compute canada | calcul canada

# Cloud Computing For Researchers

December 2016

Compute Canada is often asked about the potential of outsourcing to commercial clouds. This has been investigated as an alternative or supplement to purchasing and operating advanced research computing resources (HPC systems, cloud-focused systems, and storage).

Compute Canada has regular communication with the major providers of commercial cloud services in Canada, and has open dialogues concerning pricing, services, and characteristics of their offerings. These discussions helped lead to many of the elements of our Cloud Strategy, including commercial cloud bursting and Elastic Secure Cloud Services (described below). It has also informed some of the division of labor and user expectations for cloud support.

Currently, it is far more cost effective for the Compute Canada federation to procure and operate in-house cyberinfrastructure than to outsource to commercial cloud providers.

Compute Canada has recently performed a detailed analysis of two recent large compute cluster purchases along with storage purchases, totalling over $15M (Cdn) in capital expenditures. Comparison pricing was used from commercial offerings of a very large commercial cloud provider. Configurations were selected which most closely match CPUs, storage, and other characteristics of Compute Canada's new systems. Applicable volume discounts were applied, and US funds were converted to Canadian Funds.

The key result was that the total costs, over five years, of the commercial provider were significantly higher than the total cost of ownership of newly-purchased systems over the same five years. In addition, the commercial cloud provider could not provide the range of services required to serve a large portion of the Canadian research community.

Cloud-based costs ranged from 4x to 10x more than the cost of owning and operating our own clusters. Some components were dramatically more expensive, notably persistent storage which was 40x the cost of Compute Canada's storage.

Furthermore, the cloud provider did not offer all of the key characteristics of two of Compute Canada's new national systems ARBUTUS at the University of Victoria and CEDAR at Simon Fraser University. Compared to ARBUTUS, the cloud provider did not offer adequate local/dedicated storage. Researchers with needs for very high performance on-node storage would be dramatically underserved by the commercial provider. These include research fields such as physics, environmental science and other fields.

Compared to CEDAR, the cloud provider did not offer tightly coupled sets of nodes, with a low-latency high-performance interconnect. This would create great challenges for much of the CEDAR workload, which is dominated by parallel multi-node computation. This type of configuration is standard in data intensive

research and is the basis for a large majority of advanced research computing in areas of importance such as environmental science, complex system modelling, physics etc.

The overall cost-comparison outcome may be summarized as follows:

- For ARBUTUS, outsourced direct (capital) costs of 5 years of the closest comparable configuration was approximately 7.4x the cost of the GP1 purchase. However, the outsourced solution would not meet the requirements of the ARBUTUS workload, due to the lack of local node storage.
- For ARBUTUS, when 5 years of operational costs are added, the 7.4x decreases to approximately a 4x cost increase, to outsource versus purchase.
- For CEDAR, outsourced direct (capital) costs of 5 years of the closest comparable configuration was approximately 15.4x the cost of the CEDAR purchase, reflecting 6.6x for compute, and 40x for storage. However, the outsourced solution would not meet the requirements of the CEDAR workload, due to the lack of local node storage (for the OpenStack partition) and absence of a high performance low latency interconnect.
- For CEDAR, when 5 years of operational costs are added, the 15.4x decreases to approximately a 10.9x cost increase, to outsource versus purchase.

Commercial cloud providers often provide cost analysis or return on investment calculators, to assist in deciding whether or not to outsource to the commercial cloud. These analyses are not well-suited for Compute Canada, for several reasons:

- The Canadian cyberinfrastructure funding model, including matching, favors capital equipment;
- Member sites already own and operate data centers and other physical infrastructure; these do not need to be built, purchased or leased;
- Research computing tends to fully occupy the systems they use, and keep those systems busy all the time. ROI calculations often assume systems will be unused for much of the time, whereas research computing systems tend to be fully occupied at all times, 24x7x365;
- Compute Canada has robust and capable networks, from local area network to the regional and national network; this carries relatively little direct cost to the national platform, and no incremental costs per unit of usage. Those networks would be needed wherever cyberinfrastructure is located;
- Compute Canada's expertise is in place, to support all aspects of user need; therefore, there is no incremental cost in user support, which is otherwise often included as part of ROI analysis;
- Compute Canada provides an all-Canadian cyberinfrastructure, while most commercial cloud providers do not base their cyberinfrastructure and related support entirely within Canada.

