Computing for the LHC: operations during Run 2 and getting ready for Run 3

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

The computing infrastructure for the LHC data handling (Worldwide LHC Computing Grid - WLCG) was well prepared for Run 2 and delivered a steady data processing since the first collisions. Over the past years including the Run 1 period of the LHC the WLCG developed a unique expertise in building and operating a very large scale infrastructure for unprecedented amounts of LHC-produced data. Together with other research projects the WLCG entered the era of Data Intensive Science. The cost of Computing infrastructures and services reached the level needed for building the LHC detectors. In this contribution, we will present ongoing and planned developments to evolve the computing models towards the needs for the High Luminosity (HL) LHC, including the disappearance of the distinction between "online" and "offline", analyses of the physics performance and corresponding computing costs, and inter-operation between commercial cloud services and publicly funded e-infrastructures.

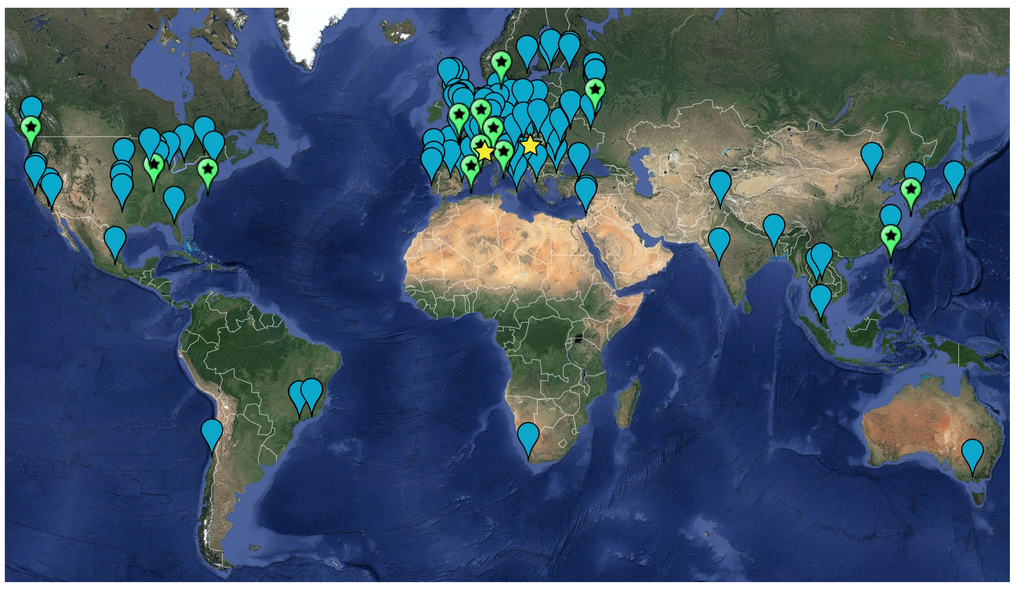

WLCG resources in 2015:
nearly 170 sites
41 countries
~ 350 000 cores
~ 500 PB of storage
> 2 million jobs/day
10 – 100 Gbit/s links

WLCG is the world's largest Computing grid.

The Worldwide LHC Computing Grid (WLCG)

WLCG is a global computing infrastructure which provides computing resources to store, distribute and analyze the data generated by the LHC experiments. It enables access to the data for all participating institutions regardless of their physical location.

Performance of WLCG during Run 1 and Run 2

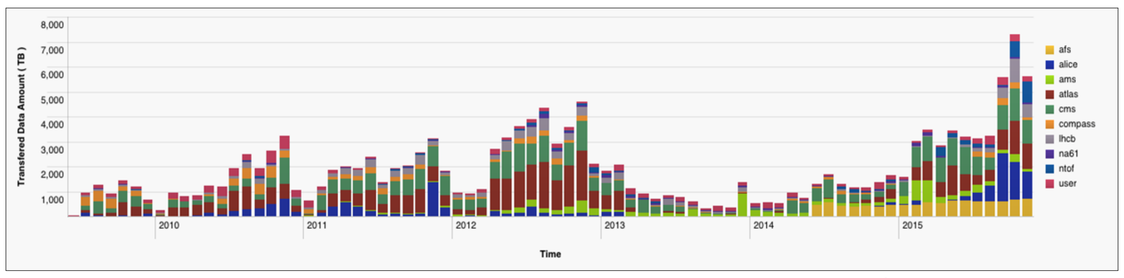

Data transfers to CASTOR tape storage at CERN Tier-0 during Run 1 and Run 2, in TB/month.
Collected during Run 1: 15 PB/2010, 23 PB/2011, 27 PB/2012.
Collected during the first year of Run 2: ~ 40 PB of data including the Heavy Ion (HI) run in the end of 2015.

