# SUSE Cloud

1.0

www.suse.com

December 03, 2012

End User Guide

# *End User Guide*

**List of Authors:** Tanja Roth, Frank Sundermeyer

# Contents

# About This Guide

SUSE® Cloud is an open source software solution that provides the fundamental capabilities to deploy and manage a cloud infrastructure based on SUSE Linux Enterprise. SUSE Cloud is powered by OpenStack, the leading community-driven, open source cloud infrastructure project. It seamlessly manages and provisions workloads across a heterogeneous cloud environment in a secure compliant, and fully-supported manner. The product tightly integrates with other SUSE technologies and with the SUSE maintenance and support infrastructure.

This guide helps cloud users to launch instances, manage volumes and track usage. Most of these tasks can either be achieved with the Web interface (the SUSE Cloud Dashboard) or the OpenStack command line tools.

Many chapters in this manual contain links to additional documentation resources. These include additional documentation that is available on the system as well as documentation available on the Internet.

For an overview of the documentation available for your product and the latest documentation updates, refer to http://www.suse.com/documentation/suse_cloud10.

# 1 Available Documentation

The following manuals are available for this product:

*Deployment Guide* (↑*Deployment Guide*)
: Gives an introduction to the SUSE® Cloud architecture and describes how to set up, deploy, and maintain the individual components.

*User Guide for Administrators* (↑*User Guide for Administrators*)
: Guides you through management of projects and users, images, flavors, and quotas with SUSE Cloud Dashboard or the command line interface.

*End User Guide* (page i)
: Describes how to launch instances, manage volumes, and track usage.

HTML versions of the product manuals can be found in the installed system under `/usr/share/doc/manual`. Find the latest documentation updates at http://www.suse.com/documentation where you can download the manuals for your product in multiple formats.

# 2 Feedback

Several feedback channels are available:

Bugs and Enhancement Requests
: For services and support options available for your product, refer to http://www.suse.com/support/.

  To report bugs for a product component, log in to the Novell Customer Center from http://www.suse.com/support/ and select *My Support* > *Service Request*.

User Comments
: We want to hear your comments about and suggestions for this manual and the other documentation included with this product. Use the User Comments feature at the bottom of each page in the online documentation or go to http://www.suse.com/documentation/feedback.html and enter your comments there.

Mail
: For feedback on the documentation of this product, you can also send a mail to `doc-team@suse.de`. Make sure to include the document title, the product version, and the publication date of the documentation. To report errors or suggest enhancements, provide a concise description of the problem and refer to the respective section number and page (or URL).

# 3 Documentation Conventions

The following typographical conventions are used in this manual:

- `/etc/passwd`: directory names and filenames
- *`placeholder`*: replace *`placeholder`* with the actual value

- `PATH`: the environment variable PATH
- `ls`, `--help`: commands, options, and parameters
- `user`: users or groups
- Alt, Alt + F1: a key to press or a key combination; keys are shown in uppercase as on a keyboard
- *File*, *File* > *Save As*: menu items, buttons
- This paragraph is only relevant for the architectures `amd64`, `em64t`, and `ipf`. The arrows mark the beginning and the end of the text block.

  This paragraph is only relevant for the architectures `System z` and `ipseries`. The arrows mark the beginning and the end of the text block.
- *Dancing Penguins* (Chapter *Penguins*, ↑Another Manual): This is a reference to a chapter in another manual.

# 4 About the Making of This Manual

This book is written in Novdoc, a subset of DocBook (see `http://www.docbook.org`). The XML source files were validated by `xmllint`, processed by `xsltproc`, and converted into XSL-FO using a customized version of Norman Walsh's stylesheets. The final PDF is formatted through XEP from RenderX.

# 1 Using SUSE Cloud Dashboard

The SUSE® Cloud Dashboard is a Web interface that allows cloud administrators and users to manage various OpenStack services. It is based on OpenStack Dashboard (also known under its codename `Horizon`).

After a short introduction to the Dashboard, learn how to execute key tasks such as creating images, launching and managing instances, and how to use volumes for persistent storage.

## 1.1 Requirements

The following requirements need to be fulfilled to access the SUSE Cloud Dashboard:

- The cloud operator has set up SUSE Cloud.
- You have a recent Web browser that supports HTML5. It must have cookies and JavaScript enabled. For using the Dashboard's VNC client, which is based on `noVNC`, your browser needs to support HTML5 Canvas and HTML5 WebSockets. For more details and a list of browsers that support `noVNC`, refer to `https://github.com/kanaka/noVNC/blob/master/README.md`, and `https://github.com/kanaka/noVNC/wiki/Browser-support`, respectively.

# 1.2 SUSE Cloud Dashboard—Overview

Learn how to log in to SUSE Cloud Dashboard and get a short overview of its Web interface.

## 1.2.1 Logging in to the SUSE Cloud Dashboard

To access the SUSE Cloud Dashboard, ask the cloud operator for the following information:

- Hostname or (public) IP address of the SUSE Cloud Dashboard. (The Dashboard is available on the node that has the `nova-dashboard server` role.)
- Username and password of the cloud administrator or cloud user with which you can log in to SUSE Cloud Dashboard.

**1** Start a Web browser and make sure that JavaScript and cookies are enabled.

**2** As a URL, enter the hostname or IP address that you got from the cloud operator.

`https://IP_ADDRESS_OR_HOSTNAME/`

---

**NOTE: Certificate Warning**

Depending on your browser and browser options, you may get a certificate warning when trying to access the URL for the first time. (In case no certificate is provided when setting up the Dashboard, SUSE Cloud uses a self-signed certificate that is not considered trustworthy by default).

In this case, verify the certificate.

To proceed anyway, you can add an exception in the browser to bypass the warning.

---

**3** On the SUSE Cloud Dashboard login screen, enter the *User Name* and *Password* and click *Sign In*.

***Figure 1.1:*** *SUSE Cloud Dashboard—Login Screen*

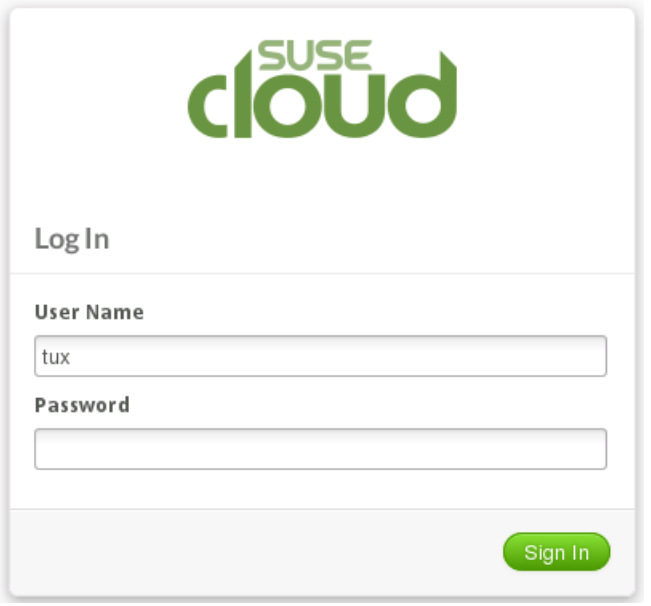


After logging in, the Dashboard's Main Screen (User's View) appears.

