

CERN Upgrades Data Center and Restarts Large Hadron Collider

August 10, 2015


Share

Tweet

Like

g+

TAKE THIS CONTENT



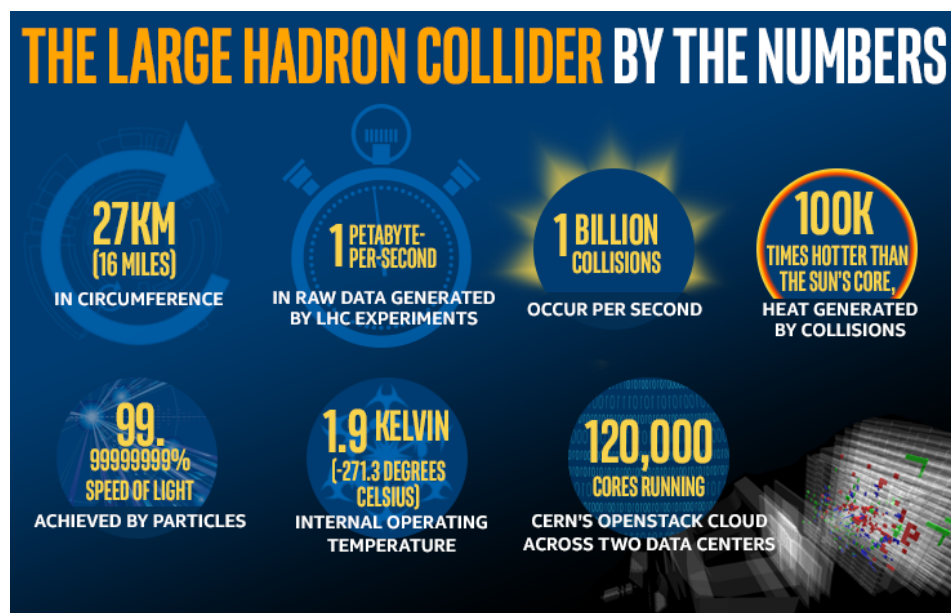
With upgraded hardware and capabilities, CERN pushes physics into new realms

In April 2015, the most powerful particle accelerator ever built—the Large Hadron Collider (LHC) at CERN, the European Organization for Nuclear Research—was fired up for its second big series of runs, this time smashing things together at nearly [twice the energy](#) of its famous first run.

That three-year run led to the discovery of the [Higgs boson](#) particle, one of the last mysteries of the [Standard Model](#) of physics—our best understanding of what stuff is fundamentally made of.

“We have a fantastic model—that we both love and hate,” said Steven Goldfarb, a physicist from the University of Michigan working on the ATLAS experiment at the LHC. “We love it because it has stood up to precision measurements for 50 years. We get more and more precise, and it stands up and stands up to our measurements. But we hate it, because it doesn’t explain the universe. So something more is needed.” Things like gravity and 95 percent of what makes up the universe, scientists say, remain unexplained, unaccounted for.

Importantly, the LHC’s first run also proved that this massive machine—a 27-kilometer (almost 17 miles) ring, 100 meters under the France-Switzerland border, consisting of 9,500 superconducting magnets kept at -271.3 degrees Celsius, which slams particles into each other at a shade under the speed of light—actually works.




A second journey: Filtering through a petabyte-per-second to improve our understanding of

the universe.

After confirming the Higgs boson, the scientists were inspired to push even harder. Teams at CERN spent the past two years upgrading magnets (which guide the powerful beams around the ring), narrowing and cranking up the energy of the beams, and improving electrical systems, cooling, and the vacuum that surrounds the beam. The CERN Data Center also expanded—almost 60,000 new cores and over 100 petabytes of additional disk storage [were installed](#). These upgrades included equipment spanning the Intel Xeon processor E5 v3 family, Intel SSDs, networking gear, and a smattering of software tools and expertise.

During the LHC's second run, CERN's researchers are casting a much wider net, pushing mind-bending frontiers including new exotic particles, dark matter, supersymmetry, extra dimensions, antimatter, and quark-gluon plasma (read more at [CERN's website](#)). It's the kind of stuff we *think* might make up some, if not all, of the universe. The increased computing requirements of the LHC's second run means it's more important than ever to optimize high-energy physics codes for new computing architectures.



Silent boom! Particles created from a proton collision stream out from the center of the CMS detector, one of seven detectors along the LHC's ring. This is one of the first images captured with a stable beam at 13 TeV. ©CERN for the benefit of the CMS collaboration.

Tara Shears, a physicist working at the University of Liverpool and a member of CERN's LHCb collaboration, told the BBC: "Instead of trying to test the truth of [the standard model], what we really want to do now is break it."

In other words, CERN's researchers are focusing on new mysteries, which means collecting vast amounts of data and filtering with great precision. They're venturing into hitherto unexplored territory.

A team of Intel experts have worked with CERN's IT teams "to implement a new infrastructure that will be more flexible and more open to modification," explains Claudio Bellini, a business development manager in the Sales and Marketing Group (he's based in Italy but spends a lot of time with CERN in Switzerland). CERN's infrastructure needs to both capture and filter live data from the LHC—up to a petabyte per second—and provide simulation and analysis of offline data for scientists all over the world.

