

CASE STUDY

High Performance Computing (HPC)
With Intel® Omni-Path Architecture



Simon Fraser University's Cedar HPC Cluster Makes Large Impact in British Columbia

From its more than 60,000 cores to a highly power-efficient data center, Cedar is turning heads in computational research, AI, and local business



Cedar highlights:

- Phase 1: 902 nodes of 2x Intel® Xeon® processor E5-2683A (launched April 2017)
- Phase 2: 640 nodes of 2x Intel Xeon Platinum 8160 processor (launched April 2018)
- Intel® Omni-Path Architecture fabric
- Supports traditional research computing, OpenStack* cloud, and Singularity* containers

Executive Summary

British Columbia's [Simon Fraser University](#) (SFU) was overdue for an HPC upgrade. In early 2017, the university installed Cedar, a 902-node dual-socket cluster built on Intel® Xeon® processor E5-2683A with 16 cores each. At installation, Cedar was Canada's most powerful academic resource for advanced research computing (ARC). Today it supports researchers in a wide variety of fields: traditional simulation research, social science research, and Artificial Intelligence (AI).

Challenge

Before 1999, at Simon Fraser University (SFU), computational researchers such as Martin Siegert ran their own simulations on self-built computers.



But, in 1999, Simon Fraser University's IT department hired Siegert to build their first centralized computational cluster. "That was when people were still building Beowulf* clusters by piecing together desktop-like systems that researchers could run simulations on," stated Siegert. "We started fairly small at that time, and we've been growing ever since."

The university's first cluster was an eight-node system. The second one was commissioned in 2002—a 96-node cluster with dual-processor servers. The third system, commissioned in 2009, was a 160-node High Performance Computing cluster with quad-core Intel Xeon processor 5430 compute nodes and InfiniBand* DDR Architecture as the interconnect. Two years later it was expanded with an additional 256 nodes of dual six-core Intel Xeon processors X5650. "In those days, it was state of the art," added Siegert. "It supported all the different research going on at the university and WestGrid."

For the next five years, researchers ran their computations on this system, while the University waited for additional funding to build a new HPC resource.

Solution

In 2016, SFU was selected by Compute Canada, the national organization for advanced research computing (ARC), to house a new national HPC system in its data center. That was the first step that led to the design of a 902-node, 1.3 petaFLOPS HPC cluster—what would become the most powerful ARC cluster in the country at its launch.

Designing Cedar

"The new system was designed to do almost everything well," commented Siegert. "Its initial name was GP2, General Purpose system #2, but we ended up naming it Cedar, after British Columbia's official tree, the Western Red Cedar. It was designed to serve researchers from all areas of science, which is why Cedar is not a homogeneous system."

Cedar was built by Scalar Decisions and Dell, and it has various types of nodes:

- Traditional compute nodes, with two 16-core Intel Xeon processor E5-2683 v4 with 4 GB memory per core, form the “workhorse” cluster on which most computing is done.
- Fat nodes with up to 3 TB of memory per node for workloads, such as bioinformatics, that are not designed for massively parallel computing.
- GPU nodes with four NVIDIA* P100 cards, used mostly for molecular dynamics and Artificial Intelligence (AI) applications.
- A 15 PB storage system serves the entire cluster.

The Intel® Omni-Path Architecture (Intel® OPA) provides the interconnect across the entire cluster and the storage system.

“Initially, the design was for a system with islands comprising 32 nodes in each island,” said Siegert. “Within each island, the network topology is non-blocking. We expected to be able to run parallel applications using up to 1,024 cores within an island.”

During the RFP process, Siegert and his colleagues realized they could use the Intel OPA network to design a far better infrastructure. “The whole Cedar system is now using what is essentially a homogenous network based on Intel OPA. We still have islands, but the network architecture results in only a 2:1 blocking factor between islands, which for most applications has no negative impact on performance. What’s more important is the latency for the applications, and that’s not affected by the level of blocking. So, we can run far larger parallel workloads—essentially across the entire approximately 30,000 cores—than what we initially had in mind.”



