

Case Study

High Performance Computing
Intel® Xeon® Gold Processors
Intel® Optane™ Persistent Memory
Intel® SSD Drives



Arbutus Opens Doors to New Cloud Computing and Machine Learning Resources for Researchers

Scientists leverage University of Victoria's next-generation cloud infrastructure to expand knowledge using non-traditional high-performance computing

"When a researcher requests an environment, what we consider their own virtual lab, we set up the network and hardware to support their work. They can then create their own virtual lab in minutes, with or without support from our services."

—Belaid Moa, Ph.D., Advanced Research Computing Specialist with the Research Computing Services department, University of Victoria

Executive Summary

[University of Victoria](#), or UVic, located on Canada's Vancouver Island, is home to over 22,000 students and hundreds of faculty/researchers. Its [Research Computing Services](#) (RCS) unit provides Advanced Research Computing (ARC) infrastructure and services to university researchers, scientists at institutions across the country, and through international collaborations. The facility hosts one of [Compute Canada's](#) ARC data centers and the Arbutus cloud, an OpenStack cloud with an emphasis on hosting virtual machines and other cloud workloads. Arbutus was designed to augment traditional large-cluster HPC workloads and support research projects that require different capabilities than traditional HPC clusters, including online machine learning/artificial intelligence, big data, and collaborative computing. Arbutus was built with Lenovo SR630, SR670, and SD530 nodes, using 2nd Gen Intel Xeon Gold processors with Intel Optane persistent memory and Intel SSDs.

Challenge

In 2015 UVic, in partnership with Compute Canada, WestGrid and the University of Sherbrooke, launched the first phase of Arbutus to enable a new generation of researchers. Unlike researchers who rely on traditional HPC clusters to run massively parallel computing jobs or large-scale, simulation-focused workloads, these investigators had different needs.

"Our existing IT services at the time did not have the infrastructure that could provide the answers to some of our researchers' advanced computing needs," [Belaid Moa, Ph.D.](#), Advanced Research Computing Specialist with the Research Computing Services unit, University Systems Department, commented. "We had HPC clusters, but researchers were in dire need for high availability collaborative platforms, customized web sites, root access, micro-services environments, and other services of cloud computing, which was rapidly becoming as important as HPC clusters and an essential ARC service for many researchers."

Thus, the private Arbutus OpenStack cloud infrastructure was built. This first phase included 7,000 CPU cores of Intel Xeon E5-2680 v4 processors across 250 nodes with on-node storage, 10 GbE networking, and 1.6 PB (4.8 PB total) of triple-redundant Ceph storage. Arbutus uses virtualization to deliver Infrastructure-as-a-Service (IaaS) resources to support researchers' diverse workloads.

Over the following four years, new research projects were launched, many of which began using emerging technology capabilities and research environments, such as machine learning (ML), artificial intelligence (AI), JupyterHub, and big data. These new projects, along with increasing demand for cloud services, required more



Lenovo and Intel are working together to accelerate the convergence of HPC and AI, creating solutions of all sizes to unlock new levels of customer insight. Through collaboration on systems and solutions, software optimizations and ecosystem enablement, their goal is to speed discovery and outcomes for the world's most challenging problems in the Exascale era and beyond. Lenovo servers, the leading system choice for the TOP500 fastest supercomputers, are powered by Intel Xeon Scalable processors and Intel's leading-edge technology for storage, memory, and software, providing the innovative foundation to quickly drive forward science and industry progress.

storage, advanced computing, and larger memory pools—leading to a larger cloud infrastructure and Arbutus phase 2.

Solution

Arbutus Phase 2 was deployed in early 2020. The new system comprises an additional 208 Lenovo ThinkSystem SR630, SR670, and SD530 nodes with 119 GB of ThinkSystem TruDDR4 memory and 1 TB of Intel Optane persistent memory (Intel Optane PMem) per node. The expansion gives UVic 7,968 more cores of 2nd Gen Intel Xeon Gold 6248 processors and Intel Xeon Gold 6130 processors to add to their cloud infrastructure. The Ceph platform was expanded 5.7 PB of SSD storage with Intel SSD S4610 drives. The cloud nodes included two new Database-as-a-Service nodes to offer dedicated high-performance structured data access via SQL.

Intel innovations in memory, storage, and processor performance offer new capabilities to UVic. Intel Optane persistent memory enables very large memory capacities per node with DRAM-like performance. Intel Optane PMem can be used as extremely large memory in memory mode, or as a persistent storage with low latency and DRAM-like access in storage over app direct mode. Using Intel Optane PMem in memory mode, each node in Arbutus Phase 2 balances the advanced computing performance of 2nd Gen Intel Xeon Gold 6248 processors with additional memory capacity.

