**CASE STUDY**
**Intel® Xeon® Processor E5 v3 Family**
**High-Performance Computing**
Data Center Efficiency

# Exploring the unknown

**CERN upgrades to the latest Intel® Xeon® processors, bringing LHC experiments onto one platform**

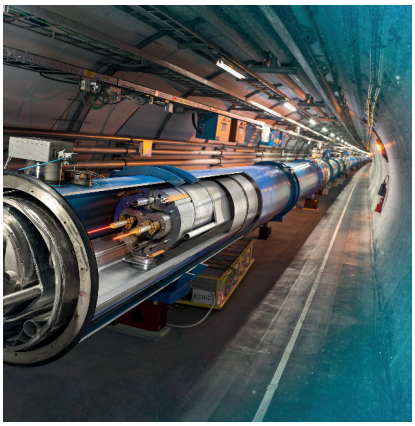

*© 2015 CERN*

CERN supplier

"The combination of processors such as the Intel® Xeon® processor E5-2630 v3 and the OpenStack cloud software enables us to meet the computing needs of the physicists for the second run of the LHC with increased flexibility and efficiency."

*Tim Bell,*
*Leader of the IT Operating Systems and Infrastructure Services Group, CERN*

CERN, the European Organization for Nuclear Research, attracts around 11,000 visiting researchers working on various experiments examining the very particles that make up our universe. After two years of shutdown for maintenance and upgrade work, it has restarted the famous Large Hadron Collider (LHC) to try and determine even more about the world we live in, with the first particle collision planned for summer 2015. During the first run, the LHC was fine-tuned to find one single particle. Now the scientific community is exploring more widely and requires a more flexible IT infrastructure. CERN has upgraded to the Intel® Xeon® processors E5 v3 family to achieve increased capacity and even better performance to deal with the unknown challenges ahead.

## Challenges

- **Computing capacity.** Huge volumes of data must be processed for each of the LHC experiments, so advanced computing capacity is essential
- **Research breadth.** During the second run of the LHC, CERN researchers are focusing on new mysteries, which means collecting vast amounts of data and filtering with great precision
- **Cost efficiency.** CERN needs to provide resources on a fixed budget, so it relies on technological improvements to achieve better results at the same cost

## Solutions

- **Latest technology.** The Intel Xeon processor E5-2600 v3 product family was tested using an industry benchmark and CERN's own analysis. According to its own internal tests, CERN discovered that performance and energy efficiency tested better than previous generations[1]
- **Unified, flexible platform.** CERN is now using the latest Intel Xeon processors across the whole organization and can integrate a mixture of other solutions with ease, which is a huge benefit
- **Performance boost.** CERN used Intel® SSD Data Center S3700 Series (Intel® SSD DC S3700 Series) to improve performance within its private virtualization environment and accelerate its block storage solution, based on Ceph*

## Impact

- **Energy efficiency.** Intel® technology helps deliver the performance and energy efficiency needed to support CERN's OpenStack solution and helps to keep power costs as low as possible
- **Easy refresh.** Using a standard platform makes it quicker and easier to refresh solutions regularly and cheaper distribution channels can be taken advantage of
- **Accelerated storage.** Thanks to Intel SSD DC S3700 Series, CERN's own tests have shown that it can at least double the input/output operations per second (IOPS) capacity and reach up to three to five times the burst performance[1]

## Unexplored horizons

Situated on the French-Swiss border, CERN's unique LHC creates close to light-speed particle collisions at four points. During these collisions, the temperature generated is around 100,000 times hotter than the heart of the sun. At each point of collision is an independent experiment, with its own data center for the immense amount of data that is captured by sophisticated sensors. These four LHC experiments must each filter all the data collected to retain relevant figures. Then, the data is transferred to CERN IT and can be accessed from all over the world from the main data center in Geneva.

Following the discovery of the Higgs boson particle, CERN is restarting the LHC to research matter from the 95 percent of our universe that is still unknown. In the first phase, protons were colliding at eight teraelectron volts (TeV), but now they have reached a much higher collision energy of 13 TeV. Now, these pioneering scientists and engineers are venturing into hitherto unexplored territory. This creates a need for technical solutions that are flexible and easy to adjust, since the range of the sensors must be finely tuned to search for a variety of matter within different planes. Because of the increased energy and rate of collisions, the scientists are also having to sift through much more data. Many data points could point to a potential discovery.

"Filtering will be heavily impacted," said Dr. Olof Bärring, leader of the Facility Planning and Procurement Section in CERN's IT Department. "We need more computing and storage capability, and more efficiency and flexibility from the IT infrastructure for this second phase. During the LHC's first three-year run, the CERN data center stored approximately 75 petabytes of data from the experiments, equivalent to about 3 million Blu-ray discs (i.e., almost 700 years of HD video). Over 100 petabytes of data are permanently archived on tape, and that is only data

# CERN sifts through a universe of data with the Intel® Xeon® processor E5 v3 family

identified as significant. The quantity of data filtered and saved is an infinitesimal fraction of the data collected. To cope with the even greater quantity of physical data being generated, filtered, and stored, we had to find an extremely reliable and powerful platform. At the same time as we were increasing power and efficiency, we also had to stick to the same budget."

