compute canada

Cyberinfrastructure Consultations

January 20-22, 2015

Outline of Today's Presentation

- Overview of Compute Canada
- The Challenge 2 Context
- Assessment of Future Needs
- Stage-1 System Types
- Example Configuration for Stage-1
- Questions & Answers

Introduction to Compute Canada

- Compute Canada's mission is to enable excellent research through delivery of Advanced Research Computing (ARC) services to Canadians.
- A national not-for-profit corporation - members are Canadian research universities and research hospitals.
- Capital and operations funding from CFI, matched by provinces, institutions and other partners.
- Currently operate 50 systems at 27 data centres across the country (about 200k cores, more than 20PB disk).
- Serve more than 2500 faculty-led research groups containing more than 8000 users in all disciplines.

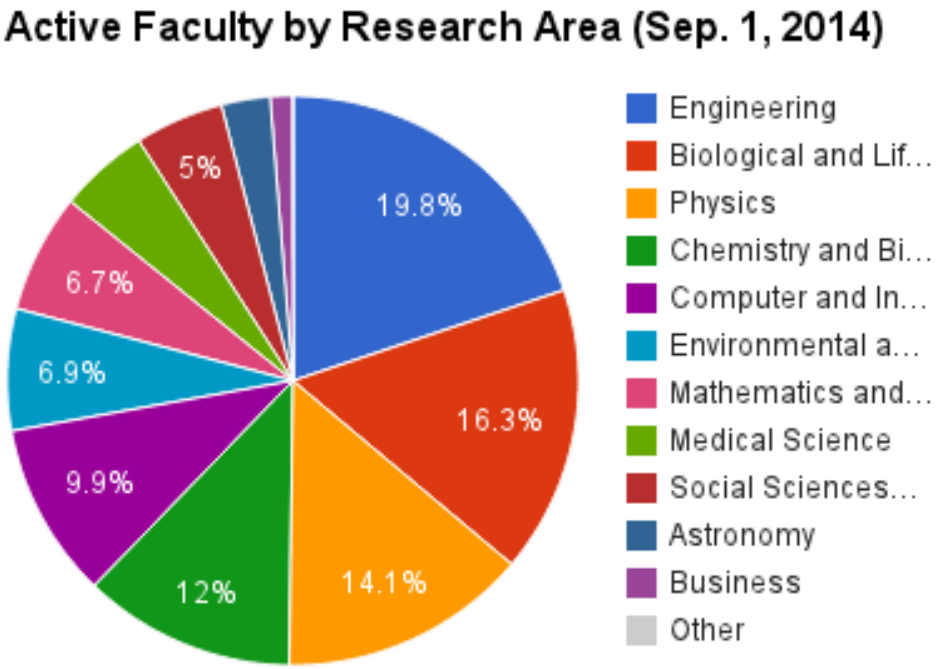

compute canada | calcul canada

Did you know that we...

- **Have over 200 staff, more than 60 PhDs, at 35 institutions?**
- Employ Database experts and run **databases** on behalf of certain research groups?
- Offer connection to 24 of our sites through the **globus.computecanada.ca** portal? You can transfer files in-and-out and share datasets with other researchers around the world with the click of your mouse.
- Employ **visualization experts** who run training workshops and work closely with researchers on their visualization needs.
- Offer a **collaboration platform** based on Vidyo - HD video conferencing, screen sharing, recording, phone bridge (vidyo@computecanada.ca).
- Launched our first **Research Platforms and Portals** competition in fall 2014? Science gateways, web portals, large collaborative science - multi-year allocations.
- Can offer **Hadoop** and other services on-demand?
- Are launching a general purpose research (OpenStack) **cloud** in the next 2 months?

Today's Consultation

- As you have just heard, Compute Canada has a responsibility in Challenge 2 to write a proposal to upgrade and modernize Canada's ARC infrastructure.
- This has been split into two stages. Today's session is focused on the stage-1, proposal due in April, 2015.
- A high-level summary document was circulated to all Compute Canada users and is available on the Compute Canada website.
- Written feedback any time: **sparc@computecanada.ca**

