# Rapid Access Cloud Guide: Part 1

## Welcome to The Cloud

Cybera operates the Rapid Access Cloud, an OpenStack cloud providing Infrastructure-as-a-Service, which allows users to set up virtual servers in any configuration desired. The virtualized cloud environment allows for the utilization of several operating systems (e.g. ubuntu, Windows) and makes it easy to scale up and down resources. It also offers object and volume storage, accommodating further use cases.

The cloud consists of two geographically distinct regions, hosted in Calgary and Edmonton. The high speed research and education network links the two regions through a 10Gbps connection providing excellent network performance and low latency. Upon registration, users are allocated a quota of 8 vCPUs, 8 GB RAM, and 500GB of storage space, and an IPv6 address for each instance.

> ⓘ The quota is shared between the two regions, but instances, volumes, etc. that are housed in the other region cannot be seen on the dashboard.

---

## Overview

This document steps through the process of using the Rapid Access Cloud dashboard to create a single Ubuntu 16.04 Linux instance accessible via ssh with a 25GB attached storage volume. This process will provide a basic understanding of the Rapid Access Cloud and how to create and manage instances and volumes. Once these steps have been completed, you are encouraged to explore the dashboard and build other instances, then move on the Advanced Guide for an understanding of some of the additional features that OpenStack and the Rapid Access Cloud provide, including information pertaining to Windows images and instructions on how to automate deployment in the cloud.

### Conventions in this document

There are step-by-step instructions provided throughout this document, however there are points along the way in which it is required to input information, such as the name of the instance or to make decisions, such as what operating system is installed in the instance.

- User inputted variables will be formatted with angle brackets:

```
<instance_name>, <security_group_name>, <description>
```

- If there are pre-determined choices, they will be formatted with parentheses:

```
(chocolate, strawberry, vanilla)
```

- Commands to be entered at a command prompt will be prefixed with a '$'
- File name paths are indicated as `/path/to/<some_file_or_command>`

## Step-by-step

# Create Account

Access to the Rapid Access Cloud is provided by third-party Federated identity providers: Google Identity Platform and Canadian Access Federation.

Users can log into either the Edmonton or Calgary region at cloud.cybera.ca.

## Google Identity Platform

If you already have an existing Google account, proceed to creating a Rapid Access Cloud account at https://rac-portal.cybera.ca/ and click Sign in using Google, otherwise create an account with Google at https://accounts.google.com/SignUp.

## Canadian Access Federation

Canadian Access Federation (CAF) permits member institutions a single sign-on (SSO) solution for access to network and network resources across Canada. Please check the list of participating institutions to see if you can use "Federated Single Sign-On" option at https://rac-portal.cybera.ca.

# Create a Key Pair

Before instances can be created, users will require a key pair that will be injected into the instance to permit access; any instance launched from an image in the Rapid Access Cloud **must** use key pairs to access the virtual machine for the first time, as password sign-on is disabled by default, key pairs being much more secure than a default password, though it does introduce extra steps. However, once a key pair is created, it can be used for any future instances that are created; further additional key pairs can be generated for different instances if there is a need to restrict access to various systems among different users.

1. Log-in to the Rapid Access Cloud dashboard at https://cloud.cybera.ca.
2. In the left-hand panel click on "Compute".
3. Click the "Key Pair" tab, then click "+Create Key Pair".
4. Enter a `<key_pair_name>`, then click "Create Key Pair". The browser will automatically download a file named `<key_pair_name>.pem`.
5. Move or save this file on your computer somewhere you will remember. It will be used when accessing instances created with this key pair.

⚠

⚠ If this key file is saved to an operating system that uses file-system permissions (Unix, Linux, BSD, OSX) then make sure the permissions are set appropriately. Typically, the `.ssh` directory permissions ought to be set to `700` (`drwx------`) and the private key (`*.pem`) should be `600` (`-rw-------`). To change the permissions of the downloaded key to `600`, do:

```
$ chmod 600 /path/to/<key_pair_name>.pem
```

🚫 The private key generated in the step above is not recoverable if it is lost. It is highly recommended that a backup of the key is made and kept safe, ideally on a separate hard drive or USB key.

