Introduction to Cloud and Virtual Machines

Venkat Mahadevan

Cloud Architect

UBC Advanced Research Computing

July 5th, 2022

Introduction

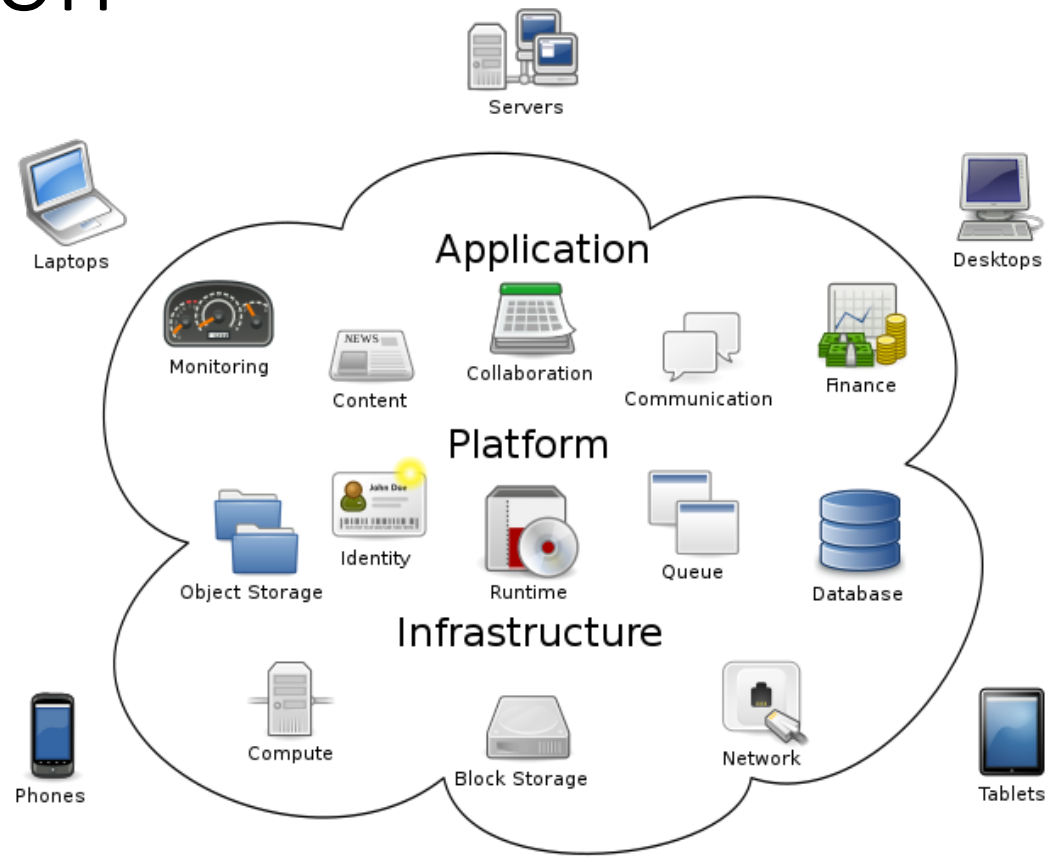

https://commons.wikimedia.org/wiki/File:Cloud_computing.svg
Author: Sam Johnston

Introduction (cont.)

- Delivers high level services and access to system resources over the Internet.
- Services: collaboration (E-mail, calendaring, etc.), web, Dropbox-like file hosting, etc.
- System resources i.e. infrastructure: compute, disk, networking, load balancing, etc.

Introduction (cont.)

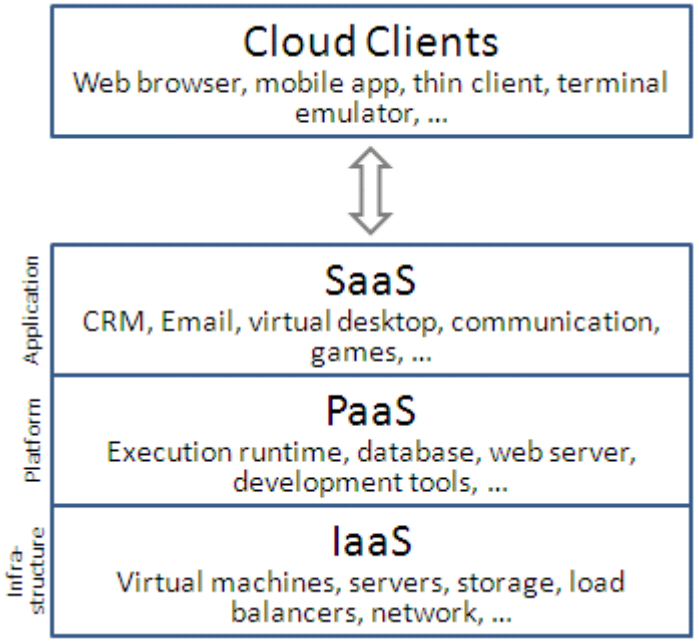

https://commons.wikimedia.org/wiki/File:Cloud_computing_layers.png

Digital Research Alliance of Canada (The Alliance) Cloud

Arbutus cloud (arbutus.cloud.computecanada.ca)

Node count	CPU	Memory (GB)	Local (ephemeral) storage	Interconnect	GPU	Total CPUs	Total vCPUs
156	2 x Gold 6248	384	2 x 1.92TB SSD in RAID0	1 x 25GbE	N/A	6,240	12,480
8	2 x Gold 6248	1024	2 x 1.92TB SSD in RAID1	1 x 25GbE	N/A	320	6,400
26	2 x Gold 6248	384	2 x 1.6TB SSD in RAID0	1 x 25GbE	4 x V100 32GB	1,040	2,080
32	2 x Gold 6130	256	6 x 900GB 10k SAS in RAID10	1 x 10GbE	N/A	1,024	2,048
4	2 x Gold 6130	768	6 x 900GB 10k SAS in RAID10	2 x 10GbE	N/A	128	2,560
8	2 x Gold 6130	256	4 x 1.92TB SSD in RAID5	1 x 10GbE	N/A	256	512
240	2 x E5-2680 v4	256	4 x 900GB 10k SAS in RAID5	1 x 10GbE	N/A	6,720	13,440
8	2 x E5-2680 v4	512	4 x 900GB 10k SAS in RAID5	2 x 10GbE	N/A	224	4,480
2	2 x E5-2680 v4	128	4 x 900GB 10k SAS in RAID5	1 x 10GbE	2 x Tesla K80	56	112

Location: University of Victoria
Total CPUs: 16,008 (484 nodes)
Total vCPUs: 44,112
Total GPUs: 108 (28 nodes)
Total RAM: 157,184 GB
5.3 PB of Volume and Snapshot **Ceph** storage.
12 PB of Object/Shared Filesystem **Ceph** storage.

The Alliance Cloud (cont.)

Cedar cloud (cedar.cloud.computecanada.ca)

Node count	CPU	Memory (GB)	Local (ephemeral) storage	Interconnect	GPU	Total CPUs	Total vCPUs
28	2 x E5-2683 v4	256	2 x 480GB SSD in RAID1	1 x 10GbE	N/A	896	1,792
4	2 x E5-2683 v4	256	2 x 480GB SSD in RAID1	1 x 10GbE	N/A	128	2,560

Location: Simon Fraser University

Total CPUs: 1,024

Total vCPUs: 4,352

Total RAM: 7,680 GB

500 TB of persistent **Ceph** storage.

The Alliance Cloud (cont.)

Graham cloud (graham.cloud.computecanada.ca)

Node count	CPU	Memory (GB)	Local (ephemeral) storage	Interconnect	GPU	Total CPUs	Total vCPUS
6	2 x E5-2683 v4	256	2x 500GB SSD in RAID0	1 x 10GbE	N/A	192	
2	2 x E5-2683 v4	512	2x 500GB SSD in RAID0	1 x 10GbE	N/A	64	
8	2 x E5-2637 v4	128	2x 500GB SSD in RAID0	1 x 10GbE	N/A	256	
8	2 x Xeon(R) Gold 6130 CPU	256	2x 500GB SSD in RAID0	1 x 10GbE	N/A	256	
3	2 x E5-2640 v4	256	2x 500GB SSD in RAID0	1 x 10GbE	N/A	120	
12	2 x Xeon(R) Gold 6248 CPU	768	2x 1TB SSD in RAID0	1 x 10GbE	N/A	480	

Location: University of Waterloo

Total CPUs: 1,368

Total vCPUs:

Total RAM: 15,616 GB

84 TB of persistent **Ceph** storage.

The Alliance Cloud (cont.)

Béluga cloud (beluga.cloud.computecanada.ca)

Node count	CPU	Memory (GB)	Local (ephemeral) storage	Interconnect	GPU	Total CPUs	Total vCPUs
96	2 x Intel Xeon Gold 5218	256	N/A, ephemeral storage in ceph	1 x 25GbE	N/A	3,072	6,144
16	2 x Intel Xeon Gold 5218	768	N/A, ephemeral storage in ceph	1 x 25GbE	N/A	512	10,240

Location: École de Technologie Supérieure

Total CPUs: 3,584

Total vCPUs: 16,384

Total RAM: 36,864 GiB

200 TiB of replicated persistent SSD **Ceph** storage.

1.7 PiB of erasure coded persistent HDD **Ceph** storage.

The Alliance Cloud (cont.)

- There is also the NextCloud service which provides 50GB of backed up Dropbox-like storage (https://docs.alliancecan.ca/wiki/Nextcloud).

The Alliance Cloud (cont.)

- The IaaS clouds are built on OpenStack.
- OpenStack is an open-source software platform for deploying clouds i.e. build your own cloud environment.
- Can work with a variety of hardware, network switches, hypervisors.

The Alliance Cloud (cont.)

- Various commercial vendors provide OpenStack:
  - SUSE
  - Redhat
  - Ubuntu
  - Mirantis

- Also exists a free implementation called OpenStack-Ansible which is in use by The Alliance:
  - https://github.com/openstack/openstack-ansible

Cloud Resources

Attributes	Compute Cloud[1]	Persistent Cloud[1]
Who can request	PIs only	PIs only
VCPUs (see VM flavours)	80	25
Instances[2]	20	10
Volumes[2]	2	10
Volume snapshots[2]	2	10
RAM (GB)	300	50
Floating IP	2	2
Persistent storage (TB)	10	
Object storage (TB)[3]	10	
Shared filesystem storage (TB)[3]	10	
Default duration	1 year[4], with 1 month wall-time	1 year (renewable)[4]
Default renewal	April[4]	April[4]

Cloud Resources (cont.)

- You can request resources via the Rapid Access Service (RAS) or Resource Allocation Competition (RAC):

- https://alliancecan.ca/en/services/advanced-research-computing/research-portal/accessing-resources

Other Free Services

- https://www.infoworld.com/article/3179785/aws-vs-azure-vs-google-cloud-which-free-tier-is-best.html

- Amazon Web Services: https://aws.amazon.com/free/; 1-2 VCPU free for 12 months (t2.micro or t3.micro instances depending on region).

