



## Where Technology Meets Medicine: SickKids High Performance Computing Data Centre

### The Hospital for Sick Children (SickKids)

Industry:	Healthcare
Stats:	15,337 patients per year
Location:	Toronto, Canada
Website:	<a href="http://www.sickkids.ca">http://www.sickkids.ca</a>



The Peter Gilgan Centre for Research and Learning brings together 2,000 scientists and research staff from different scientific disciplines and a variety of clinical perspectives, to accelerate discoveries, new knowledge and their application to child health. The facility physically connects SickKids science, discovery and learning activities to its clinical operations.

### Background

On an average day, 285 patients come through the doors of The Hospital for Sick Children (SickKids) in Toronto, and over the course of a year, over 15,000 will be admitted for treatment. We are all familiar with the role that doctors, nurses and hospital staff play, providing invaluable care to our loved ones. What we do not see are the legions of researchers and scientists working behind the scenes to analyze patient data, looking for trends that will ultimately direct treatment programs and improve patient care.

Biology and genomics have become a data science. Research discovery and innovation have moved from wet lab work to a world of high performance computing (HPC). The SickKids Centre for Computational Medicine is working to better understand the makeup of a human organism, analyzing terabytes of data to predict the minute differences from one person to the next. The only way to determine the nuances that differentiate us from one another is to analyze this data with algorithms. HPC has become essential, allowing researchers to discover the tiny distinctions that are otherwise impossible to find.

### A new era of personalized medicine

The “one size fits all” approach to medicine is changing. The notion of personalized medicine, made possible by the growing understanding of the human genetic makeup and high performance computing, is allowing for customized care, fitted to a person’s unique genetic makeup. “We are at the beginning of personalized medicine. Our ability to customize treatments is getting better and better. Your treatment plan is not optimized just based on your genome - it will be optimized based on your life factors, on the environment in which you live, and on your body type. All of this together can really improve and guide therapies,” said Dr. Michael Brudno, Director of the Centre for Computational Medicine at SickKids and Associate Professor of Computer Science at the University of Toronto.

*“We are at the beginning of personalized medicine. Our ability to customize treatments is getting better and better.”*



## Objectives

Architect and implement a best of breed HPC data centre

## Solution

- 7,000 cores, 14,000 compute threads
- 2 petabytes disk storage capacity, 4.5 petabytes tape storage capacity
- InfiniBand network
- Scalar project management

## Results

- Five-fold improvement in HPC capacity
- Reduced cluster reboot time from 4 to 6 hours down to 20 minutes
- Architected to scale as requirements grow with the ability to add compute and storage resources in a modular way

## Solution

### Best of breed high performance computing

In order to make personalized medicine a reality, SickKids has built the Peter Gilgan Centre for Research and Learning which has brought SickKids scientists and research staff together under one roof. The state-of-the-art laboratory and learning facilities house the HPC data centre which is comprised of best of breed components across the compute systems, network and storage infrastructure.

Scalar was awarded the data centre project after rigorous analysis through a detailed RFP process. “We felt we could do this differently. We created alignment between a number of technical vendors to create a sophisticated architecture that contained the absolute best technology in each area. Our role was to aggregate the technologies together and make it seamless for SickKids. Our mantra was no compromise on any one component,” said Neil Bunn, P. Eng - Chief Architect, Infrastructure and Cloud Solutions, Scalar Decisions.

“We went from 1,200 to approximately 7,000 cores. Storage capacity has gone up drastically, from 550 terabytes to 2 petabytes. We have moved from an Ethernet network to an InfiniBand network which is considerably faster and meets the requirement of 80 gigabytes per second from the compute nodes to the storage,” said Jorge Gonzalez-Outeirino, Ph.D., Facility Manager of the Centre for Computational Medicine at SickKids. The system is able to do 107 trillion calculations per second - one of the largest systems dedicated to health research. “We can now process 5 times the work in the same amount of time,” said Brudno.

A management software layer was added to seam and blend all of the different pieces of technology into a single interface, making it easy to manage the large, complex cluster. “In the old system, when we had a problem, it could take anywhere from 4 to 6 hours to reboot and reinstall the cluster. Now, after a shutdown, it only takes us 20 minutes to reinstall the nodes and reestablish services,” said Gonzalez-Outeirino.

After the architecture was finalized, implementation took place over a period of 3 months, followed by a 5-month period

***The system is able to do 107 trillion calculations per second - one of the largest systems dedicated to health research.***



**Compute Nodes**  
SGI® C-Series

**Storage**  
EMC® Isilon® X400  
System, SpectraLogic  
T950 Tape System

**Network**  
Mellanox® InfiniBand®  
SX6036 with Ethernet  
Gateway Functions and  
SX1024 Ethernet Switches

**Management,  
Provisioning &  
Monitoring**  
Bright Cluster Manager®

**Scheduling and  
Workload Management**  
Adaptive Computing  
Moab® HPC Suite

**Project Management**  
Scalar certified PMI  
Project Manager, single  
point of call solution

of acceptance testing. The move to the new building posed unique challenges - the physical data centre was not yet complete and there was no network connectivity in place. “Scalar worked closely with us through these challenges. We are talking about a big system that has a lot of potential for things to go wrong, but the process was very smooth,” said Gonzalez-Outeirino.

#### Looking Forward

#### High Performance Computing for Health Sciences (HPC4Health)

The next phase of this project sees SickKids and University Health Network (UHN) Princess Margaret Cancer Centre partnering to build a pilot IT infrastructure that will provide researchers and clinicians with secure cloud-computing services, that will also satisfy personal health information privacy requirements. The pilot project is called High Performance Computing for Health Sciences, also known as HPC4Health.

“Together, with Princess Margaret Cancer Centre we designed HPC4Health, a resource-sharing model to increase efficiencies not only to operating budgets, but also to help clinicians analyze patient tests faster,” said Brudno.

HPC4Health will initially be available to hospitals in Toronto, but will be expanded to include healthcare organizations throughout Ontario and eventually across Canada. This health node will provide significant savings to all partners offering more access to high performance computing without the cost of building and operating their own IT infrastructure.

“What we are doing here is truly making a difference and is going to have positive outcomes for patients in a very short period of time, and lead to scientific discoveries in many fields,” concluded Brudno.

For more information, please visit <http://www.scalar.ca>

For the latest news, visit our blog at <http://blog.scalar.ca>

*“What we are doing here is truly making a difference and is going to have positive outcomes for patients in a very short period of time, and lead to scientific discoveries in many fields.”*