The Alliance Cloud Connect Pilot

Call for Participation for Researchers

Digital Research Alliance of Canada | Alliance de recherche numérique du Canada

The Alliance Cloud Connect Pilot (ACCP)

Call for Participation for Researchers

Version 0.9 - June 2024

---

1. Introduction

The Digital Research Alliance of Canada (the “Alliance”) is excited to announce *The Alliance Cloud Connect Pilot (*“ACCP”*)*, a unique opportunity for vendors, institutions, individuals (associated with institutions) and researchers to participate in the development and testing of a cutting-edge initiative that will revolutionize how Canadian researchers access Community Cloud and key Commercial Cloud Services.

As part of our commitment to fostering innovation and facilitating seamless cloud services to Canadian researchers, we invite Canadian Researchers to submit a response for consideration to this Call for Participation (“CfP”).

2. Objectives

The vision of the ACCP is to: leverage an enhanced and rich suite of cloud services, provide an on-demand approach to accessing cloud resources and develop a long-term strategy for using cloud computing to benefit all Canadian researchers. The ACCP aims to build national cloud research capacity and deliver on the Alliance’s mandate to meet the increasing DRI resource demand from the Canadian research community.

The ACCP will consist of three (3) separate calls for participation (“CfP”) that will target:

1. Commercial Cloud Providers
2. Platform Developers
3. Canadian Researchers (this Call)

The “Commercial Cloud Provider” and “Platform Developer” CfPs will develop a unified pilot platform (“Platform”) for accessing Commercial Cloud Service Providers (“CSP”). This CfP aims to engage Canadian Researchers in the testing and refinement of the services and capabilities developed as part of the ACCP.

In this CfP, the Alliance is inviting researchers from diverse sectors to work together in an open and collaborative context to identify ways to facilitate researchers' use of cloud frameworks.

Note that this CfP is not a funding call or a procurement for services; it is an invitation to researchers to gain early access to and influence the development of the ACCP platform and the services being developed.

The overall timeline for the ACCP is provided below, with Researcher involvement (this CfP) beginning active participation in September of 2024 and ending in March 2025.

Stage	Description	Time Frame
1 - CfP: Identify possible participants	Call for participation that describes the ACCP and its components. A letter of intent from respondents indicating their relevance and capacity to contribute to one or more components. Once submitted, a committee will filter out inappropriate or poorly fitting submissions.	April 2024
2 - Kick-off meeting & workshops (1/2 to 1 day)	A workshop with accepted respondents (participants) to explain the project, its goals, components, specific details and to answer any questions. Participants will interact with the Alliance and other participants with the goal of one or more participants forming teams that will design, build and operate (for the period of the pilot) each component. Each team will be co-led by an Alliance representative and a team member. Workshop meetings follow to refine each component proposal.	June 2024
3 - Component Proposals	Each team (per component) will prepare a component proposal including architecture, technology stack, time frame and budget, including operations and support of the component during the pilot.	June 2024
4 - Component Sub-project Implementation (Agile/iterative)	Teams build components, with a focus on integration with ACCP front end component #1 ("Pane of glass"), IAM (component #6) and reporting (component #7).	July - Aug 2024
5 - Deployment & Testing (*pre-Researcher*) (Agile/iterative)	In-cloud testing and deployment of components, individually and integrated. Infrastructure-as-code (IaC) deployment only. Alliance and select end users (researchers) will be involved in acceptance testing.	July - Aug 2024
6 - Pilot Operations and Researcher Evaluation	Deployment into production of all components. Teams responsible for operations and support during this time.	Sept 2024 - Feb 2025

3. Background

3.1. The Alliance

The Alliance is funded by the Government of Canada through Innovation, Science and Economic Development Canada (ISED) via the Digital Research Infrastructure Strategy. The Government of Canada is promoting the broader adoption of cloud computing to leverage the benefits in delivering faster, more adaptable and scalable solutions to meet the growing demand for digital resources in the Canadian research community.[1]

In 2022, ISED gave the green light to the Alliance's 2023-25 Multi-Year Funding Proposal and National Cloud Strategy, endorsing a unified vision, strategy and action plan for Canada's Digital Research Infrastructure (DRI) Ecosystem.

