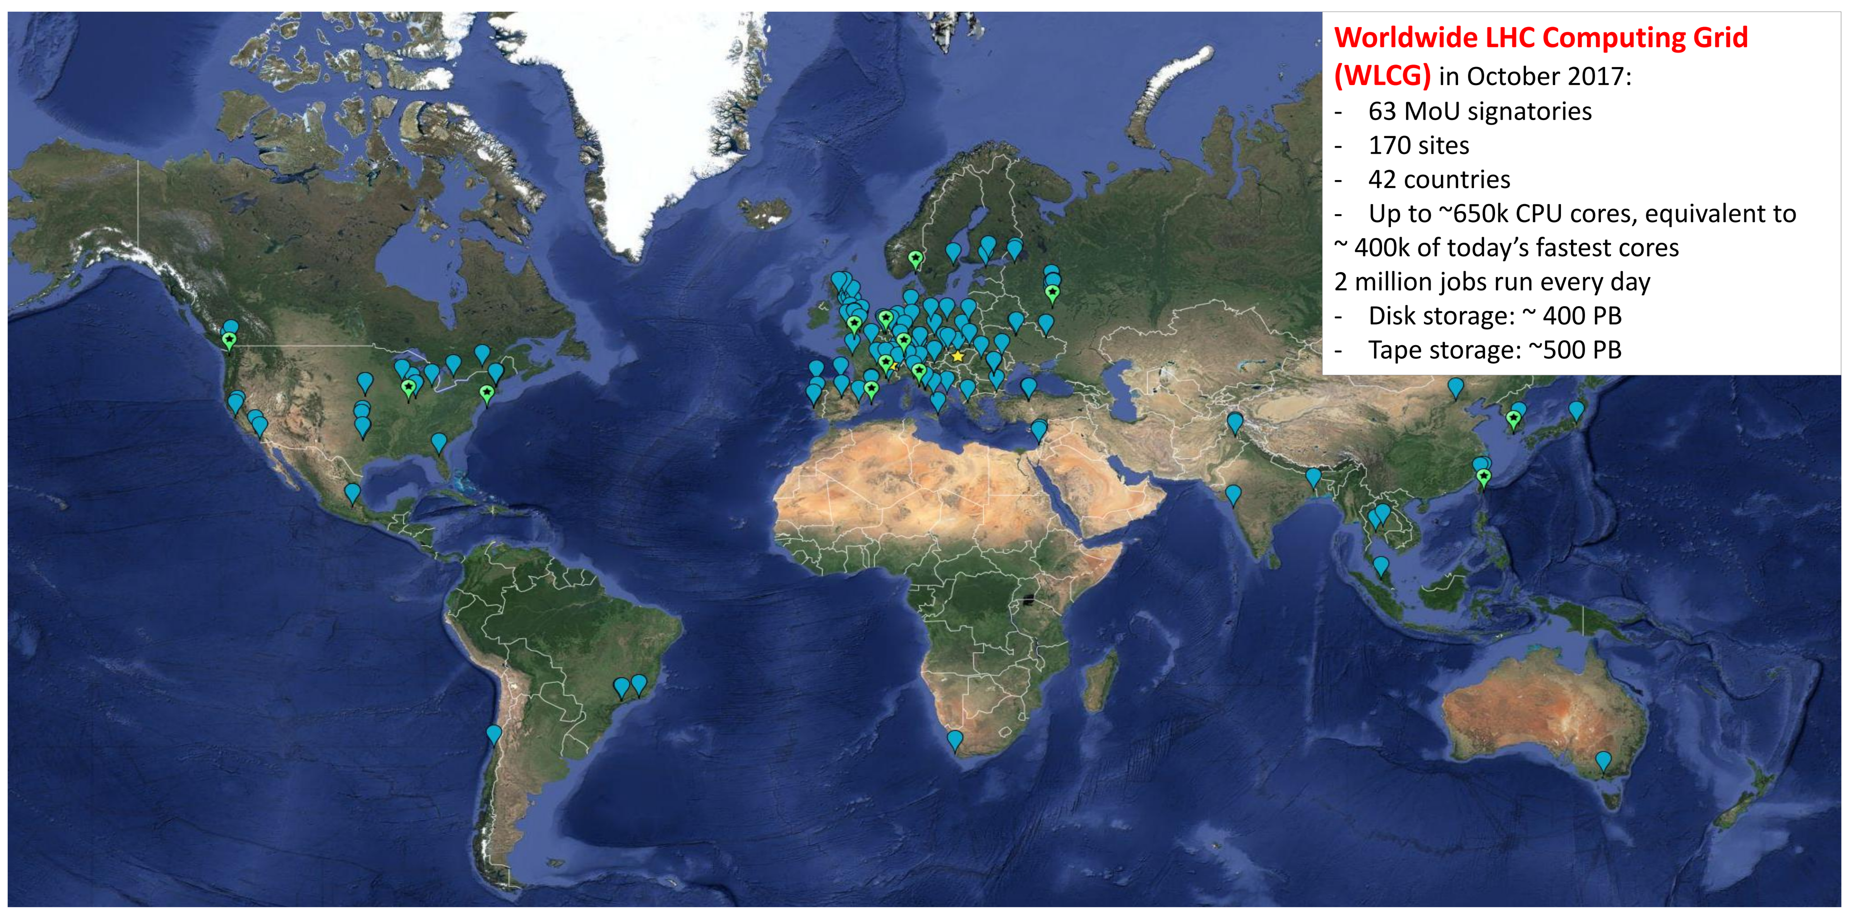

Worldwide LHC Computing Grid (WLCG) in October 2017:
- 63 MoU signatories
- 170 sites
- 42 countries
- Up to ~650k CPU cores, equivalent to ~ 400k of today's fastest cores
2 million jobs run every day
- Disk storage: ~ 400 PB
- Tape storage: ~500 PB

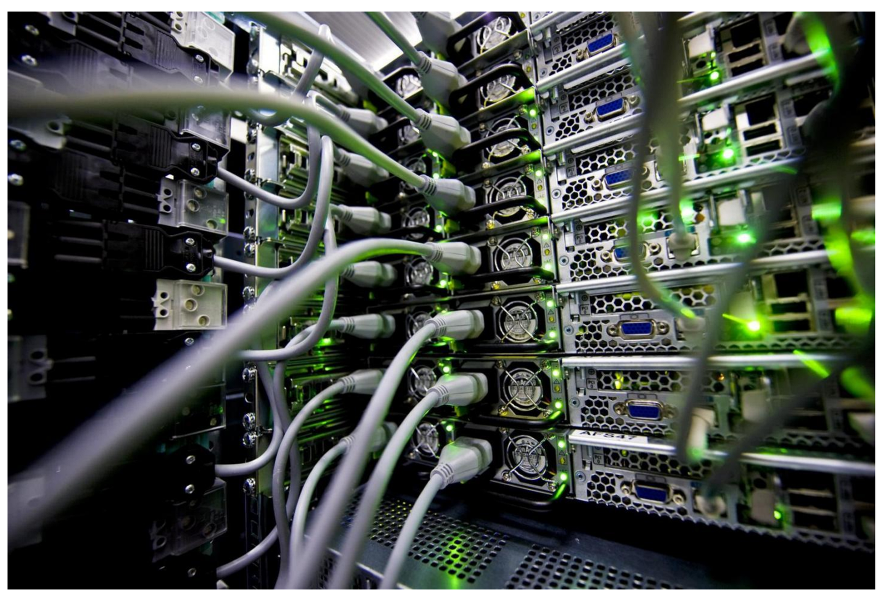

Data servers at the CERN Tier-0 computing center send one copy of the LHC raw (primary) data to local tape and distribute the second copy over the tape storage of Tier-1 centers worldwide. The links connecting CERN with Tier-1s are steadily being upgraded to capacities typically up to 100 Gb/s for the next few years.