Federating innovation in Digital Research Infrastructure accessibility

> *As Canada's innovation agenda continues to evolve, it is essential that our advanced computing sector move in lock-step. In that regard, together with Compute Canada, we are committed to ensuring that researchers have the access and expertise needed to utilize the increasingly essential infrastructure that underpins their work.*
>
> - Nizar Ladak, Compute Ontario CEO
> Compute Canada Annual Report 2015-2016

The purpose of Canada's Digital Research Infrastructure (DRI) is to be at the forefront of scientific advances by providing Canadian researchers with the digital tools they need to perform and innovate at an international level. When it comes to Advanced Research Computing (ARC), most of the technical innovation currently comes from individual initiatives from host sites or from the implementation of already existing popular research software developed internationally. We propose a nationally concerted effort for innovation, which would propel DRI to new heights of research software development and availability.

Current issues

Canadian DRI accessibility is lacking

Significant advances have been made by the Compute Canada Federation in teaching the basics of research computing[1] and in adopting an agnostic approach to making popular research software available to users[2]. However, the burden of grasping advanced computing concepts and adapting research software to work with DRI falls primarily on the researchers, their staff, and their students. This effectively curtails what researchers can accomplish with DRI and creates a potential for wasted resources in a computing ecosystem with limited resources. The steep learning curve can be in part attributed to software tools designed for users with a technical background that is unrepresentative of Canadian DRI reality and ambitions. Compute Canada currently leverages software that were designed by and for laboratories in the United States like LLNL (Slurm), Argonne (Globus), OpenSFS (Lustre), LANL (ParaView) and NASA (OpenStack), but contributes little to those software and has yet to offer homegrown efforts centered on Canadian researchers' needs.

In addition, researchers interested in experimenting directly with the national DRI are invited to use default allocations on national clusters as a sandbox[3]. However, they are subject to being blamed for mistakes that would cause resource wastage and to being penalized by the national support staff for infringing the terms of use[4]. Researchers who do not have the time or expertise to develop their own software solutions are mainly constrained by the Federation's ability to make existing software easily accessible to all users regardless of their technical skills in a timely manner. As an example, a request was made via Compute Canada's support system in 2017 for the availability of RStudio, a graphical data science tool [OTRS #08318], and it took three years to reach general availability [OTRS #085712], despite a solution being already available and promoted internally throughout those years[5–7] .

Researchers with access to qualified personnel and who are interested in developing their own accessibility software can do so outside of the Compute Canada environment via Compute Canada's Research Platform and Portal (RPP) competition. While the competition has success stories like syzygy[8] and has helped maintain the success of platforms like CBRAIN[9] and GenAP[10], the program has several shortcomings. Once RPPs are awarded, researchers are left to fend for themselves, as mentioned in the RPP competition: "[researchers are expected to] be able to develop, operate and manage the proposed portal or platform with minimal support from the Compute Canada Federation."[11] In addition to the lack of support, there is also a lack of sustainability in maintaining successful RPPs outside of repeated applications to RPP competitions. Leaving the burden of sustainability of scientific gateways, platforms, and portals to the researchers is a missed opportunity to grow the research enterprise and the expertise within the federation, as well as to keep software alive that could benefit other communities.

DRI innovation relies on individual initiatives and lacks a sustainability plan

DRI host sites or institutions are currently expected to be the primary source of innovation that would promote access to their systems but local initiatives are not consistently encouraged and there are inconsistencies in their promotion, sharing, and acceptance across the federation. When initiatives are not rapidly fruitful, the scarcity of human resources and the lack of a federated support model hinder early developments and doom them to failure.

The absence of a federated structure for local innovation becomes a risk when the diversity of ARC services launched by Compute Canada is not in line with the diversity of needs presented by the DRI user community. In the last 6 years, Compute Canada publicly announced the availability of three national file storage and transfer services: a Globus Portal for file transfer[12], an object storage system[1], and a NextCloud File hosting service[13]. The object storage system has yet to become available to users. In addition, two similar user support systems were launched: a Research support ticketing system[1] and a Bioinformatics helpdesk[14]. While these solutions are required by some users, they cover only a small portion of the spectrum of user needs.