- Data egress is typically charged.

Hands-On

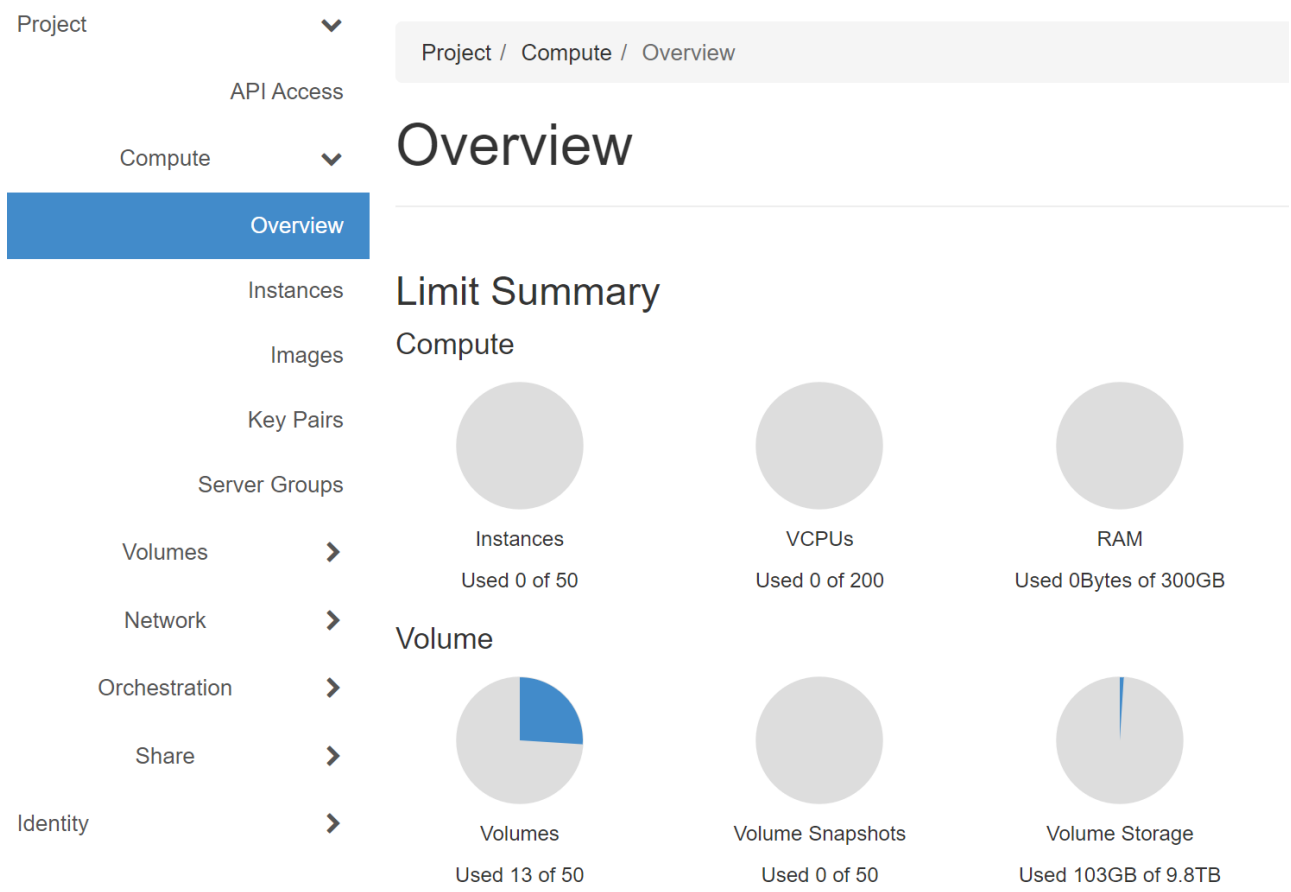
Project
API Access
Compute
Overview
Instances
Images
Key Pairs
Server Groups
Volumes
Network
Orchestration
Share
Identity
Project / Compute / Overview
Overview
Limit Summary
Compute
Instances
Used 0 of 50
VCPUs
Used 0 of 200
RAM
Used 0Bytes of 300GB
Volume
Volumes
Used 13 of 50
Volume Snapshots
Used 0 of 50
Volume Storage
Used 103GB of 9.8TB

Create SSH Key Pair and Download Private Key

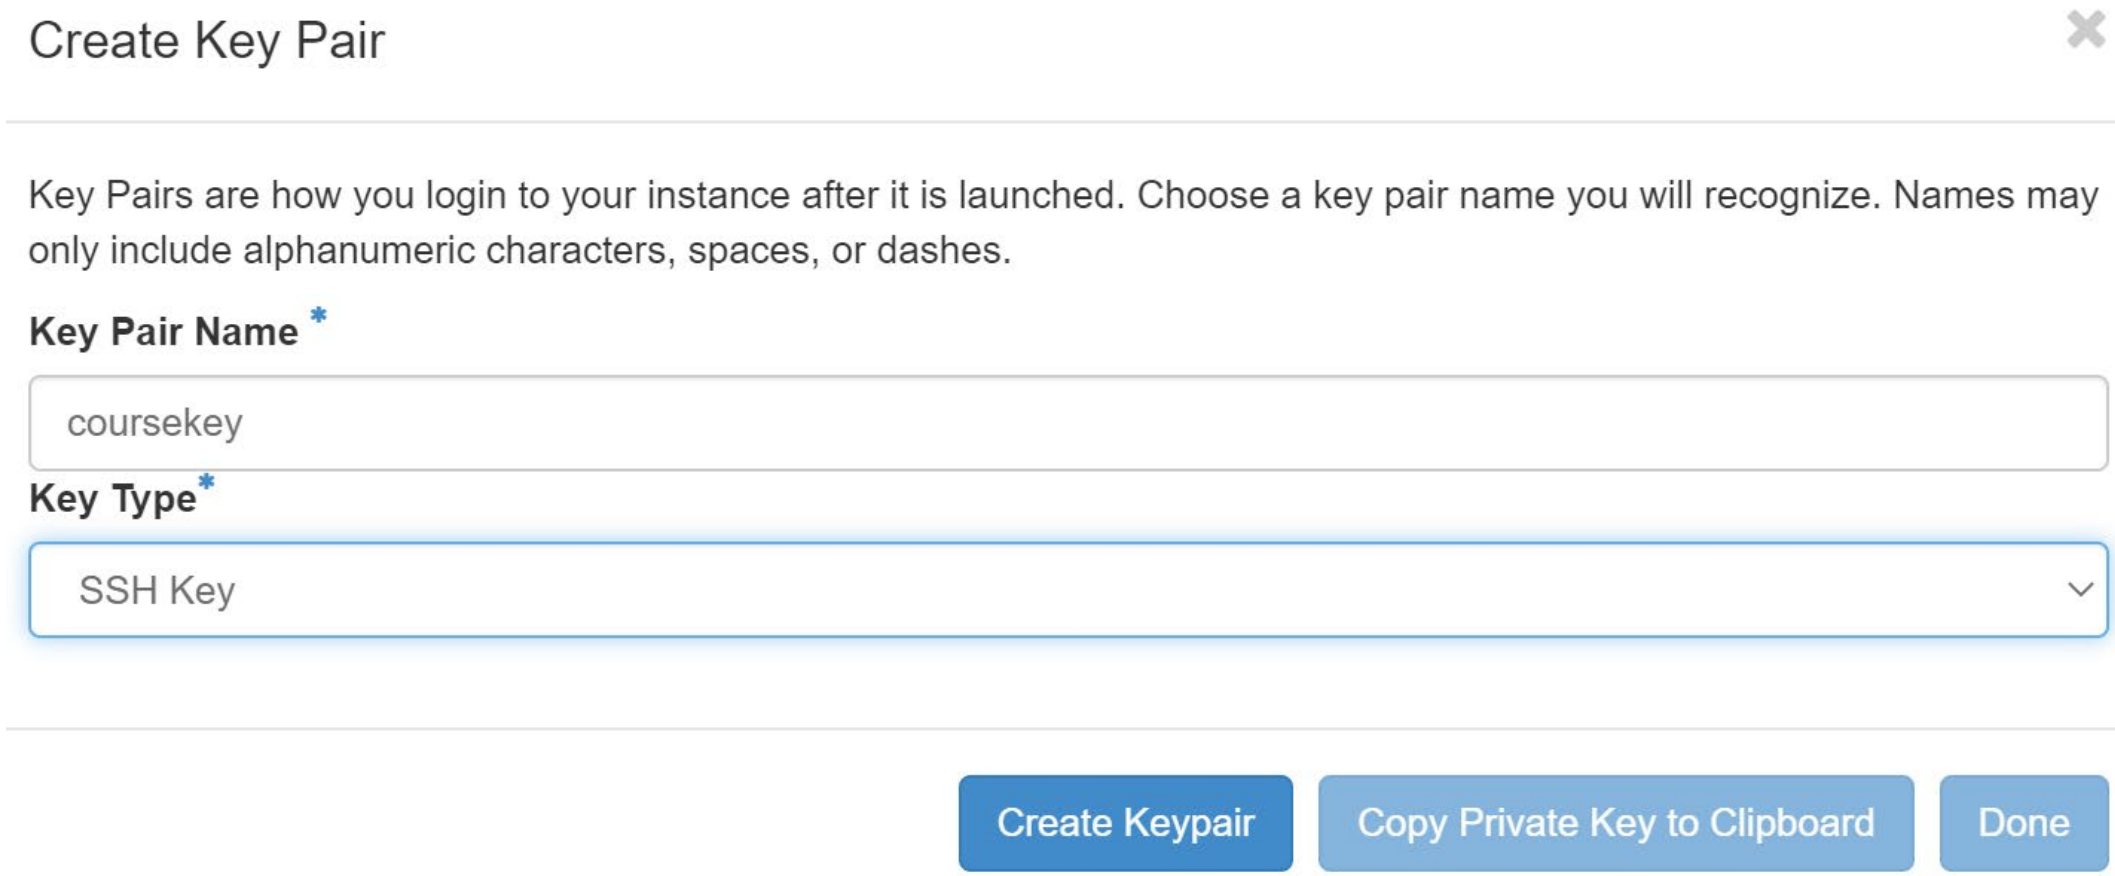
Create Key Pair
Key Pairs are how you login to your instance after it is launched. Choose a key pair name you will recognize. Names may only include alphanumeric characters, spaces, or dashes.
Key Pair Name *
coursekey
Key Type*
SSH Key
Create Keypair
Copy Private Key to Clipboard
Done

