Case Study
Intel® Xeon® Processor E5-2650 v2 Product Family
High-Performance Computing
Cloud Computing

Seeing the universe through the clouds

CERN upgrades to the latest Intel® Xeon® processors to support its new cloud environment

"Physicists can now request a virtual machine in 15 minutes rather than previously, where they would be waiting weeks for physical hardware"

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

At CERN, the European Organization for Nuclear Research, physicists and engineers are probing the fundamental structure of the universe. They use the world's largest and most complex scientific instruments to study the fundamental particles that make up matter. The organization's famous Large Hadron Collider (LHC) creates close to light-speed particle collisions that have led to the groundbreaking discovery of the Higgs boson particle. However, 96 percent of our universe is still unknown and the challenges ahead for the scientific community are striking.

Challenges

- **Regular refresh.** CERN needs to upgrade its data center resources regularly to keep up with calculation requirements
- **Peak performance.** It needed to ensure it has sufficient compute capacity and performance to support the next run of the LHC
- **Greater agility.** CERN wanted to increase flexibility while maintaining staffing levels

Solutions

- **Latest technology.** The Intel® Xeon® processor E5-2650 v2 product family was tested using an industry benchmark and CERN's own analysis. Its performance and energy efficiency tested better than previous generations of the processor
- **Two phases.** CERN deployed 880 CPUs to support current requirements, with an extra 2,200 to be deployed during summer 2014, ahead of the next LHC run planned for 2015
- **Virtual environment.** The project will include these new resources as virtual machines into CERN's OpenStack cloud infrastructure

Impact

- **Simple switching.** CERN's internal tests have shown that researchers can start up virtual machines in just 15 minutes
- **Flexible support.** Computing resources can be assigned to whichever activity is the top priority at a given time
- **Efficient use.** All servers can be used at all times, reducing power waste

Big changes for small particles

During the periods that the collider is switched on, CERN's two primary (Tier 0) data centers must be able to handle the huge spike in computing capacity and processing power that is needed to analyze the new results as they come in.

When the LHC is shut down to undergo sensitive upgrade and maintenance work, researchers are kept busy carrying out highly complex, data-intensive calculations and simulations to analyze the collected data and make discoveries that push forward our understanding of the universe.

Both use cases require powerful computing resources that must be regularly updated to keep them at the peak of performance. "We carry out regular technology refreshes to make sure we have optimum resources to meet our experiments' requirements," explains Dr. Olof Bärring, leader of facility planning and procurement section, IT department CERN.

With the next switch-on of the LHC about a year away, the team at CERN wanted to ensure it was ready. "We're expecting twice the energy from collisions next time we run the LHC, and this will drive a significant increase in the volume of data we'll be handling," says Tim Bell, leader of the IT operating systems and infrastructure services group, CERN. In addition to its usual hardware update, the CERN team wanted to move to a private cloud-based model in its main data centers to boost its efficiency and build in the capacity to expand its total compute capabilities as needed. "With a cloud environment, we can simply add virtual machines to create a larger computing resource without having to hire new staff to look after physical servers," adds Bell. "This means we can grow while keeping staff costs stable."

This meant it had an additional requirement for its new server technology. Besides the usual performance and energy efficiency benefits, it needed to support the organization's new OpenStack cloud environment.

Scientific groundbreaker implements new cloud platform powered by Intel® Xeon® processor E5-2650 v2 product family

Careful evaluation

CERN always runs performance benchmark tests prior to any new server implementation to ensure that any new technology introduced to its environment meets its stringent standards. It uses the HEP-SPEC06* benchmark, which is based on the widely used industry-standard SPEC* CPU2006 benchmark suite. This enables the team to configure their applications to test the capabilities of new server technologies. "With Intel® technologies, we see an increase in performance per processor each time, so it was no surprise that the latest Intel Xeon processor E5-2650 v2 product family also de-livered a performance improvement," says Bärring.

The team also conducts its own total cost of operation (TCO) calculation for any new processors, which assesses their power consumption characteristics. It found that the latest Intel Xeon processors performed 10 to 15 percent better in this calculation than the previous generations of the technology. "TCO analysis was fundamental to select the new Intel server platform as the preferred solution in our tenders," says Bärring.

Following these tests, CERN implemented 880 CPUs of the Intel Xeon processor E5-2650 v2 product family in its two Tier 0 data centers, with support from OEMs E4 Computer Engineering SpA and Action S.A.. The platform will support the computing needs of CERN and the organizations with which it collaborates during the end of the LHC's upgrade period. A further 2,200 CPUs will be added as a second phase in the run-up to it being switched on again.

The CERN scientists will use these OpenStack cloud virtual machines to run their simulations and calculations moving forward.

More flexible resources

The flexibility of the OpenStack platform makes it quick and easy for researchers to switch virtual machines from one application to another. "We now have an elastic pool of resources that enables scientists to adjust workloads depending on their needs, without having to raise a ticket and wait for the IT department to address it for them," says Bell. "It's now possible to start applications on a virtual machine in just 15 minutes, whereas before it could take up to several weeks to obtain the physical resources. This means they have a lot more time in which to carry out calculations, enabling them to drill deeper into the data and pull out more insights than was possible previously."

This gives the CERN team much greater flexibility to focus on the research that's most important at any given time. For example, a scientist planning to speak at an upcoming physics conference may need to focus on data analysis instead of simulation for a week or so ahead of his presentation to gather the results to illustrate what he will be saying. The cloud framework allows him to switch from simulation to analysis quickly, while the high-performance Intel Xeon processors provide the power to generate a large volume of in-depth results.

When the LHC is switched on, Data Acquisition and Filtering (DAQ) servers handle the data generated by the detectors. When the LHC is shut down, DAQ servers are not active and can be allocated to support ongoing simulation work. Then, when they are needed for a specific task, they can be quickly reassigned. This helps maximize the use of all available resources, which is particularly important for CERN, since all servers need to be kept switched on at all times to reduce condensation build-up in the underground data centers. The new cloud-based model ensures that all machines can be put to good use at all times, thereby reducing power waste.

Lessons Learned

CERN is world renowned as a driver of innovation and scientific breakthroughs, such as the discovery of the Higgs boson particle in 2012. To maintain this reputation, it needs to put a lot of thought and dedication into the upkeep of its technology resources and to apply its spirit of innovation, as well as following tried and tested best practices. It is for this reason that it has collaborated closely with the industry, including Intel, over many years to continuously drive innovation in its solutions. With this particular project, it has combined its annual technology refresh with a bigger shift to the cloud, creating a new computing environment that will deliver the flexibility, capacity and performance essential to the next round of LHC analysis and breakthroughs.

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

0614/JNW/RLC/XX/PDF 330701-001EN