## 1.2.2 Main Screen (User's View)

The top-level row of the main screen shows the username with which you are logged in. It also allows you to access the *Settings* or to *Sign Out* of the Web interface.

---

**NOTE: Available Functions**

The visible tabs and functions in the Dashboard depend on the access permissions of the user that is logged in. They are defined by roles.

---

If you are logged in as a user, the main screen only shows the *Project* tab. This shows details for the projects (or `tenants`) that you are a member of.

***Figure 1.2:*** *SUSE Cloud Dashboard—Project Tab*

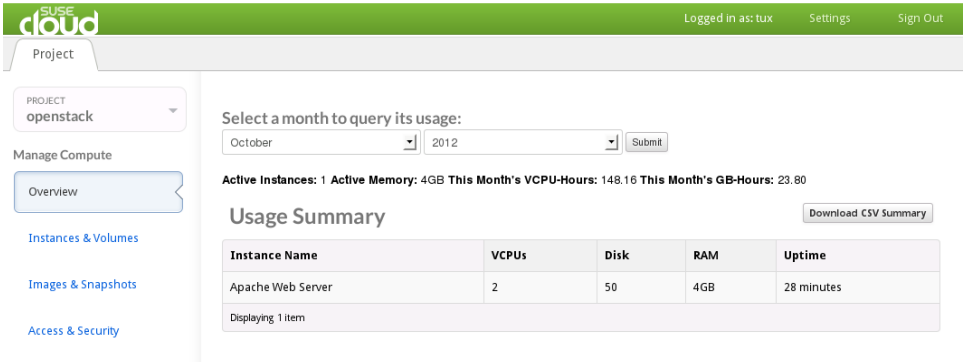


Select a *Project* from the drop-down list on the left-hand side to access the following categories:

*Overview*
: Shows basic reports on the project.

*Instances & Volumes*
: Lists instances and volumes created by users of the project. From here, you can terminate, pause, or reboot any instances or connect to them via VNC.

*Images & Snapshots*
: Lists images and snapshots created by users of the project, plus any images that are publicly available.

*Access & Security*
: Allows to allocate or release floating IP addresses, manage security groups and keypairs.

# 1.3 Managing Images

User permissions to manage images are defined by the cloud operator during setup of SUSE Cloud. Image upload and management may be restricted to cloud administrators or cloud operators only.

After uploading an image to Nova, it cannot be changed any more (“golden image”).

Whereas nearly all key tasks can either be executed from the SUSE Cloud Web interface or from the command line, images can only be uploaded with a command line tool, `glance image-create`. For details, refer to Section "Adding Images" (Chapter 2, *Using OpenStack Command Line Interfaces*, ↑*User Guide for Administrators*).

# 1.4 Launching Instances

Instances are virtual machines that run inside the cloud. To start an instance, a virtual machine image must exist that contains the following information: which operating system to use, a username and password with which to log in to the instance, file storage etc. The cloud contains a pool of such images that have been uploaded to Glance and are accessible to members of different projects.

## 1.4.1 Key Parameters

When starting an instance, you need to specify the following key parameters:

Flavor
: In OpenStack, flavors define the compute, memory, and storage capacity of `nova` computing instances. To put it simply, a flavor is an available hardware configuration for a server. It defines the "size" of a virtual server that can be launched.

  For more details and a list of default flavors available, refer to Section "Managing Flavors" (Chapter 1, *Using SUSE Cloud Dashboard*, ↑*User Guide for Administrators*).

Keypair
: Keypairs are SSH credentials that are injected into images when they are launched. For this to work, the image must contain the `cloud-init` package.

  Create at least one keypair per project. If you already have generated a keypair with an external tool, you can import it into OpenStack. The keypair can be used for multiple instances belonging to that project.

  For details, refer to Section 1.5.1, "Creating or Importing Keys" (page 11).

Security Group

In SUSE Cloud, security groups are used to define which incoming network traffic should be forwarded to instances. Security groups hold a set of firewall policies (security group rules).

For details, refer to Section 1.5.2, “Configuring Security Groups and Rules” (page 12).

If needed, you can assign a floating (public) IP address to a running instance and attach a block storage device (`volume`) for persistent storage. For details, refer to Section 1.5.3, “Managing IP Addresses” (page 16) and Section 1.7, “Managing Volumes” (page 24).

## 1.4.2 Booting From Volumes

You can start an instance directly from one of the images available in Glance or from an image that you have copied to a persistent volume before. For the preparation of the volume, refer to Procedure 1.1 (page 6). When booting an image from a volume, the procedure is basically the same as when launching an instance from an image in Glance, except for some additional steps.

***Procedure 1.1:*** *Creating and Preparing the Volume*

To be able to boot an instance from a volume, create the volume and copy an image to it:

1. Create a volume as described in Procedure 1.10, “Creating or Deleting Volumes” (page 24). Its size must be big enough to store an unzipped image.

2. Create an image with SUSE Studio or SUSE Studio Onsite. For details, refer to Section “Building Images with SUSE Studio” (Chapter 2, *Using OpenStack Command Line Interfaces*, ↑*User Guide for Administrators*).

3. Launch an instance as described in Procedure 1.2, “Launching an Instance” (page 8).

4. Attach the volume to the instance as described in Procedure 1.11, “Attaching Volumes to Instances” (page 25).

5. Assuming that the attached volume is mounted as `/dev/vdb`, use one of the following commands to copy the image to the attached volume:

- For a raw image:

  ```
  cat IMAGE >/dev/null
  ```

  (alternatively, use `dd`)

- For a non-raw image:

  ```
  qemu-img convert -O raw IMAGE /dev/vdb
  ```

- For a `*.tar.bz2` image:

  ```
  tar xfjO IMAGE >/dev/null
  ```

6 As only *detached* volumes are available for booting, detach the volume. For details on how to do so, refer to Procedure 1.11, "Attaching Volumes to Instances" (page 25), Step 9.

7 For booting an instance from the volume, continue with Procedure 1.2, "Launching an Instance" (page 8).

## 1.4.3 Launching an Instance

You can start an instance directly from one of the images available in Glance. In that case, SUSE Cloud will create a local copy of the image on the respective Compute Node where the instance will be started.

---

**NOTE: Launching Instances from a Volume**

Alternatively, you can start an instance from an image that has been copied to a persistent volume. In that case, the instance will be booted from the volume (provided by nova-volume) via iSCSI.

For preparation details, refer to Procedure 1.1, "Creating and Preparing the Volume" (page 6).

