

Cloud Service

Book with links and materials with Cloud user manual, tutorials and training.
Includes Virtualization, security considerations for VMs, and resource accounting

- [Cloud Tutorials](#)
- [Security recommendations for cloud](#)
 - [Security recommendations for the Cloud](#)
- [How to check resource accounting](#)
- [Time synchronization](#)

Cloud Tutorials

Openstack Tutorials

An Openstack Horizon/Dashboard tutorial can be found [in these slides](#)

The Openstack CLI tutorial is being maintained in a github repository [tutorial-oscli](#)

Older versions of the Openstack Tutorials:

- Openstack Command Line Tools [materials](#)
- Openstack Graphical User Interface [materials](#)
- Openstack EC2 API [materials](#)

Security recommendations for cloud

Security recommendations for cloud users, in particular when instantiating and managing VMs

Security recommendations for the Cloud

A non exhaustive list of security recommendations

These recommendations target specially user and administrators of cloud computing resources (VMs), as well as interaction through the Openstack dashboard and CLI.

- default security groups, leave only ssh (22) open and nothing else, other rules should be to create another "Create Security Group". Each security group should be as minimal as possible. When too many ports are added to the Default, like the 80 and 443 these are immediately available on the VM, and soon that the service comes up it can be immediately attacked.
- The first thing when instantiating a VM is to immediately update / upgrade the operating system and reboot to get the latest kernel, before installing other SW. In fact, update / upgrade regularly all VMs.
- Install fail2ban (see below)
- When testing or developing a given service such as wordpress, it's better to have only one private IP and make an ssh tunnel from another VM with a public IP, until the service is properly configured and secured.
- Deploy the service endpoints with https/certificate, if possible only port 443 should be opened in the security groups.
- Only associate the new security group (for example 443 and / or 80) after the service is configured properly. Avoiding windows of opportunity while developpeing/testing the service.
- If necessary, restrict the range of IPs that can be connected from outside.

Updates, installation and configuration of fail2ban

On the following use either `yum` for Cento7 or `dnf` for Centos8, for Ubuntu also available, update and reboot the VM:

```
sudo -s
dnf -y update
dnf -y install epel-release
shutdown -r now
```

Reboot the VM

```
dnf -y install fail2ban
```

The fail2ban configuration files are located in the `/etc/fail2ban/` directory and filters are stored in the `/etc/fail2ban/filter.d/` directory (the filter file for sshd is `/etc/fail2ban/filter.d/sshd.conf`).

The global configuration file for the fail2ban server is `/etc/fail2ban/jail.conf`, however, it is not recommended to modify this file directly, as it will probably be overwritten or improved in case of a package upgrade in the future.

As an alternative, it is recommended to create and add your configurations in a `jail.local` file or separate `.conf` files under the `/etc/fail2ban/jail.d/` directory. Note that configuration parameters set in `jail.local` will override whatever is defined in `jail.conf`.

For this article, we will create a separate file called `jail.local` in the `/etc/fail2ban/` directory as shown.

```
vi /etc/fail2ban/jail.local
```