## Fathoming data

The huge amount of data produced by the LHC experiments – in the range of 1 PB every second – requires significant processing power for both filtering and reconstruction. CERN tends to refresh its IT infrastructure on a cyclical basis. In preparation for the LHC restart, the IT team was looking to acquire a fresh platform to serve across the different projects, with the aim of doubling capacity to meet expected data rates. So, the team ran a proof of concept comparing identical configurations, which demonstrated to CERN that a new server based on a S2600KP board running the Intel Xeon processor E5-2630 v3 is approximately 20 percent more efficient than the previous server platform based on a S2600JP board running an Intel Xeon E5-2650 processor v2. It also provides the same performance level on the HEP-SPEC06 benchmark, so CERN IT are benefiting from the upgrade to the Intel Xeon processor E5-2600 v3 product family.

CERN IT and the LHC Experiments teams had the opportunity to test and validate the Intel Xeon processor E5-2680 v3 prior to launch and recognized a notable improvement in performance and performance per watt. The scientists who selected this processor are able to process more data in a shorter timeframe, increasing productivity and enabling a greater range of data to be analyzed more quickly. Total power consumption of the system is a key part of CERN IT's adjudication criteria for the tender. Results of an internal total cost of operation (TCO) evaluation that measured power consumption characteristics helped the team to select this platform. "Our contract adjudication included a TCO analysis that was fundamental in choosing the new Intel Xeon processor E5 family-based server platform," said Bärring.

CERN is also using Intel SSDs to accelerate its distributed storage solution based on Ceph. "The amount of data collected is so massive that it is impossible to store on SSDs," said Tim Bell, leader of the IT Operating Systems and Infrastructure Services Group at CERN. "So, we are not using them for saving LHC physical data but as an efficient way to accelerate the read and write capabilities of the Ceph journal, our block storage solution for cloud infrastructure based on OpenStack. Using this method we can look to double the IOPS capacity and potentially reach up to five times our previous burst performance." CERN is using Intel SSDs with bcache to increase the IOPS for specific virtual machines where read/write operations on local spinning disks could represent a bottleneck.

## Optimized solution

CERN is a unique research center, but CERN IT prefers to run an optimized standard infrastructure for maximum flexibility. "We have to be careful to invest granted funds wisely," said Bärring. "Thankfully, we do not have to rely on custom-built, high-end solutions but can take full advantage of impressive industry-standard, high-volume servers based on x86 64-bit architectures such as the Intel Xeon processors. Use of standard solutions helps to reduce costs and allows for a quicker and easier regular refresh of the solutions so we can make the most of better memory and stronger networking. Plus, the IT team can get the best out of commercial off-the-shelf Intel technology thanks to how the platform is optimized. It is fantastic that manufacturers like Intel give us early access to technologies that make standard solutions work smarter."

CERN is required to run competitive tenders, and prioritizes tenders from Europe. Being able to choose a flexible platform was very important. The majority of the deal was handled by Megware, Bios-IT, and Transtec; Megware offered the latest Intel PCSD server platform powered by the Intel Xeon processor E5-2630 v3 and E5-2680 v3, Transtec used a Quanta server running on the Intel Xeon Processor E5-2650 v2, and Bios-IT delivered a Supermicro Server running on the Intel Xeon Processor E5-2650 v2. "With Intel technology, we can use a mixture of solutions with no problem, and this flexibility is a huge benefit," said Bärring.

## Strong relationship

Intel has worked closely with CERN throughout the search for the Higgs boson and beyond, providing ongoing support and consultancy. "We rely heavily on technology so we can continue to innovate and provide outstanding services, and this is something that Intel provides us with. Without new lines of processors and platforms that generate increased performance and power at a reasonable cost, we would not be able to deliver the services we can today and must in the future," said Bärring.

**Lessons Learned**

CERN is one of the world leaders in pushing boundaries but, as research expands into different areas, researchers need more compute capacity and storage capability, as well as greater efficiency and flexibility from a computing infrastructure. CERN IT is impressed by the improved scaling of the latest high-performance computing data processors and estimates it will be able to collect 50 times more data by 2020. This upgrade is a crucial step towards achieving that performance level.

Some years back, CERN also purchased licenses for Intel compiler and Intel software development tools used to optimize high energy physic applications for a better usage of the deployed Intel processors.

The whole solution is implemented by CERN IT in a private cloud based on the Icehouse* version of OpenStack powered by industry standard systems including Intel technology, with one single cloud across two physical data centers. CERN OpenStack cloud is currently made up of 120,000 cores in 4,700 hypervisors. There are around 12,000 virtual machines. The OpenStack framework allows CERN IT and LHC experiments to benefit from a more flexible infrastructure, allowing the scientists to quickly schedule the appropriate workload thereby reducing inactivity time. The additional capacity added to this OpenStack infrastructure boosts the overall efficiency of the cloud computing infrastructure.

This framework also allows CERN to benefit from others' research and contribute improvements that people in the community (companies, researchers, and developers) can use. "Ever since the start of the World Wide Web, CERN has been strongly involved in open-source collaboration with other research centers around the world," said Bell. "With our OpenStack framework, other researchers can use the same software as CERN, so the extreme computing challenges of the LHC contribute to the global knowledge base of mankind."

Find the solution that's right for your organization. View **success stories from your peers**, learn more about **server products for business** and check out the **IT Center**, Intel's resource for the IT Industry.

[1] Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to http://www.intel.com/performance.



0415/JNW/RLC/XX/PDF

332426-001EN