ⓘ **About key pairs**

Key pairs are a set of mathematically generated strings, one is the *private key* and the other is the *public key*. The key pairs that are used in the Rapid Access Cloud are ssh keys generated by the OpenStack dashboard, keeping the public key to be injected as needed into new instances, and the private key is the *.pem file automatically downloaded by the browser in the steps above. If you already have a key pair suited for use, [import that key](#) by following the steps in the Advanced Guide.

A detailed explanation of public-key cryptography is out of the scope of this document; [this](#) will help you understand it better.

---

# Security Groups

Security groups are the policies that deny access to the network ports of an instance. Security groups are therefore **firewalls** for instances, with a set of default policies that block all access to each port from any source, including the computer you are using to access the Rapid Access Cloud dashboard. Before an instance can be accessed the appropriate ports will need to be opened and a source IP address or range of addresses will need to be configured.

There is a default security group that cannot be deleted, however it can have rules added and removed from it. Additional security groups can also be created depending on need. When a new Rapid Access Cloud account is created, the default security group has four rules. The Egress rules, traffic going out from the instance, is permitted to go out by default. The Ingress rules, traffic going in to the instance, is denied by default since it does not specify any network as seen in the Remote IP Prefix. Thus, a few rules are required to permit basic access.

🚫 We strongly advise against highly permissive security groups. Allowing access from any source (i.e. 0.0.0.0/0 or ::/0) means that anyone in the world can access your instance. We advise limiting the traffic to your instance to the smallest possible CIDR.

Allowing access to all ports (port range 1 : 65535) means that anyone from the allowed CIDR can access all services on your instance - even services which should only be internal.

We monitor for security groups rules which are providing open access to specific ports. If we discover and open access rule, it will be deleted. An "open access" security group rule is one where the CIDR is set to either 0.0.0.0/0 for IPv4 or ::/0 for IPV6. The security group rule deletion will be applied for the following ports:

- 42, 1512: WINS
- 88: Kerberos
- 135, 137, 138, 139: NetBIOS/SMB/Samba
- 389: LDAP
- 445, 3268: Active Directory
- 3389: Remote Desktop
- 5985, 5986: WinRM
- 9200: Elasticsearch
- 27017: MongoDB

You can still configure access to these ports, but they must not be in an open-access nature. Please use a CIDR of a specific IP address when creating security group rules for these ports.

For best practices, please read [this informative blog on OpenStack Security.](#)

## Modify the default security group

⚠

1. Log-in to the Rapid Access Cloud dashboard at https://cloud.cybera.ca.
2. In the left-hand panel under "Network", click "Security Groups".
3. Click the "Manage Rules" button on the right hand side associated with the "default" security group. The list initially has four rules, however we are going to add rules that:
   a. permit ICMP for ping and traceroute, from any IPv4 or IPv6 address
   b. permit SSH from any IPv4 or IPv6 address
4. Click "+Add Rule" in the top right. We are going to be adding four rules. For each rule input the values, then click the blue "Add" button. Note, the first and third rules are for IPv4 access, while the second and fourth are for IPv6:

Rule: `(All ICMP)`

Remote: `(CIDR)`

CIDR: `0.0.0.0/0`

Rule: `(All ICMP)`

Remote: `(CIDR)`

CIDR: `::/0`

Rule: `(SSH)`

Remote: `(CIDR)`

CIDR: `0.0.0.0/0`

Rule: `(SSH)`

Remote: `(CIDR)`

CIDR: `::/0`

5. Verify the new security group rules in the "default" security group.

### ICMP

Adding rules for the ICMP protocol will allow you to "ping" your instance from the outside world. ICMP rules can be added in one of two ways:

1. Choosing "All ICMP". This will allow all ICMP traffic to your instance.
2. Specifying a specific ICMP type. This involves setting a "Type" and optional "Code". See here for a reference.

---

# Instances

Instances are the virtual machines that run in the Rapid Access Cloud, and they are provisioned with a set of specifications not unlike traditional bare-metal hardware with processors, memory and storage being the primary configurable elements. The Rapid Access Cloud utilizes **flavors** or pre-configured templates that determine the number of virtual CPU (vCPU), available memory (RAM) and disk space. There are six flavours to choose from with the details of each available during the instance creation process.