Launch Instance of a Virtual Machine

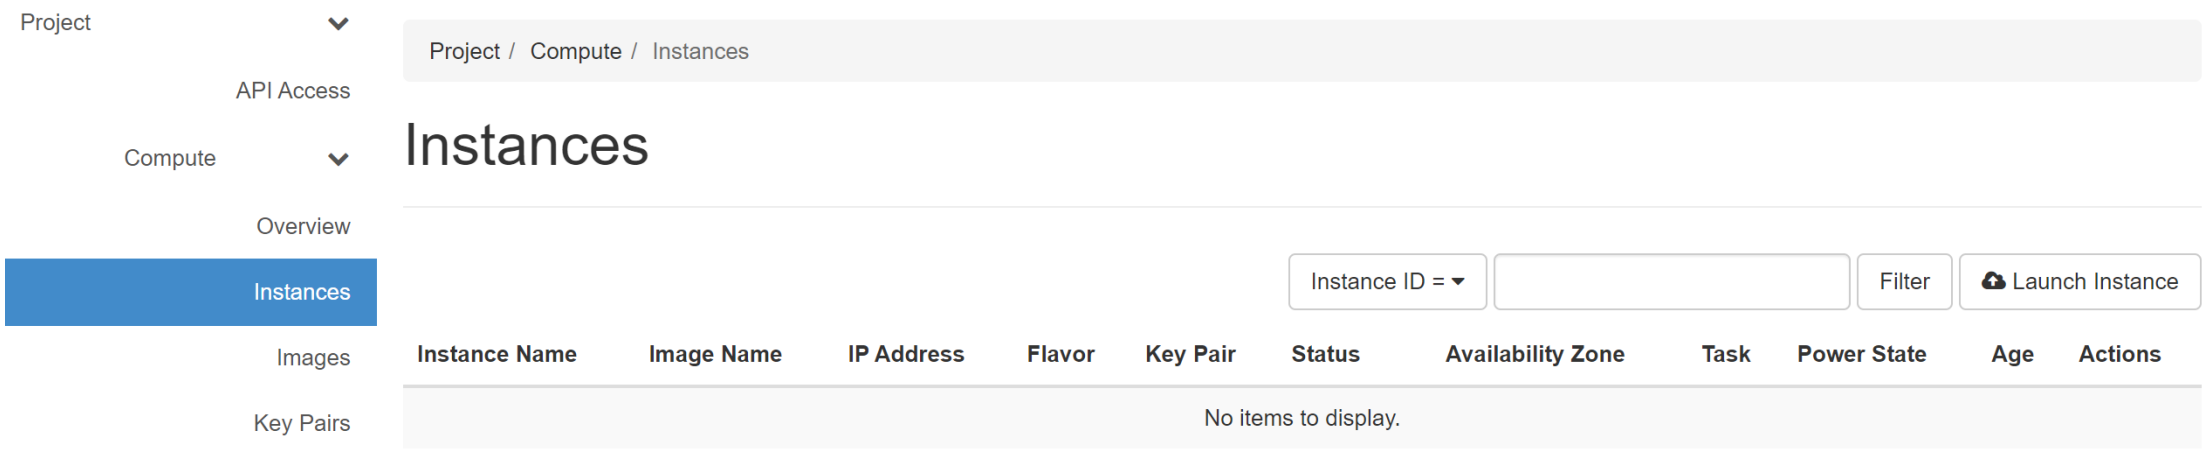

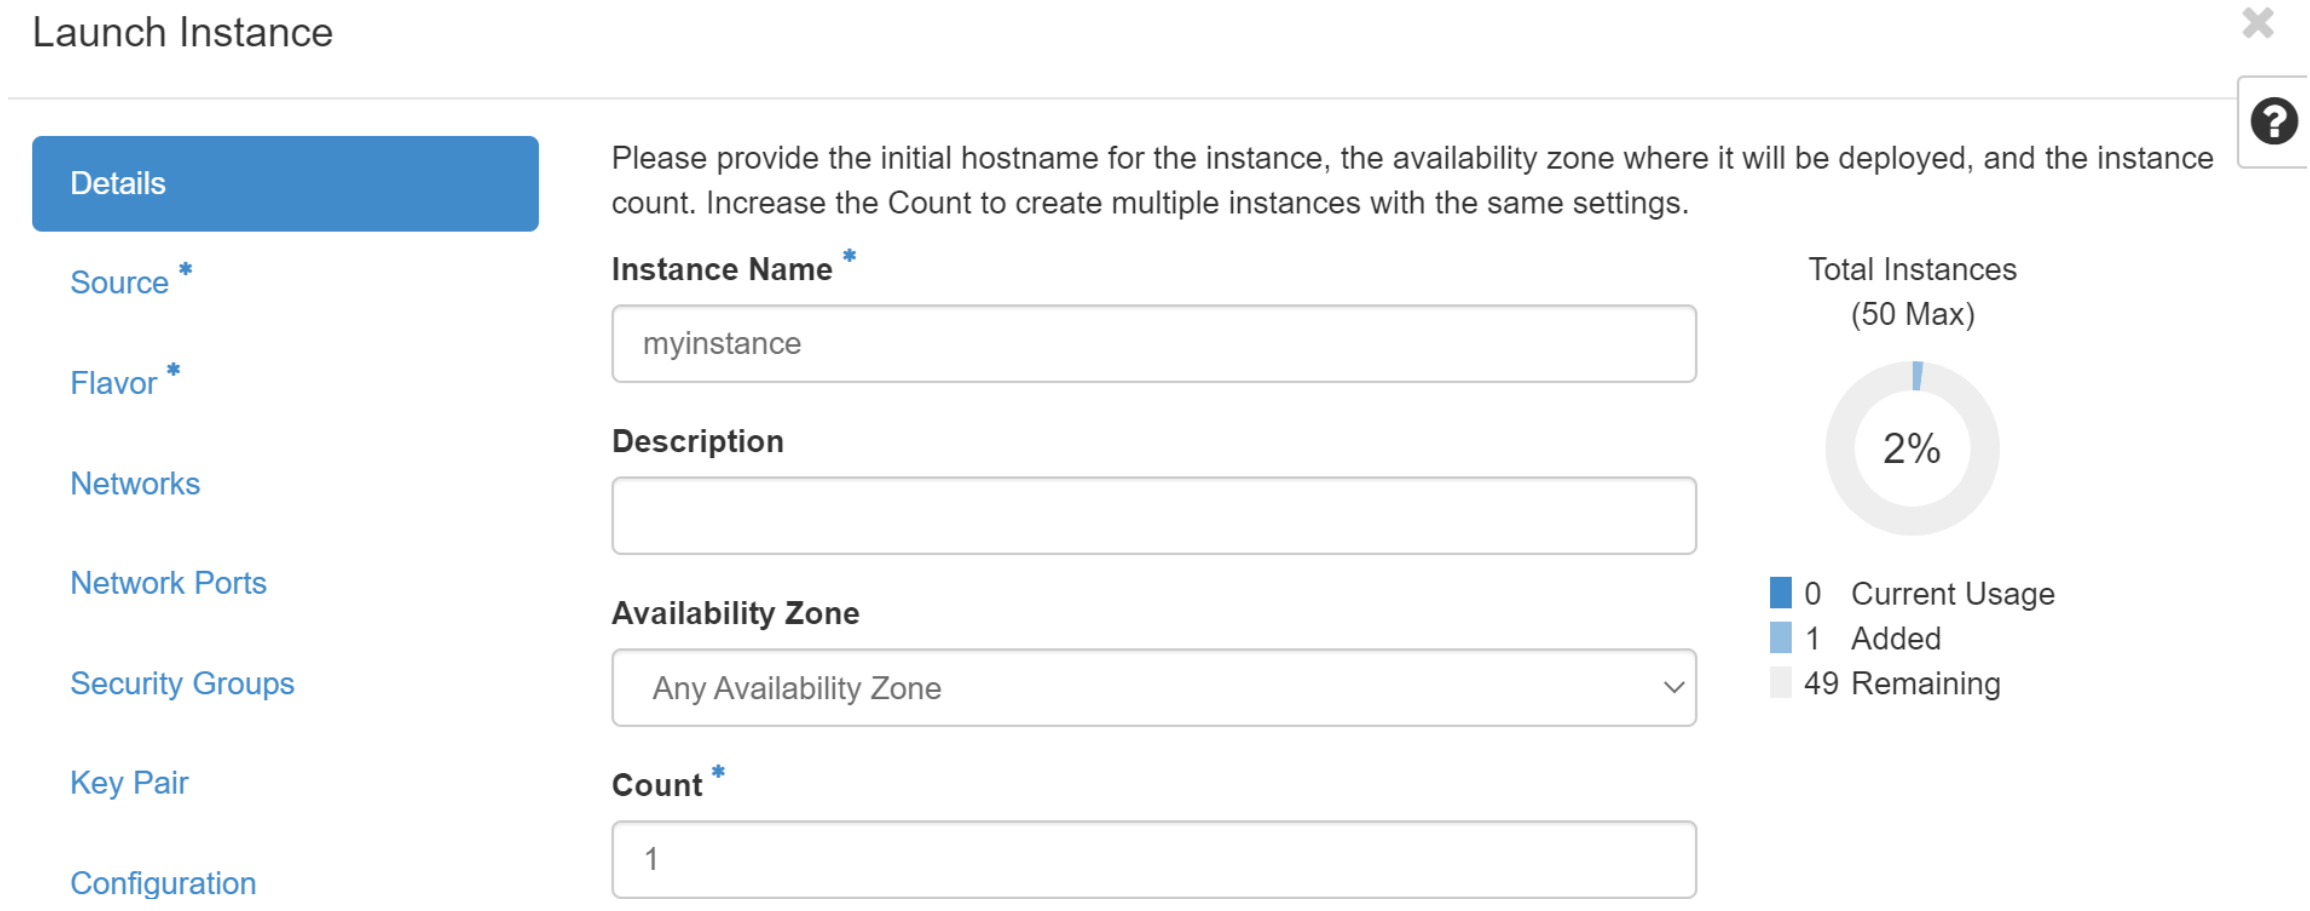

Digital Research Alliance of Canada

SFU SIMON FRASER UNIVERSITY

UBC Advanced Research Computing

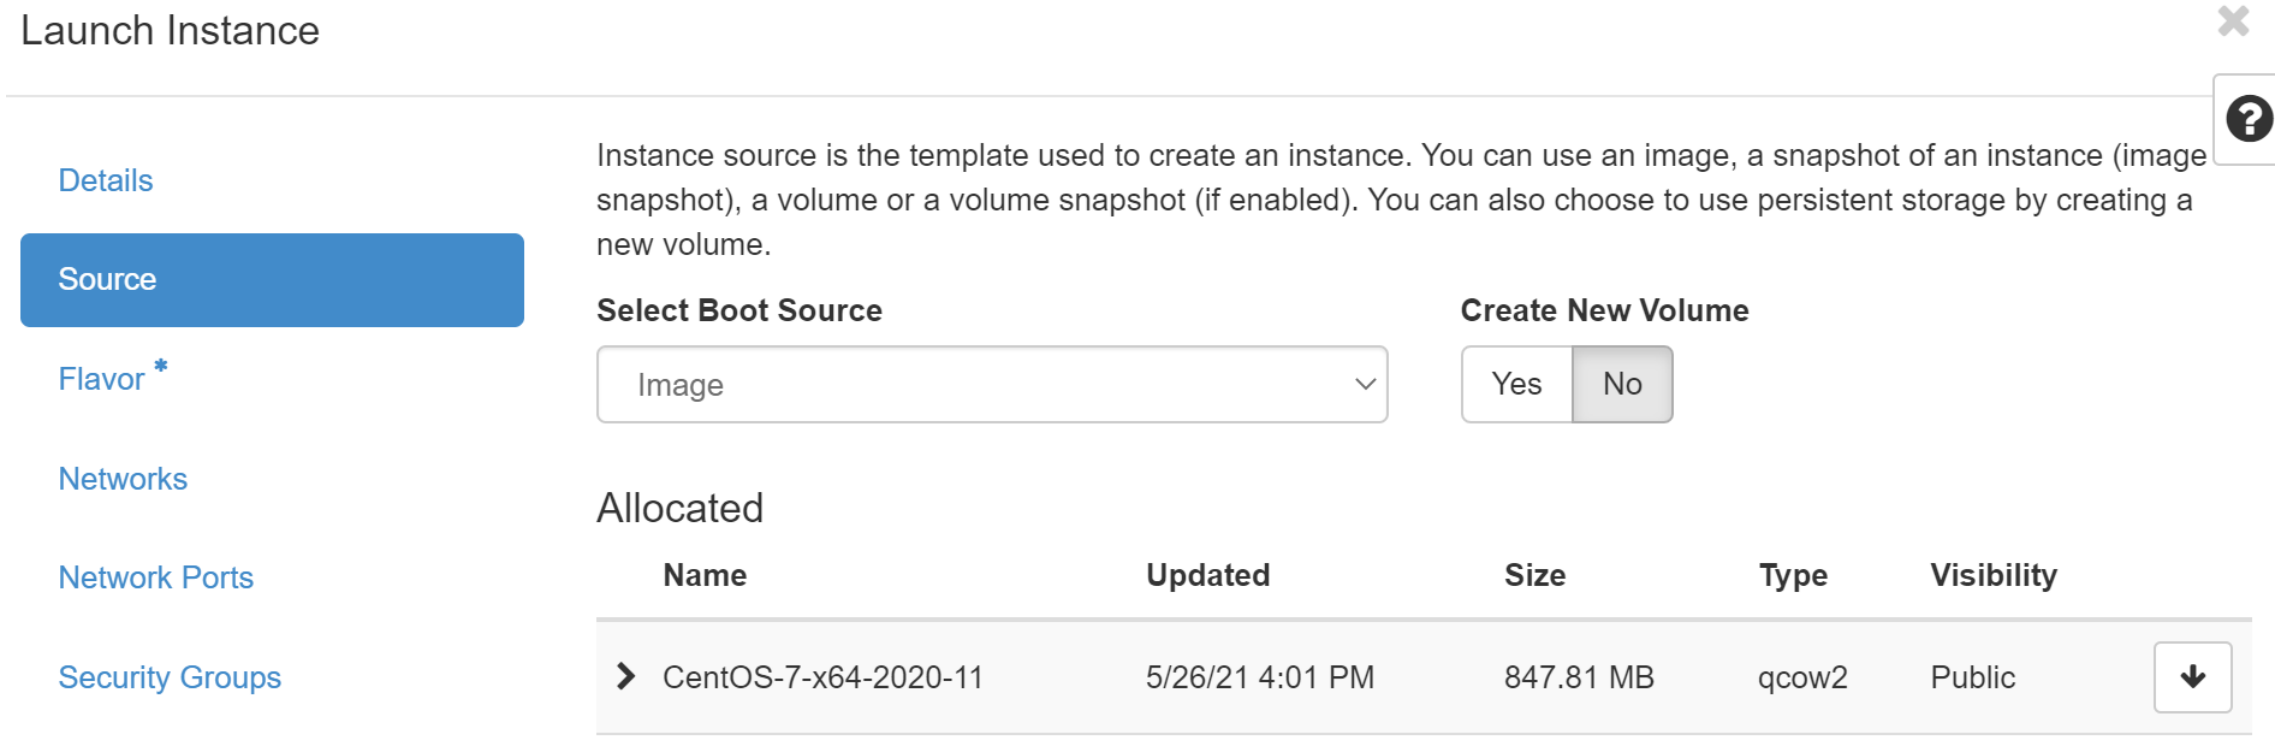