Compute Canada continuously monitors pricing and services, and interacts with other research computing providers, as well as commercial cloud providers, to evaluate potential benefits of shifts in the balance of service. Procurement opportunities for the Challenge 2 systems are, in fact, open to any provider who can meet the needs of the specified systems and services - including cloud providers.

Compute Canada is engaging in collaboration with cloud providers, in order to achieve benefits of cloud bursting for those users who desire it. The goal is decreased time to solution for users: by accessing commercial cloud resources, users may be able to get access to more resources, or faster access to resources, than through Compute Canada systems. Commercial clouds may offer services or features not available from Compute Canada. Costs for such commercial cloud uses would be borne by the users, not by the national platform.

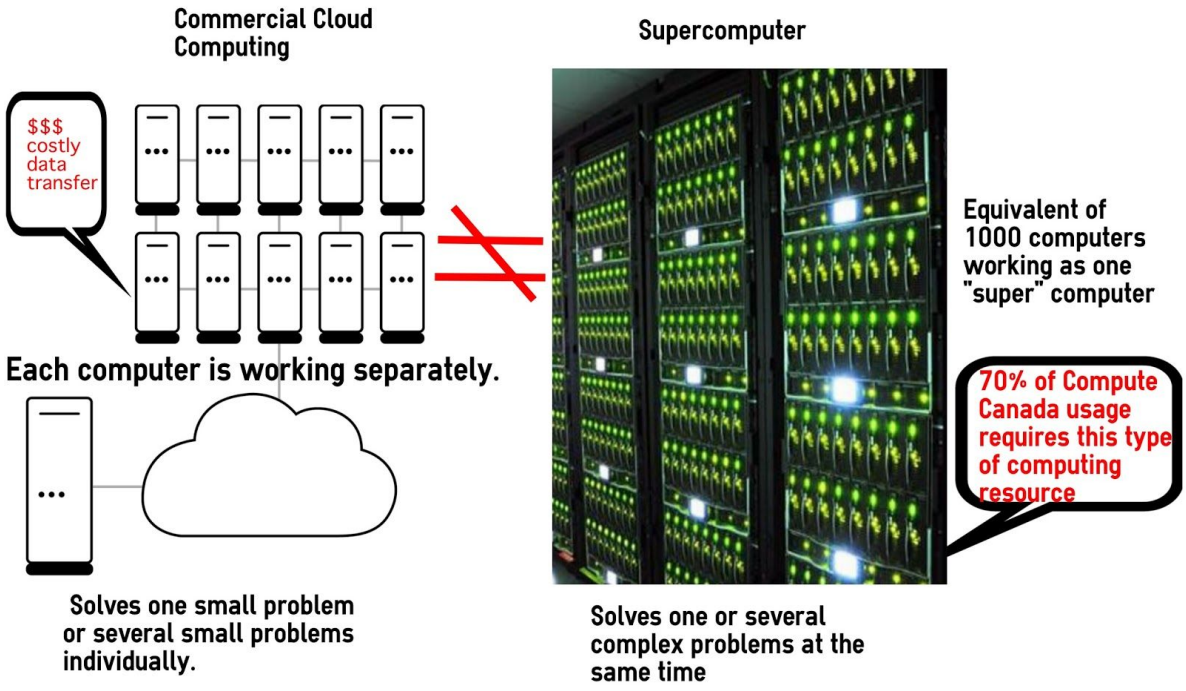


## Overview of the Compute Canada Cloud Services

Since 2015, Compute Canada has provided cloud services to the Canadian academic research community. Services delivered to date have primarily been directed at users experienced with systems administration and support. The Compute Canada cloud is classified as a multi-site, multi-tenant private cloud, built upon the OpenStack platform. This platform is a standard in science cloud computing.

