

# Cloud Infrastructure Documentation

## Bare Metal Specs

The cluster comprise 8 nodes each of which has the following specs:

<b>Model</b>	Supermicro SYS-6018U-TRT+
<b>CPU</b>	2 x Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20GHz
<b>RAM</b>	16 x 16GB DDR4 @2133MHz
<b>Disks</b>	<ul style="list-style-type: none"><li>- 4 x Seagate ST3000NM0005-1V4 (3TB, 3.5" HDD)</li><li>- 1 x SuperMicro SSD (64GB, Disk On Module)</li><li>- Intel Corporation DC P3600 SSD (400GB NVMe)</li></ul>
<b>Network interfaces</b>	Intel Corporation Ethernet Controller 10-Gigabit X540-AT2

<b>DNS</b>	<b>Public IP (1G)</b>	<b>Private IP (10G)</b>
cloud1.cslab.ece.ntua.gr	147.102.4.97	10.0.0.1
cloud2.cslab.ece.ntua.gr	147.102.4.98	10.0.0.2
cloud3.cslab.ece.ntua.gr	147.102.4.99	10.0.0.3
cloud4.cslab.ece.ntua.gr	147.102.4.100	10.0.0.4
cloud5.cslab.ece.ntua.gr	147.102.4.101	10.0.0.5
cloud6.cslab.ece.ntua.gr	147.102.4.102	10.0.0.6
cloud7.cslab.ece.ntua.gr	147.102.4.103	10.0.0.7
cloud8.cslab.ece.ntua.gr	147.102.4.104	10.0.0.8

## Operating System

Each host runs **Ubuntu 16.04.3 LTS**, with Linux Kernel **4.4.0-97-generic x86\_64 GNU/Linux**

# Networking/VPN Setup

Each node accesses two networks:

1. One **publicly** accessible network (147.102.4.0/24) that utilizes 1G links
2. One **private** network (10.0.0.0/16) that facilitates VM connectivity and utilizes 10G links

By default, each allocated VM has access only to the **private** network. In order for a user to connect to their VMs, they must establish a VPN connection with the the cluster's VPN server that acts as an endpoint between the VMs and the outer world. The VPN is located at cloud-gw.cslab.ece.ntua.gr and it is a VM that accesses both (1) and (2) networks. The VPN server is configured to provide a *tun* interface, i.e., it acts as a router rather than a bridge (so as to achieve isolation and enhance security). cloud-gw runs **OpenVPN 2.3.10**.

For further networking details, see Neutron from the Openstack Setup section.

# Hypervisor Setup

The cluster runs **Qemu 2.10.0** with **KVM** support.

# Storage Setup

The cluster is configured to persist data to a **CEPH 12.2.0 (Luminous)** cluster. Below you can find a brief summary of CEPH's setup.

The cluster comprises:

1. 8 x OSD nodes (cloud1-cloud8)
2. 1 x monitoring node (cloud1)
3. 1 x manager node (cloud1)

Each OSD node, i.e., actual data holder, utilizes a RAID5 array for persisting its data. Each RAID5 array comprises 4 x 3TB disks, formatted to XFS (according to CEPH's instructions).

There exist 2 storage pools:

- vms
- images

Each storage pools is used to store a specific type of data. The *vms* pool persists the root devices of the VMs deployed to the Openstack cluster. The *images* pools, is utilized by Openstack image service in order to persist the cluster's images. The data transmission utilizes both network interfaces between different OSDs: The *public* network (i.e., 1G) is utilized for writing the master data replica, heartbeat between the OSDs and the monitor, etc. The *private* network (i.e., 10G) - also called cluster network in CEPH's language - is utilized for replicas transfer. More information regarding CEPH's network architecture can be found [here](#).

# Openstack Setup

The cluster runs the latest **Openstack Pike version**. Below you can find a description for each Openstack module configuration:

Role	Hosts	Services
controller	cloud1.cslab.ece.ntua.gr	nova-api nova-conductor nova-novncproxy nova-consoleauth nova-scheduler  glance-api glance-registry  neutron-dhcp-agent neutron-server neutron-linuxbridge-agent neutron-metadata-agent  rabbitmq mariadb
compute	cloud1.cslab.ece.ntua.gr	nova-compute  neutron-linuxbridge-agent
	cloud2.cslab.ece.ntua.gr	
	cloud3.cslab.ece.ntua.gr	
	cloud4.cslab.ece.ntua.gr	
	cloud5.cslab.ece.ntua.gr	
	cloud6.cslab.ece.ntua.gr	
	cloud7.cslab.ece.ntua.gr	
	cloud8.cslab.ece.ntua.gr	

## Nova

Nova is Openstack's compute service, which is responsible to run the VMs, schedule the VM placement, etc. The cluster nodes are differentiated to two roles: the *controller* node and the *compute* nodes. The above table describes the services that are executed to each cluster node.

The version utilized is **Nova 16.0.0**. The compute nodes are configured to use **libvirt 3.6.0**, that, in turn, is configured to spawn VMs using qemu-kvm as described above. Moreover, nova is configured to persist the VM's root devices to the *vms* storage pool.

## Glance

Glance is Openstack's image service, used to manage the original images, VM snapshots, etc. **Glance 15.0.0** is utilized. It only runs in the controller node and it persists its data in the *images* storage pool. Moreover, according to CEPH-Openstack integration best practices, only raw images are uploaded (not qcow or other zipped formats).

## Neutron

Neutron is Openstack's networking service. **Neutron 11.0.1** is utilized. Neutron provide two configuration options: the [provider](#) and the [self-service](#) networks. The provider network utilizes the hosts' network (bridging) whereas the self-service network enables the dynamic creation of new networks (routing). In our case, we chose the provider network. Specifically, neutron is configured to bridge the *private* host interface to the VM network and, hence, all the VMs belong to the same network. Any time a VM is launched, neutron's DHCP service assigns an IP to it, based on the available DHCP pool. In order to make the *private* network publicly accessible, we have configured cloud1 (10.0.0.1) to be the network's gateway and have enabled packet forwarding and NAT.

Name	Pool	Gateway	Usage
provider	10.0.1.0 - 10.0.10.255	10.0.0.1	Default VM network
external	147.102.4.105 - 147.102.4.110	147.102.4.200	Used by admin VMs/ VMs that need to be publicly accessible

## Other Services

1. **RabbitMQ 3.5.7**: Used by different Openstack modules for communication
2. **MariaDB 10.0.31**: Used by different Openstack modules to persist their data and metadata.

Both for security and performance reasons, all internal Openstack services listen to the private host interfaces.