



BCNET2019

Filling the Void: The Case for Local Digital Research Infrastructure (DRI)

Roman Baranowski, Ryan Thomson, Venkat Mahadevan, Wade Klaver
UBC Advanced Research Computing



Overview

- The Canadian Digital Research Infrastructure (DRI) landscape has evolved rapidly over the past 10 years.
- There is increasing recognition that DRI is vital to keeping Canadian research institutions competitive on the global stage.
- Major growth in DRI usage across all disciplines (including humanities & social sciences, natural sciences, and engineering).





Overview

- Data sizes and computational requirements continue to grow.
- Increase in demand for High Performance Computing (HPC), large scale data storage, and on-demand computing, including Cloud computing.






Overview

- Investments in DRI have traditionally come from a variety of funding sources including national agencies, provincial, and institutional funds.
- The Canadian Foundation for Innovation (CFI) provides significant funding for large scale DRI.
- Push to centralize CFI funded DRI at national sites under the Compute Canada umbrella.



Centralization

compute | **calcul**
canada | canada



Compute Canada, in partnership with regional organizations [ACENET](#), [Calcul Québec](#), [Compute Ontario](#) and [WestGrid](#), leads the acceleration of research and innovation by deploying state-of-the-art advanced research computing (ARC) systems, storage and software solutions.

<https://www.computecanada.ca/about/>

BCNET 2019



National Sites

- University of Victoria (BC): Arbutus Cloud
- Simon Fraser University (BC): Cedar HPC
- University of Waterloo (Ontario): Graham HPC
- University of Toronto (Ontario): Niagara HPC
- École de Technologie Supérieure (Quebec): Béluga HPC
- Sick Kids Hospital (Ontario): HPC4Health





Benefits of Centralization

- Aggregate computing power at national sites with trained support staff.
- Reduce costs (e.g. data centre space, volume purchasing, staffing).
- Improve utilization and platform consistency.





Benefits of Centralization

“In theory there is no difference between theory and practice; in practice there is.”

– *Quote overheard at a Computer Science conference*



Things Fall Apart



BCNET 2019

There Simply Isn't Enough





Problems Meeting Demand

- <https://www.computecanada.ca/research-portal/accessing-resources/resource-allocation-competitions/rac-2019-results/>
- For 2019, Compute Canada was able to meet 40.3% of the demand for CPU core years.
- Only 20.3% of GPU requests were able to be fulfilled for 2019.





Problems Meeting Demand

- Some researchers who received reasonable allocations in previous years, received very low allocations for 2019.
- Young researchers often find it difficult to get their research off the ground due to lack of resources.
- Research is in jeopardy!





General Purpose HPC

- 3 national sites run General Purpose (GP) HPC systems.
- There are challenges running different types of workloads on GP systems.
- Traditional HPC workloads e.g. Physics generally run without issue.





General Purpose HPC

- Other kinds of workloads, notably Bioinformatics, which can read/write to many small files, can bring parallel filesystems to their knees.
- One size fits all approach is hard to manage and unpredictable.
- Highly complex scheduling policies (by core vs by node): very hard to debug.





CFI Funding Conditions

- CFI is imposing conditions on grants that mandate DRI infrastructure usually be contributed and integrated with national systems.
- In practice, this has proven to be an extremely time consuming and costly process for both national host sites and contributing institutions.
- Integrated systems represent a one-size-fits-all model that may not be suitable for all researchers.





CFI Funding Conditions

- Case A: PI with a CFI grant with conditions at University A wanted to buy server hardware; told to integrate with national system at University B.
- From date of funding availability to computing on contributed hardware was in excess of 1 year.





CFI Funding Conditions

- Case B: PI wants to apply for computing resources. Should they request everything in hardware form or use Compute Canada resources (which have no guarantee) or a combination of both?
- The CFI provides funding for Compute Canada.
- Let not thy left hand know what thy right hand doeth.





Data Privacy and Security

- Many kinds of research data, i.e. health or indigenous data, have governance requirements enforcing strict privacy, security and curation controls.
- Such data typically cannot be stored on national systems and are often mandated to be held at the institution performing the research.
- Various other jurisdictional laws and data sharing agreements can apply.





Institutional Impacts

- Virtually all research fields have become digital in some respect.
- Top researchers will go where they can access the resources they need. Worse, they will go elsewhere if they cannot.
- Globally leading institutions are making significant investments in local DRI





Present Day, Present Time

- Institutions and funding agencies must be under no illusion that centralized resources can fulfill the needs of the community.
- Institutional investments in local DRI must be increased to maintain relevance on the global stage.
- Investments are needed for the training of support staff and the formation of Advanced Research Computing teams.





Present Day, Present Time

- UBC is in the process of deploying HPC and large-scale storage systems on-premises.
- Priority focus on new Faculty and researchers for whom the national platform is not suitable.
- Serve as a gateway to national systems.





Turning the Tide

- UBC Advanced Research Computing (ARC) team is expanding to meet the needs of researchers.
- UBC continues to make investments to retain its status as a leading research institution.



BCNET[→]2019