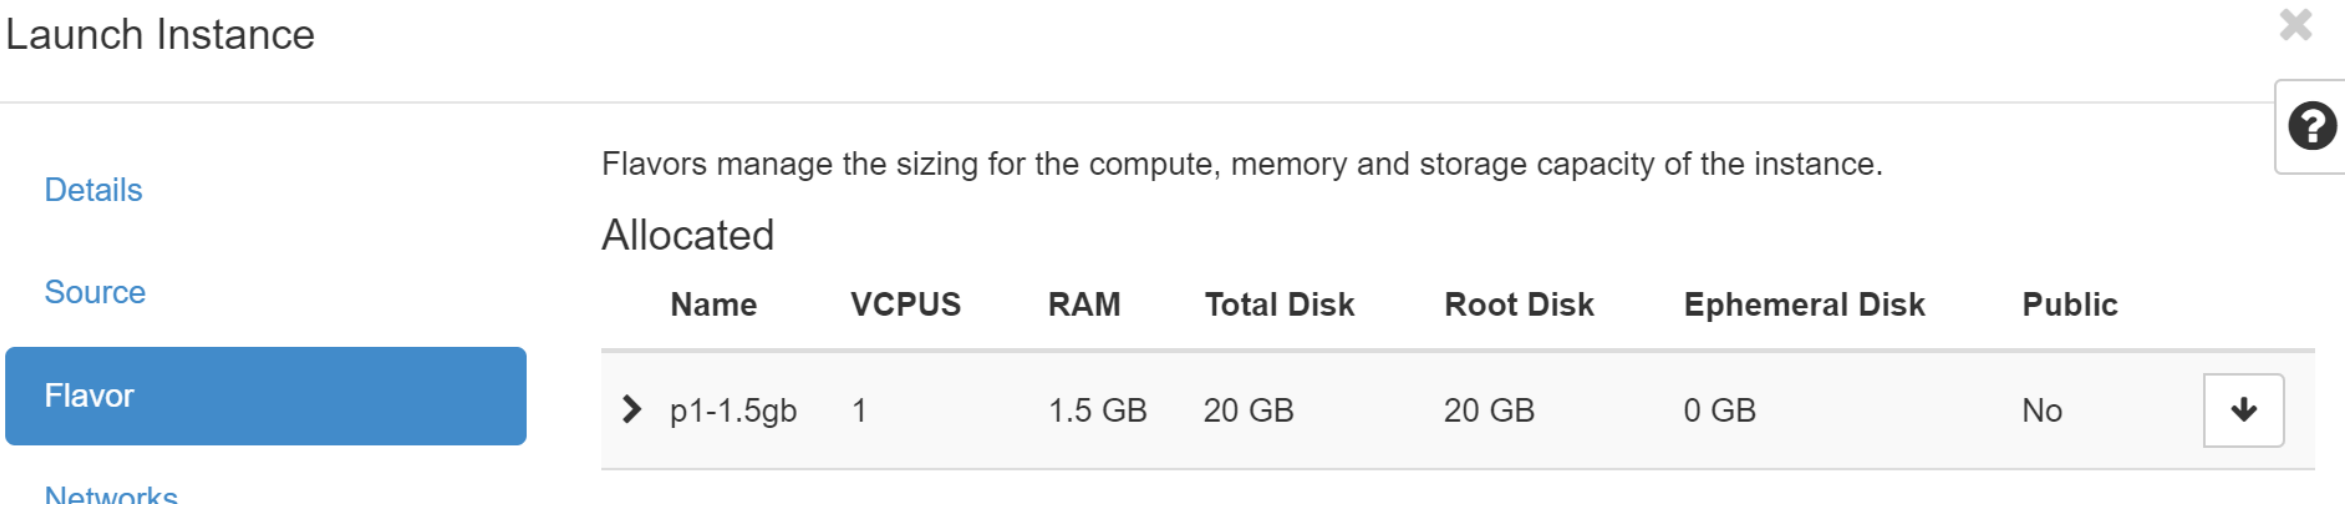

Digital Research Alliance of Canada

SFU SIMON FRASER UNIVERSITY

UBC Advanced Research Computing

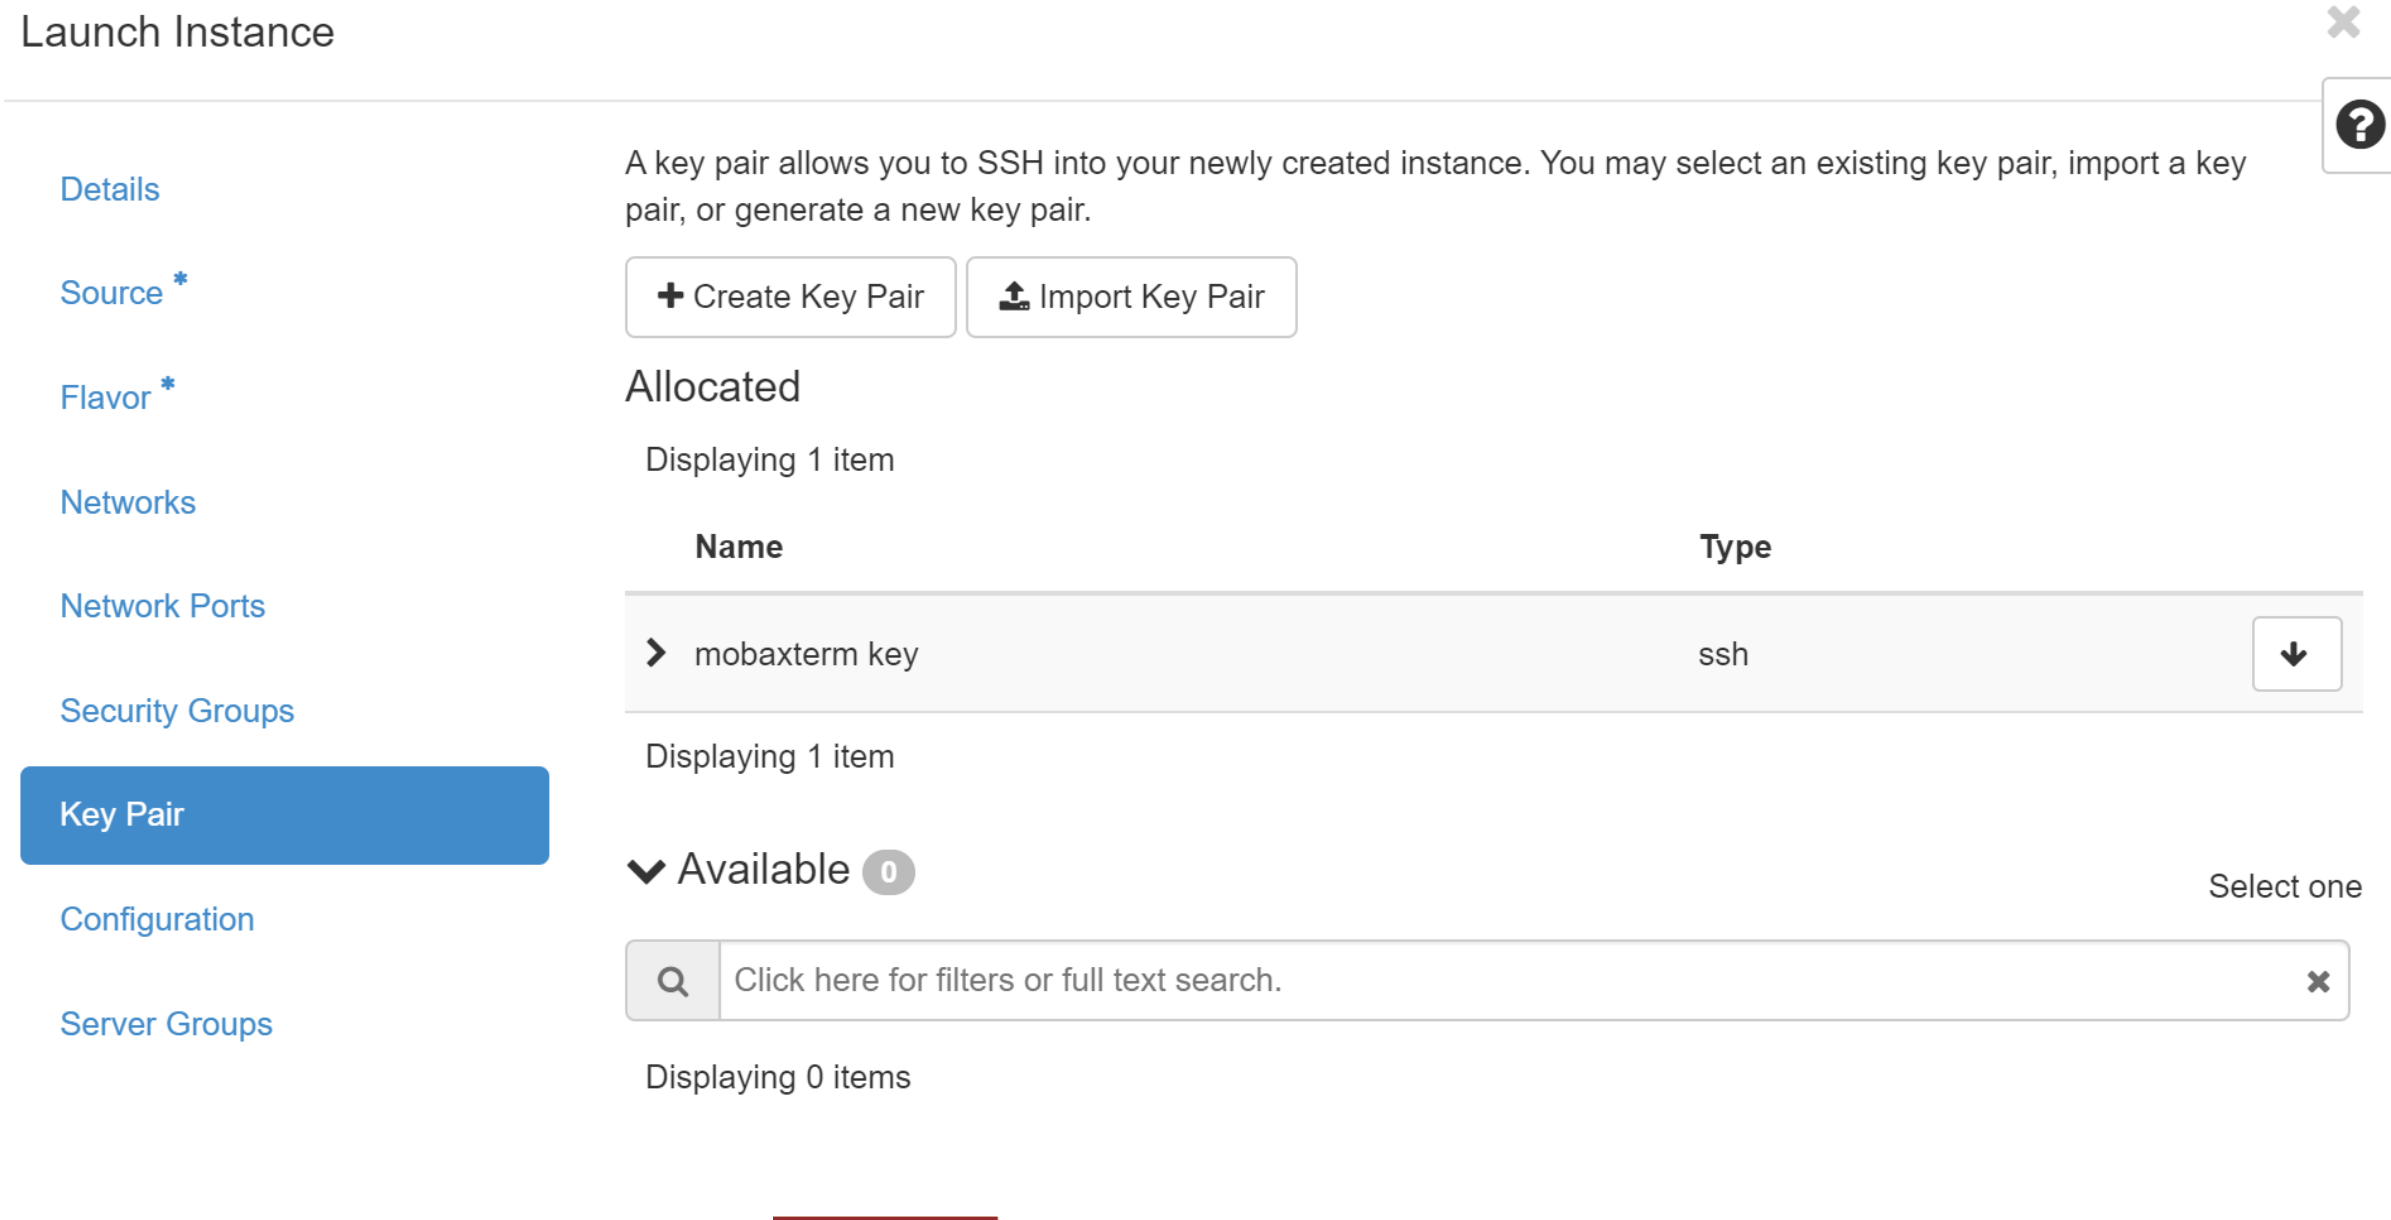

Digital Research Alliance of Canada

SFU SIMON FRASER UNIVERSITY

UBC Advanced Research Computing

Launch the Instance

- Click Launch Instance to launch the virtual machine. Make sure to note the name of your instance.
- OpenStack will boot the VM and insert the SSH key into it.
- Once the VM is booted, we can try to access it remotely.
- But need to configure security and public networking first.