The core concept of the Compute Canada cloud is **federation**. The Compute Canada cloud is a federation of multiple sites. Federation is defined, for these purposes, as follows:

1. Unified login. Users access cloud resources utilizing the same username, password, UID, GID and other identifying characteristics, across the Compute Canada federated cloud.

2. Portability of workload. Through use of the OpenStack platform, the exact same virtual machine will run identically on any cloud system.
3. Similar user experience. In addition to OpenStack, other aspects of the cloud systems will be essentially the same. This includes available filesystems and software.
4. Centralized support. All cloud support is accessed via the national helpdesk (i.e., cloud@computecanada.ca), with local/regional support delivered within the existing Compute Canada service model.
5. Local or regional support, where available. On-site or regional support is an important part of the national platform, and cloud support is part of this offering.

# Compute Canada Cloud Strategy 2016-2018

Over the next year, Compute Canada will complete a multi-region cloud, with improved access among the federated sites. Portability of user workloads will be even greater, and there will be capabilities for automated load balancing among sites. Until then, Compute Canada's cloud federation will be achieved through common practices across multiple cloud sites. This will deliver the federated cloud characteristics listed above, but without operating as a single multi-region cloud.

| |
|---|
| Target date for multi-region cloud: Coincident with Challenge 2 Stage 2 cloud resources, approximately mid-2017. |

Another direction that may lessen the need for dedicated cloud resources is increased capabilities for scheduling new HPC resources. Current time-sensitive workloads for ingestion of experimental or observational data may be tolerant of batch-mode post-processing. In such cases, it may be more efficient to utilize HPC resources as part of the workflow, rather than keeping the entire workflow on dedicated cloud resources.

| |
|---|
| Target date for workflows spanning cloud & HPC resources, including data movement: To be developed as appropriate for Challenge 1 Stage 1 service development, late 2016 & beyond. |

# On-shoring of Data and Computation

It is the intention of the Compute Canada cloud federation that all on-premises data and computation is within the confines of the Compute Canada federation and, therefore, will not leave or transit any national boundaries. This includes authentication data, backups, storage, and all other data traffic or exchange.


## Commercial Cloud Bursting

Compute Canada plans to make it easier for users to seamlessly move their workload between commercial clouds and the CC cloud. This is in recognition that the quantity of available Compute Canada resources, for all types of computation, is not currently adequate to meet all user demand.

Some Compute Canada users may wish to mainly utilize commercial cloud resources, and then periodically expand their computation or storage to Compute Canada systems. Other users may primarily make use of Compute Canada resources, and may occasionally want to expand to commercial cloud resources.

It is not the intention of Compute Canada to pay for commercial cloud resources in such situations. However, it is in the interest of all parties to assure ease of access, bi-directionally, for computation and data access. Therefore, Compute Canada has proposed to CFI that a portion of the Challenge 2 Stage 2 effort be devoted to integrating bi-directional cloud bursting with HPC and cloud resources. Initial investigation of these capabilities is ongoing, but large-scale development and support will depend on CFI's assessment of the Challenge 2 Stage 2 proposal.

| |
|---|
| Target date for bi-directional cloud bursting as a mainstream service offering: Part of Challenge 2 Stage 2 activity, from mid-2017. Initial investigation is ongoing in mid- to late-2016. |

## Elastic Secure Cloud

Elastic Secure Cloud (ESC) concepts were proposed to CFI for Challenge 2 Stage 2. ESC will build on the federated cloud model, by adding better capabilities for isolation of data and computation in cloud partitions. For ESC, a cloud-based project would be allocated one or more cloud nodes or virtual machines. Those nodes would have access to data partitions that are not available to other cloud tenants. If desired, ESC tenants could utilize their own authentication and authorization services, rather than Compute Canada's. Additional assurances and procedures, such as for certifications of systems administrators, encrypted or off-site backups, or isolation of network traffic, will be developed. Elasticity in the secure cloud would be achieved by changing the quantity or configuration of allocated resources to particular projects, over time. Secure cloud concepts are already in place at several Compute Canada member institutions, and will form a basis for ESC in the cloud federation.