The Alliance Researcher Council's updated cloud priorities for 2023:[2]

1. Provide a flexible model for commercial cloud computing that facilitates transparent access or by providing in-kind credits to researcher groups to directly purchase cloud compute, storage or services from a vendor of their choice.

2. Ensure that Alliance-managed access to commercial clouds is done so that compute, storage and services are provided in a manner that is vendor-agnostic.

3. Consider commercial clouds for opportunistic computing when Alliance resources are oversubscribed or offline.

ISED has made $5 million available to the Alliance for the ACCP project, which, if leveraged fully, would require a match of approximately $2 million. This amount is for the entire project, including the CSP CfP.

3.2. Current State

The "Cloud," defined as those services provided via the major CSPs and second-tier providers, has created a mature Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS) ecosystem over the last two decades. This ecosystem has been growing at approximately 15% year over year

[1] Government of Canada's Cloud Adoption Strategy: 2023 Update, https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/cloud-services/cloud-adoption-strategy-2023-update.html.
[2] Alliance Researcher Council: Meeting the Digital Research Infrastructure Needs of the Canadian Research Community - Update on Cloud Computing, https://zenodo.org/doi/10.5281/zenodo.7974805

since 2006, and now serves all levels of business, enterprise and government. Higher education in Canada and the United States has lagged in this area, although interest is accelerating as pricing becomes more competitive, along with specialized programs and economies of scale.

The Alliance is one of the key providers of Digital Research Infrastructure (DRI) services in Canada. It is a member-based organization of post-secondary institutions, research hospitals, colleges and related organizations, with 45 primary members and 88 associated members spanning the country. Computational services include both traditional High Performance Computing (HPC) as well as Community Cloud on-premise infrastructure, with over 21,000 active research users and 5,600 active principal investigators. The demand for Alliance DRI computational services is continually increasing, with the Alliance only able to meet a fraction of the demand from current researchers due to limits in capacity. For example, in 2023 the Alliance resource allocation process[3] was only able to meet 46% of the requested computation for HPC (216K core years allocated out of a requested 468K) and 76% of requested Cloud resources (27K virtual CPU years allocated out of a requested 35K). The ability to meet the needs of Canadian researchers in terms of specialized hardware is even more challenging, with only 20% of requested GPU resources allocated (2K GPU years allocated out of 10K requested). In some cases, Canadian researchers are directly utilizing Commercial Cloud to meet their research needs. The Alliance envisions hybrid Commercial and Community Cloud as a mechanism to help meet the needs of Canadian researchers, including but not limited to (from Alliance National Cloud Strategy):[4]

- enabling access to cloud-based, container-orchestration workflow systems;
- supporting shared scheduling of HPC-like workloads;
- enabling interactive and collaborative workloads;
- providing alternative open source tools;
- supporting rapid access to virtual HPC;
- offering collaborative portals and publishing; and
- accessing innovative AI tools and algorithms.

4. Project Overview

The goal of this project is to create and operate a pilot platform that streamlines and enhances the process of Canadian researchers accessing and administering community and commercial cloud providers operating in Canada. This platform aims

[3] Alliance 2022-2023 Annual Report, https://www.alliancecan.ca/sites/default/files/2023-08/annual_report_2022-2023.pdf, Page 13
[4] Alliance National Cloud Strategy 2023, https://zenodo.org/doi/10.5281/zenodo.10214474

to provide better access to compute, software and data through a centralized and user-friendly interface, ensuring efficiency, security and compliance for the Canadian research community seeking cloud services.

The objective of *this* CfP is to identify potential Canadian researchers who will work with the CSP and Platform teams to test and assess the ACCP design and implementation.

The intent is to have initial production versions of the CSP and Platform solutions available for Researcher testing (as per this CfP) from September 2024 to March 2025.

Pursuant to the 2021 Researcher Needs Assessment and the Researcher Council Priorities Report, five (5) Use Cases were developed. The Use Cases provide examples that detail the specific functional and technical requirements of the desired platform and can be found here.