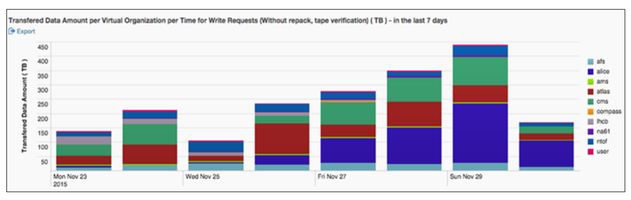

Data transfers from experiments to CERN during one week of the HI run in 2015, in TB/day.
Data rates ~8 GB/s sustained for 8 hours, peaks ~10.5 GB/s.
Up to 450 TB/day (previous max was ~220TB/day).

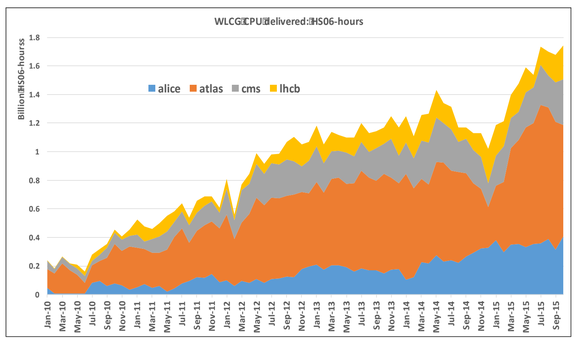

Delivery of CPU resources during Run 1 and Run 2, in billion HepSpec06. Ramp-up almost by 1/3 in 2015.

What is different in Run 2 in comparison with Run 1

- More flexible computing models: less specific roles of different Tiers, more even distribution of tasks (reconstruction, simulation, analysis).
- Use of HLT farms of the experiments for data processing during LHC technical stops and machine development periods.
- Use of clouds and other opportunistic and commercial resources.
- Data popularity analysis and clean-up of unused data have become even more important.
- Network access to data: moving data around is expensive (disks needed). Access any data from any site without the need to copy it first. However, most of the processing still needs to be performed inside the centers holding the data or nearby via fast networks.
- Further simplification and consolidation of middleware services and their clients.
- Enhanced experiment software performance:
  - -- Simulation faster up to factor 2.
  - -- Performance of reconstruction improved significantly, up to 4 times.
  - -- Multicore jobs: more efficient memory usage.
- Other optimizations:
  - -- Reduced number of (re-)processing passes.
  - -- New optimized analysis data formats.
  - -- More organized analysis processing workflows ("analysis trains").
  - -- Faster calibration and alignment (ALICE), online calibration (LHCb)
- New transatlantic networking via ESNet (links 100 Gb/s).
- New sites joining.
- Growth at existing sites.

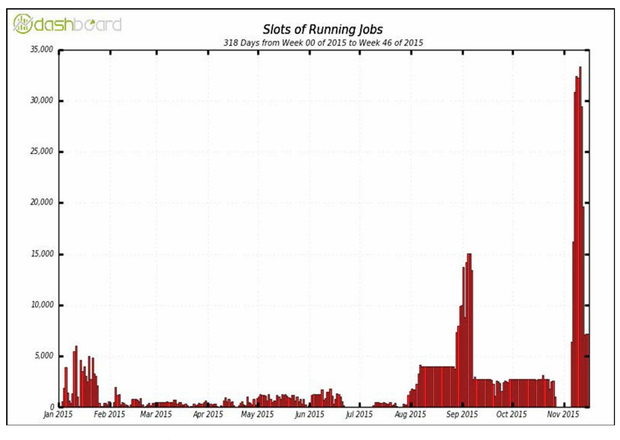

Usage of the ATLAS HLT farm during 2015
maximum: 33 372, average: 2 259.

