# Expanding knowledge with high-performance computing in the cloud.

The **University of Victoria** uses Lenovo ThinkSystem solutions, powered by 2nd Gen Intel® Xeon® Scalable processors, to enable new types of research projects such as improving synthetic organs, with cloud-based artificial intelligence and machine learning.


**Lenovo Infrastructure Solutions for The Data-Centered**

intel XEON GOLD inside™

Lenovo

1

## Background

The University of Victoria (UVic) is one of Canada's leading research universities. Its campus on Vancouver Island is home to more than 22,000 students and hundreds of researchers, as well as some of the country's leading research institutes.

UVic's Research Computing Services (RCS) unit hosts one of Compute Canada's Advanced Research Computing (ARC) data centers, which provides high-performance computing (HPC) infrastructure and services not only to research teams at the university itself, but also at other institutions across Canada and around the world.

In addition to these HPC services, RCS also provides Arbutus, an OpenStack cloud solution that supports research projects with non-traditional computing requirements, such as online machine learning, big data analytics, and collaborative computing.

2

## Challenge

The origins of Arbutus go back to 2015, when UVic and its partners, Compute Canada, WestGrid, and the University of Sherbrooke, saw an opportunity to innovate by creating an infrastructure that would empower a new generation of researchers. Recognizing that, while traditional HPC clusters are extremely valuable for research projects that depend on massively parallel computing jobs or large-scale simulations, there was also a significant demand for more flexible, general-purpose cloud computing environments.

Belaid Moa, Ph.D. Advanced Research Computing Specialist with the RCS unit, University Systems Department, recalls: “Our existing IT services did not have the infrastructure that could provide answers to some of our researchers’ advanced computing needs. We had HPC clusters, but researchers were in dire need of high-availability collaborative platforms, customized web sites, root access, micro-services environments, and other cloud computing services.”

The Arbutus OpenStack cloud infrastructure was designed to meet these needs, and UVic delivered an initial deployment comprising 7,000 CPU cores distributed across 250 nodes with on-node storage, as well as 1.6 PB of triple-redundant Ceph storage. This infrastructure was capable of supporting thousands of virtual machines, each of which could be configured for the precise needs of individual research projects.

Over the next four years, many research teams launched projects running on Arbutus, which proved especially useful in emerging technology areas such as artificial intelligence and machine learning (AI/ML), as well as hosting JupyterHub workspaces that enable teams of researchers to collaborate on big data analysis. The success of these projects created even more demand for storage, memory, and advanced computing resources, which prompted UVic to start working on a second phase of the Arbutus initiative.

## Why Lenovo? Innovations in HPC and AI convergence.

UVic recognized the work that Lenovo and Intel are doing to accelerate the convergence of HPC and AI/ML technologies, and realized that the combination of Lenovo servers, 2nd Gen Intel® Xeon® Scalable processors and Intel® Optane™ persistent memory would create an ideal platform for Arbutus Phase 2. In particular, the Intel® Optane™ persistent memory would be extremely useful for AI/ML and big data workloads, as it gives the option of either massively expanding the amount of memory available per node or providing very fast persistent storage for applications.

The fact that Lenovo servers are currently the leading system choice for the TOP500 fastest supercomputers also gave the university confidence in selecting Lenovo as its server infrastructure provider.

## Designed for deep learning, built for big data.

Deployed in early 2020, Arbutus Phase 2 adds a further 208 Lenovo ThinkSystem SR630, SR670 and SD530 nodes, and nearly 8,000 more cores to the cluster with 2nd Gen Intel® Xeon® Scalable processors. Each node features an average of 119 GB of Lenovo ThinkSystem TruDDR4 Memory plus 1 TB of Intel® Optane™ persistent memory.

The new Lenovo and Intel infrastructure enables the second generation Arbutus cloud platform to offer higher performance data processing capabilities, with greatly increased memory capacity per node. As a result, UVic is able to run many more virtual machines per server, which means it can support more researchers and more projects than ever before.

The platform is well-equipped to support AI/ML workloads thanks to advancements made in the Intel® Xeon® processor architecture such as Intel® Deep Learning boost and software-specific features such as Intel® Optimization for TensorFlow and Intel® Distribution for Python. These allow researchers to compile their AI/ML code to take advantage of the full power of the 2nd Gen Intel® Xeon® Scalable processors.

The use of Intel® Optane™ in App Direct Mode increases the CPU utilization and utility of each server. In-memory databases, in-memory analytics frameworks, and ultra-fast storage applications are examples of workloads that benefit from using App Direct Mode. Using App Direct Mode also provides for substantially faster system restarts eliminating costly downtime.

Furthermore, in App Direct Mode data is encrypted using a key stored on the module in a security metadata region, which is only accessible by the Intel® Optane™ persistent memory controller. The modules are locked at power loss, and a passphrase is needed to unlock and access the data. If a module is repurposed or discarded, a secure cryptographic erase and DIMM over-write operation keeps data from being accessed.

3

## Results

Phase 2 of the Arbutus initiative has enabled UVic Research Computing Services to support a much larger number of researchers and projects, while providing a more comprehensive and flexible range of cloud services to complement its traditional supercomputing resources. For example, many AI/ML jobs can be run in parallel on smaller clusters of nodes over longer periods of time, allowing research teams to ramp up much more quickly than if they had to wait for an available window on a larger cluster.

Moa explains: “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. 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.”

Dennis K. Hore, Ph. D., Professor of Chemistry and Computer Science at UVic, adds: “Most of my projects over the past 15 years have revolved around studying how molecules interact with surfaces. For example, my team studies how proteins in the human body interact with plastics such as catheters, stents, sutures, and artificial organs, trying to get at the molecular basis of biocompatibility.”

UVIC

- **208 Lenovo ThinkSystem SR630, SR670, and SD530 nodes with Lenovo ThinkSystem TruDDR4 Memory**
- **7,969 cores of Intel® Xeon® Gold 6248 Processors and Intel® Xeon® Gold 6130 Processors**
- **1 TB of Intel® Optane™ persistent memory per average node**

## Drug checking to reduce overdoses and save lives.

Using Arbutus, Professor Hore and his team have launched a project that uses big data and machine learning to help improve the lives of people using non-prescription street drugs.

Professor Hore comments: “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 and use that data 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 people to make informed decisions on the use of their substances, according to their composition and strength.”

Initiatives such as these demonstrate how UVic’s researchers are harnessing Lenovo and Intel technologies to advance material, chemical, and medical sciences, and make a positive contribution to society, with benefits for healthcare, public safety, and other fields.

“Built on cutting-edge Lenovo and Intel technologies, the Arbutus platform has become as important as HPC clusters and an essential ARC service for many researchers.”

**Belaid Moa Ph.D.**
Advanced Research Computing Specialist, Research Computing Services,
University of Victoria

# What will you do with Lenovo HPC solutions?

The Data-Centered advance groundbreaking research with Lenovo smarter infrastructure solutions, powered by Intel®.

**Explore HPC Solutions**