Future DRI State

Improved accessibility to resources for researchers of all backgrounds and computational skills

We see future clusters equipped with true sandboxed spaces that could be used without impacting the performance of production clusters. These sandboxes would work as incubators for both local and global development and testing, in addition to a dedicated system that would accept targeted requests when research groups lack the proper resources to develop customized solutions.

A structured approach to innovation, starting from local initiatives to long-term support by the federation

We propose a federated team with a global perspective on all fields of research that currently need access to our resources and that target needs specific to Canadian researchers for which solutions might not yet have been implemented internationally. Such a federated effort will require a team of full-time staff that can take ownership of promising initiatives before they get dropped by the local teams. The team’s role would be to create and maintain a clear structure that would manage all levels of software accessibility, from supporting local initiatives to federal adoption and long-term support after the initial adoption phase.

How to bridge the gap

An incubation process with a dedicated team and incentives

We propose the creation of a federated team that reviews and maintains open-source code developed internally, by local initiatives, or by research groups. The service is supported by development clusters, either virtual or physical. The group’s responsibilities encompass DRI on clusters and in the cloud, either within the federation or commercial cloud. The innovation structure starts with a local sandboxing phase, then an incubating phase, which also accepts “injections” or projects coming from local initiatives but still in need of some incubation (Fig. 1). Then, some early adopter sites validate the project. Once the project has “graduated”, the federation expects the bulk of the sites to adopt the solution, followed by the late adopters and those who have voiced concerns about the solution and have made change requests. Financial incentives are provided by the stakeholders during every phase leading to the graduation to stimulate local initiatives and encourage natural adoption. Those financial incentives are a function of the willingness of each team or host site to adopt national solutions.

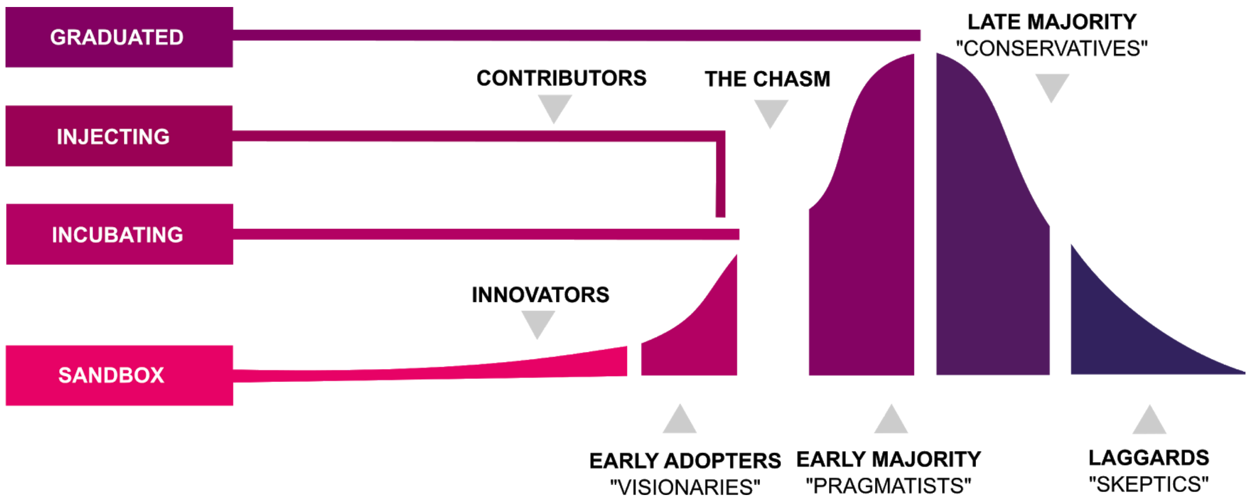

Fig 1. Detailed innovation structure and process
Sandbox: initiatives for new services developed by any DRI actor or **innovators**. **Incubating**: actors and the DRI team work on a plan to bring the project to an early majority and offer the infrastructure-as-code required to deploy the initiative nationally and in collaboration with the instigators. Promotion at the national level to gain traction and seek feedback. **Injecting**: any initiative from **contributors** outside the Canadian DRI ecosystem and brought to the team for incubation. The **chasm** preventing software from going from the early-developmental stage to the graduation process should be overcome by the efforts of a federated team. **Graduation**: projects selected by the DRI group receive approval from the **early majority**, become national solutions, and are expected to be adopted by the **late majority**. All sites share their progress during reports and it is expected that some solutions will raise concern in some **skeptic** sites. Those sites should provide a justification in quarterly progress reports and when the team makes a valid argument, changes would be made to the solution by the DRI team. Figure adapted from CNCF[15].

