

# CASE STUDY

## SciNet

Moab Adaptive HPC Suite™ delivers energy efficiency for one of the world's greenest data centers



Powered by Moab™

### KEY FACTS

#### Overview

The University of Toronto's SciNet Consortium provides resources to the University's scientific researchers and ten affiliate research hospitals by combining industry-leading innovations in data-center design.

#### Challenge

SciNet had to show it could meet strict budgets for ongoing energy use while providing multi-OS provisioning to meet the needs of its wide range of users. SciNet was looking for a very green, energy-efficient system in a very energy-efficient facility.

#### Solution

- Energy-aware, stateless, on-demand multi-OS provisioning
- Moab Adaptive HPC Suite™ and xCAT provisioning software from IBM
- 4,000-server IBM System x iDataPlex supercomputer system
- 30,000 Intel Xeon 5500 cores with a theoretical peak of 306 tflop/s

#### Results

A state-of-the-art data center that saves enough energy to power more than 700 homes yearly. On-demand provisioning allows users to make their OS choices part of their automated job templates. SciNet always has several different flavors of Linux running simultaneously.



### OVERVIEW

The University of Toronto's SciNet is one of seven regional consortia of universities and colleges across Canada with mandates to provide high-performance computing (HPC) resources to their own academic researchers, other users across the country, and support for international collaboration.

SciNet provides resources to the University's scientific researchers and ten affiliate research hospitals by combining industry-leading innovations in data-center design.

### CHALLENGE

SciNet was one of seven consortia universities that shared a \$120 million budget for the acquisition of new infrastructure. However, very little money was available for operations. Unusually, this had to include power, so the overriding issue was total cost of ownership (TCO).

SciNet's research activities required two different HPC system architectures. The first is a tightly coupled capability system with 104 IBM Power 6 nodes, 3,328 cores @ 4.7 GHz, 14 TB of RAM, and a theoretical peak of 60 tflop/s.

*"Why should we pay for cooling when it's so cold outside? Toronto is pretty cold for at least half of the year.*

*"We could have bought a humongous pile of cheap x86 boxes but couldn't power, maintain, or operate them in any logical way.*

*"We needed to partner with people who understood the problem. None of us had done something at this scale before."*

Dr. Daniel Gruner, Chief Technology Officer of Software for SciNet

The second system is a very large x86-based commodity cluster suitable for a large variety of serial and parallel applications.

In addition, the university didn't have space or power availability for a data center requiring a minimum of 1.5 MW on campus or in downtown Toronto and needed to optimize use of both infrastructure and operating budget. A significant research area for SciNet users is climate change and global warming, so creating one of the greenest data centers in the world was of key importance in this project.

The Power 6 and x86 compute clusters would be optimized for different types of computing, but they would eventually need to be intelligently managed as a single shared pool of resources. SciNet had to show it could meet strict budgets for ongoing power use while providing energy-aware multi-OS provisioning capabilities to meet the needs of its wide range of users.



## SOLUTION

The data center is located off campus in a newly renovated commercial complex and has a dedicated 4 MW power feed. Total floor space is 12,000 square feet, which is subdivided into a separate machine room (3,000 square feet), electrical space, mechanical room, and offices.

Energy efficiency was a key consideration in the design of the data center as well as in the selection of the computer equipment and power and cooling infrastructure. A cooling tower allows the facility to operate completely in "free-cooling" mode when the outdoor temperature is less than 4° C (39° F) and helps to lower cooling costs at all other times.

- 4,000-server IBM System x iDataPlex supercomputer
- 30,000 Intel Xeon 5500 cores, 60 TB of RAM—a theoretical peak of 306 tflop/s
- 10 times more powerful than any other system in Canada
- No. 16 on the June 2009 Top500 list

The IBM System x iDataPlex server is specifically designed for data centers that require high performance, yet are constrained by floor space, power, and cooling infrastructure. This system provides up to five times the compute density offered by competitive products and a unique water-cooled technology—IBM's Rear Door Heat Exchanger—that extracts more heat than the systems actually generate.

Moab Adaptive HPC Suite, combined with open-source xCAT software from IBM, delivers on-demand multi-OS provisioning capability that allows

*"How do we run it? Green, stateless on-demand provisioning. Moab and xCAT."*

Dr. Daniel Gruner

*"This was the first Moab-xCAT implementation on a scale of this size for an adaptive data center. We were pushing the envelope; we expected delays and issues. These large, exotic, and ambitious projects require customer and vendors cooperating together. It takes a partnership to deliver one of the greenest data centers in the world."*

Egan Ford, Executive IT Specialist, IBM

users to choose their preferred OSES on login or to make their OS choices part of their automated job templates. Applications requiring several different flavors of Linux always run simultaneously on the system. Moab and xCAT provide:

- On-demand multi-OS, stateless image provisioning
- Automatic health checks and amber-light avoidance
- Power monitoring
- Green mode: nodes powered off when idle, powered on and provisioned as needed to be ready when required
- Intelligent automation of extreme scale, heterogeneous systems

Scalability is critical because of the need to run a single job, a single application, across the entire system. To run a single job across thirty thousand cores or four thousand servers is something that very few people have ever done.

## FUTURE OPPORTUNITIES

Looking to the future, greater efficiencies can be obtained by combining the Moab schedulers running on both systems. This will enable fair-share scheduling on both systems with central accounting and reporting of unified usage statistics.

From the energy-efficiency perspective, the ability to measure

power usage during runtime and generate "power profiles" for jobs enables Moab to schedule jobs where they run best and most cheaply. To further reduce costs, power-hungry jobs can be scheduled when power is cheaper.

Another option is to schedule based on the spot price of electricity and to charge users accordingly as the price of power varies.

## RESULTS

Faced with extremely diverse usage requirements, Adaptive Computing's Moab technology provides SciNet users and administrators the ability to consolidate and virtualize resources across heterogeneous systems, dynamically allocate and manage applications, optimize service levels, and reduce operational costs.

Moab Adaptive HPC Suite reduces power consumption through intelligent, workload-aware power management and temperature-aware balancing.

When combined with additional energy-efficient technologies, including dynamic provisioning software that automatically turns off processors not in use, the state-of-the-art data-center design at the University of Toronto saves enough energy to power more than seven hundred homes yearly.

*"The most important thing about this data center is a Power Usage Effectiveness (PUE) value of 1.16. A typical data center with a PUE of 1.5 uses half a watt for cooling for every watt used in computation. That's a lot of power and, in our view, a lot of waste. For us the difference between a PUE of 1.5 and a PUE of 1.16 is about 700 KW or about \$500,000 a year. We wanted to be at least as good as Google."*

Dr. Daniel Gruner

**To learn more** about how Moab technology can improve HPC, data center, or cloud computing, contact us—

North America, Latin America, Asia-Pacific . . . . . +1 (801) 717-3700  
 U. S. toll-free. . . . . +1 (888) 221-2008  
 EMEA (Europe, Middle East, Africa) . . . . . +44 (0) 1483 243578  
 Email . . . . . info@adaptivecomputing.com  
 Web site . . . . . www.adaptivecomputing.com

