

CERN Is Betting Big on Exascale

By Oliver Peckham

April 1, 2021

The European Organization for Nuclear Research (CERN) involves 23 countries, 15,000 researchers, billions of dollars a year, and the biggest machine in the world: the Large Hadron Collider. Even with so much organizational and mechanical firepower behind it, though, CERN and the LHC are outgrowing their current computing infrastructure, demanding big shifts in how the world's biggest physics experiment collects, stores and analyzes its data. At the 2021 EuroHPC Summit Week, Maria Girone, CTO of the CERN openlab, discussed how those shifts will be made.

The answer, of course: HPC.

The Large Hadron Collider – a massive particle accelerator – is capable of collecting data 40 million times per second from each of its 150 million sensors, adding up to a total possible data load of around a petabyte per second. This data describes whether a detector was hit by a particle, and if so, what kind and when.

Currently, this data load is handled by the Worldwide LHC Computing Grid (WLCG), a network of systems spanning 167 sites across 42 countries and comprising a million CPU cores, along with an exabyte of storage. Much of this infrastructure is old, however, and much of it is running on old code – Python, in many cases over 20 years old. (Girone characterized this legacy code as “challenging to optimize.”)



“There will be a gap of resources”

Complicating matters, the LHC is far from static. It's been almost 13 years since its inauguration, and the LHC is currently in “Run 2” of a roughly 30-year operation plan. Looming on the horizon is a new regime, the “High-Luminosity LHC,” that will dramatically increase the effective resolution (and data generation) of the behemoth machine beginning with Run 4 in 2028. “In a way,” Girone said, “it is like moving from looking for a needle in a haystack to producing many more needles.”

Somewhat more imminently, Run 3 (scheduled for 2022) will introduce substantial upgrades to the LHC “beauty” experiment (LHCb) and the A Large Ion Collider Experiment (ALICE), causing each of the projects to produce ten times more data.

Two of CERN's major projects – ATLAS (A Toroidal LHC Apparatus) and CMS (Compact Muon Solenoid) – expect that the HL-LHC upgrades will cause sixfold to tenfold increases in computing power requirements, along with threefold to fivefold increases in disk space requirements (already, the combined volume of LHC data is approximately an exabyte).

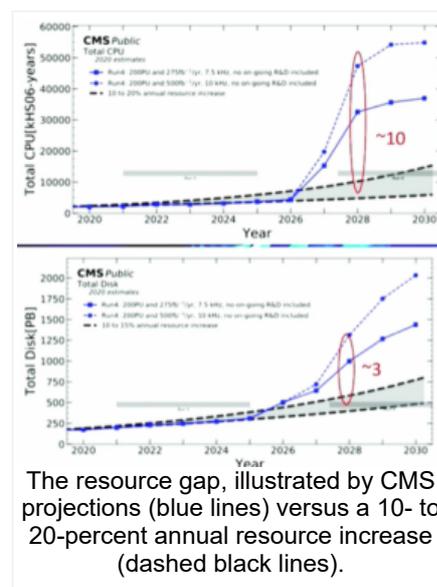
The current pace of expansion is not sufficient to meet either of the anticipated needs.

“There will be a gap of resources,” Girone said. “This resource gap is motivating an ambitious program of R&D to modernize and optimize our software to adapt the code to hardware accelerators and high-performance computing, to reduce the storage footprint and to use more efficient techniques like AI and ML in the experiment workflows.”

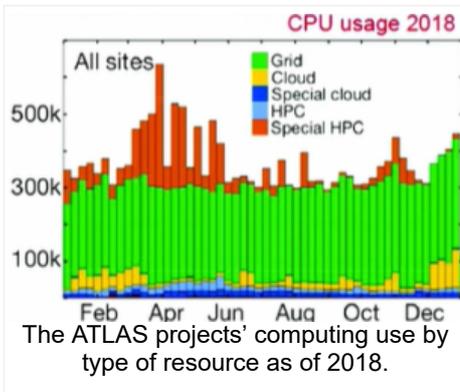
To bridge the gap, CERN is looking to HPC for answers. Girone noted that supercomputers are expected to become ten times more powerful in the time it will take for the HL-LHC to become operational – and furthermore, that supercomputers are taking advantage of the AI, machine learning and heterogeneous architectures that CERN is quickly realizing will be crucial to its future operations. “The advent of heterogeneous hardware, that is used heavily in high-performance computing, has shaken also our landscape, and today all experiments are exploring GPUs for accelerated event reconstruction simulation,” Girone said, adding, “We believe that high-performance computing can play a critical role in the success of the HL-LHC.”

Supercomputers, of course, are already a part of CERN's operations: Girone said that ATLAS, for instance, uses supercomputing for around 10 percent of its computing needs, stemming from a large number of diverse systems and facilities (CSCS, OLCF, ALCF, NERSC, TACC, LRZ, etc.). CMS, similarly, has used a high-energy physics resource provisioning service called HEPCloud (on which it once used 100K cores simultaneously) and CINECA's Marconi system (peak usage of 22K cores).

“The experiment progress at individual sites is good and very encouraging,” Girone said, “but you can also see that it still represents a small fraction of the one million-core operation of the WLCG. ... If high-performance computing is going to make a serious dent in the resource gap, we need to expand the types of hardware we can use and increase the number of supercomputing sites.”



The resource gap, illustrated by CMS projections (blue lines) versus a 10- to 20-percent annual resource increase (dashed black lines).



“An exascale project for an exascale problem”

Given the timeline for HL-LHC and its eventual needs, CERN sees strong promise in the exascale era. To that end, CERN launched a project last summer in collaboration with GÉANT (a major European research network), PRACE (the Partnership for Advanced Computing in Europe) and SKAO (the organization leading the Square Kilometer Array, which will be the largest radio telescope on Earth when completed).

The project consists of a number of pillars, mostly centered around developing expertise in heterogeneous hardware and developing proofs of concept for benchmarking, data access, authorization and authentication. Girone called the collaboration “an exascale project for an exascale problem.”

Already, the collaboration has borne fruit: the researchers successfully ran a benchmarking suite for multiple

architectures at scale in real-world HPC sites. CERN is also testing AI and machine learning techniques throughout its data pipeline and is working to modernize its data transfer network to reach 10 PB of data processing per day in order to demonstrate feasibility on the path to exascale.

Girone didn't name names when it came to future collaborations and supercomputer integrations, but said that such partnerships were “very important” and that “engagement with high-performance computing is very important to the future of high-energy physics computing.” However, the host of the discussion, Matej Praprotnik, hinted that tests on some of the [new and forthcoming EuroHPC systems](#) may already be underway.

