

Welcome to Avitohol Cloud User Guide



- Avitohol Cloud Location and Login
- Configure access and security for instances
 - Add a rule to the default security group
 - Add a key pair
 - Import a key pair
 - Allocate a floating IP address to an instance
- Launch and manage instances
 - Launch an instance
 - Connect to your instance by using SSH
 - Track usage for instances
 - Create an instance snapshot
 - Manage an instance
- Create and manage networks
 - Create a network
 - Create a router
 - Create a port

Avitohol Cloud Location and Login

Before you begin you need to obtain access for Avitohol Cloud Console. Once your access is granted navigate to the [Avitohol Cloud Console](#). Then follow the sections of this document based on your use case.

Configure access and security for instances

Before you launch an instance, you should add security group rules to enable users to ping and use SSH to connect to the instance. Security groups are sets of IP filter rules that define networking access and are applied to all instances within a project. To do so, you either add rules to the default security group Add a rule to the default security group or add a new security group with rules.

Key pairs are SSH credentials that are injected into an instance when it is launched. To use key pair injection, the image that the instance is based on must contain the cloud-init package. Each project should have at least one key pair. For more information, see the section Add a key pair.

If you have generated a key pair with an external tool, you can import it into OpenStack. The key pair can be used for multiple instances that belong to a project. For more information, see the section Import a keypair.

✔ Note

A key pair belongs to an individual user, not to a project. To share a key pair across multiple users, each user needs to import that key pair.

When an instance is created in OpenStack, it is automatically assigned a fixed IP address in the network to which the instance is assigned. This IP address is permanently associated with the instance until the instance is terminated. However, in addition to the fixed IP address, a floating IP address can also be attached to an instance. Unlike fixed IP addresses, floating IP addresses are able to have their associations modified at any time, regardless of the state of the instances involved.

Add a rule to the default security group

This procedure enables SSH and ICMP (ping) access to instances. The rules apply to all instances within a given project, and should be set for every project unless there is a reason to prohibit SSH or ICMP access to the instances.

This procedure can be adjusted as necessary to add additional security group rules to a project, if your cloud requires them.

✔ Note

When adding a rule, you must specify the protocol used with the destination port or source port.

1. Log in to the dashboard.
2. Select the appropriate project from the drop down menu at the top left.
3. On the [Project](#) tab, open the [Network](#) tab. The [Security Groups](#) tab shows the security groups that are available for this project.
4. Select the default security group and click [Manage Rules](#).
5. To allow SSH access, click [Add Rule](#).
6. In the [Add Rule](#) dialog box, enter the following values:
 - **Rule:** SSH
 - **Remote:** CIDR
 - **CIDR:** 0.0.0.0/0

✔ Note

To accept requests from a particular range of IP addresses, specify the IP address block in the [CIDR](#) box.

7. Click [Add](#).
8. To add an ICMP rule, click [Add Rule](#).
9. In the [Add Rule](#) dialog box, enter the following values:
 - **Rule:** All ICMP
 - **Direction:** Ingress
 - **Remote:** CIDR
 - **CIDR:** 0.0.0.0/0
10. Click [Add](#).

Instances will now accept all incoming ICMP packets.

Add a key pair

Create at least one key pair for each project.

1. Log in to the dashboard.
2. Select the appropriate project from the drop down menu at the top left.
3. On the [Project](#) tab, open the [Compute](#) tab.
4. Click the [Key Pairs](#) tab, which shows the key pairs that are available for this project.
5. Click [Create Key Pair](#).
6. In the [Create Key Pair](#) dialog box, enter a name for your key pair, and click [Create Key Pair](#).
7. The private key will be downloaded automatically.
8. To change its permissions so that only you can read and write to the file, run the following command:

```
$ chmod 0600 yourPrivateKey.pem
```

✔ Note

If you are using the Dashboard from a Windows computer, use PuTTYgen to load the *.pem file and convert and save it as *.ppk. For more information see the [WinSCP web page for PuTTYgen](#).

9. To make the key pair known to SSH, run the **ssh-add** command.

```
$ ssh-add yourPrivateKey.pem
```

Import a key pair

1. Log in to the dashboard.
2. Select the appropriate project from the drop down menu at the top left.
3. On the [Project](#) tab, open the [Compute](#) tab.