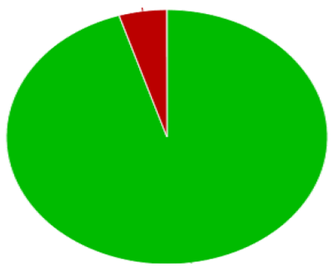

ALICE: Access to data/ Data popularity in 2015
(green accessed, red not accessed).

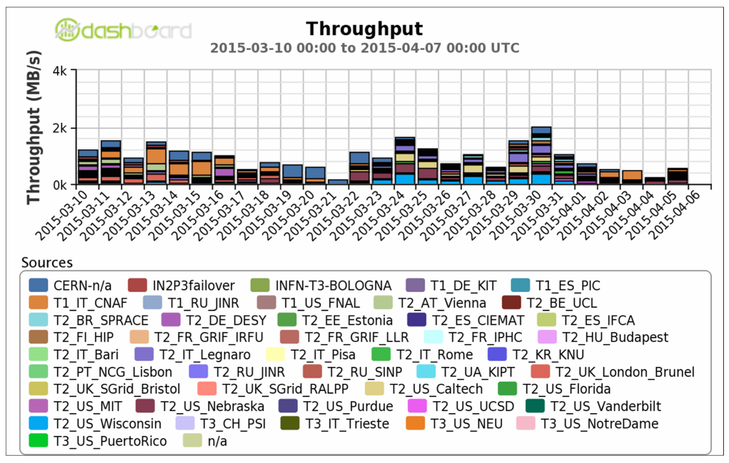

CMS: scale testing campaign to validate the project "Any data Anytime Anywhere (AAA)".

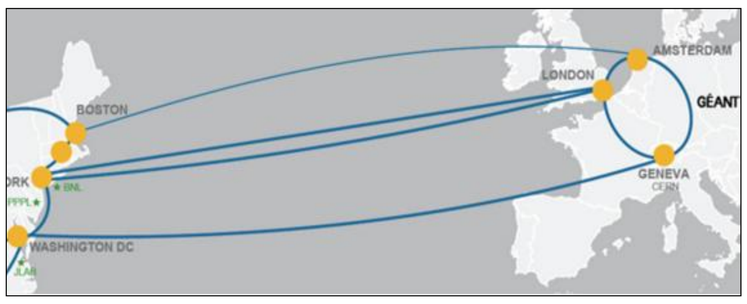

Schema of a new transatlantic network connection: links 100 Gb/s.

Longer term prospects for the LHC computing

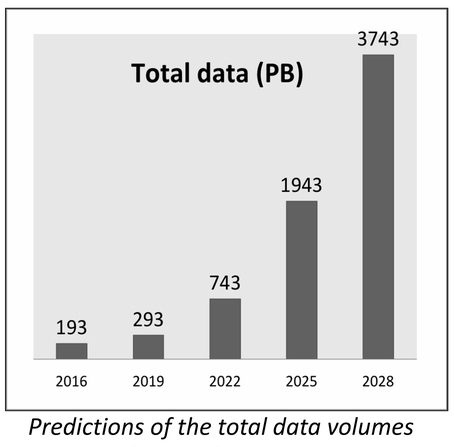

Predictions of the total data volumes generated beyond 2016.

The basic question: What physics can be done with the computing we can afford?
Depends on evolution of technology and costs of HEP computing.

- During the LEP era, computing was affordable on commodity platforms.
- In the LHC era: simple calculation shows that the global yearly cost of WLCG hardware is approximately 100 million CHF/$/Euro.
- 5-year cost is almost the same as the construction cost of ATLAS or CMS detectors.
- What it includes: CPU, disks, tapes, networking, electricity, computing facilities, sys-admins.
- The costs of facilities and power motivates an idea that commercially provisioned computing may soon be more cost effective for HEP (huge scale facilities located in regions with cheap power and cooling).
- There are also other important developments to meet the challenges coming with the Run 3 era:
  - -- The continued optimization of the reconstruction and physics code to get more performance out of evolving traditional CPU architectures (x86_64).
  - -- A reimplementation of such code to work on new architectures that will be less expensive (e.g. ARM, GP-GPU).