Partners from Compute Canada and WestGrid tour the new SFU Data Center in Burnaby, Canada where the Cedar supercomputer is located. (Photo by Greg Ehlers courtesy of Simon Fraser University)

Expanding with Intel® Xeon® Scalable Processors

In 2018, Cedar got an upgrade. “After one year of operation,” added Siegert, “we found we had a larger need in Cedar’s compute section. So, we recently purchased an additional 30,720 cores consisting of 640 dual-socket nodes with 24-core Intel Xeon Platinum 8160 processors.” That expansion is larger than the original configuration, creating a cluster with more than 60,000 cores of Intel Xeon processors. With the larger 48-core nodes, it is also possible to run larger shared memory applications on a single node, giving researchers larger resources on which to run more massively parallel workloads.

Results

Cedar serves a wide range of scientific research, such as large molecular dynamics simulations done in chemistry for drug design, which is almost exclusively done on the computer first. Bioinformatics workloads are used to study complex problems in genome analysis and protein folding. Material science researchers also run large scale simulations. Other areas include social science and AI. “We’ve seen a huge growth in artificial intelligence and deep learning applications,” commented Siegert, “such as in natural language processing.” The system is available to all faculties and researchers across the country from all disciplines. Consequently, SFU has a huge spread of applications going across all sciences and the arts and social sciences.

One interesting project that Siegert says stands out was criminology research that analyzed real police data. “The data we received were raw, and the first thing we needed to do was de-identify, or anonymize, the data. Then we created the databases for the researchers, and they ran analyses on those data.”

According to Siegert, one AI research group is using Cedar to build an English to French translator. They are using AI on Cedar’s GPU nodes to train their translator to improve its algorithms. “They’ve actually won several competitions on their translator program,” added Siegert. The Intel OPA interconnect supports all multi-node GPU-based applications.

Containers and Clouds

A requirement for Cedar was for it to run an OpenStack cloud. Researchers needed an infrastructure where they could stand up their own environments, load their own operating systems and applications, and run their workloads. To accommodate these users, Siegert and his colleagues partitioned 128 nodes to run OpenStack. They added 10 gigabit Ethernet* for OpenStack to these nodes along with Intel OPA, so the nodes could be dynamically reassigned to run in the cloud or in the larger HPC cluster.

Cedar runs on the CentOS* 7 distribution of Linux*. But several researchers at SFU work on the CERN Atlas project, processing massive amounts of data from the Large Hadron Collider (LHC). That project runs its codes on CentOS 6. To stand up CentOS 6 nodes on Cedar, the Atlas team developed a method of running Singularity* containers on the cluster. The Intel OPA fabric supports the entire cluster, both the containerized workloads and traditional applications.

Data Center Impact

“Installing a system the size of Cedar at the University had quite a ripple effect,” commented Siegert. “Researchers, who in previous years had tried to run their applications on smaller systems, started using Cedar. It made a large difference, and others noticed, so we were able to attract researchers from other areas that we had not worked with before.”

That was just the beginning. People also noticed what SFU was doing with their new data center.

To house Cedar, SFU built an entirely new data center with energy efficiency as a key requirement. It has a Power Utilization Efficiency of 1.07—only seven percent of consumed power goes to utilities, e.g., cooling pumps.



Cedar is the most powerful academic supercomputer in Canada and will support researchers across the nation.
(Photo courtesy Simon Fraser University)

Research groups across the city realized that SFU had a very large and energy efficient data center. And they wanted to co-locate some of their applications in the facility. According to Siegert, the response has been somewhat overwhelming. "It's clearly more attractive for companies in the city to house their systems here than far away in another efficient data center."

Find out more about Cedar at www.sfu.ca/sfunews/stories/2017/04/canadas-most-powerful-academic-supercomputer-will-launch-at-sfu.html.

Find out more about Intel® Omni-Path Architecture at www.intel.com/omnipath

Summary

After several years of running older technology, Simon Fraser University was able to install a powerful supercomputer for academic research in Canada in 2017. Since then, Cedar has supported many new users, expanding research from traditional areas, such as physics and chemistry, to social sciences and AI. The data center's power-efficient design even attracted other companies to co-locate their own hardware in the data center to save on operational costs. In April 2018 Cedar doubled in size to more than 60,000 cores, creating a very powerful, Intel Xeon processor-based cluster with the Intel OPA fabric for academic research across Canada.



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at <https://www.intel.com/content/www/us/en/products/network-io/high-performance-fabrics.html>.

Intel, the Intel logo, Xeon and Phi are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.

© 2018 Intel Corporation