McGill University expanding the frontiers of science

With an IBM iDataPlex and QLogic solution

Overview

The need

McGill University needed a more powerful high-performance computing solution for advanced research that was flexible enough for a variety of workloads and scalable enough for future expansion plans.

The solution

McGill deployed 1,200 IBM System x® iDataPlex® dx360 nodes using energy efficient Intel® Xeon® processors with QLogic TrueScale QDR InfiniBand switches and adapters.

The benefit

The solution is Canada's third-most powerful high-performance computing cluster for academic research, enabling scientists to conduct advanced research and providing the school with a foundation for future expansion.

Want to know more about star development? That means studying rotating celestial objects called pulsars through observational analysis. How do airplane wings accumulate ice deposits in various climates? It's a critical question for the aviation industry, but too dangerous to test in the real world. Ever wondered how air, rain, soil and vegetation collectively shape local ecosystems? Understanding these relationships in detail may lead to critical insights about the environment.

These complex questions now have a powerful answer at McGill. Using a high-performance computing (HPC) cluster from IBM that leverages QLogic switches and adapters, leading researchers from around the world are converging on the campus in Montreal, Canada, to reveal fundamental insights about our world.

In today's research environment, extremely powerful supercomputers are mandatory to sift through gargantuan data sets, explains Bryan Caron, Ph.D., director of business operations for CLUMEQ, Calcul Quebec and Compute Canada at McGill University.

"We look at these systems as a tool for discovery," says Caron. "With the volumes of data that we're looking at these days, supercomputing facilities are really key components of that entire workflow."

Momentum grows for supercomputing capabilities in Canada

McGill has long been a leading research center in the province of Quebec. Founded in 1821, it is now Canada's leading post-secondary institution with two campuses, 11 professional schools, 300 programs of study and more than 36,000 students, including 8,300 graduate students. These students come from more than 150 countries around the world and make up 20 percent of the student body.

"We wanted something that was modular and scalable to a large number of cores, both for the computers as well as for the storage. It also had to be commodity hardware, and I think we found a good balance with IBM iDataPlex."

—Bryan Caron, Ph.D., director of business operations, CLUMEQ and Calcul Quebec, Compute Canada, McGill University

McGill is one of several regional institutions that formed a high-power computing consortium in 2001 called CLUMEQ—an acronym for the various schools involved. The CLUMEQ organization is evolving with other HPC centers in the province to form Calcul Quebec, and works together with Compute Canada, Canada's national high-performance computing organization, to support leading edge research. HPC clusters were established at McGill and elsewhere in Quebec, establishing the region as a Canadian supercomputing hotbed.

Over time, the existing HPC facilities could not keep up with the demands being placed upon them. McGill and the CLUMEQ consortium needed a new HPC solution that could leverage the processing power necessary to facilitate the efficient capture, storage, search, sharing, analysis and visualization of vast amounts of research data, says Caron, a former physicist.

"The pure computational need exceeded the capability or the capacity within those facilities," adds Caron. "What we needed was a modular, flexible and scalable HPC cluster that could serve as an important resource for multidisciplinary research efforts across Canada."

Specific requirements for a flexible, scalable solution

The key was finding a general research computing platform that could provide the right balance for a variety of unique HPC workload requirements, from fast processors for rapid numerical calculations to a large memory footprint to aid in the analysis of large data sets. The solution also had to be modular and scalable to a large number of cores. "I think we found a good balance with IBM iDataPlex," says Caron.

With funding from the Canadian Foundation for Innovation (CFI), a government organization which promotes the development of world-class research and technology investment in Canada, McGill worked with IBM on an $17.6 million contract for a new HPC cluster. The first phase of the solution encompasses 1,200 IBM System x iDataPlex dx360 nodes using energy efficient Intel Xeon 5600 processors.

The resulting solution has 14,400 available cores, each with between two and six gigabytes of addressable memory, arranged in three partitions to handle different workloads. One partition is designated for general serial processing jobs, another is for large addressable memory jobs of up to one terabyte, and a third handles applications which need high-bandwidth storage partitions, explains Caron.

Solution components

Hardware

- IBM System x® iDataPlex® dx360 class
- Intel® Xeon® 5600 processors
- QLogic TrueScale QDR InfiniBand switches and adapters

Software

- IBM General Parallel File System (GPFS™)
- xCAT
- Linux
- ScaleMP

Underlying the solution are QLogic TrueScale QDR InfiniBand switches and adapters. The network includes two core QLogic switches which unify all of the iDataPlex nodes across the various HPC functions. QLogic host channel adapters reside within the iDataPlex nodes to provide communications capabilities for workloads and make the file system accessible to all nodes as well.

"Those are key for us," says Caron. "We're very happy with QLogic for things such as the adaptive and distributed routing capabilities and the tools that they provide. Working with QLogic has been very good."

On the software side, the solution runs IBM General Parallel File System (GPFS™), Linux, ScaleMP, the xCAT open source provisioning system and other specialized applications to deliver capabilities for rapid provisioning and management capabilities in the HPC environment.

Living up to expectations—with room for growth

Caron says the solution has performed exactly as was predicted by IBM. As a shared resource within the Compute Canada infrastructure, the solution has already been in high demand from researchers around the country. Within months of the phase one rollout, the cluster saw workloads capable of achieving 95 percent peak utilization, says Caron.

The demand is understandable given the solution's performance characteristics. The solution was ranked No. 83 in the November 2011 edition of the Top500 list of most powerful supercomputers worldwide.[1] According to McGill, it is Compute Canada's third-fastest HPC system.

"Having this facility come online to address these computational needs is very important," says Caron. "At the same time, it has been a very important tool to attract researchers from across Canada, because they know they have this available to get their scientific work done."

Looking ahead, Caron says planning is underway for a phase two expansion of the cluster, which will double the number of available cores and increase the storage capacity from two petabytes to four. Says Caron, "The deployed IBM iDataPlex solution is certainly a strong foundation on which to build in the future."

For more information

Contact your IBM representative or IBM Business Partner, or visit us at: **ibm.com**/systems/x/hardware/idataplex

For more information about QLogic, visit: www.qlogic.com

For more information about McGill University, visit: www.mcgill.ca

IBM Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
March 2012

[1] Top500 Supercomputing Sites, November 2011 ranking. (http://www.top500.org/lists/2011/11).

Please Recycle

XSC03108-USEN-01