4.1. Key Features

- **Seamless Integration:** The Platform should support easy integration with leading commercial and community cloud providers, allowing users to manage and deploy resources effortlessly.
- **Security and Compliance**: Robust security measures and compliance features must be implemented to meet industry standards and regulatory requirements in Canada.
- **User-Friendly Interface**: The Platform should have an intuitive and user-friendly interface, enabling both technical and non-technical users to navigate and utilize its features effectively.
- **Scalability**: Allow researchers to combine the power of an agile, elastic and pay-per-use cloud experience that can accommodate the evolving needs of the researchers and adapt to changing requirements.

4.2. General Principles

Figure 1 captures the basic architecture for the Minimum Viable Product (MVP) for the Platform. The Platform is divided into three categories, each containing one or more components, indicated in red borders in the diagram. As a platform that targets supporting researchers, researchers are expected to interact with many of the components of the Platform.

4.3. Minimum Viable Product (MVP)

The MVP document provides a description of the features required for the development of the Cloud Connect Platform MVP, as well as a list of desirable

features. The MVP document was used to inform the development of the ACCP through the Provider and Platform CfPs.

The Platform components are expected to run on one or more of the cloud service providers. Deployment of the Platform components onto the CSPs is required to use Infrastructure-as-Code (IaC). The MVP document contains details regarding the IaC requirements.

The Platform is made up of six components, in four categories.

1. **'Single Pane of Glass' Portal (*Component #1*):** A web interface (portal) providing researchers with a single integrated view across all of the Platform services (2, 3, 4 below) and pass-through links to the CSPs' web consoles. The development of this component will require integration with the below components and links to the CSPs' web consoles. The portal will also provide Alliance and community admins with integrated administration views of all the components.

2. **Researcher-Facing Managed Services:** A set of Alliance and/or community-managed, CSP-hosted, researcher-facing specialized compute services, which will allow researchers to deploy their own instances of these services. These services (components), for the purposes of the pilot, are Jupyter **(*Component #2*)**, Galaxy **(*Component #3*)** and Magic Castle **(*Component #4*)**. The user interfaces for the deployment of these services by researchers will be integrated into the portal. Additional services (Others - ***Component #5***) will be considered where relevance, cost and community support justify their inclusion in the pilot.

3. **Identity and access management (IAM) (*Component #6*):**
   a) Delegates identity management to a federated authority; and
   b) Is the authority for access management and propagates access management policy to the cloud providers.

   Access management in this context is primarily the ability to invoke, manage, view, etc. CSPs' services/infrastructure. Researchers will be able to sign on with their existing credentials and access cloud resources as per their role in their projects. An administrative view will allow Alliance and community admins to manage accounts, CSP access, etc.

4. **A multicloud reporting and management dashboard (*Component #7*):** A component allowing researchers and Alliance and community administrators to monitor, control and optimize cloud costs, infrastructure and services. The dashboard will offer a unified view across the cloud providers, supporting roll

ups by researcher, project, cloud provider, etc. The dashboard will support reporting and alerting for cost management, intrusion and anomaly detection, compliance, security policy, network issues, performance issues and access policy issues.