4. Click the [Key Pairs](#) tab, which shows the key pairs that are available for this project.
5. Click [Import Key Pair](#).
6. In the [Import Key Pair](#) dialog box, enter the name of your key pair, copy the public key into the [Public Key](#) box, and then click [Import Key Pair](#).

The Compute dashboard registers the public key of the key pair.

The Dashboard lists the key pair on the [Key Pairs](#) tab.

Allocate a floating IP address to an instance

When an instance is created in OpenStack, it is automatically assigned a fixed IP address in the network to which the instance is assigned. This IP address is permanently associated with the instance until the instance is terminated.

However, in addition to the fixed IP address, a floating IP address can also be attached to an instance. Unlike fixed IP addresses, floating IP addresses can have their associations modified at any time, regardless of the state of the instances involved. This procedure details the reservation of a floating IP address from an existing pool of addresses and the association of that address with a specific instance.

1. Log in to the dashboard.
2. Select the appropriate project from the drop down menu at the top left.
3. On the [Project](#) tab, open the [Network](#) tab.
4. Click the [Floating IPs](#) tab, which shows the floating IP addresses allocated to instances.
5. Click [Allocate IP To Project](#).
6. Choose the pool from which to pick the IP address.
7. Click [Allocate IP](#).
8. In the [Floating IPs](#) list, click [Associate](#).
9. In the [Manage Floating IP Associations](#) dialog box, choose the following options:
 - The [IP Address](#) field is filled automatically, but you can add a new IP address by clicking the [+](#) button.
 - In the [Port to be associated](#) field, select a port from the list.

The list shows all the instances with their fixed IP addresses.

10. Click [Associate](#).

✔ Note

To disassociate an IP address from an instance, click the [Disassociate](#) button.

To release the floating IP address back into the floating IP pool, click the [Release Floating IP](#) option in the [Actions](#) column.

Launch and manage instances

Instances are virtual machines that run inside the cloud. You can launch an instance from the following sources:

- Images uploaded to the Image service.
- Image that you have copied to a persistent volume. The instance launches from the volume, which is provided by the cinder-volume API through iSCSI.
- Instance snapshot that you took.

Launch an instance

1. Log in to the dashboard.
2. Select the appropriate project from the drop down menu at the top left.
3. On the [Project](#) tab, open the [Compute](#) tab and click [Instances](#) category.

The dashboard shows the instances with its name, its private and floating IP addresses, size, status, task, power state, and so on.

4. Click [Launch Instance](#).
5. In the [Launch Instance](#) dialog box, specify the following values:

[Details](#) tab

Instance Name

Assign a name to the virtual machine.

Availability Zone

By default, this value is set to the availability zone given by the cloud provider (for example, us-west or apac-south). For some cases, it could be nova.

✔ Note

The name you assign here becomes the initial host name of the server. If the name is longer than 63 characters, the Compute service truncates it automatically to ensure dnsmasq works correctly.

After the server is built, if you change the server name in the API or change the host name directly, the names are not updated in the dashboard.

Server names are not guaranteed to be unique when created so you could have two instances with the same host name.

Count

To launch multiple instances, enter a value greater than 1. The default is 1.

Source tab

Instance Boot Source

Your options are:

Boot from image

If you choose this option, a new field for **Image Name** displays. You can select the image from the list.

Boot from snapshot

If you choose this option, a new field for **Instance Snapshot** displays. You can select the snapshot from the list.

Boot from volume

If you choose this option, a new field for **Volume** displays. You can select the volume from the list.

Boot from image (creates a new volume)

With this option, you can boot from an image and create a volume by entering the **Device Size** and **Device Name** for your volume. Click the **Delete Volume on Instance Delete** option to delete the volume on deleting the instance.

Boot from volume snapshot (creates a new volume)

Using this option, you can boot from a volume snapshot and create a new volume by choosing [Volume Snapshot](#) from a list and adding a [Device Name](#) for your volume. Click the [Delete Volume on Instance Delete](#) option to delete the volume on deleting the instance.

Image Name

This field changes based on your previous selection. If you have chosen to launch an instance using an image, the [Image Name](#) field displays. Select the image name from the dropdown list.