To boot an instance from the volume, follow Procedure 1.2, "Launching an Instance" (page 8). Especially note the following steps:

- Step 4: To be able to select from which volume to boot, launch an instance from an arbitrary image. The image you select there will *not* be booted. It will be replaced by the image on the volume that you choose during the next steps.

  In case you want to boot a *Xen* image from a volume, note the following requirement: The image you launch in Step 4 needs to be of the same type (`fully virtualized` or `paravirtualized`) as the one on the volume.

- Step 5f: Select the *Volume or Volume Snapshot* to boot from. Enter a *Device Name* (`vda` for KVM images, `xvda` for Xen images).

---

***Procedure 1.2:*** *Launching an Instance*

**1** Log in to SUSE Cloud Dashboard.

**2** If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.

**3** Click the *Images & Snapshot* category. The Dashboard shows the *Images* that have been uploaded to Glance and are available for this project.

**4** Select an image and click *Launch*.

**5** In the window that opens, specify the following:

**5a** Enter a *Server Name* that will be assigned to the virtual machine.

**5b** From the *Flavor* drop-down list, select the "size" of the virtual machine to launch.

**5c** Select a *Keypair*. For details, refer to Section 1.5.1, "Creating or Importing Keys" (page 11). In case an image uses a static `root` password or a static key set (neither is recommended), you do not need to provide a keypair on starting the instance.

**5d** In *Instance Count*, enter the number of virtual machines to launch from this image.

**5e** 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 should be forwarded to instances. For details, refer to Section 1.5.2, “Configuring Security Groups and Rules” (page 12). If you have not created any specific security groups, you can only assign the instance to the default security group.

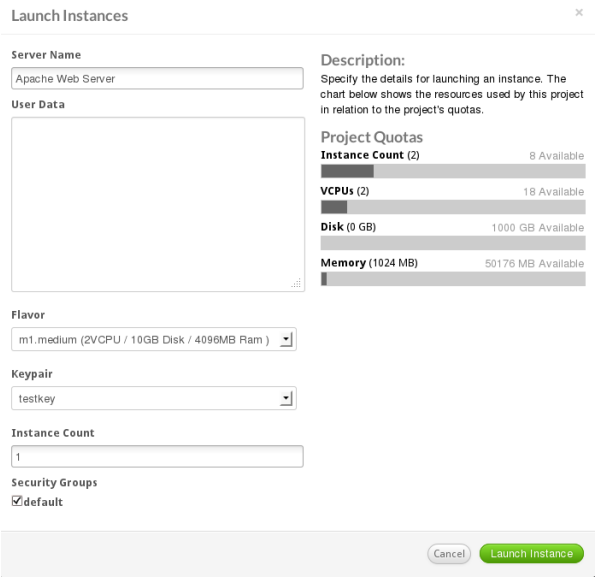


**5f** If you want to *Boot From Volume*, click the respective entry to expand its options. Set the options as described in Launching Instances from a Volume (page 7).

**5g** Click *Launch Instance*. The instance will be started on any of the Compute Nodes in the cloud.

After you have launched an instance, switch to the *Instances & Volumes* category to view the *Instance Name*, its (private or public) *IP address*, its *Size*, its *Status*, *Task*, and *Power State*.

***Figure 1.3:*** *SUSE Cloud Dashboard—List of Launched Instances*

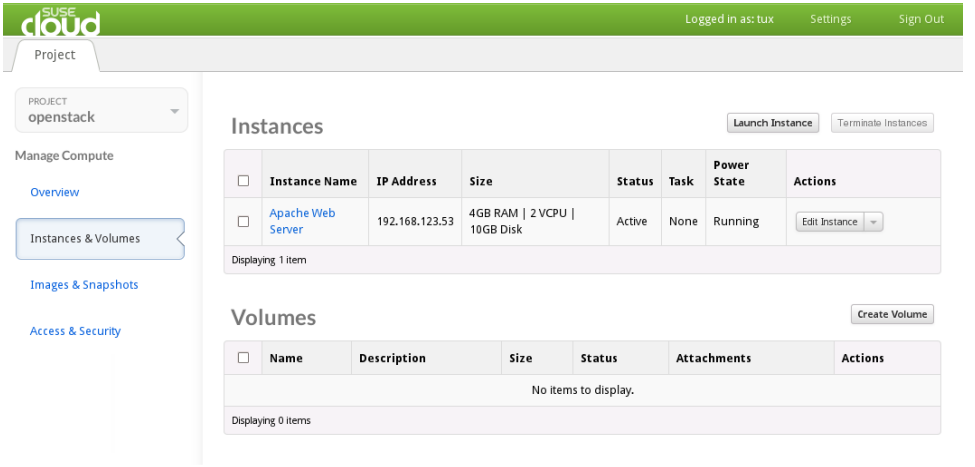


If you did not provide a keypair on starting and have not touched security groups or rules so far, by default the instance can only be accessed from inside the cloud via VNC at this point. Even pinging the instance is not possible. To change this, proceed with Section 1.5, “Configuring Access to the Instances” (page 10).

# 1.5 Configuring Access to the Instances

Access to an instance is mainly influenced by the following parameters:

- keypairs
- security groups and rules
- IP addresses

For SSH access to an instance, you usually need to provide a keypair at launch time. The security rules need adjustment, too, since the default rules block access to SSH ports and prevent pinging an instance. To make the instance also accessible from outside the cloud, assign a floating (public) IP address.

## 1.5.1 Creating or Importing Keys

Keypairs are SSH credentials that are injected into images when they are launched. For this to work, the image must contain the `cloud-init` package.

Create at least one keypair per project. If you already have generated a keypair with an external tool, you can import it into OpenStack. The keypair can be used for multiple instances belonging to that project.

1. Log in to SUSE Cloud Dashboard.
2. If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.
3. Click the *Access & Security* category. The Dashboard shows *Floating IPs*, *Security Groups*, and *Keypairs* for the selected project.
4. To import a keypair that you have generated with an external tool:
   - **4a** Click *Import Keypair*.
   - **4b** In the window that opens, enter a name for the keypair and copy the public key into the respective input field.
   - **4c** Confirm your choice.
5. To create a new keypair:
   - **5a** Click *Create Keypair*.
   - **5b** In the window that opens, enter a name for the keypair and confirm your choice.

     OpenStack generates a keypair and provides the private key for download as a `*.pem` file.
   - **5c** Save the `*.pem` file locally and change its permissions so that only you can read and write to the file:

     ```
     chmod 600 MY_PRIV_KEY.pem
     ```

The public key of the keypair is registered at the Nova database. The Dashboard lists the keypair in the *Access & Security* category as shown in Figure 1.4, “SUSE Cloud Dashboard—Keypairs” (page 12).