| |
|---|
| Target date for elastic secure cloud as a CC-wide service offering: Part of Challenge 2 Stage 2 activity, from mid-2017. |

## Long-Term Visions of Cloud/HPC Integration and HPC as a Service

Compute Canada closely tracks developments in the research computing industry, as well as developments in other jurisdictions. In the long run, industry forecasts anticipate that the current distinctions between high performance computing (HPC) batch-style computations and OpenStack-style cloud platforms will decrease, and that the two will eventually merge.

One step towards that merger will be for allocation of virtual *clusters*, not just virtual machines (VMs), on cloud platforms. A virtual cluster, whether initiated via OpenStack or an HPC batch scheduler, would give a user the ability to develop a customized software environment similar to how VMs are addressed with today's Infrastructure as a Service (IaaS)*. The difference would be for the workflow to encompass parallel computations - including efficient utilization of high performance interconnect networks and parallel filesystems.

Compute Canada anticipates that increasing numbers of users or user projects will desire such customized environments. As with today's VMs, the virtual cluster choices would be persistent for a given user, and could be applied across different physical systems.

Also, as with today's federated cloud offerings, Compute Canada would provide the underlying software environment to support the virtual clusters.

## Infrastructure as a Service (IaaS)

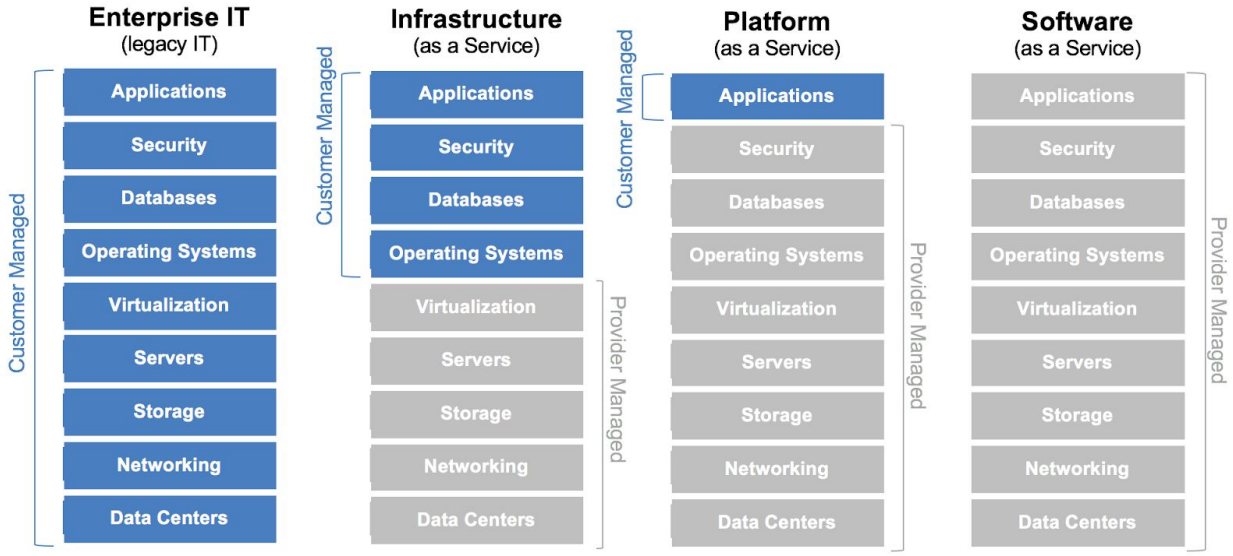


Credit: The enterprise Cloud Blog
https://mycloudblog7.wordpress.com/2013/06/19/who-manages-cloud-iaas-paas-and-saas-services/