Configuring Remote Access

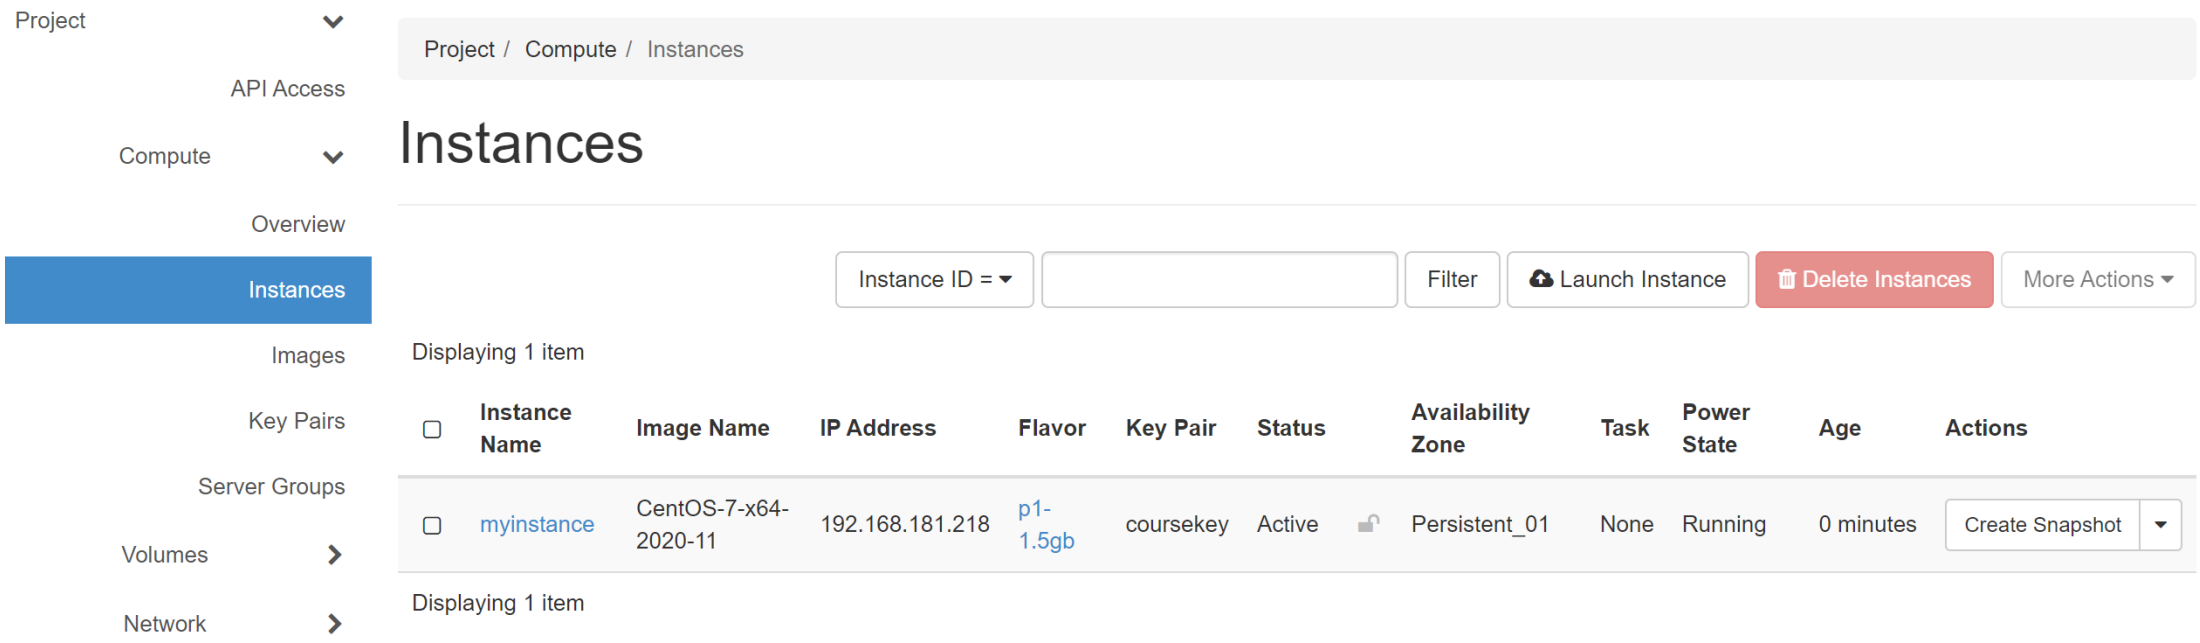

Digital Research Alliance of Canada

SFU SIMON FRASER UNIVERSITY

UBC Advanced Research Computing

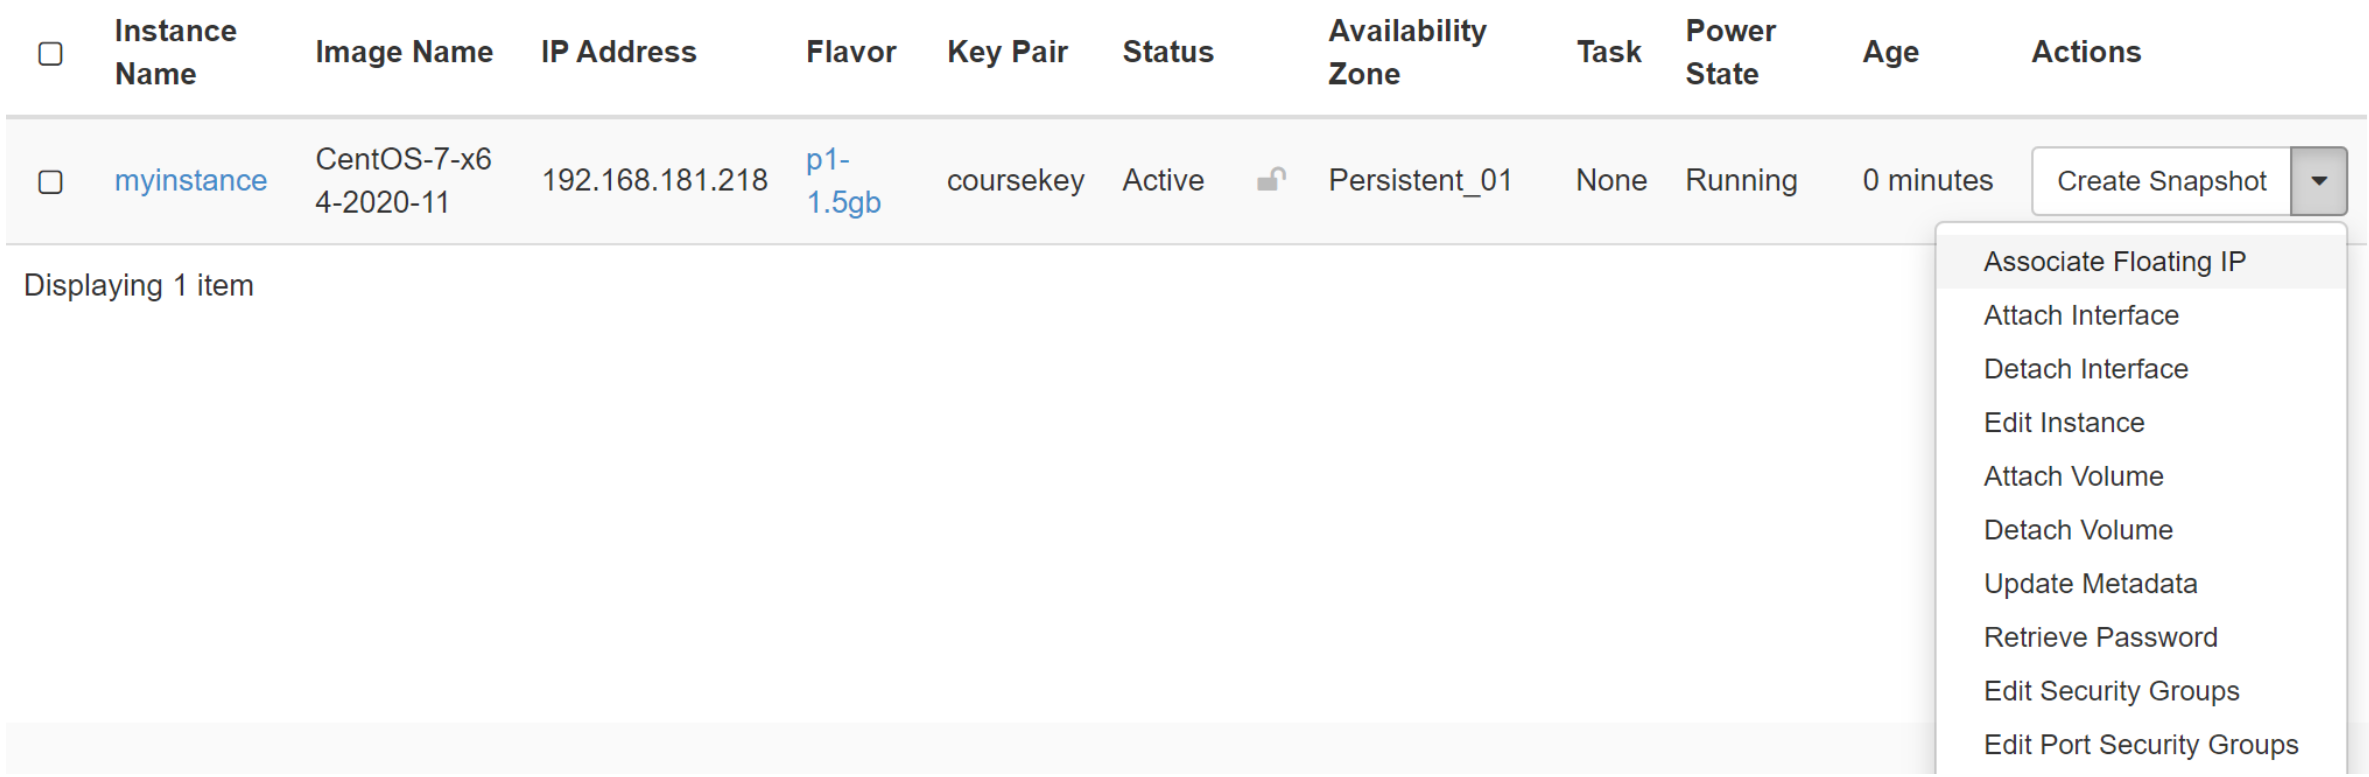
Instance Name
Image Name
IP Address
Flavor
Key Pair
Status
Availability Zone
Task
Power State
Age
Actions
myinstance
CentOS-7-x64-2020-11
192.168.181.218
p1-1.5gb
coursekey
Active
Persistent_01
None
Running
0 minutes
Create Snapshot
Displaying 1 item
Associate Floating IP
Attach Interface
Detach Interface
Edit Instance
Attach Volume
Detach Volume
Update Metadata
Retrieve Password
Edit Security Groups
Edit Port Security Groups

Manage Floating IP Associations

IP Address *

206.12.88.135

Port to be associated *

myinstance: 192.168.181.218

Select the IP address you wish to associate with the selected instance or port.