***Figure 1.4:*** *SUSE Cloud Dashboard—Keypairs*

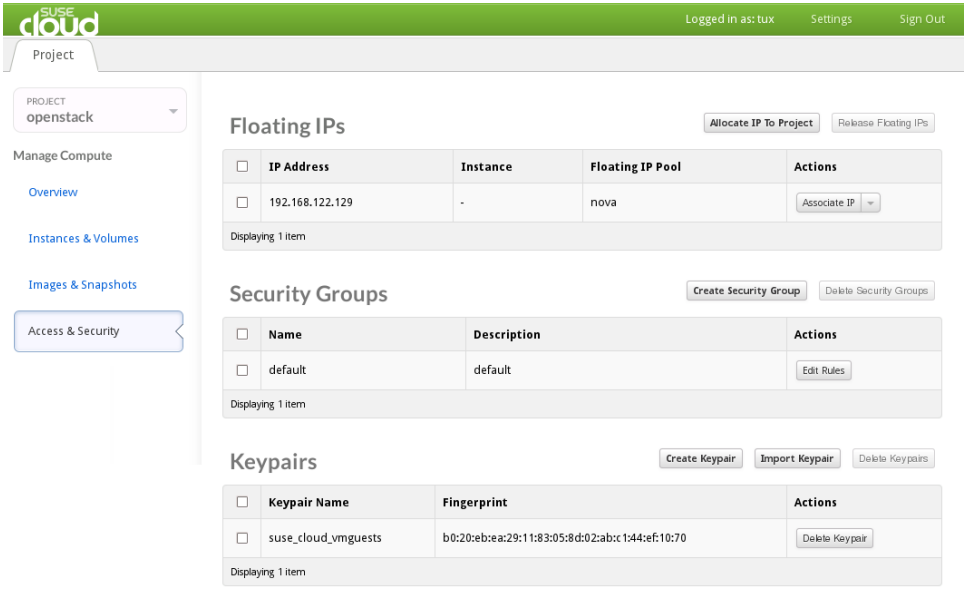


## 1.5.2 Configuring Security Groups and Rules

In SUSE Cloud, security groups are used to define which incoming network traffic should be forwarded to instances. Security groups hold a set of firewall policies (security group rules).

### 1.5.2.1 Security Groups

When launching an instance, you need to define which security groups it should belong to. A default security group is available for each project. It allows all network traffic from other members of this group and discards traffic from other IP addresses and groups.

Multiple security groups for a project can be defined, with each group holding a different set of firewall policies. This is useful if you have groups of instances that should differ in firewall configuration (for example, front-end and back-end servers). An instance can be assigned to multiple security groups.

***Procedure 1.3:*** *Creating or Deleting Security Groups*

1. Log in to SUSE Cloud Dashboard.
2. If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.
3. Click the *Access & Security* category. The Dashboard shows *Floating IPs*, *Security Groups*, and *Keypairs* for the selected project.
4. To create a new security group:
   - **4a** Click *Create Security Group*.
   - **4b** In the window that opens, enter a *Name* and *Description* for the group and confirm your changes.
5. To delete one or multiple security groups:
   - **5a** Activate the check boxes in front of the groups that you want to delete.
   - **5b** Click *Delete Security Groups* and confirm your choice in the pop-up that appears.

     A message on the Web page shows if the action has been successful.

---

**NOTE: Deleting Security Groups**

The default security group for a project cannot be deleted.

If another group cannot be deleted, it is because it is still assigned to a running instance.

---

### 1.5.2.2 Security Group Rules

You can adjust rules of the default security group as well as rules of any other security group that has been created. As soon as the rules for a group are modified, the new rules are automatically applied to all running instances belonging to that security group.

Adjust the rules in a security group to allow access to instances via different ports and protocols. This is necessary to be able to access instances via SSH, to ping them, or to allow UDP traffic (for example, for a DNS server running on an instance).

Rules in security groups are specified by the following parameters:

Source of traffic
: Decide whether to allow traffic to instances only from IP addresses inside the cloud (from other group members) or from *all* IP addresses.

Protocol
: Choose between TCP (for SSH), ICMP (for pings), and UDP.

Destination Port on Virtual Machine
: Define a port range. To open a single port only, enter the same value twice. ICMP does not support ports. In that case, enter values that define the codes and types of ICMP traffic to be allowed.

If no further security groups have been created, any instances are automatically assigned to the default security group (if not specified otherwise). Unless you change the rules for the default group, those instances cannot be accessed from any IP addresses outside the cloud.

***Procedure 1.4:*** *Configuring Security Group Rules*

1. Log in to SUSE Cloud Dashboard.

2. If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.

3. Click the *Access & Security* category. The Dashboard shows *Floating IPs*, *Security Groups*, and *Keypairs* for the selected project.

4. Select the security group to modify, then click *Edit Rules*. The window that appears shows which rules have already been configured.

**5** To allow SSH access to the instances:

**5a** Set *IP Protocol* to `TCP`.

**5b** Enter the value `22` in both *From Port* and *To Port*.

**5c** To allow access from *all* IP addresses (specified as IP subnet in CIDR notation as `0.0.0.0/0`), leave the other fields unchanged.

Alternatively, allow only IP addresses from other security groups to access the specified port. In that case, select the desired security group from the *Source Group* drop-down list.

**5d** Confirm your changes to add the rule.

**6** To allow pinging the instances:

**6a** Set *IP Protocol* to `ICMP`.

**6b** Enter the value `-1` in both *From Port* and *To Port*. This allows access to all codes and all types of ICMP traffic, respectively.

**6c** To allow access from *all* IP addresses (`0.0.0.0/0`), leave the other fields unchanged.

Alternatively, allow only members of other security groups to ping instances. In that case, select the desired security group from the *Source Group* drop-down list.

**6d** Confirm your changes to add the rule.

**7** To allow access via UDP port (for example, for a DNS server running on a VM):

**7a** Set *IP Protocol* to `UDP`.

**7b** Enter the value `53` in both *From Port* and *To Port*.

**7c** To allow access from *all* IP addresses (`0.0.0.0/0`), leave the other fields unchanged.

Alternatively, allow only IP addresses from other security groups to access the specified port. In that case, select the desired security group from the *Source Group* drop-down list.

**7d** Confirm your changes to add the rule.

**8** To delete one or multiple security group rules:

**8a** Select the security group to modify, then click *Edit Rules*. The window that appears shows which rules have already been configured.

**8b** Select the rule or rules to remove.

**8c** Click *Delete Rules* and confirm your choice.

***Figure 1.5:*** *SUSE Cloud Dashboard—Adding Security Group Rules*

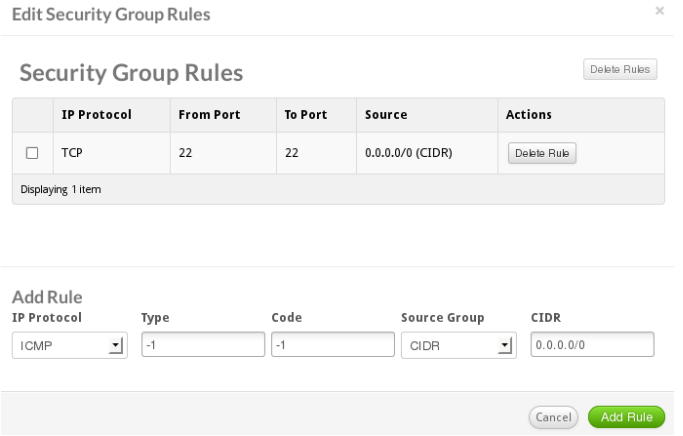


# 1.5.3 Managing IP Addresses

Each instance can have two IP addresses: a private (fixed) IP address and a public (floating) one. Private IP addresses are used for communication between instances, and public ones are used for communication with the outside world. When an instance is launched, it is automatically assigned a private IP address, which stays the same until

the instance is explicitly terminated. (Rebooting the instance does not have an effect on the private IP address.)