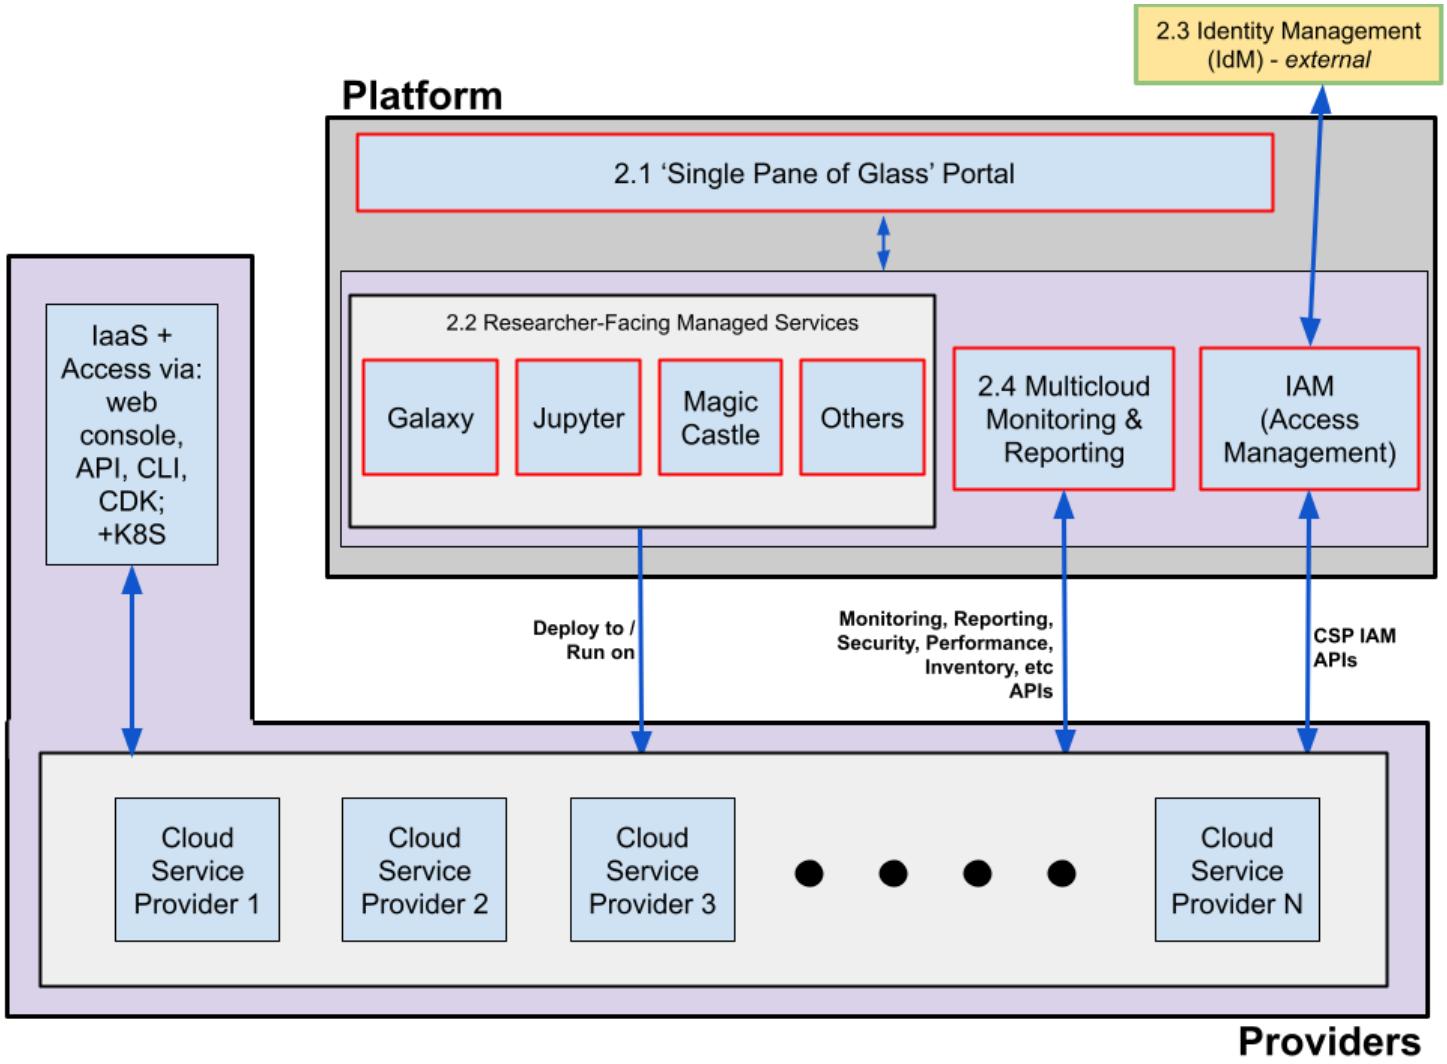

Figure 1: ACCP MVP

5. Scope of Participation

5.1. Participant Eligibility

The ACCP initiative provides a unique opportunity for collaboration and partnership between the Alliance, academia, health care, private industry, researchers and other partners to accelerate cloud computing across the Canadian DRI ecosystem.