## Launch an instance

1. Log-in to the Rapid Access Cloud dashboard at https://cloud.cybera.ca.
2. In the left-hand panel under "Compute", click "Instances".
3. Click on "Launch Instance" in the top right corner.
4. In the "Details" tab, specify the following parameters. **Do not yet click "Launch"**:

Availability Zone: `(nova)`

Instance Name: `<your_instance_name>`

Flavor: `(m1.small)`

Instance Count: `1`

Instance Boot Source: `(Boot from image)`

Image Name: `Ubuntu 16.04`

> ⚠ The following Instance Boot Source options are NOT supported:
>
> - Boot from volume
> - Boot from image (creates a new volume)
> - Boot from volume snapshot (creates a new volume)
>
> If you select one of those options, it will result in an error.

5. Click "Access & Security" tab within the Launch Instance field, and select the `<key_pair_name>` created earlier in the tutorial. The default security group should be checked as well.

6. Click the "Launch" button. It should take the instance less than 2 minutes to launch; progress can be monitored in the "Status" column and should say "Active" when ready.

## About Flavors

OpenStack uses 'flavors' to define the compute, memory, and storage capacity of computing instances. To put it simply, a flavor is an available hardware configuration for a server. The flavors available in the Rapid Access Cloud allow for a broad deployment of virtual machines given the default quotas available to users. The flavors are in two classes: m1 and g1.

- **M1** class flavors are General Purpose Instances. This family provides a balance of compute, memory, and network resources, and it is a good choice for many applications.

| Flavor | VCPUs | Root Disk (GB) | RAM (MB) | Swap (MB) |
|---|---|---|---|---|
| m1.tiny | 1 | 5 | 512 | 512 |
| m1.micro | 1 | 5 | 1,024 | 1,024 |
| m1.small | 2 | 20 | 2,048 | 2,048 |
| m1.medium | 2 | 40 | 4,096 | 4,096 |
| m1.large | 4 | 40 | 8,192 | 4,096 |
| m1.xlarge | 8 | 40 | 16,384 | 4,096 |

- **G1** instances are intended for general-purpose GPU compute applications. Use cases include machine learning, rendering, and other server-side GPU compute workloads. (see GPU-Enabled Instances for more on how to launch GPU instances.)

| Flavor | VCPUs | Root Disk (GB) | RAM (MB) | Swap (MB) |
|---|---|---|---|---|
| g1.medium.auto-destruct | 8 | 40 | 8,192 | 4,096 |
| g1.large.auto-destruct | 8 | 40 | 32,768 | 4,096 |
| g1.xlarge.auto-destruct | 16 | 40 | 32,768 | 4,096 |

## Accessing instances

Having an instance up and running is one thing, and perhaps just a single Linux 'sandbox' to run some code is all that is needed, however the real power of computers, virtual or otherwise, is in connectivity, and that means networks. The instances in the Rapid Access Cloud can be connected to and accessed in a variety of ways, permitting users to create an environment with multiple instances networked together in the same way a bare-metal environment can be built, but in this case with virtual machines providing the routing, switching and other network functions along with the expected servers running applications on top of operating systems like Linux and Windows.

## IP addresses

Once an instance is provisioned, it is automatically given two addresses: a private IPv4 address, and a public IPv6 address.

The **private address**, called a *fixed-ip* in OpenStack, is not publicly routable (that is, not reachable from the public internet; though see the Making the most of a single IPv4 address section for more on this) and is used by the OpenStack application stack to provision the instance. The private address can also be used to communicate between instances without the need of routing traffic out on the internet and back again. The private addresses are assigned from a range of 10.1.0.0 - 10.2.254.254.