With the high performance and much larger memory per node, UVic can run many more virtual machines per server to support the growing base of researchers, especially in their needs for persistent workloads that must run 24/7 to support their projects. Advances in Intel Xeon processor architecture with Intel® Deep Learning boost (Intel® DL boost), and software specific to deep learning, such as Intel® Optimizations for TensorFlow and Intel® Distribution for Python, help speed ML tasks when codes are compiled for the 2nd Generation Intel Xeon Scalable processors.

Result

With Arbutus phase 2, UVic's Research Computing Services can support many more researchers across the country with more capable cloud—and even smaller-scale HPC—computing resources. While the facility continues to support large traditional supercomputing workloads with its big HPC clusters, researchers can also run smaller parallel jobs on Arbutus and ramp them up more quickly than waiting for a window on larger machines.

“When a researcher requests an environment, what we consider their own virtual lab, we set up the network and hardware to support their work,” Moa explained. “They can then create their own virtual lab in minutes, with or without support from our services.”

According to Moa, Arbutus allows users to choose from different ML environments, such as TensorFlow, PyTorch, Julia, Pandas, scikit-learn, and Apache Spark. These environments rely on Conda distributions. The Conda distribution uses the Intel MKL for low-level operations when using packages, such as numpy, scipy, and scikit-learn. In the future, UVic will install Intel Optimizations for TensorFlow and Intel Distribution of Python.

“Some virtual labs are even running small-scale HPC workloads, such as GROMACS, the molecular dynamics software used for studying things like the SARS-CoV-2 virus,” concluded Moa.

[Professor Dennis K. Hore, Ph. D.](#) is familiar with GROMACS and the capabilities of leveraging the cloud for research. He is a researcher and professor in the Chemistry and Computer Science departments at UVic with a team of 25 researchers working across [15 different projects](#).

“Most of my projects over the past 15 years revolved around studying how molecules interact with surfaces,” Professor Hore stated. “As an example, there are a lot of plastics used in the human body: catheters, stents, sutures, artificial organs, and others. My team studies how proteins interact with them, trying to get at the molecular basis of biocompatibility.”

But over the last three years his team has launched a project that combines chemical analyses with big data and machine learning using Arbutus to help improve the lives of people using non-prescription street drugs.

“At three different sites across Victoria,” Hore added, “we work anonymously with people to inform them about the makeup of drugs they bring in voluntarily for analysis. We run a host of chemical analyses on their samples using state-of-the-art analytical instruments. We then use the data we collect along with chemical libraries and databases to build machine learning algorithms and applications. One of the goals of the program is to convey information that enables



Drug testing is provided to the public by the Vancouver Island Drug Checking project, in collaboration with Health Canada and the University of Victoria. Photo credit: Jay Wallace



For infrared absorption spectroscopy measurements, a small quantity of sample is placed on a crystal through which IR light is reflected. Photo credit: Jay Wallace

people to make informed decisions on the use of their substances, according to their composition and strength.”

This is just the tip of the iceberg, according to Hore. In addition to the human benefits, the data, and computing science his project offers, there are potential applications for remote healthcare. The learnings and applications from this research could lead to the development of portable devices and kiosks that can quickly and interactively analyze chemical compounds. This remote analysis, using online interactive machine learning, can then provide insight to the potential effects of the sample and guidance to those seeking its analysis.

“The project began 4 years ago after a harm-reduction pharmacist contacted me to do quality control on one of the prescription drugs he was dispensing,” Hore explained. “While he had been acquiring the particular drug for years from the same manufacturer, his customers were telling him it was affecting them differently from past usage. He wanted analysis of the drug with concentrations of its constituents. This is another potential application of the science.”

That request led to the Drug Checking project that is helping build new knowledge for social and chemical sciences using advanced computing technologies with cloud computing infrastructures built on Intel® architecture. The implications of the research could result in improving healthcare, public safety, and other fields.

Solution Summary

At University of Victoria, demand by researchers for more and different types of cloud computing resources prompted the University’s Research Computing Facility to expand their existing Arbutus infrastructure. Arbutus phase 2 added nearly 8,000 new cores with advanced Intel Xeon Scalable processors, Intel Optane persistent memory, and Intel SSD S4610 drives. The larger, more advanced Arbutus is used for a wide range of computing, including web services, AI/ML, and big data.

Where to Get More Information

Learn more about [U Vic’s Research Computing Services](#).

Find out more about [Arbutus](#).

Explore the capabilities of the [2nd Generation Intel Xeon Scalable processors](#) with integrated Intel® Deep Learning Boost capabilities for accelerated AI inferencing.

Discover [Lenovo HPC solutions](#).

Solution Ingredients

- 208 Lenovo ThinkSystem SR630, SR670, and SD530 nodes with ThinkSystem TruDDR4 memory
- Nearly 8,000 cores of 2nd Gen Intel Xeon Gold 6248 Scalable processors and Intel Xeon Gold 6130 processors
- 1 TB of Intel Optane persistent memory per node
- Intel SSD S4610 drives



¹ Top500.org June, 2020 list

Intel technologies may require enabled hardware, software or service activation.
No product or component can be absolutely secure.

Your costs and results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.