We propose that the group follows and enforces some major guidelines:

1. When code is approved for national distribution, it must be adopted by members of the federation
2. If a group or host site refuses to adopt the national solution, a justification must be provided in a progress report.
3. If the report makes a valid argument for the refusal of the solution, changes are made to the national solution to comply with the requests of the group or host site
4. Financial incentives are provided by the stakeholders to stimulate local initiatives and promote the general adoption of national solutions
5. In return, the federated team agrees to be responsible for the long-term maintenance of the national solutions, to fix bugs, and to keep high standards in the quality of the code.

Authors

Julie Faure-Lacroix is a Science Liaison Agent at Calcul Québec. She is developing data visualization and cluster accessibility tools aiming at non-traditional and new users. She has won several awards, including Compute Canada's Outstanding Achievement award in 2017.

Félix-Antoine Fortin is the lead of Université Laval's new research software development team. He has been developing and proposing open-source solutions for Canadian DRI accessibility since 2015 and was awarded a Compute Canada award of excellence for his work in 2018.

Florent Parent is the lead of Calcul Québec at Université Laval team. Since its inception in 2009, he has been enabling his team to innovate by providing local support, resources, and mentorship that materialized in a large amount of national and international talks, awards, and contributions.

Darren Boss is team lead for Middleware Infrastructure at Compute Canada. He has demonstrated strong advocacy of federated innovations over the years, leading and maintaining the Compute Canada identity provider, deploying and improving continuous integration infrastructure, and proposing a novel Kubernetes-as-a-service solution to the federation.

Drew Leske is a senior programmer analyst with Compute Canada and a member of the Middleware Infrastructure Team, the Automation and Configuration Management (ACME) team, and the Infrastructure Operations National Team (IONT). Drew has worked closely with researchers and staff to work with existing services and advocate for additional services for these communities. Drew leads the FRAK Burst Enablement project.

References

1. Compute Canada. Compute Canada Annual Report 2015-2016. https://www.computecanada.ca/ar/en/1.html.
2. Boissonneault, M., Oldeman, B. E. & Taylor, R. P. Providing a Unified Software Environment for Canada's National Advanced Computing Centers. in *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (learning)* 1–6 (ACM, 2019). doi:10.1145/3332186.3332210.
3. Compute Canada. Access Policy | Compute Canada. https://www.computecanada.ca/research-portal/accessing-resources/access-policy/.
4. Compute Canada. Terms of Use | Compute Canada. https://www.computecanada.ca/research-portal/information-security/terms-of-use/.
5. Fortin, F.-A. Jupyterhub: spearheading interactive computing update. (2018).
6. Fortin, F.-A. Jupyterhub: spearheading interactive computing. (2017).
7. Fortin, F.-A. Jupyter and Magic Castle: spearheading interactive computing. (2019).

8. Compute Canada. Compute Canada and PIMS Launch Jupyter Service for Researchers. https://www.computecanada.ca/featured/compute-canada-and-pims-launch-jupyter-service-for-researchers/.
9. Sherif, T. *et al.* CBRAIN: a web-based, distributed computing platform for collaborative neuroimaging research. *Front. Neuroinformatics* **8**, (2014).
10. GenAP. https://galaxyproject.org/use/genap/.
11. Compute Canada. Research Platforms and Portals (RPP) Competition | Compute Canada. https://www.computecanada.ca/research-portal/accessing-resources/resource-allocation-competitions/rpp/.
12. Compute Canada. Compute Canada Annual Report 2014-2015. https://www.computecanada.ca/ar2015/en/1.html.
13. Compute Canada. Nextcloud - CC Doc. https://docs.computecanada.ca/wiki/Nextcloud.
14. Compute Canada. Compute Canada Launches Bioinformatics Support Resource. https://www.computecanada.ca/featured/bioinformatics-2/.
15. Cloud Native Computing Foundation. Graduated and incubating projects. *Cloud Native Computing Foundation* https://www.cncf.io/projects/.