At the limits of physics—and the distributed data center

CERN's work is really two monster projects: running the Large Hadron Collider and the experiments, and running the compute environment to capture, store, analyze, and distribute the data that comes from the LHC.


Data from collisions are collected and stored at the CERN Data Centre. But this is precious, hard-to-generate data, so CERN uses the Worldwide LHC Computing Grid (WLCG) to copy it elsewhere. This technology enables collaborating scientists around the globe use their own supercomputers and workstations to take a crack at the data, too.

Founded in 2002, the WLCG is a [four-tier](#) distributed computing infrastructure that's the largest of its kind on the planet. The grid ensures the data is safe, and dramatically boosts the available compute power to dig into it. In the U.S., for example, [18 different labs](#) and universities will participate in tiers of the grid for this phase of LHC runs.

After data is captured and stored at the CERN Data Centre—that's tier 0—data is piped over dedicated networks to 13 [tier 1 centers](#) around the world—from the Karlsruhe Institute of Technology (just a hop from [Intel's site](#)) to Taiwan's national academy [in Taipei](#) to Fermilab outside Chicago. These centers then distribute data to 155 tier 2 sites at universities and scientific institutes, which then feed to countless more local tier 3 clusters and PCs.

(Prefer a quick visual tour? Watch how protons harvested from hydrogen gas end up deconstructed as

petabytes of data in CERN's animation below.)



Public-private partnership meets 'computing challenges of the LHC'

If there's an emerging data center technology—like silicon photonics, software-defined [networks](#) and infrastructure, new coprocessors—it's likely to be put to the test somewhere in CERN's grid.


It's no accident. CERN runs a public-private partnership called [CERN openlab](#), created in 2001 "to develop the innovative IT systems needed to cope with the unprecedented computing challenges of the LHC." Intel has been a CERN openlab partner since its inception, contributing to a variety of hardware and software challenges over that time.

Bellini says that one goal was to use the collaboration as a test bed for technologies, receiving engineering feedback on "alpha prototype platforms and software solutions" from demanding use cases. Those ranged from testing early versions of the Intel Xeon Phi coprocessor to using Intel tools to optimize software on new instruction set features.

Now entering what CERN openlab calls its [fifth phase](#), the partnership is "moving into a new project-based collaboration," Bellini explains. Other partners in CERN openlab's fifth three-year phase are Oracle, Huawei, and Siemens, as well as other contributors and associates. CERN is also expanding the effort beyond the needs of the LHC to include other research institutions and topics including data analytics, healthcare, and entrepreneurship. Intel just [signed on](#) with CERN to join this phase on several projects, such as:

- **Modernizing the code:** The increased computing requirements of the LHC's second run means it's more important than ever to optimize high-energy physics codes for new computing architectures. There are 4 different use cases:
 1. [Geant software simulation toolkit](#), widely used in space, particle, and medical research, with aims of up to a 5X performance boost, thanks to experts and tools from Intel's Software and Services Group; this use case is implemented in collaboration with Intel IPCC program
 2. FairRoot, (in collaboration with GSI) a widely used HEP experiment framework used by a large number of collaborations. Key parts of the framework will be re-designed to make optimal use of the Xeon Phi architecture
 3. Cx3D (in collaboration with School of Computing Science – Newcastle University), a leading cell growth simulation code. This code is currently written in Java but will be re-written in C++ for this project. The C++ rewrite will be combined with a re-engineering of the data structures to make optimal use of the Intel CPU vector instruction and of the Xeon Phi architecture.
 4. Several Beams Department Injector Simulation codes. These codes are written in a mix of FORTRAN, C/C++ and Python. Some use already MPI. Additional optimization possibilities will be researched.
 5. Inside this project Intel is also delivering on site workshops addressing latest Intel s/w tools and training on code vectorization technologies.
- **HTC: High Throughput Computing aims at exploring developing the new data acquisition chain for the LHCb detector.** The data rate and volume will increase by one order of magnitude between now and 2020; this three-year project (2015-2017) will focus on building and assessing prototypes based on the latest Intel technology (manycore & networking), as well as assessing the performance of the continuously evolving acquisition software developed by CERN in parallel.
- **Hardware and software proof-of-concepts** that span from high-throughput computing for future LHC runs, to bringing real-time analytics to the LHC control system, to using Intel [Rack Scale Architecture](#) for a software-defined data center;
- **An entrepreneurship project** to support and accelerate IT startups, helping researchers go from

vetting concepts to launching commercial products and services.



“CERN openlab played a vital part in [the Higgs boson] discovery,” said Rolf Heuer, CERN Director-General, in openlab’s 2012 report. “Without the ability to sift, record, and process the huge volumes of data generated by the Large Hadron Collider, it would have been impossible for the experiments to sort the wheat—a few hundred Higgs candidates—from the chaff among the trillions of particle collisions that took place.” ©CERN

Related Articles:

1. [Green Supercomputer Crunches Big Data in Iceland](#)
2. [Solar Car Design Powered by High Performance Computing](#)
3. [Tech Levels Playing Field for Filmmakers](#)
4. [U.S. Postal Service Using Supercomputers to Stamp Out Fraud](#)

Tweet

Like 

© Intel Corporation [Terms of Use](#) [Trademarks](#) [Privacy](#)



[About](#) [Contact](#)