Cancel Associate

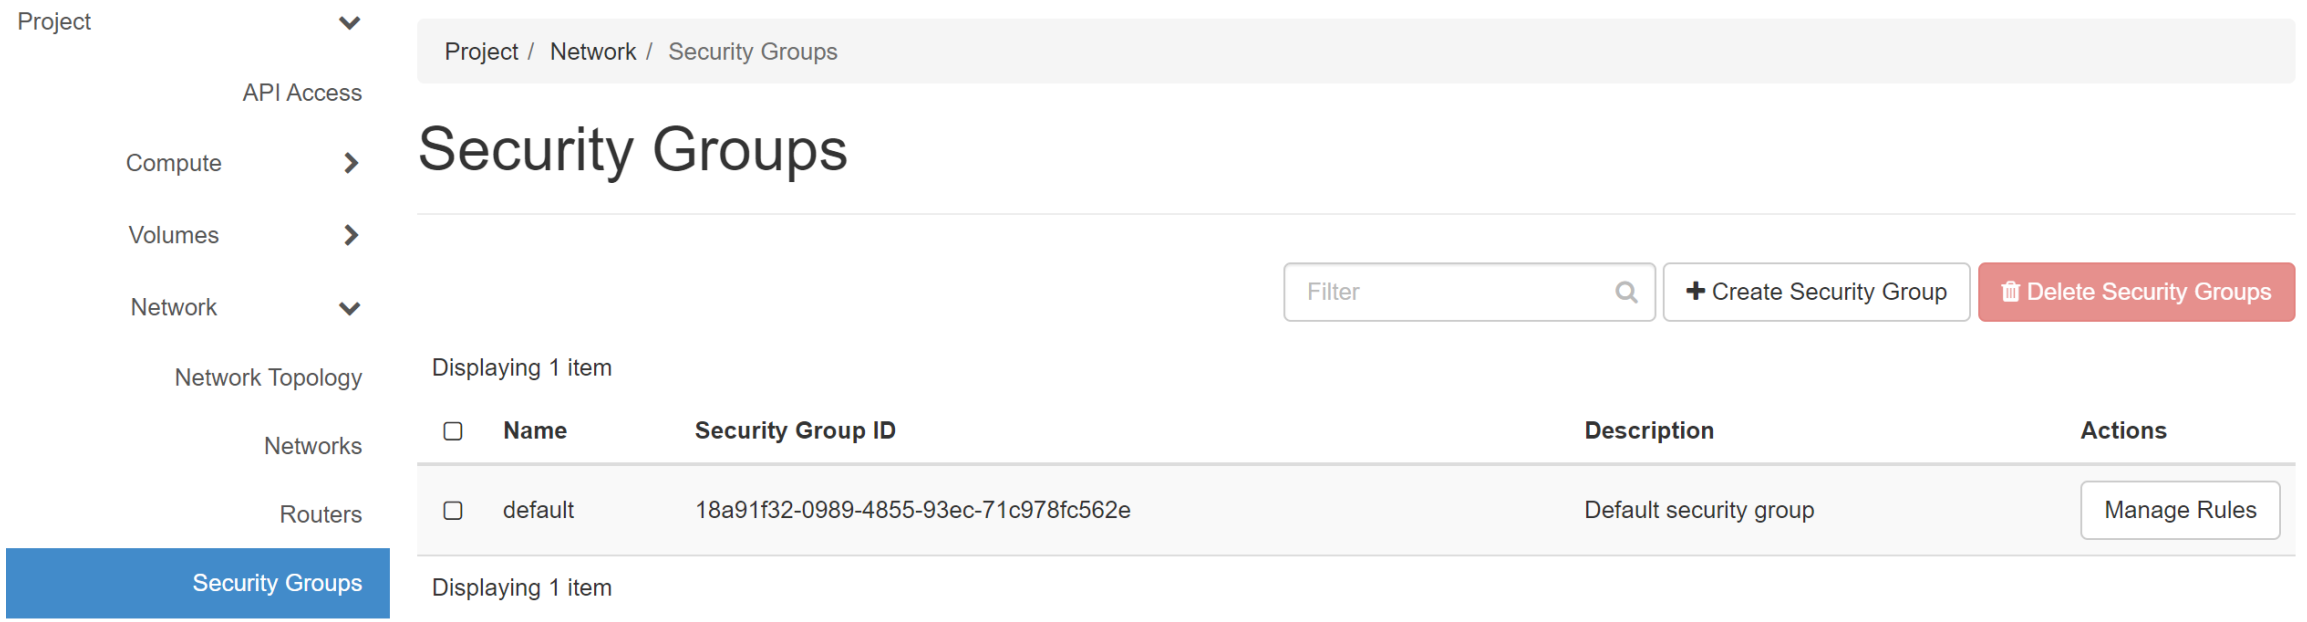

This can be ignored, it has already been pre-configured but is here for informational purposes

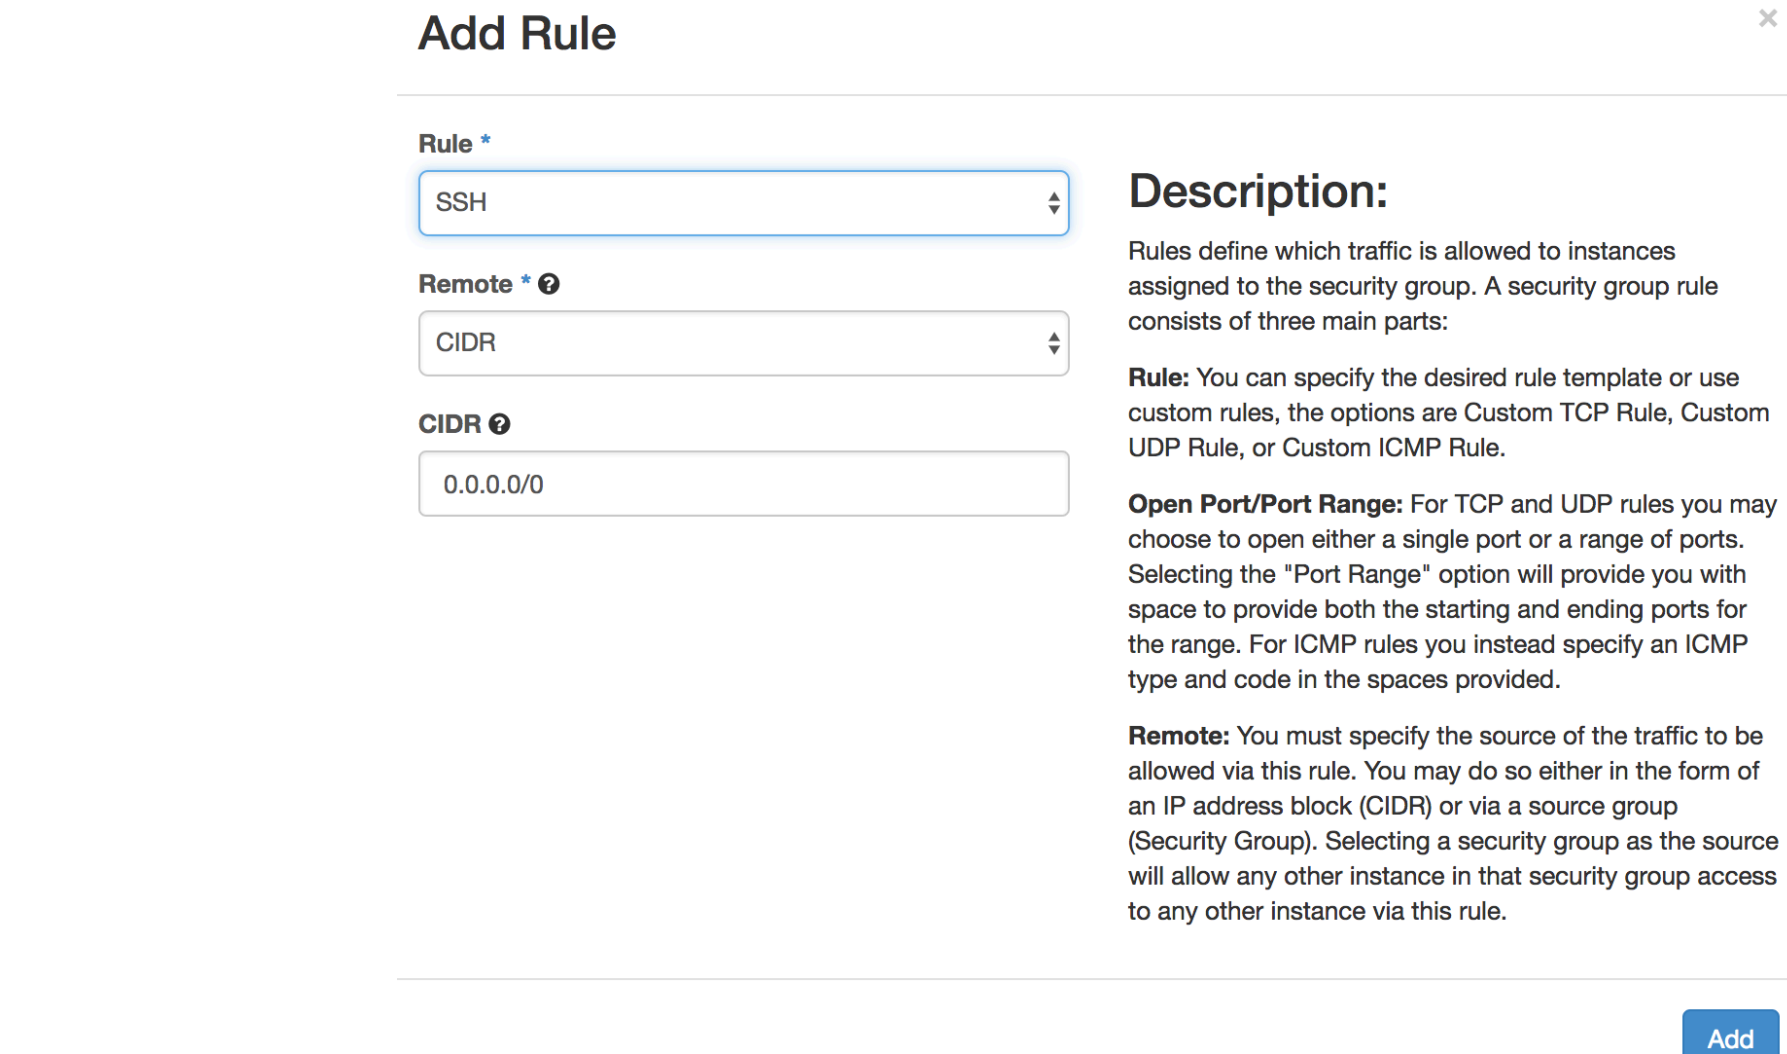

This can be ignored, it has already been pre-configured but is here for informational purposes