Once the file is open, copy and paste the following configuration in it. The `[DEFAULT]` the section contains global options and `[sshd]` contains parameters for the sshd jail.

```
[DEFAULT]
ignoreip = 192.168.56.2/24
bantime = 21600
findtime = 300
maxretry = 3
banaction = iptables-multiport
backend = systemd

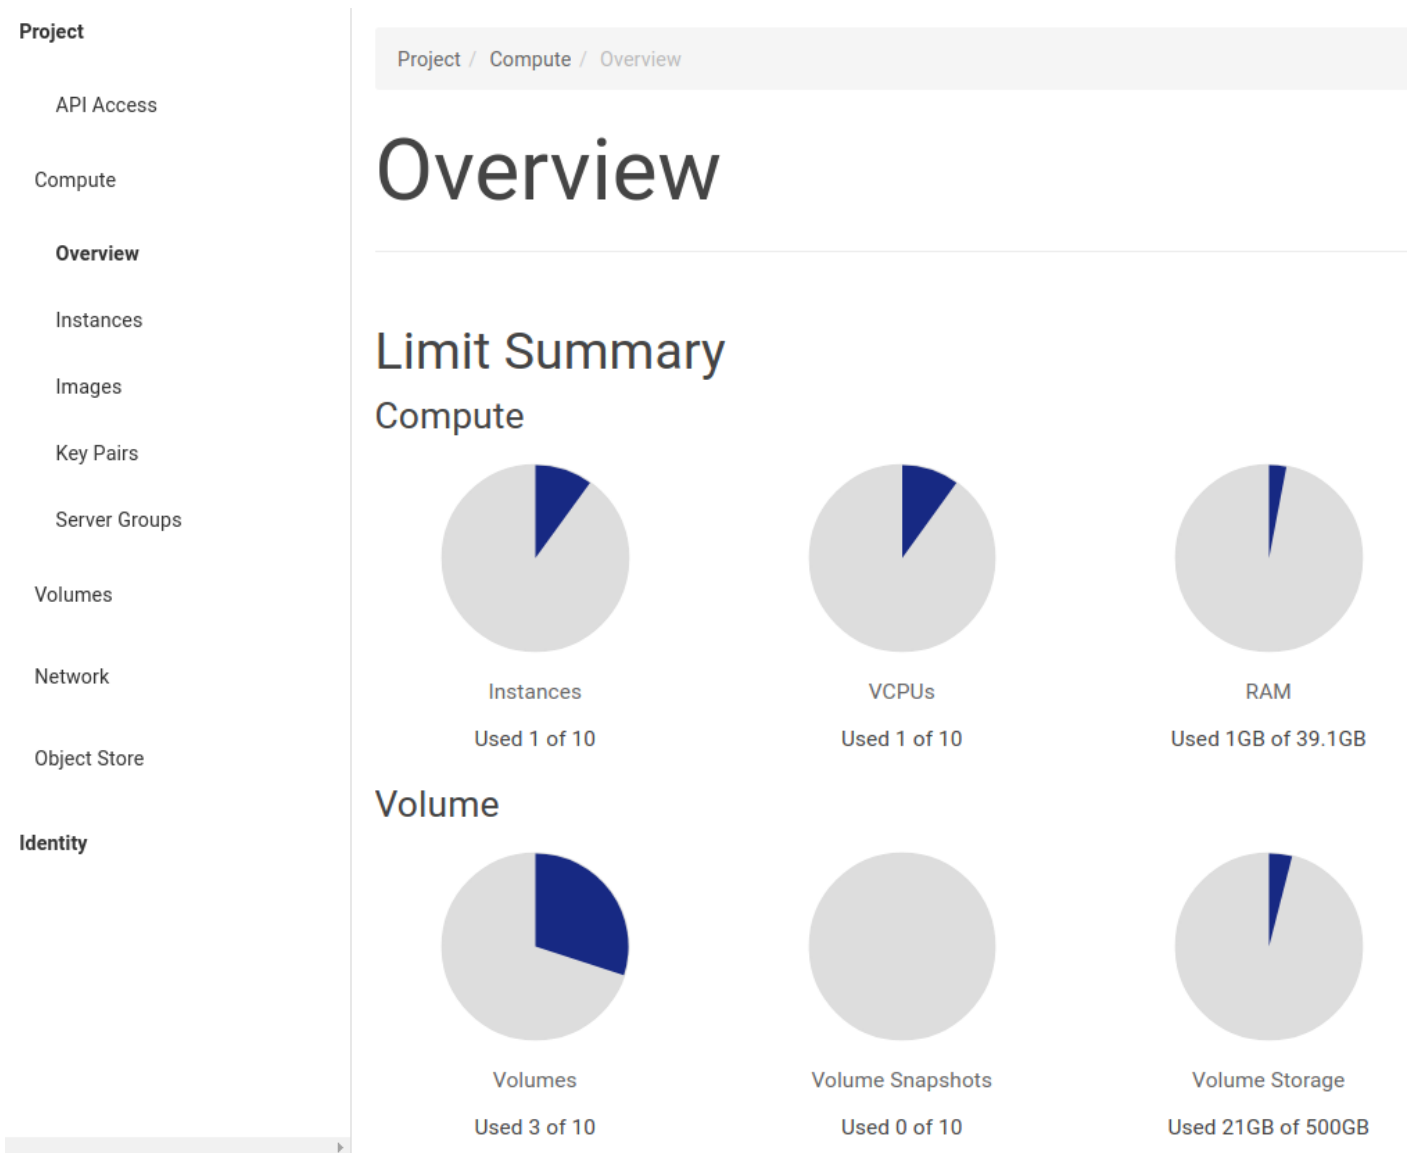
[sshd]
enabled = true
```

Enable and start the service:

```
systemctl start fail2ban  
systemctl enable fail2ban  
systemctl status fail2ban
```

How to check resource accounting

In the Openstack dashboard, select "Overview" on the left panel:



Scroll down to see the "Usage Summary" and select the initial date, for example January 1st 2021, and end date (the default is the current day):

Project

- API Access
- Compute
- Overview
- Instances
- Images
- Key Pairs
- Server Groups
- Volumes
- Network
- Object Store
- Identity

Security Groups

Used 1 of 10

Security Group Rules

Used 5 of 100

Networks

Used 1 of 100

Ports

Used 5 of 500

Select a period of time to query its usage:
The date should be in YYYY-MM-DD format.

2021-01-18 → 2021-01-19 SUBMIT

Active Instances: 1
Active RAM: 1GB
This Period's VCPU-Hours: 35.23
This Period's GB-Hours: 352.32
This Period's RAM-Hours: 36077.13

Usage

Displaying 1 item

Instance Name	VCPU	Disk	RAM	Age
UBuntu	1	10GB	1GB	3 days

Displaying 1 item

After you click "submit" you will see the summary table with resource usage, shown below:

Usage Summary

Select a period of time to query its usage:

The date should be in YYYY-MM-DD format.

2021-01-01 → 2021-01-19 SUBMIT

Active Instances: 1

Active RAM: 1GB

This Period's VCPU-Hours: 187.00

This Period's GB-Hours: 1869.95

This Period's RAM-Hours: 191483.23

Time synchronization

Recommendations for NTP configuration

For stratus.ncg.ingrid.pt use the INCD provided NTP servers:

```
server ntp01.ncg.ingrid.pt  
server ntp02.ncg.ingrid.pt
```

To use other NTP servers configure them as "peer" not "server" like:

```
peer ntp02.oal.ul.pt  
peer ntp04.oal.ul.pt  
peer ntp01.fccn.pt  
peer ntp02.fccn.pt
```