Quotes from CFI Draft Call For Proposals - Challenge 2, Stage-1
Stage 1: Up to $15 million will be provided for the upgrading and modernization of the computational and data storage capacity of the pan-Canadian advanced research computing platform. **As the managers of this platform, Compute Canada will be invited to submit a proposal on behalf of the advanced research computing community**; Stage 1: The renewal of the pan-Canadian advanced research computing platform will be conducted in two stages. For Stage 1, the CFI invites Compute Canada, on behalf of its member institutions, to propose three distinct options **for the capabilities and services that will enable leading-edge research and address the most pressing immediate needs.** This proposal will focus on the upgrading and modernizing of the computational and data storage capabilities managed by Compute Canada.

compute canada | calcul canada

Context - Pressing Needs

- This is a **tremendous opportunity** to improve the ARC resources available to Canadian researchers!
- In stage-1 alone we cannot greatly expand the total computational power available to researchers.
- Aging infrastructure means that the most pressing and immediate needs are to de-fund some very inefficient older systems and replace them with more efficient (and more stable) systems.
- **By the end of stage-2, nearly all current Compute Canada systems will have been defunded and replaced with modern infrastructure.**
- Of course, the world has changed since the current national platform was built. **The new systems will be designed for modern workloads.**

Context - Consolidation

- While we operate systems in 27 data centres today, we intend to greatly reduce this number in the future.
- **By 2018, Compute Canada expects to have large shared systems in 5-10 data centres.**
- We have separated the discussion of what we will buy from where it will be located. **Site selection is proceeding through an independent process involving our members.**
- Wherever the final systems are located, they will serve all Canadian researchers.
- **Compute Canada is committed to maintaining research support personnel on campuses, close to researchers, not only where the hardware is located.**

The Changing Needs of Researchers

- This consultation follows several other steps to assess the needs of the research community going forward:
  - User needs survey (2013)
  - Strategic plan consultation sessions (2013/2014)
  - SPARC white paper call (2014)
- We have also analyzed our internal usage data and responses to our annual resource allocation competition.
- Our Advisory Council on Research (ACOR) has provided feedback.

Compute: Increasing Demand

White Paper	Predicted Increase from Current to 2020
Numerical Relativity	3x
Subatomic Physics	3x
Materials Research	5x
Canadian Genome Centres	8x
Canadian Astronomical Society	10x
Theoretical Chemistry	12x

- Also:
  - Clear need for accelerators.
  - Clear need for mix of memory sizes.

Storage: Increasing Demand

Subatomic Physics Storage Requirements (IPP+CINP White Paper)								
	2014	**2015**	**2016**	**2017**	**2018**	**2019**	**2020**	**2021**
Disk (PB)	12.9	14.9	19.4	22.6	26.5	30.4	37.0	43.9
Tape (PB)	5.5	7.2	10.4	13.7	16.0	23.4	30.9	40.7
Total (PB)	**18.4**	**22.1**	**29.8**	**36.3**	**42.5**	**53.8**	**67.9**	**84.6**

- **Far more dramatic growth in genomics - 450PB of disk by 2020.**
- Many big data projects on the horizon drive the storage needs far beyond our current infrastructure.

Security

- **We have seen increasing demand (medical, social science, industry) to handle research datasets with special privacy limitations.**
- Two challenges for Compute Canada:
  - Security policy framework
  - Infrastructure to allow efficient security protection, auditing, etc.
- We have appointed a Director of Information Security and are currently working on a new security policy framework which will apply across all CC sites.
- **This capital renewal will allow us to re-design our networking and physical security for the new systems.**
- This increased effort will help to ensure the security of existing users, and also is intended to give confidence to future users, including from government and industry.
- Increased security, when needed, will also facilitate highly available (HA) systems.

How the Resources are Served

- Command-line logins and batch job submission: typical for “traditional” high performance computing.
- This refresh gives us a chance to improve ease of access. Expect:
  - Common login
  - Better support for virtualization, cloud
  - Better support for portals and gateways
  - Better ability to utilize a mix of resources, across different sites
    - The mix of resources might involve different systems for pre- and post-processing, computation, analysis and visualization
- Compute Canada is already engaged with numerous new approaches to serve users better:
  - Data analytics (e.g., Hadoop)
  - Databases, Web services
  - Analytics and visualization

Types of Systems

Large Parallel (LP): a system optimized for running large message passing (e.g., MPI) jobs, focused on serving applications using 512 cores and more in any single parallel job. This type of system will have a high-speed interconnect and a relatively homogeneous set of nodes with relatively low requirements on memory/node.

General Purpose (GP): a system optimized for running a wide range of applications including serial computation and parallel applications spanning a relatively small number of nodes. This type of system will be comprised of a heterogenous set of nodes (eg. some with large memory, some with GPUs) and will be well-suited to data-intensive applications.