The **public address** given to the instance automatically is an IPv6 address. The ability to connect to an instance via IPv6 will be limited by the network the connection is coming *from*; unfortunately many schools, workplaces, or home internet providers do not have IPv6 capable networks. Use the following tools to determine if your network has the ability to route IPv6 traffic:

- Test-IPv6.com
- IP6.me

Cybera has confirmed that most users who use TELUS or Shaw, University of Alberta users (via WiFi), and Concordia users have IPv6 connectivity. If you fall into one of these groups, please use the above links to verify and if at all possible, we ask that you use IPv6 to connect to your instance in order to help us conserve IPv4 addresses.

## Why IPv6?

IPv4 addresses are limited to 4.3 billion addresses and the world is quickly approaching the limit and obtaining new ones will become impossible. IPv6 uses 128-bit addresses instead of IPv4's 32-bit addresses, removing this limit and providing better security features over IPv4.  IPv6 is not a new technology as it has been available for over a decade, though it has been slow to implement due to legacy concerns across multiple industries.

If you're more curious about IPv6 itself we recommend reading the Wikipedia article on IPv6.

## IPv6 Considerations

One thing to note about using IPv6 addresses is that some tools require a different format when attempting to connect directly via the IP address. In general we recommend to use our helper DNS record or use a DNS record in place of the IPv6 address.

Please note that of the most common utilities on macOS and Linux to copy files to and from instances has a different syntax to use raw IPv6 addresses instead of IPv4 addresses. You will need to include and escape square brackets around the IPv6 address.

```
scp myfile ubuntu@\[2001:db8:0:1\]:myfile
```

# Public IPv4 addressing

Rapid Access Cloud accounts are able to also have public IPv4 addresses allocated to them, however due to demand the default quota is 0. This is referred to as a *floating-ip* in OpenStack and can be associated with only one instance at a time.

This limitation on IP address availability can be overcome, and in fact can make for a more robust and secure cloud environment in some cases. Please see the Making the most of a single IPv4 Address section for solutions to this problem.

---

# Requesting, Allocating and associating floating IPs

If the default public IPv6 address is not sufficient for your use case, public IPv4 addresses are available. By default a new Rapid Access Cloud account does not have a quota permitting a public IPv4 address. Please see the IP address section above for testing whether IPv6 is a limitation for access to your resources.

## Requesting a floating IP address

To request a floating IP address you can fill out a form on the dashboard by:

1. Log-in to the Rapid Access Cloud dashboard at https://cloud.cybera.ca.
2. In the left-hand panel under "RAC", click "Quota Change".
3. Request one floating IP and fill out the reason and then press "Submit a Quota Change Request"

rac-admin will be notified and we will process your quota change request as soon as possible.

## Allocating a floating IP address

If your project does have have quota for a floating IP, the floating IP is not allocated to your automatically. You will need to allocate a floating IP to your project. Given the scarcity of the addresses for Cybera (and indeed, the world), **addresses allocated to projects that have gone unused for three months will be reclaimed**, however you are welcome to allocate an IP address again if needed.

1. Log-in to the Rapid Access Cloud dashboard at https://cloud.cybera.ca.

2. In the left-hand panel under "Compute", click "Instances".
3. Click the Action drop-down button on the right-hand side and select "Associate Floating IP".
4. Click on the "+" sign next to "Select an IP address".
5. There is only one pool of addresses available (`nova`) and the quota shows only one IP address from that pool, so simply click "Allocate IP".
6. After the IP address has been allocated, click the "Associate" button the the right hand side. Under the Instances summary, your Instance should now have three IP addresses, including a publicly accessible IPv4 address.

### Associating the floating IP address to other instances

Given the ease of creating and destroying instances, along with the possibilities of changing needs, moving the IPv4 address around may be required. Once the address is allocated to a project, it can be assigned to any instance associated with the project. When an instance is destroyed, the associated IP address remains allocated to the project. Simply follow the steps above to associate the IP address with a new instance, omitting the steps for allocating the address (steps 4 and 5).

## Windows instances

Cybera does provide Windows images in the Rapid Access Cloud, in the form of Windows Server 2016. You will need to provide your own valid license from Microsoft, and follow additional steps covered in the Advanced Guide.

---

# Log in

After creating a key pair, modifying the default security group, and clicking through the process of launching an instance and assigning a floating IP, you are now able to log in. The basics are the same regardless of where you are logging in *from*, but like many things it is the details that matter.