A pool of floating IPs is available in OpenStack Nova, as configured by the cloud operator. You can allocate a certain number of these to a project—the maximum number of floating IP addresses per project is defined by the quota. From this set, you can then add a floating IP address to an instance of the project. Floating IP addresses can be dynamically disassociated and associated with other instances of the same project at any time.

***Procedure 1.5:*** *Allocating Floating (Public) IPs to a Project*

Before you can assign a floating IP address to an instance, you first need to allocate floating IPs to a project.

**1** Log in to SUSE Cloud Dashboard.

**2** If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.

**3** Click the *Access & Security* category. The Dashboard shows a list of *Floating IPs*, *Security Groups*, and *Keypairs* for the current project.

**4** To allocate a floating IP address to the current project:

**4a** Click *Allocate IP to Project*.

**4b** In the window that opens, select a *Pool* out of which to take the IP address.

**4c** Click *Allocate IP*.

The Dashboard shows the allocated IP addresses in the *Access & Security* category. The first IP shown for an instance is the private IP address; the second one is the floating IP address.

**5** To release one or multiple floating IP addresses from a project:

**5a** Activate the check boxes in front of the IP addresses that you want to release.

**5b** Click *Release Floating IPs*. The IP addresses are put back into the pool of IP addresses that are available for all projects. If an IP address is currently as-

signed to a running instance, it will automatically be disassociated from the instance.

***Procedure 1.6:*** *Assigning Floating (Public) IP Addresses to Instances*

After floating IP addresses have been allocated to the current project, you can assign them to running instances. One floating IP address can be assigned to only one instance at a time.

**1** Log in to SUSE Cloud Dashboard.

**2** If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.

**3** Click the *Access & Security* category. The Dashboard shows a list of *Floating IPs*, *Security Groups*, and *Keypairs* for the current project.

**4** To assign an IP to an instance:

**4a** Select an IP address and click *Associate IP*.

**4b** In the window that opens, select the *Instance* to associate the IP with and confirm your choice.

In the *Access & Security* category, the list of *Floating IPs* shows the ID of the instance with which the IP has been associated. The instance is now publicly available under the respective floating IPs address (provided you have also configured the security group rules for the instance accordingly). For details, refer to Section 1.5.2, “Configuring Security Groups and Rules” (page 12).

**5** To remove a floating IP address from an instance:

**5a** Click the *Access & Security* category.

**5b** Select the IP address to remove.

**5c** Click *Disassociate IP* and confirm your change.

***Figure 1.6:*** *SUSE Cloud Dashboard—Instance IPs*


# 1.6 Managing Instances

The following are typical tasks for managing instances:

- Accessing instances from remote
- Viewing logs
- Creating instance snapshots to preserve a certain disk state of an instance
- Using instance snapshots as base for new images
- Rebooting or terminating instances
- Pausing or suspending instances
- Tracking instance usage

## 1.6.1 Viewing Instance Logs

1 Log in to SUSE Cloud Dashboard.

2 If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.

3 Click the *Instances & Volumes* category.

4 Select the instance and from the *Actions* drop-down list, select *View Log*.

Alternatively, click the instance's name and switch to the *Log* tab that opens.

The Dashboard shows the output of the instance's serial console. To make use of this feature, the respective image must have set the serial console correctly in GRUB.

## 1.6.2 Accessing Instances from Remote

The Dashboard's built-in VNC client allows you to access instances at any time.

***Procedure 1.7:*** *Accessing an Instance via VNC*

1 Log in to SUSE Cloud Dashboard.

2 If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.

3 Click the *Instances & Volumes* category.

4 Select the instance to access and from the *Actions* drop-down list, select *VNC Console*.

Alternatively, click the instance's name and switch to the *VNC* tab that opens.

5 When establishing the first connection, you might be prompted by your browser to trust a certificate before you can see the VNC screen.

6 To display a larger VNC screen, use the link *Click here to show only VNC*.

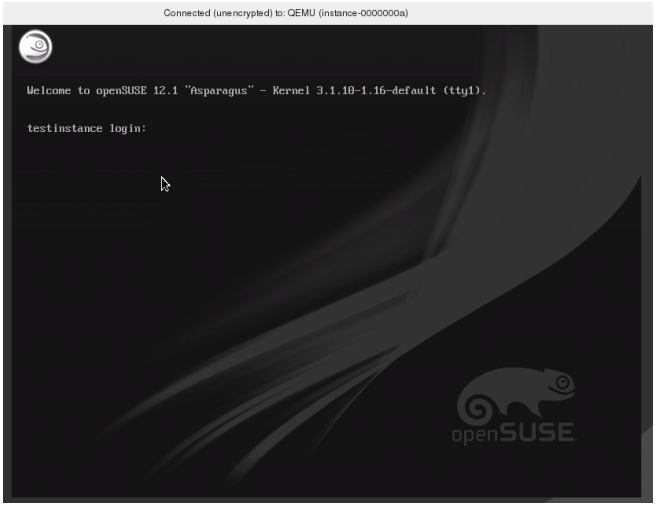


**7** To leave the large VNC screen, use the back button of the browser.

To access an instance via SSH, the following requirements need to be fulfilled:

- `sshd` must be running inside the virtual machine.
- Port `22` must be open in the virtual machine's firewall.
- The security group which the instance is assigned to, must be configured to allow SSH access.
- To allow SSH access from outside the cloud, a floating IP address must be assigned to the instance.
- You must know the private or public IP address of the instance.

## 1.6.3 Using Instance Snapshots

Instance snapshots preserve the disk state of a running instance. You can launch a new instance from a snapshot or use a snapshot to create a new image based upon the snapshot. Ephemeral disks are not included in any snapshots.

***Procedure 1.8:*** *Creating Instance Snapshots*

1. Log in to SUSE Cloud Dashboard.
2. If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.
3. Click the *Instances & Volumes* category.
4. Select the instance of which to create a snapshot. From the *Actions* drop-down list, select *Snapshot*.
5. In the window that opens, enter a name for the snapshot and confirm your changes. The Dashboard shows the new *Instance Snapshot* in the *Images & Snapshot* category.
6. To launch a new instance from the snapshot, select the snapshot and click *Launch*. Proceed with launching an instance as described in Procedure 1.2, "Launching an Instance" (page 8).