LP in More Detail

- Some of our largest consumers of CPU cycles require a large number of cores in a single location, with fast interconnect (and message-passing) between nodes. For example:
  - Computational fluid dynamics
  - Astrophysics
  - Materials Science
  - Ocean and Atmospheric Modelling
- Some of the systems slated for de-funding current serve this community and that functionality must be replaced. Need at least one such system in stage-1.
- Relatively homogeneous system, relatively low memory/core.
- Scheduling jobs which use 512-cores or more on a general-purpose system is complex. Policies on LP should favour large parallel jobs.

GP in More Detail

- GP systems will serve a wide range of users who do serial computation, moderate parallel jobs, big data analytics, etc.
- Mix of technologies:
  - both low and high memory nodes
  - some nodes with accelerators (eg. GPUs)
  - support for virtualization and containers
  - at least 2 different security “zones”
- Support full workflow at a single site (eg. one step needs GPUs, another needs high-memory nodes - share the same storage).
- Focus GPUs at GPs since few currently use GPUs from multiple nodes in parallel.
- Several large GP users require at least two sites (redundancy).

Sample System Configuration

System	LP	GP2	GP3	GP1
Cores	30k+	16k+	16k+	10k+
Fast storage	3PB+	4PB+	4PB+	2PB+
Mid tier storage	2PB+	4PB+	4PB+	2PB+
GPUs	4 - for vis.	768	256	4 - for vis 32- alt. arch
Minimum Expected Power Draw (base kW, before overhead of cooling and other auxiliary systems)	540 kW	430 kW	360 kW	200 kW
Estimated Cash Purchase Price (net of vendor in-kind) ($million)	9.27	8.28	7.13	4.97

Timeline

Date	Task
January 2015	Consultations
February 2015	Refine hardware plan, site selection committee reports
March 2015	Produce full draft proposal, present final site selection to members
April 2015	Approval by board, submission to CFI
June 2015	Potential approval by CFI board
July 2015	Begin writing RFP
September 2015	Issue RFP
2016	Install and commission new systems
first half of 2017	migrate users, de-fund replaced systems

Summary

- Challenge 2 provides a tremendous opportunity for the research community to gain access to new shared ARC infrastructure.
- Compute Canada is working on a proposal to meet the most pressing community needs in a stage-1 proposal due in April.
- Your input is welcome today and any time via **sparc@computecanada.ca**

Question: Competitiveness

1. Does this plan provide you with the resources you need to remain competitive in your field over the next few years?

1. Are there near-term science opportunities that this refresh should be tuned to enable? In your opinion, would the existing plan enable researchers to take advantage of those opportunities?

Question: General Hardware

1. Based on this draft plan, what (if any) hardware you require will be completely missing from the Compute Canada platform in 2017?

1. One of the hardest optimization questions for Compute Canada is the balance between CPU, GPU, storage, parallel computing, serial computing, etc. Should the balance of funding in the current draft plan be shifted? (eg. for your research you will require more GPUs).

Question: General Future Needs

1. What is your most pressing hardware need for 2017? Compute, storage, GPUs, high memory nodes, etc.?

1. What type of services are you looking for? VMs, interactive systems, secure data, archiving, etc.?

Question: What is Compute Canada's Role?

Compute Canada is the national organization to provide and support advanced research computing systems and services.

Within the context of the CFI challenges, Compute Canada has been designated a role for partnering with and ultimately supporting needs for advanced research computing and other services, for the Challenges' successful bidders.

Question: What has Compute Canada done to understand users?

In 2014, Compute Canada embarked on the SPARC effort. This resulted in statements from numerous organizations concerning current and future needs for advanced research computing and other requirements for the use of national-level cyberinfrastructure.

This has helped to start a dialog with CFI and other stakeholders concerning the needs for the Challenge and other future activities.

More input is needed, including as part of the Challenge solicitation. Please do share your science stories, infrastructure needs, hopes, and goals via the SPARC email address.

How can you help?

Let decision makers in your institution know that advanced research computing is essential to your research efforts

Ensure the President and VPR are aware of how digitally intensive research is transformative and that requires sustainable predictable funding models.

Mention the resources you use in your publications and presentations

Write your MP and share news of your research and how you use advanced research computing

Work with your communications department on success stories highlighting the advances in digitally intensive research