Because the instance built according to this document is a Ubuntu Linux distribution, the default username is **ubuntu**. For each of the other Linux distributions available in the Rapid Access Cloud, the username follows the same scheme: centos, debian, and fedora.

## IPv6 Access

As mentioned in the Accessing Instances section, connecting to an instance via IPv6 is preferred over using a Floating IP. If you do have a floating IP address associated with your instance you can use your floating IP instead of the IPv6 address or use the DNS name autogenerated as seen in the instructions below.

To make IPv6 addresses more friendly, instances are given an automated DNS name that maps to your IPv6 address and if applicable your floating IP address. The autogenerated domain can be found on your instance's details page under the Metadata section. (eg. 12345.yyc.cybera.ca)

## From Linux, UNIX, or BSD (including OS X)

The simplest way to use ssh from these operating systems is:

1. Open a terminal and enter:

```
$ ssh -i /path/to/<key_pair_name>.pem ubuntu@<ipv6_address_or_dns_name>
```

2. Answer 'yes' to the following question:

```
The authenticity of host '<ipv6_address> (<ipv6_address>)' can't be established.
RSA key fingerprint is e5:de:ad:c3:be:ef:b2:ba:be:a1:ba:dc:af:ea:ce:d4.
Are you sure you want to continue connecting (yes/no)?
```

3. You are now logged in and will then be presented with the Message Of The Day and a shell prompt:

```
--------------------------
Cloud Image Helper Scripts
--------------------------
To enable automatic updates please run:
/usr/local/bin/enableAutoUpdate
To install the latest OpenStack tools please run:
/usr/local/bin/installOpenStackTools
To use the local software update proxy please run:
/usr/local/bin/localSUS
To remove this message from your message of the day please run:
sudo rm /etc/motd
ubuntu@<your_instance_name>:~$
```

# From Windows

There are plenty of applications that allow ssh access from within Windows, none of which are bundled *with* Windows (at this time). You are welcome to use any that fit your needs, however puTTY is a widely used and well supported suite of SSH utilities that includes key management and generation, an scp client and the ssh client itself.

In order to connect to your instances using puTTY, you first need to convert your private key to a puTTY compatible format:

### Convert OpenStack private key to puTTY compatible key

1. Launch PuTTYGen, installed as part of the PuTTY suite.
2. Click Conversions from the "PuTTY Key Generator" menu and select Import key.
3. Navigate to the OpenStack private key (*.pem) used in the instance you would like to connect to and click "Open".
4. Under Actions / Save the generated key, select Save private key.
5. Choose an optional passphrase to protect the private key.
6. Save the private key to the desktop as a *.ppk.

### Connect to instance with PuTTY

1. Open the main PuTTY application.
2. Enter the remote server IP address under the "Session" category in left-hand panel.
3. Navigate to the "Connection" category, then "Data".
4. Under Login details, enter the username to log in with. For ubuntu instances on the Rapid Access Cloud, the default is `ubuntu`.
5. In the "Connection" category in the left-hand menu, select "SSH"  and then "Auth".
6. Click "Browse…" under "Authentication parameters / Private key file" for authentication.
7. Locate the *.ppk private key you generated above and click "Open".
8. Navigate back to the "Session" category to name the session and then click "Save".
9. Click "Open" to log into the remote server with key pair authentication.

---

# Volumes

You now have a running instance in the cloud. This flavor we used is m1.small, and it has 20 GB of root storage. The nature of the instances are ephemeral, so any data left in an instance when it is destroyed is gone for good. It makes much more sense to create a data volume that is independent of the instance, and can be attached and reattached as needed, much like attaching a USB disk to the computer.

Initially, in order to take advantage of volumes, a two step process of creating and then attaching must be followed. It is also likely the volume will need to be formatted before it can be used.

## Create a volume

1. Log-in to the Rapid Access Cloud dashboard at https://cloud.cybera.ca.
2. In the left-hand panel under "Compute", click "Volumes".
3. Click "+Create Volume" on the right hand side.
4. In the Create Volume screen enter the following values:

   Volume Name: `<your_volume_name>`

   Description: `<your_description>`

   Type: **LVM**

   Size (GB): `25`

   Availability Zone: `(Any availability zone)`

5. Click the blue "Create Volume" button and after a few moments a 25GB volume will be ready for attachment to an instance.