Instance Snapshot

This field changes based on your previous selection. If you have chosen to launch an instance using a snapshot, the [Instance Snapshot](#) field displays. Select the snapshot name from the dropdown list.

Volume

This field changes based on your previous selection. If you have chosen to launch an instance using a volume, the [Volume](#) field displays. Select the volume name from the dropdown list. If you want to delete the volume on instance delete, check the [Delete Volume on Instance Delete](#) option.

[Flavor](#) tab

Flavor

Specify the size of the instance to launch.

✔ Note

The flavor is selected based on the size of the image selected for launching an instance. For example, while creating an image, if you have entered the value in the [Minimum RAM \(MB\)](#) field as 2048, then on selecting the image, the default flavor is m1.small.

[Networks](#) tab

Selected Networks

To add a network to the instance, click the [+](#) in the [Available](#) field.

[Network Ports](#) tab

Ports

Activate the ports that you want to assign to the instance.

[Security Groups](#) tab

Security Groups

Activate the security groups that you want to assign to the instance.

Security groups are a kind of cloud firewall that define which incoming network traffic is forwarded to instances.

If you have not created any security groups, you can assign only the default security group to the instance.

[Key Pair](#) tab

Key Pair

Specify a key pair.

If the image uses a static root password or a static key set (neither is recommended), you do not need to provide a key pair to launch the instance.

[Configuration](#) tab

Customization Script Source

Specify a customization script that runs after your instance launches.

[Metadata](#) tab

Available Metadata

Add Metadata items to your instance.

6. Click [Launch Instance](#).

The instance starts on a compute node in the cloud.

✔ Note

If you did not provide a key pair, security groups, or rules, users can access the instance only from inside the cloud through VNC. Even pinging the instance is not possible without an ICMP rule configured.

You can also launch an instance from the [Images](#) or [Volumes](#) category when you launch an instance from an image or a volume respectively.

When you launch an instance from an image, OpenStack creates a local copy of the image on the compute node where the instance starts.

For details on creating images, see [Creating images manually](#) in the *OpenStack Virtual Machine Image Guide*.

When you launch an instance from a volume, note the following steps:

- To select the volume from which to launch, launch an instance from an arbitrary image on the volume. The arbitrary image that you select does not boot. Instead, it is replaced by the image on the volume that you choose in the next steps.

To boot a Xen image from a volume, the image you launch in must be the same type, fully virtualized or paravirtualized, as the one on the volume.

- Select the volume or volume snapshot from which to boot. Enter a device name. Enter `vda` for KVM images or `xvda` for Xen images.

✔ Note

When running QEMU without support for the hardware virtualization, set `cpu_mode="none"` alongside `virt_type=qemu` in `/etc/nova/nova-compute.conf` to solve the following error:

```
libvirtError: unsupported configuration: CPU mode 'host-model'
for ``x86_64`` qemu domain on ``x86_64`` host is not supported by hypervisor
```

Connect to your instance by using SSH

To use SSH to connect to your instance, use the downloaded keypair file.

✔ Note

The user name is `ubuntu` for the Ubuntu cloud images on TryStack.

1. Copy the IP address for your instance.
2. Use the **ssh** command to make a secure connection to the instance. For example:

```
$ ssh -i MyKey.pem ubuntu@10.0.0.2
```

3. At the prompt, type `yes`.

It is also possible to SSH into an instance without an SSH keypair, if the administrator has enabled root password injection. For more information about root password injection, see [Injecting the administrator password](#) in the *OpenStack Administrator Guide*.

Track usage for instances

You can track usage for instances for each project. You can track costs per month by showing meters like number of vCPUs, disks, RAM, and uptime for all your instances.

1. Log in to the dashboard.
2. Select the appropriate project from the drop down menu at the top left.
3. On the [Project](#) tab, open the [Compute](#) tab and click [Overview](#) category.
4. To query the instance usage for a month, select a month and click [Submit](#).
5. To download a summary, click [Download CSV Summary](#).