***Procedure 1.9:*** *Basing an Image on a Snapshot*

1. Log in to SUSE Cloud Dashboard.
2. If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.
3. Click the *Images & Snapshots* category.
4. Select the snapshot and from the *Actions* drop-down list, select *Edit*.
5. In the window that opens, enter the image properties. For more information, refer to Section "Managing Images" (Chapter 1, *Using SUSE Cloud Dashboard*, ↑*User Guide for Administrators*).
6. Click *Update Image*.

   Dashboard shows the newly created image in the list of images in the *Images & Snapshots* category. If you delete the snapshot, upon which the image is based, the image will be deleted as well. If you delete the image, the snapshot upon which it is based, will be deleted as well.

## 1.6.4 Pausing, Suspending, Rebooting, or Terminating Instances

For maintenance reasons, you can pause or suspend images—provided they are running on KVM or Xen. Pausing or suspending avoids the consequences that come with terminating an instance.

If you pause an instance, the content of the virtual machine is stored to memory (RAM) and the image is kept running in a "frozen" state. When suspending an instance, the content of the virtual machine is stored to disk, and memory and VCPUs are freed.

---

**WARNING: Terminating Instances: Risk of Data Loss**

*Terminating* an instance has the following consequences:

- All data on the image's root disk and ephemeral disks are destroyed. To prevent that, use volumes and attach them to an instance for persistent storage.
- If a floating IP address was assigned to that instance, the IP address is disassociated from that image. However, it is still available in the pool of allocated IP addresses for the current project.

---

1. Log in to SUSE Cloud Dashboard.
2. If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.
3. Click the *Instances & Volumes* category.
4. Select the instance that you want to put out of the running state. From the *Actions* drop-down list, select the respective action.

## 1.6.5 Tracking Usage

Use the Dashboard's *Overview* category to track usage of instances per project. This allows you to track costs per month by showing metrics like number of VCPUs, disks, RAM, and uptime of all your instances.

If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab. Select a month and click *Submit* to query the instance usage for that month. The Dashboard also allows to download a CVS summary.

***Figure 1.7:*** *SUSE Cloud Dashboard—Usage Overview*

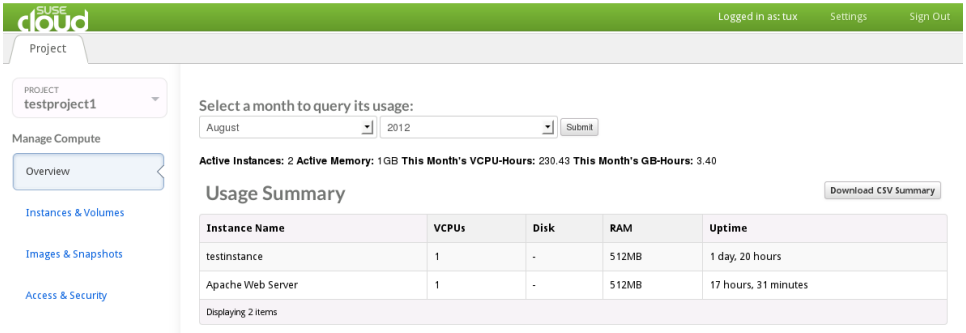


# 1.7 Managing Volumes

Volumes are block storage devices that can be attached to instances. They allow for persistent storage as they can be attached to a running instance (or detached and attached to another instance at any time). In contrast to the instance's root disk, the data of volumes is not destroyed when the instance is terminated.

***Procedure 1.10:*** *Creating or Deleting Volumes*

1. Log in to SUSE Cloud Dashboard.
2. If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.
3. Click the *Instances & Volumes* category.

4 To create a volume:

   4a Click *Create Volume*.

   4b In the window that opens, enter a name to assign to a volume, a description (optional), and define the size in GB.

   4c Confirm your changes.

      The Dashboard shows the volume in the *Instances & Volumes* category.

5 To delete one or multiple volumes:

   5a Activate the check boxes in front of the volumes that you want to delete.

   5b Click *Delete Volumes* and confirm your choice in the pop-up that appears.

      A message on the Web page shows if the action has been successful.

After having created one or multiple volumes, you can attach them to instances. A volume can only be attached to one instance at a time. View the *Status* of a volume in the *Instances & Volumes* category of the Dashboard: the volume is either `available` or already `In-Use`.

***Procedure 1.11:*** *Attaching Volumes to Instances*

1 Log in to SUSE Cloud Dashboard.

2 If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.

3 Click the *Instances & Volumes* category.

4 Select the volume to add to an instance and click *Edit Attachments*.

5 In the window that opens, select an instance to attach the volume to.

6 Enter a *Device Name* under which the volume should be accessible on the virtual machine.

7 Confirm your changes. The Dashboard shows the instance to which the volume has been attached and the volume's device name.

8 Now you can log in to the instance, mount the disk, format it, and use it.

If the instance is running the latest SUSE Linux Enterprise Server SP2 Kernel, it is not necessary to reboot the virtual machine to make the device appear. Otherwise load the `acpiphp` module manually:

```
modprobe acpiphp
```

9 To detach a volume from an instance:

**9a** Select the volume and click *Edit Attachments*.

**9b** In the window that opens, click *Detach Volume* and confirm your changes.

A message on the Web page shows if the action has been successful.

***Procedure 1.12:*** *Creating Volume Snapshots*

1 Log in to SUSE Cloud Dashboard.

2 If you are a member of multiple projects, select a *Project* from the drop-down list at the top of the tab.

3 Click the *Instances & Volumes* category.

4 Select the volume of which to create a snapshot.

5 From the *Actions* drop-down list, select *Create Snapshot*.

6 In the window that opens, enter a *Snapshot Name* and a *Description*.

7 Confirm your changes. The Dashboard shows the new *Volume Snapshot* in the *Images & Snapshots* category.

# Using OpenStack Command Line Interfaces

# 2

The OpenStack project provides a variety of command line tools with which you can manage the services within your cloud and automate tasks by using scripts. Each of the core OpenStack components has its own command line tool.

## 2.1 OpenStack Commands—Overview

The following command line tools are available for the respective services' APIs:

`keystone`
: For managing users and projects. Provided by the `python-keystoneclient` package.

`nova`
: For managing instances and flavors. Provided by the `python-novaclient` package.

`glance`
: For managing images. Provided by the `python-glanceclient` package.

`swift`
: For managing the object store. Provided by the `python-swiftclient` package.

All of them have tab completion.

Help and detailed information about the individual commands and their arguments are available with

```
COMMAND help
```

For help on subcommands, use

```
COMMAND help SUBCOMMAND
```

For example: `glance help` or `glance help image-create`

# 2.2 OpenStack RC File

To set the necessary environment variables for the OpenStack command line tools, you need to download and source an environment file, `openrc.sh`. It is project-specific and contains the credentials used by OpenStack Compute, Image, and Identity services. You can download it from the SUSE Cloud Dashboard (either as user `admin` or as any other user).