## Attach a volume

1. Log-in to the Rapid Access Cloud dashboard at https://cloud.cybera.ca.
2. In the left-hand panel under "Compute", click "Volumes".
3. Click the Action drop-down button on the right-hand side and select "Manage Attachments".
4. Under "Attach to Instance" select the instance the volume is be attached to, then click the blue "Attach Volume" button.
5. After a few moments, the volume will be attached. Take note of the "Attached to" column on the summary screen, it will list where it is attached like "/dev/sdc".

## Use a volume

1. [Log into](#) your instance via ssh.
2. Format the volume:

```
$ sudo mkfs.ext4 /dev/sdc
```

> ⓘ   Attached volumes will typically be assigned device names in sequential order (i.e. /dev/sdc, /dev/sdd, /dev/sde, etc.)

3. List all disks from within the instance with:

```
$ sudo fdisk -l
```

> ⓘ   /dev/sda and /dev/sdb are the system volumes that make up the instance.

4. Create a mount point for the volume:

```
$ sudo mkdir /mnt/<mount_point_name>
```

5. Mount the volume device to the mount point:

```
$ sudo mount /dev/sdc /mnt/<mount_point_name>
```

6. Permissions may need to be changed on the new volume, as they are initially set to root:

```
$ sudo chown ubuntu:ubuntu /mnt/<mount_point_name>
```

Additional information on how to use other volume types is available in our Advanced Guide under [Volume Types](#).

---

# Backups

The Cybera Rapid Access Cloud is offered on a best-effort basis and it is the users' responsibility to ensure appropriate backups of all their data are made. While Cybera's record for uptime and data-loss is very good, we strongly encourage all of our users to ensure regular backups are made.

## Instance Snapshots

Snapshots are point-in-time copies of your instances. You can snapshot your instance and then either download that snapshot as a backup or use that snapshot as a "cookie cutter" template to build new instances based off of it. This has the advantage that you will not have to reconfigure your virtual machine after launching from the snapshot, as opposed to launching from the image. Also, you can use this to resize your instance: if you want your instance to have different resources assigned to it (e.g. CPU or RAM) you can re-launch from the snapshot according to your preferred configuration (note that the root disk cannot be smaller than in the original instance).

To snapshot your instance:

1. Log-in to the Rapid Access Cloud dashboard at [https://cloud.cybera.ca](https://cloud.cybera.ca).
2. In the left-hand panel under "Compute", click "Instances"
3. For the instance you want to snapshot, **make sure it is in a shut-off state**.
4. Click on the "Create snapshot" button.
    a. Select a name to save the snapshot as
5. OpenStack will then navigate to the Images section. The snapshotting process can take considerable time depending on the size of the instance.

You should verify that the snapshotting process has worked following completion. Check the size of the snapshot as it should be at least several hundred MB in size. Also consider test launching an instance from the snapshot in order to ensure the process concluded successfully.

### Exporting a Snapshot

If you would like to take an instance snapshot to a different cloud, or a different account within the Rapid Access Cloud, you can use the OpenStack cli tool to export the snapshot image.

1. In a terminal, with the openstack cli tools installed, and the openrc file sourced, type: **openstack image list**

2. In the list of returned images take the ID for the snapshot you have crated and run: **openstack image save --file export.qemu $id** # Where $id is the ID from the command above.

The above will save your image as export.qemu in the folder you ran the command. You can then upload that image to a new account or another cloud.

---

# Object storage

Users may also use the Object Storage service to back-up their data. See this document for more information.

---

# Conclusion

Congratulations! You now have enough knowledge of Cybera's Rapid Access Cloud and OpenStack to deploy instances, manage networking and data on and amongst them. This encompasses a majority of use cases for cloud and may very well be enough for your needs. However, the true potential of cloud architecture is realized once you move beyond manual creation and automate the creation and destruction of instances, while maintaining data integrity in volumes. These topics and more are covered in the Advanced Guide.