Create an instance snapshot

1. Log in to the dashboard.
2. Select the appropriate project from the drop down menu at the top left.
3. On the [Project](#) tab, open the [Compute](#) tab and click the [Instances](#) category.
4. Select the instance from which to create a snapshot.
5. In the actions column, click [Create Snapshot](#).
6. In the [Create Snapshot](#) dialog box, enter a name for the snapshot, and click [Create Snapshot](#).

The [Images](#) category shows the instance snapshot.

To launch an instance from the snapshot, select the snapshot and click [Launch](#). Proceed with launching an instance.

Manage an instance

1. Log in to the dashboard.
2. Select the appropriate project from the drop down menu at the top left.
3. On the [Project](#) tab, open the [Compute](#) tab and click [Instances](#) category.
4. Select an instance.
5. In the menu list in the actions column, select the state.

You can resize or rebuild an instance. You can also choose to view the instance console log, edit instance or the security groups. Depending on the current state of the instance, you can pause, resume, suspend, soft or hard reboot, or terminate it.

Create and manage networks

The OpenStack Networking service provides a scalable system for managing the network connectivity within an OpenStack cloud deployment. It can easily and quickly react to changing network needs (for example, creating and assigning new IP addresses).

Networking in OpenStack is complex. This section provides the basic instructions for creating a network and a router. For detailed information about managing networks, refer to the [OpenStack Networking Guide](#).

Create a network

1. Log in to the dashboard.
2. Select the appropriate project from the drop down menu at the top left.
3. On the **Project** tab, open the **Network** tab and click **Networks** category.
4. Click **Create Network**.
5. In the **Create Network** dialog box, specify the following values.

Network tab

Network Name: Specify a name to identify the network.

Shared: Share the network with other projects. Non admin users are not allowed to set shared option.

Admin State: The state to start the network in.

Create Subnet: Select this check box to create a subnet

You do not have to specify a subnet when you create a network, but if you do not specify a subnet, the network can not be attached to an instance.

Subnet tab

Subnet Name: Specify a name for the subnet.

Network Address: Specify the IP address for the subnet.

IP Version: Select IPv4 or IPv6.

Gateway IP: Specify an IP address for a specific gateway. This parameter is optional.

Disable Gateway: Select this check box to disable a gateway IP address.

Subnet Details tab

Enable DHCP: Select this check box to enable DHCP.

Allocation Pools: Specify IP address pools.

DNS Name Servers: Specify a name for the DNS server.

Host Routes: Specify the IP address of host routes.

6. Click **Create**.

The dashboard shows the network on the **Networks** tab.

Create a router

1. Log in to the dashboard.
2. Select the appropriate project from the drop down menu at the top left.
3. On the **Project** tab, open the **Network** tab and click **Routers** category.
4. Click **Create Router**.
5. In the **Create Router** dialog box, specify a name for the router and **External Network**, and click **Create Router**.

The new router is now displayed in the **Routers** tab.

6. To connect a private network to the newly created router, perform the following steps:
 1. On the **Routers** tab, click the name of the router.
 2. On the **Router Details** page, click the **Interfaces** tab, then click **Add Interface**.
 3. In the **Add Interface** dialog box, select a **Subnet**.

Optionally, in the **Add Interface** dialog box, set an **IP Address** for the router interface for the selected subnet.

If you choose not to set the **IP Address** value, then by default OpenStack Networking uses the first host IP address in the subnet.

The [Router Name](#) and [Router ID](#) fields are automatically updated.

7. Click [Add Interface](#).

You have successfully created the router. You can view the new topology from the [Network Topology](#) tab.

Create a port

1. Log in to the dashboard.
2. Select the appropriate project from the drop-down menu at the top left.
3. On the [Project](#) tab, click [Networks](#) category.
4. Click on the [Network Name](#) of the network in which the port has to be created.
5. Go to the [Ports](#) tab and click [Create Port](#).
6. In the [Create Port](#) dialog box, specify the following values.

[Name](#): Specify name to identify the port.

[Device ID](#): Device ID attached to the port.

[Device Owner](#): Device owner attached to the port.

[Binding Host](#): The ID of the host where the port is allocated.

[Binding VNIC Type](#): Select the VNIC type that is bound to the neutron port.

7. Click [Create Port](#).

The new port is now displayed in the [Ports](#) list.