

SciNet HPC - Niagara

SciNet
www.scinet.utoronto.ca
University of Toronto
Toronto, Canada

December 7, 2017

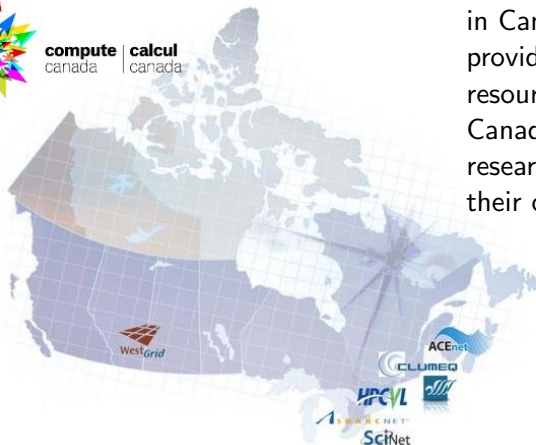
SciNet is ...

... a consortium for High-Performance Computing consisting of researchers at U. of T. and its associated hospitals.

One of 4 consortia in Canada that provide HPC resources to Canadian academic researchers and their collaborators.



compute | calcul
canada | canada



General Purpose Cluster (GPC)



General Purpose Cluster (GPC)

- 3864 nodes with two 2.53GHz quad-core Intel Xeon 5500 (*Nehalem*) x86-64 processors
- 16 GB RAM per node
- 16 threads per node
- 1:1 DDR (840 nodes) and 5:1 QDR (3024 nodes) Infiniband Interconnect
- 306 TFlops (261 HPL)
- #16 on the June 2009 *TOP500* supercomputer sites
- #179 on the June 2014 list, #3 in Canada

New HPC Compute Infrastructure in Canada

- CFI Stage 1 investment of \$75 million
 - 4 new sites
 - 1 cloud, 2 general purpose, 1 large-parallel
 - various storage
- CFI Stage 2 investment of \$50 million
 - 1 new general purpose, 1 expanded general purpose
 - more storage

Arbutus - Univ. Victoria - Westgrid

- 248 nodes (2x14 core Intel Broadwell)
- 10G Ethernet
- OpenStack Cloud
- 1.6 PB Storage
- September 2016

Cedar - SFU - Westgrid

- 27,696 cores (Intel 2x16 core Broadwell)
- 584 Nvidia P100 GPU's
- 2:1 blocking Intel Omni-Path
- Multiple memory configs (128G – 1.5TB)
- 4 PB Scratch, 10PB Project
- July 2017

Cedar Part2 - SFU - Westgrid

- X cores
- Intel Omni-Path
- Multiple memory configs
- July 2018

Graham - Waterloo - SHARCNet

- 33,376 cores (Intel 2x16 core Broadwell)
- 320 Nvidia P100 GPU's
- 8:1 EDR/FDR Infiniband
- Multiple memory configs (128G – 3TB)
- 3.6 PB Scratch, 10PB Project
- July 2017

GP4 - McGill - CQ

- X cores
- X GPUs
- X PB Scratch, X PB Project
- Sept 2018

Design Criteria

- \$17M
- Designed for Large (MPI) Jobs
- High Performance Parallel Filesystem
- Competitive RFP Process (Jan-July 2017)
 - Real HPC Codes (7) Benchmarks
 - Energy Efficiency (Flops/Watt)
 - Network Design
 - Deployment Plan
- 11 original bidders, 5 shortlisted
- Lenovo design chosen

Niagara - UofT - SciNet

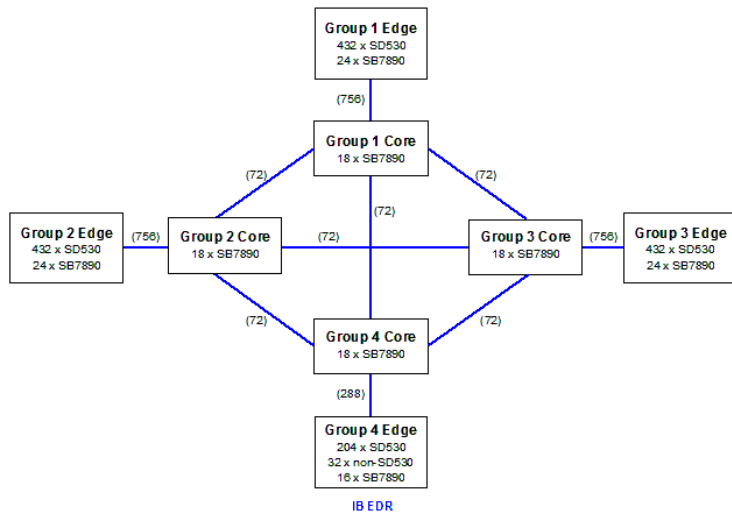
- 1500 Lenovo Nodes (2x20 core Intel Skylake @ 2.4GHz)
- 60,000 cores total
- 192GB Ram per node
- EDR Infiniband (Dragonfly+)
- 5PB Scratch, 5 PB project (GPFS)
- 256 TB Burst Buffer (GPFS)
- Designed for large parallel workloads (ie big MPI jobs)
- February/March 2018
- a peak Rmax of 4.61PF (GPC: 312TF, BGQ: 839TF, Graham: 2.6PF, Cedar: 3.7PF).

Lenovo SD530 Node

- Intel Skylake 6148 Gold (2.4 GHz, AVX512)
- 192GB Ram (150 GB/s Memory Bandwidth)
- 3TFlops/node
- 100 Gb/s EDR IB
- stateless (diskless)

Dragonfly+

- Cheaper and much more scalable than Fat-tree
- Requires only small edge switches
- Adaptive Routing
- Congestion Control
- Cray "Aries" network
- new for Infiniband



Storage

- 12PB traditional disk backed GPFS storage
- 70-90 GB/s R/W

Burst Buffer

- A 256TB Excelero burst buffer (NVMe fabric) (GPFS)
- 160 GB/s R/W
- very high IOPs performance

User Configuration

- CC sign-on (similar to Graham/Cedar)
- Software Modules
- Slurm Scheduler
- by-node scheduling
- Designed for large jobs (512+ core)


GPC vs Niagara

	GPC	Niagara
Deployed	2009	2018
Cabinets	45	21
CPU	Intel "Nehalem"	Intel "Skylake"
	SSE4.2	AVX512
Cores/node	8	40
RAM/node	16GB	192GB
Cores	30,240	60,000
Network	5:1 QDR	Dragonfly+ EDR
Storage	2 PB	10-12PB
OS	Centos 6	Centos 7
Rmax	312 TF	4.61 PF
Power	1000 kW	685 kW









Deployment Timeline

- Decommission TCS - September 2017
- Reduce GPC (1/2 Size) - November 2017
- Niagara online - February 2018
- Full GPC Decomission - April 2018

Migration Details

- HPSS Archive (Tape) stays as is
- Existing \$HOME, \$SCRATCH, \$PROJECT will migrate
- 2017 TCS and GPC RAC allocations migrate to Niagara
- 2018 RAC allocations are for Niagara
- Codes will need to be recompiled

<https://wiki.scinet.utoronto.ca/wiki/>
<https://docs.computecanada.ca>