***Procedure 2.1:*** *Downloading the OpenStack RC File*

1. Log in to the SUSE Cloud Dashboard.
2. In the top-level row of the main screen, click *Settings* > *OpenStack Credentials*.
3. Select the project for which you want to download the OpenStack RC file, click *Download RC File* and save the file.
4. Copy the `openrc.sh` file to the machine on which you want to execute OpenStack commands (for example, uploading an image with the `glance` command).
5. On any shell that you want to execute OpenStack commands from, source the `openrc.sh` file for the respective project:

   ```
   source openrc.sh
   ```

   You will be prompted for an OpenStack password.
6. Enter the OpenStack password of the user who downloaded the `openrc.sh` file.

With sourcing the file and entering the password, environment variables are set for that shell. They allow the commands to communicate to the OpenStack services running in the cloud.

## 2.3 Managing Images

User permissions to manage images are defined by the cloud operator during setup of SUSE Cloud. Image upload and management may be restricted to cloud administrators or cloud operators only.

After uploading an image to Nova, it cannot be changed any more ("golden image").

Whereas nearly all key tasks can either be executed from the SUSE Cloud Web interface or from the command line, images can only be uploaded with a command line tool, `glance image-create`. For details, refer to Section "Adding Images" (Chapter 2, *Using OpenStack Command Line Interfaces*, ↑*User Guide for Administrators*).

## 2.4 Launching Instances

Instances are virtual machines that run inside the cloud. To start an instance, a virtual machine image must exist that contains the following information: which operating system to use, a username and password with which to log in to the instance, file storage etc. The cloud contains a pool of such images that have been uploaded to Glance and are accessible to members of different projects.

Upon start of an instance, you need to specify the following key parameters:

Flavor

In OpenStack, flavors define the compute, memory, and storage capacity of `nova` computing instances. To put it simply, a flavor is an available hardware configuration for a server. It defines the "size" of a virtual server that can be launched.

For more details and a list of default flavors available, refer to Section "Managing Flavors" (Chapter 2, *Using OpenStack Command Line Interfaces*, ↑*User Guide for Administrators*).

Keypair
: Keypairs are SSH credentials that are injected into images when they are launched. For this to work, the image must contain the `cloud-init` package.

  Create at least one keypair per project. If you already have generated a keypair with an external tool, you can import it into OpenStack. The keypair can be used for multiple instances belonging to that project.

  For details, refer to Section 2.5.1, “Creating or Importing Keys” (page 32).

Security Group
: In SUSE Cloud, security groups are used to define which incoming network traffic should be forwarded to instances. Security groups hold a set of firewall policies (security group rules).

  For details, refer to Section 2.5.2, “Configuring Security Groups and Rules” (page 33).

If needed, you can assign a floating (public) IP address to a running instance and attach a block storage device (`volume`) for persistent storage. For details, refer to Section 2.5.3, “Managing IP Addresses” (page 36).

Before you can launch an instance, you need to look up a few parameters, for example, which images, flavors, and security groups are available. Proceed as follows:

***Procedure 2.2:*** *Launching an Instance*

1. On a shell, source the OpenStack RC file. For details, refer to Section 2.2, “OpenStack RC File” (page 28).

2. Look up the available flavors:

   ```
   nova flavor-list
   ```

   Memorize the ID of the flavor that you want to use for your instance.

3. Look up the available images:

   ```
   nova image-list
   ```

   Memorize the ID of the image that you want to boot your instance from.

4. Look up the available security groups:

```
nova secgroup-list
```

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

5 Look up your keypair's name (for SSH access) and memorize it:

```
nova keypair-list
```

6 Now you have all the parameters at hand for starting an instance. Do so with the following command:

```
nova boot --flavor FLAVOR_ID --imageIMAGE_ID --key_name KEY_NAME \
--security_group NAME_OF_SEC_GROUP NAME_FOR_INSTANCE
```

The command returns a list of instance properties, including the `status` of the instance. The status `BUILD` indicates that the instance has started, but is not yet online.

7 Check if the instance is online:

```
nova list
```

This command lists all instances of the project you belong to, including their ID, their name, their status, and their private (and if assigned, their public) IP addresses. If your instance's status is `ACTIVE`, the instance is online.

To refine the search, run `nova help list` to view the available options for the command.

If you did not provide a keypair on starting and have not touched security groups or rules so far, by default the instance can only be accessed from inside the cloud via VNC at this point. Even pinging the instance is not possible. To change this, proceed with Section 2.5, “Configuring Access to the Instances” (page 31).

# 2.5 Configuring Access to the Instances

Access to an instance is mainly influenced by the following parameters:

- keypairs

- security groups and rules
- IP addresses

For SSH access to an instance, you usually need to provide a keypair at launch time. The security rules need adjustment, too, since the default rules block access to SSH ports and prevent pinging an instance. To make the instance also accessible from outside the cloud, assign a floating (public) IP address.

## 2.5.1 Creating or Importing Keys

Keypairs are SSH credentials that are injected into images when they are launched. For this to work, the image must contain the `cloud-init` package.

Create at least one keypair per project. If you already have generated a keypair with an external tool, you can import it into OpenStack. The keypair can be used for multiple instances belonging to that project.

In case an image uses a static `root` password or a static key set (neither is recommended), you do not need to provide a keypair on starting of the instance.

***Procedure 2.3:*** *Creating or Importing Keys*

Use the `nova keypair-add` command to generate a new keypair, or to upload an existing public key.

1. To generate a new keypair, execute the following commands:

   ```
   nova keypair-add KEY_NAME > MY_KEY.pem
   chmod 600  MY_KEY.pem
   ```

   The first command generates a new keypair named *`KEY_NAME`*, writing the private key to the file *`MY_KEY`*`.pem` and registering the public key at the Nova database. The second command changes the permissions of the file *`MY_KEY`*`.pem` so that only you can read and write to it.

2. If you already have generated a keypair, with the public key located at `~/.ssh/id_rsa.pub`, you can upload the public key with the following command:

   ```
   nova keypair-add --pub_key ~/.ssh/id_rsa.pub KEY_NAME
   ```

The command registers the public key at the Nova database and names the keypair *KEY_NAME*.

3 Check if the uploaded keypair appears in the list of availabe keypairs:

```
nova keypair-list
```

## 2.5.2 Configuring Security Groups and Rules

In SUSE Cloud, security groups are used to define which incoming network traffic should be forwarded to instances. Security groups hold a set of firewall policies (security group rules).

### 2.5.2.1 Security Groups

When launching an instance, you need to define which security groups it should belong to. A default security group is available for each project. It allows all network traffic from other members of this group and discards traffic from other IP addresses and groups.

Multiple security groups for a project can be defined, with each group holding a different set of firewall policies. This is useful if you have groups of instances that should differ in firewall configuration (for example, front-end and back-end servers). An instance can be assigned to multiple security groups.