This CfP invites Canadian Researchers who currently use or have a need for DRI services (either Community Cloud or ARC services) to participate in the development, testing and assessment of the ACCP Platform. In order to test the full scope of the Platform, a range of researcher needs and experiences are desirable in this testing phase. Researchers are welcome who:

1. Have a range of experience using the cloud:
    a. General users
        i. Working as a consumer of the services provided by either the Alliance (e.g., SaaS such as Galaxy or Jupyter) or other research platforms (e.g., CanFar or iReceptor).
    b. Advanced users
        i. Working at the services and infrastructure layer via web console, API, CLI and Cloud Development Kit (CDK).
        ii. Working at the IaaS (Infrastructure-as-a-Service), PaaS (Platform-as-a-Service), FaaS (Function-as-a-service) and SaaS (Software-as-a-Service) services layer.
2. Require a range of computational resources:
    a. Provisioning and managing VMs (e.g., Repositories and Web Portals).
    b. Running Software-as-a-Service (e.g., Jupyter, Galaxy).
    c. Running compute workloads in the cloud.
    d. Bursting ARC compute workloads into the cloud (e.g., Magic Castle).
3. Want to be involved in the pilot as:
    a. Early adoption users of the services developed as part of the platform.
    b. Test users of the full-featured (but still beta) services provided by the platform.

The ACCP Researcher CfP is open to all Canadian researchers. Proposals will be screened based on their applicability for the testing of the ACCP Platform Service Components and the capacity of the Alliance ACCP team to support the number of proposals received. It is hoped that all proposals will be able to be supported, but this will depend on the number of researcher submissions and the capacity of the ACCP team. If proposals are screened due to capacity limitations, proposals will be

assessed based on their applicability to testing the components of the ACCP Platform Service Components. If the ACCP project has support capacity and/or has not received proposals that test all components, the CfP will be left open for proposal submissions during the duration of the testing period.

Researchers should note that the ACCP project is a pilot project. Researchers should be aware that all services being offered through this CfP should be considered in a "beta" phase of development. Although the Alliance will make best efforts to work with the CSP and Platform teams to maintain high availability and reliability of the Platform and its services, *critical* production research work *should not* be undertaken through this project.

5.2 What does the ACCP provide to research participants?

The goal of this CfP is to engage the Canadian research community in the testing and refinement of the services and capabilities developed through the ACCP. This CfP gives the research community an opportunity to be part of a cutting-edge initiative that will transform the long-term strategy for cloud computing in Canadian research. Through this CfP, researchers will be able to get early access to the ACCP platform services and provide feedback on the services at a critical time in their development. The integration of cloud services across multiple cloud providers has the potential to dramatically increase the productivity of DRI in Canada. At the same time, the ecosystem is complex, and researcher feedback is essential to the pilot project's success.

Specific benefits to researchers in participating in this CfP are:

- early access to ACCP services and capabilities;
- ability to influence the direction of ACCP through interaction with the ACCP development team;
- interaction with experts from industry (CSP and Platform) and the Alliance;
- limited access to compute cycles and storage; and
- potential access to specialized resources (advanced GPUs, large memory VMs) or services (e.g. AI/ML).

All research participants will be invited to a Researcher Kickoff Workshop where the ACCP platform services will be presented. The workshop will be used to determine which participants are interested in working with which platform services and to determine what, if any, gaps there are in service testing. This workshop will also consider researcher use cases as a mechanism to test the integration of the ACCP Service Components.

5.3 What is expected of research participants?

The Alliance and its collaborators on the ACCP value the feedback of the research community. This feedback will be critical in ensuring that the services developed will provide real value to researchers. It is recognized that researchers' time is valuable, and as such the Alliance has attempted to maximize value and minimize researcher effort in participation in the ACCP. To participate in the ACCP as an early access researcher, it is expected that researchers will:

- apply (through response to this CfP) for advanced access to the ACCP Platform;
- participate in the ACCP Researcher Kickoff Workshop;
- actively participate in using and testing the ACCP Platform for their research workflows;
- provide representation from their research group at ACCP working group meetings;
- provide feedback and engage with the development teams as applicable; and
- participate in the ACCP Platform user survey(s).