

# The Faculty of Science Research Cloud Governance Document

This is a document to outline the policies and procedures to govern the Faculty of Science research computing cloud.

## Purpose of the Research Cloud

Borrowing from the CERN/LHC computing model research computing can be broadly divided into three tiers of resources:

**Tier I:** These are large-scale clusters of thousands of machines which are shared between many researchers. They are suited for very large processing jobs but may require the user to wait a day or more for the output due to the shared nature of the resource. They also typically require jobs to be adapted to a common operating system environment. Examples include Jasper or Checkers from Westgrid and Guillimin from Calcul Quebec.

**Tier II:** These are clusters at a local institute designed to provide immediately available computing power for medium-scale storage and jobs. The operating system environment is controlled by the research group which means that it is hard to share the resource with other researchers when not in use.

**Tier III:** This is the researcher's own desktop or laptop which they use for analysis.

The Faculty of Science Research Cloud is designed as a tier II resource which will use virtualization to enable researchers to have full control over their research computing environment while allowing them to benefit from sharing resources and using centrally managed hardware. The aim is to merge many of the different local research clusters scattered throughout the Faculty into a single resource with reductions in cost and hardware maintenance being passed back to the researchers along with the benefit of being able to share unused resources.

The resources will be administered by a combination of IST and Compute Canada. IST will manage the physical hardware providing power, cooling and network access as well as taking care of warranty replacements, new machine installation etc. Compute Canada will manage the virtualization software which will enable researchers to startup and shutdown machines, clone existing machines, manage virtual networks and storage volumes etc.

Compute Canada also manage the tier I resources which many researchers use and some of these resources are also switching over to using virtualization. The research cloud will be managed using the same software and configured in a compatible way. This will enable future integration between the faculty research cloud and the national Compute Canada cloud so that researchers will not need to reconfigure, recompile etc. jobs to be able to take advantage of the

large scale, tier I resources Compute Canada can offer. It is also hoped that this approach may work with other resource providers such as Cybera, Amazon and IST.

## **Guiding Principles**

Research computing support requires an agile, flexible system which can adapt to the ever changing needs of researchers. Cloud computing is a well established technology but its use to provide interactive, low latency tier II computing, rather than large-scale low latency tier I computing, is new. As a result there is little experience in the best way to run such a cluster, particularly in a research environment, which means that this document is a work in progress and not set in stone.

We anticipate that the governance of the cluster will change over time as we better learn how researchers will use the cluster. However the broad principles which will be used to determine when, and how, to adapt the governance of the cluster are:

1. Resources purchased by researchers will be available for their use at anytime and will remain available for as long as the underlying physical hardware remains functional and practical to operate i.e. they will mimic the ownership characteristics of a physical server machine.
2. Researchers should be afforded the maximum freedom possible to use their resources as they see fit. Restrictions should only be imposed where they are required to:
  - a. Prevent undue interference with other researchers' use of the cloud
  - b. Maintain the security of the cloud
  - c. Ensure that cloud resources are shared fairly
  - d. Ensure compliance with both the law and university policies
3. Researchers will have no responsibility for any hardware maintenance. Warranty replacements, firmware updates etc. will all be handled by the cloud management team.
4. Researchers are responsible for the operating systems and software which run on the virtual machines and this can be handled by their existing research computing support if available/desired.

## **Governance Committee**

The Governance of the Research Cloud will be under the aegis of the Information Technology Oversight Committee (ITOC). This committee has representatives from every science department. Researchers with concerns, requests for information, suggestions for improvements etc. should contact their department's ITOC representative.

## **Research Groups and Departments**

Groups of researchers may merge some or all their resources into a single resource pool and may then collectively decide the permissions and policies governing use of that pool within the limitations of the OpenStack software that manages the cloud.

Departments may purchase resources in the cloud and the disposition of the resources will be determined by the department chair within the limitations of the OpenStack software managing the cluster. Departments should bear in mind that the cloud is aimed at research; if they wish to use it for teaching or administrative purposes they should discuss it with ITOC before purchase.

## **Support Level**

The cloud management team will provide support for all aspects of the cloud administration. This includes, but is not limited to: ensuring quotas are accurate, fixing any issues with the launching or shutting down of virtual machines, managing the storage, managing permissions etc.

Several stock virtual machine (VM) images will be provided for use by researchers. Researchers using these machine images will have kernel/driver-level support i.e. ensuring that the VM can access cluster storage etc.

Researchers may also develop and use their own VM images. In this case support will be limited to ensuring that the virtual machine launches and has the correct virtual hardware configuration: installation of drivers etc. will be the responsibility of the researcher.

Support for the physical hardware will be provided by IST and the machines will reside in a data centre which they manage. The OpenStack installation and configuration will be managed by Compute Canada.

## **Downtime Policy**

Occasionally servers in the cloud will need to be shutdown for maintenance and/or replacement or will go down due to power or hardware failure. For scheduled downtime a minimum of one week's notice will be provided to all users of the cloud. This notice will include a brief explanation of the need for the intervention along with the systems affected and the length of the downtime. During such downtimes administrators may, if necessary, restrict resource quotas to ensure smooth functioning of the cloud.

For unscheduled downtime due to equipment failure or other unexpected causes the priorities of the support team will be:

1. Act to minimize the impact of the failure on cloud users.

2. Communicate to the cloud users about the nature of the downtime including the cause, the ongoing impact on the cloud resources, the expected time frame before things will be fully restored and any possible workarounds to assist users function while the issue is being addressed.
3. Maintain regular communication updates to users until the situation is fully resolved.

## **Access to External Resources**

The Research Cloud will act as an interface to external, large scale computing resources such as those provided by academic programs such as Compute Canada as well as commercial offerings such as Amazon. The nature and degree of integration possible will depend on both the level of interest as well as the evolving nature of cloud technologies. ITOC will welcome contact from any parties interested in offering services to researchers as well as researchers who would like to have better access to an external service.

## **Security**

Access to the OpenStack software used to configure and launch virtual machines, manage disk storage etc. will be through Compute Canada user IDs. These are available to all researchers in Canada and so, if desired, researchers may allow access to colleagues outside the institute.

It is the responsibility and obligation of the researcher to ensure that virtual machines within the cloud are properly secured. This includes using secure authentication mechanisms to determine authorized users as well as ensuring the operating system is kept patched and up to date.

## **Resource Sharing**

Resources purchased by researchers will always be available for their use (subject to any cloud downtime for maintenance or serious hardware failure). Any idle resources may be temporarily used by any researcher on the cloud subject to the policies and procedures established.

## **Backup**

A backup system will be implemented for the cloud. Researchers wishing to have storage backed-up will need to purchase sufficient storage to store the backed up data archive as well as the live copy. The amount of storage required will depend on the details of the technical implementation, the backup schedule and the frequency with which the researcher's data changes. Details of the backup system will be provided and researchers will then need to make their own determination whether the available options provide sufficient protection for their needs.