Predictions for the LHC raw data taking beyond 2016:
- Run 3 (2020-2022): ~150PB/year
- Run 4 (2023-2029): ~600PB/year
- Peak rates of ~80GB/s

The challenge for the coming years: to be able to evolve the system while fitting within constrained budgets without compromising the physics output.*

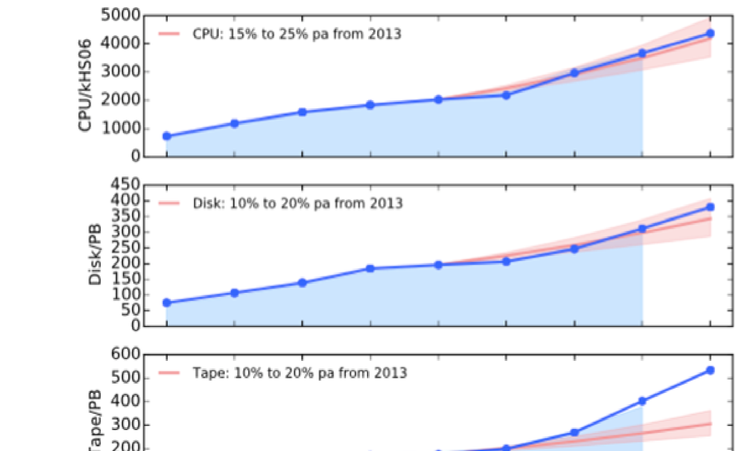

Expected growth of the computing resources requirements of the LHC experiments.

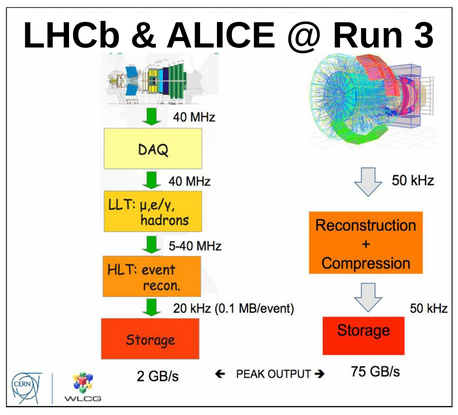

Data throughput to CERN envisaged by LHCb and ALICE for Run 3.

WLCG and usage of clouds (scientific and commercial)

Many investigations of usage of cloud resources by experiments, CERN, and other sites.
CERN: multiple procurement projects to gain experience with markets and usage models.
Can commercial cloud services supplement or (partially) replace HEP-owned and managed facilities?

Helix Nebula - the Science Cloud: EC project in Europe (together with other sciences).
A promising way of working with industry to keep up with the ever-growing amounts of the LHC data.
- Evaluation of the use of cloud technologies for LHC data processing.
- Integration of cloud computing resources with experiment software and services.

HELIX NEBULA THE SCIENCE CLOUD

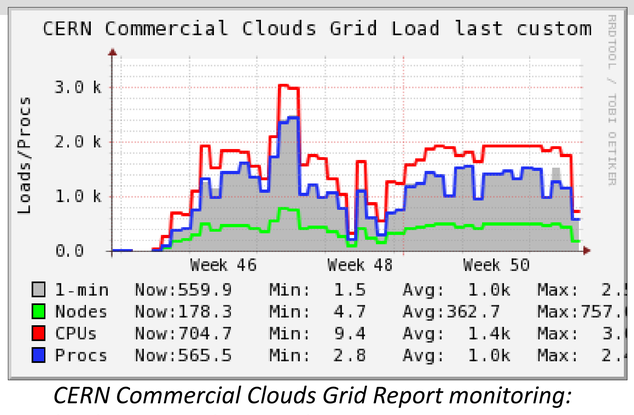

CERN Commercial Clouds Grid Report monitoring: load report in the period 1.11.2015 - 19.12.2015.

* *Ian Bird @ Computing Resource Review Board; CERN, 27.10.2015*