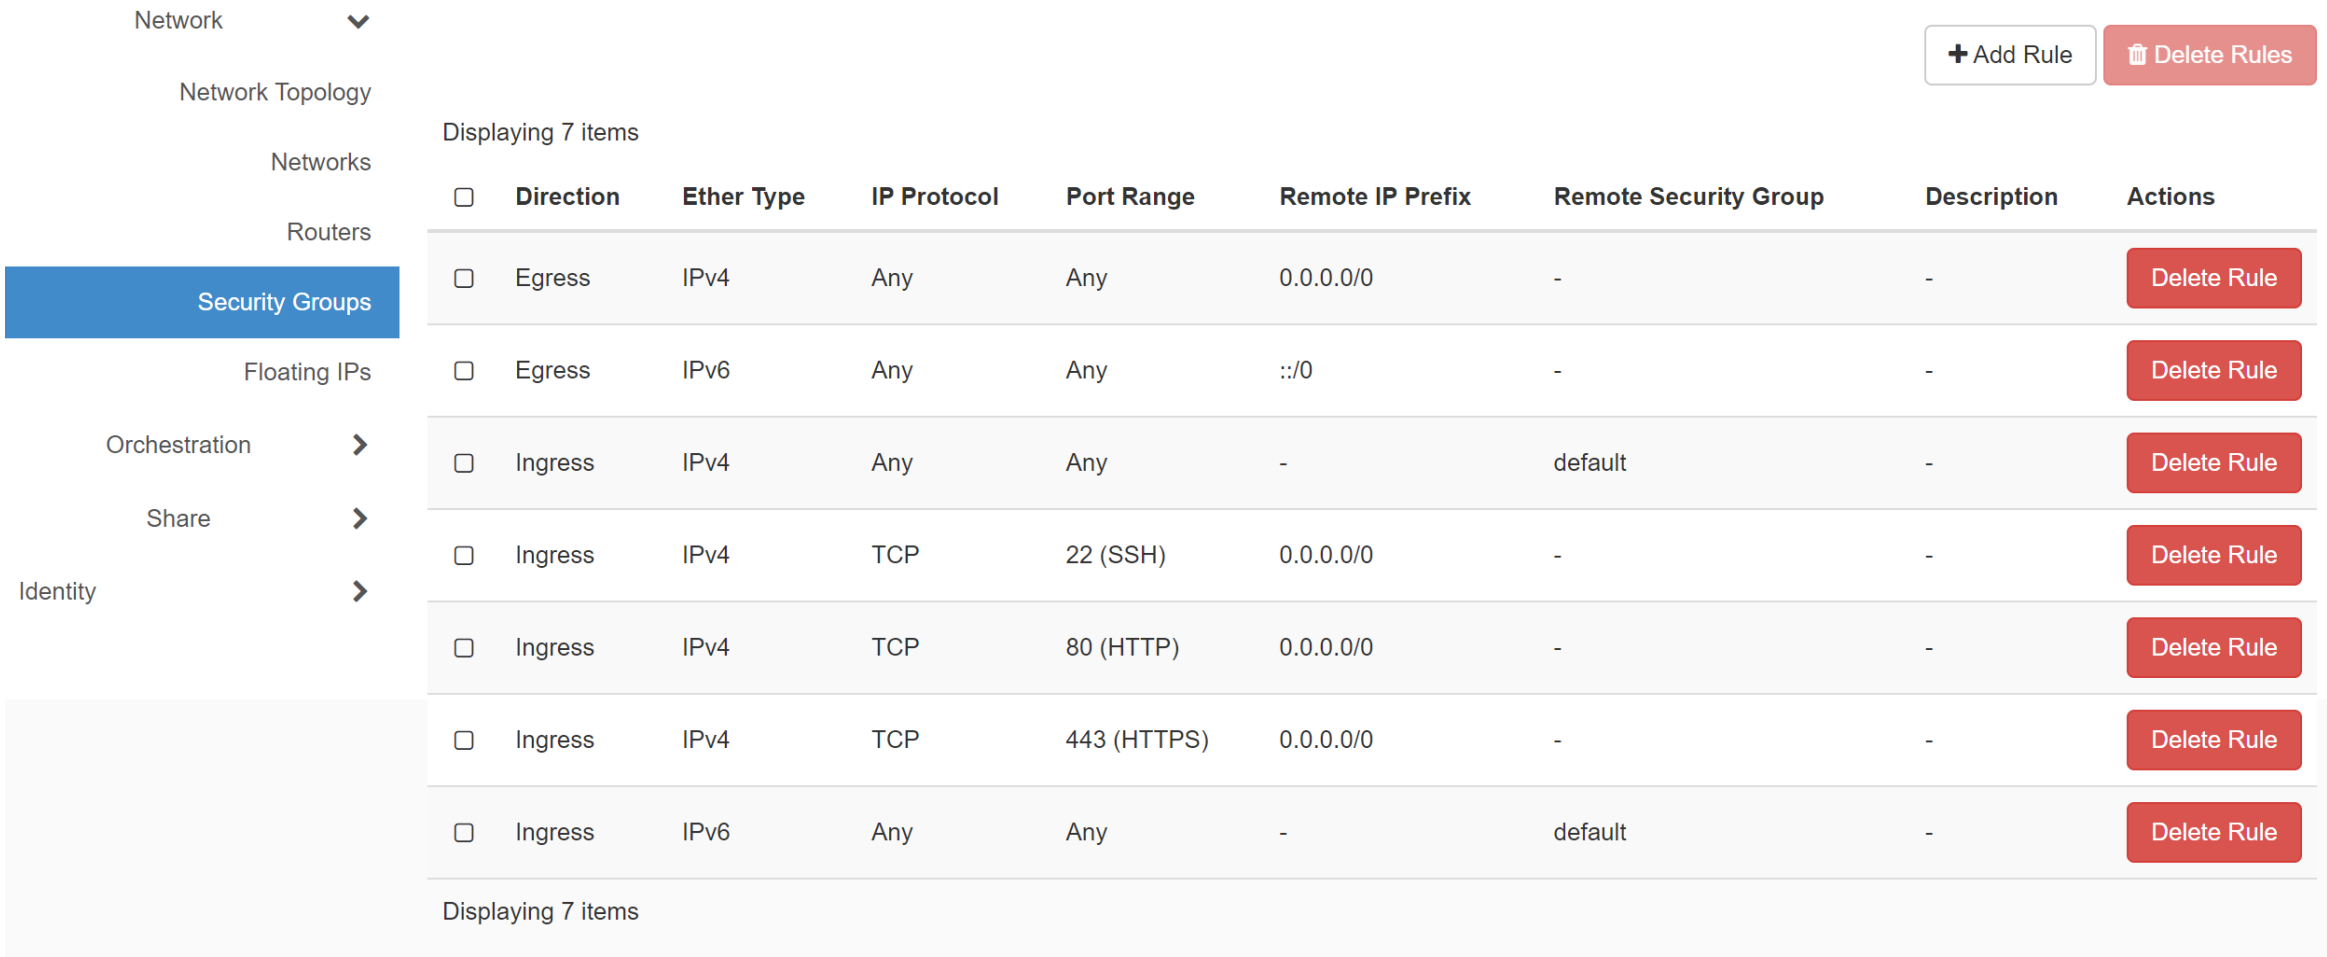

This can be ignored, it has already been pre-configured but is here for informational purposes

Connect to the Instance via SSH

```
ssh -i <key>.pem centos@<public ip>
```

If using MobaXTerm, see:

https://docs.alliancecan.ca/wiki/Connecting_with_MobaXTerm#Using_a_Key_Pair

If using Windows Subsystem for Linux, you may need to do:

```
chmod 600 <name of private key file>
```

Installing RStudio

```
sudo yum install epel-release -y
sudo yum install R -y
<< will take a while >>
```

```
sudo yum install wget -y
wget https://download2.rstudio.org/server/centos7/x86_64/rstudio-server-rhel-1.4.1717-x86_64.rpm
sudo yum install rstudio-server-rhel-1.4.1717-x86_64.rpm -y
```

```
sudo systemctl status rstudio-server.service
sudo systemctl enable rstudio-server.service
```

Add Security Rule

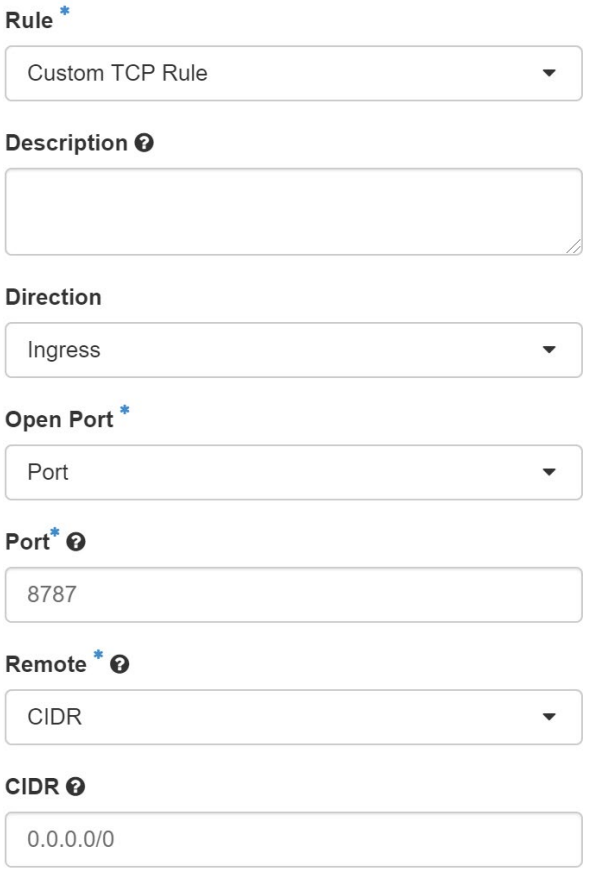

This can be ignored, it has already been pre-configured but is here for informational purposes

Add User

```
sudo useradd rstudiouser

sudo passwd rstudiouser
```

Done

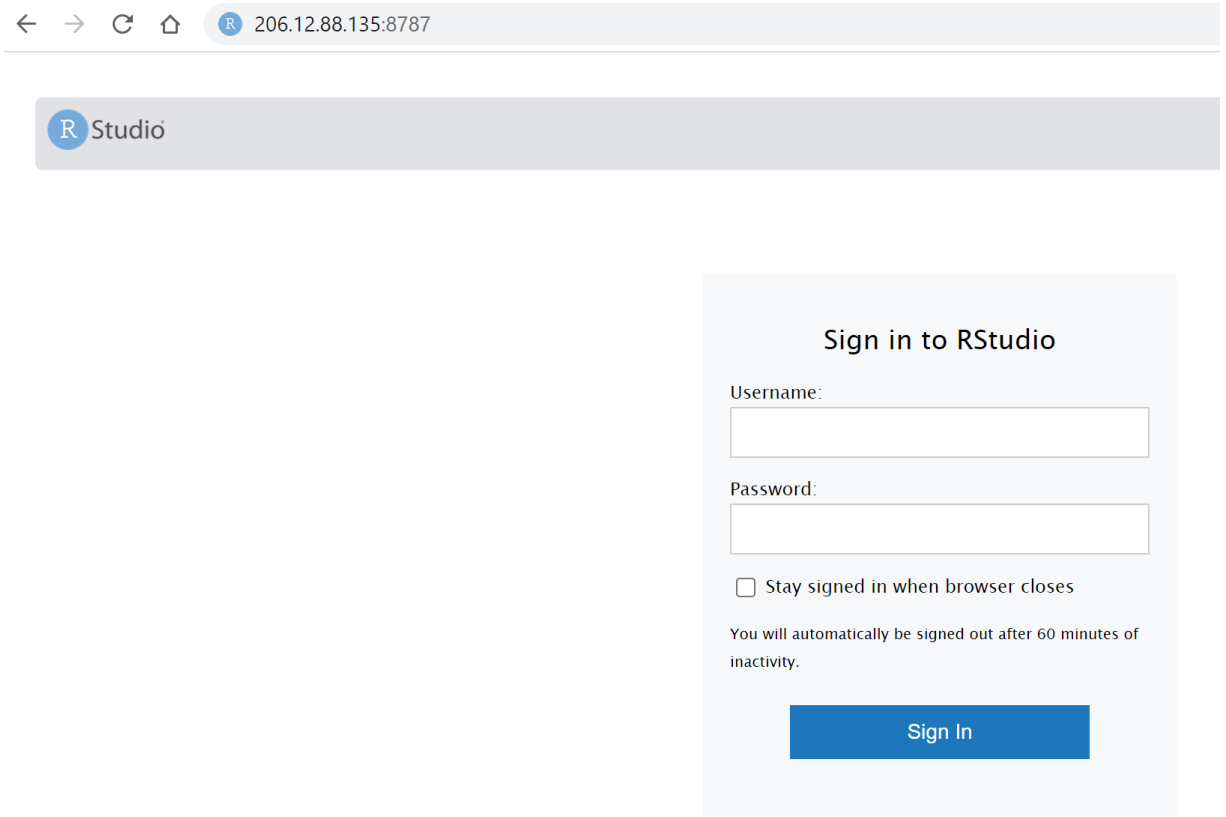
206.12.88.135:8787
RStudio
Sign in to RStudio
Username:
Password:
Stay signed in when browser closes
You will automatically be signed out after 60 minutes of inactivity.
Sign In

Maintaining Your Instance

- Install updates to the OS, e.g. for CentOS do “yum -y update”.
- Install application updates regularly for RStudio and other applications.

Resources

- The Alliance Cloud
  - https://docs.alliancecan.ca/wiki/Cloud_RAS_Allocations
  - https://docs.alliancecan.ca/wiki/Cloud_Quick_Start

- UBC Advanced Research Computing
  - https://www.arc.ubc.ca