Security groups can be managed with the `nova secgroup-*` commands, provided by the `python-novaclient` package.

Listing Security Groups

```
nova secgroup-list
```

Lists all security groups for the current project, including the groups' descriptions.

Creating a Security Group

```
nova secgroup-create SEC_GROUP_NAME GROUP_DESCRIPTION
```

Creates a new security group with the specified name and description.

Deleting a Security Group

```
nova secgroup-delete SEC_GROUP_NAME
```

Deletes the specified group.

**NOTE: Deleting Security Groups**

The default security group for a project cannot be deleted.

If another group cannot be deleted, it is because it is still assigned to a running instance.

### 2.5.2.2 Security Group Rules

You can adjust rules of the default security group as well as rules of any other security group that has been created. As soon as the rules for a group are modified, the new rules are automatically applied to all running instances belonging to that security group.

Adjust the rules in a security group to allow access to instances via different ports and protocols. This is necessary to be able to access instances via SSH, to ping them, or to allow UDP traffic (for example, for a DNS server running on an instance).

Rules in security groups are specified by the following parameters:

Source of traffic
: Decide whether to allow traffic to instances only from IP addresses inside the cloud (from other group members) or from *all* IP addresses.

Protocol
: Choose between TCP (for SSH), ICMP (for pings), and UDP.

Destination Port on Virtual Machine
: Define a port range. To open a single port only, enter the same value twice. ICMP does not support ports. In that case, enter values that define the codes and types of ICMP traffic to be allowed.

If no further security groups have been created, any instances are automatically assigned to the default security group (if not specified otherwise). Unless you change the rules for the default group, those instances cannot be accessed from any IP addresses outside the cloud.

***Procedure 2.4:*** *Configuring Security Group Rules*

Modify security group rules with the `nova secgroup-*-rule` commands. Proceed as follows:

**1** On a shell, source the OpenStack RC file. For details, refer to Section 2.2, "OpenStack RC File" (page 28).

**2** Look up the existing rules for a security group:

```
nova secgroup-list-rules SEC_GROUP_NAME
```

**3** To allow SSH access to the instances:

**3a** Either from *all* IP addresses (specified as IP subnet in CIDR notation as `0.0.0.0/0`):

```
nova secgroup-add-rule SEC_GROUP_NAME tcp 22 22 0.0.0.0/0
```

**3b** Alternatively, you can allow only IP addresses from other security groups (`source groups`) to access the specified port:

```
nova secgroup-add-group-rule --ip_proto tcp --from_port 22 \
--to_port 22 SEC_GROUP_NAME SOURCE_GROUP_NAME
```

**4** To allow pinging the instances:

**4a** Either from *all* IP addresses (specified as IP subnet in CIDR notation as `0.0.0.0/0`):

```
nova secgroup-add-rule SEC_GROUP_NAME icmp -1 -1 0.0.0.0/0
```

This command allows access to all codes and all types of ICMP traffic, respectively.

**4b** Alternatively, you can allow only members of other security groups (`source groups`) to ping instances:

```
nova secgroup-add-group-rule --ip_proto icmp --from_port -1 \
--to_port -1 SEC_GROUP_NAME SOURCE_GROUP_NAME
```

**5** To allow access via UDP port (for example, for a DNS server running on a VM):

**5a** Either from *all* IP addresses (specified as IP subnet in CIDR notation as `0.0.0.0/0`):

```
nova secgroup-add-rule SEC_GROUP_NAME udp 53 53 0.0.0.0/0
```

**5b** Alternatively, you can allow only IP addresses from other security groups (`source groups`) to access the specified port:

```
nova secgroup-add-group-rule --ip_proto udp --from_port 53 \
--to_port 53 SEC_GROUP_NAME SOURCE_GROUP_NAME
```

**6** To delete security group rules, you need to specify the same arguments that you used to create the rule. For example:

To delete the security rule that you created in Step 3a (page 35):

```
nova secgroup-delete-rule SEC_GROUP_NAME tcp 22 22 0.0.0.0/0
```

To delete the security rule that you created in Step 3b (page 35):

```
nova secgroup-delete-group-rule --ip_proto tcp --from_port 22 \
--to_port 22 SEC_GROUP_NAME SOURCE_GROUP_NAME
```

## 2.5.3 Managing IP Addresses

Each instance can have two IP addresses: a private (fixed) IP address and a public (floating) one. Private IP addresses are used for communication between instances, and public ones are used for communication with the outside world. When an instance is launched, it is automatically assigned a private IP address, which stays the same until the instance is explicitly terminated. (Rebooting the instance does not have an effect on the private IP address.)

A pool of floating IPs is available in OpenStack Nova, as configured by the cloud operator. You can allocate a certain number of these to a project—the maximum number of floating IP addresses per project is defined by the quota. From this set, you can then add a floating IP address to an instance of the project. Floating IP addresses can be dynamically disassociated and associated with other instances of the same project at any time.

Before you can assign a floating IP address to an instance, you first need to allocate floating IPs to a project.

After floating IP addresses have been allocated to the current project, you can assign them to running instances. One floating IP address can be assigned to only one instance at a time.

Floating IP addresses can be managed with the `nova *floating-ip-*` commands, provided by the `python-novaclient` package.

Listing Pools with Floating IP Addresses

```
nova floating-ip-pool-list
```

Lists the name of all pools that provide floating IP addresses.

Allocating a Floating IP Address to the Current Project

```
nova floating-ip-pool-list
```

The output of the command shows the freshly allocated IP address. If there is more than one pool of IP addresses available, you can also specify the pool from which to allocate the IP address (optional):

```
floating-ip-create POOL_NAME
```

Listing Floating IP Addresses Allocated to the Current Project

```
nova floating-ip-list
```

Lists all floating IP addresses that have been allocated to the current project. If an IP is already associated with an instance, the output also shows the instance's IP, the instance's fixed IP address and the name of the pool that provides the floating IP address.

Releasing a Floating IP Address from the Current Project

```
nova floating-ip-delete FLOATING_IP
```

The IP address is put back into the pool of IP addresses that are available for all projects. If an IP address is currently assigned to a running instance, it will automatically be disassociated from the instance.

Assigning a Floating IP Address to an Instance

```
nova add-floating-ip INSTANCE_NAME_OR_ID FLOATING_IP
```

To associate an IP address with an instance, one or multiple floating IP addresses must have been allocated to the current project. Check this with `nova`

`floating-ip-list`. In addition, you need to know the instance's name (or ID). To look up the instances that belong to the current project, use the `nova list` command.

After assigning the IP with `nova add-floating-ip`, the instance is now publicly available under the respective floating IP address (provided you have also configured the security group rules for the instance accordingly). For details, refer to Section 1.5.2, “Configuring Security Groups and Rules” (page 12).

Removing a Floating IP Address from an Instance

```
nova remove-floating-ip INSTANCE_NAME_OR_ID FLOATING_IP
```

To remove a floating IP address from an instance, you need to specify the same arguments that you used to assign the IP.