

Virtuozzo

Virtuozzo Cloud 1.0

Self-Service Guide

7/9/2026

Table of contents

About this guide	5
Managing notifications	6
Managing API credentials	9
Managing compute resources	11
Managing virtual machines	11
Supported guest operating systems	11
Creating virtual machines	13
Connecting to virtual machines	18
Setting a password inside virtual machines	19
Managing virtual machine power state	20
Attaching ISO images to virtual machines	21
Reconfiguring virtual machines	22
Monitoring virtual machines	26
Shelving virtual machines	27
Rescuing virtual machines	27
Managing guest tools	30
Troubleshooting virtual machines	32
Deleting virtual machines	33
Managing Kubernetes clusters	33
Supported Kubernetes and Fedora CoreOS versions	33
Creating and deleting Kubernetes clusters	35
Managing Kubernetes worker groups	38
Updating Kubernetes clusters	41
Using persistent volumes for Kubernetes pods	43
Managing volume snapshots in Kubernetes	48
Creating external load balancers in Kubernetes	50
Using network policies in Kubernetes	52
Assigning Kubernetes pods to specific nodes	56
Enabling GPU support for Kubernetes nodes	56
Monitoring Kubernetes clusters	57
Managing images	60
Uploading images	61
Creating volumes from images	62
Preparing templates	63
Managing volumes	68

Creating and deleting volumes	68
Attaching and detaching volumes	69
Resizing volumes	70
Changing the storage policy for volumes	71
Creating images from volumes	71
Cloning volumes	72
Managing volume snapshots	73
Transferring volumes between projects	75
Managing virtual networks	76
Managing VPN connections	79
Creating VPN connections	80
Editing VPN connections	84
Restarting and deleting VPN connections	85
Managing virtual routers	85
Traffic flow and address translation	86
Managing router interfaces	87
Managing static routes	89
Managing floating IP addresses	91
Managing security groups	92
Creating and deleting security groups	92
Managing security group rules	93
Changing security group assignment	94
Managing load balancers	94
Creating load balancers	95
Managing balancing pools	98
Editing and deleting load balancers	103
Monitoring load balancers	103
Managing backups	104
Creating backup plans	104
Managing volumes in backup plans	106
Editing and deleting backup plans	107
Creating and deleting backups manually	108
Restoring volumes from backups	108
Restoring virtual machines from backups	109
Managing SSH keys	110
Managing migrations from external platforms	111
Supported migration configurations	112

VMware migration requirements	114
Managing source clouds	115
Preparing the Windows migration worker template	117
Creating migrations	119
Viewing migration details	121
Replicating migrations	122
Deploying migrations	123
Editing and deleting migrations	125
Managing S3 resources	127
Enabling access to S3 storage	127
Managing access keys	127
Managing buckets	128
Creating and deleting buckets	128
Managing bucket policies	129
Managing files and folders	146

About this guide

This guide is intended for project administrators and explains how to manage compute and S3 resources using the self-service panel.

Managing notifications

The notification center stores and shows notifications about recent tasks of the current user in the management panel. Notifications are displayed only for tasks performed during the current user session and cleared out when the user logs out.

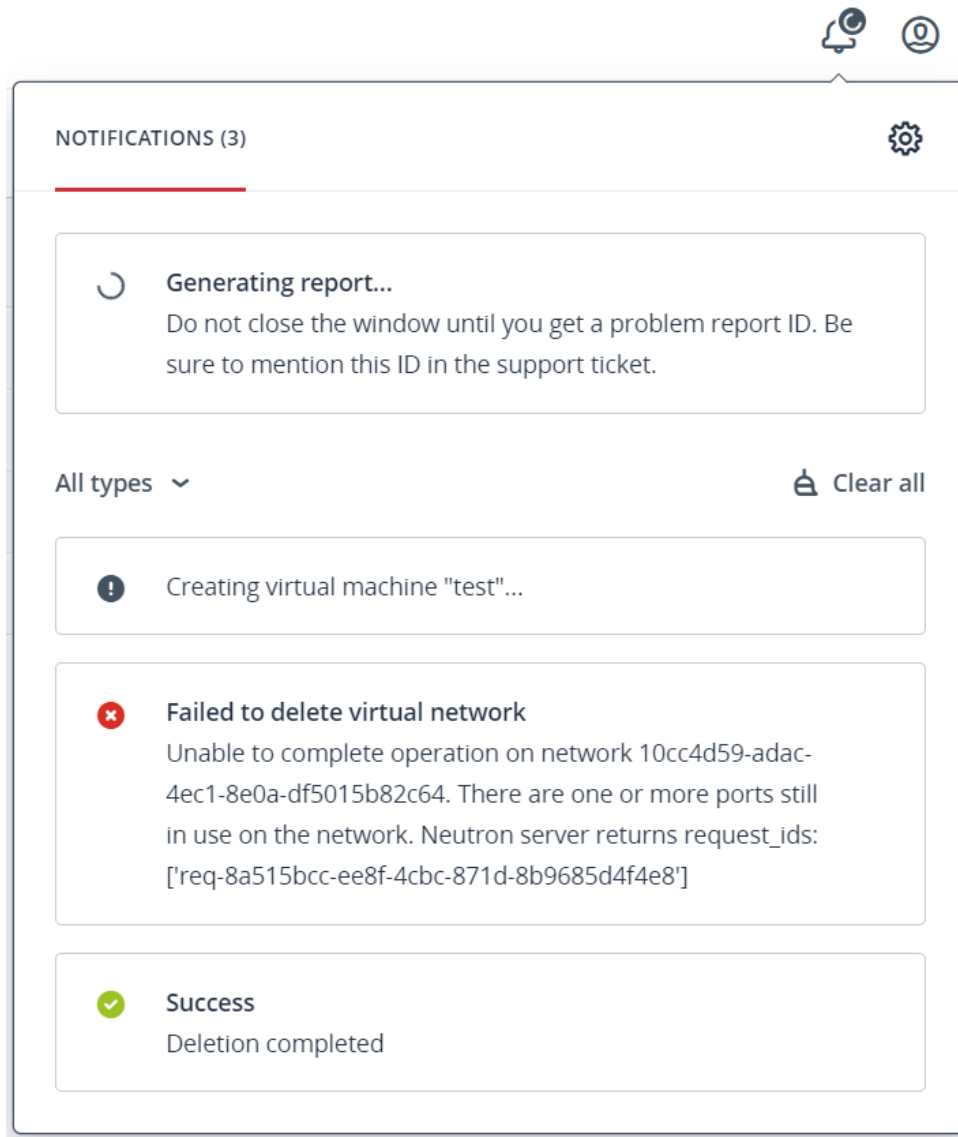
A user is informed about each task by a pop-up notification in the bottom right corner of the screen. The same notification also appears in the notification center. After the pop-up window is closed, the notification is available in the notification center.

The following table describes all of the supported notification types:

Notification type	Icon	Description	Retention period of a pop-up window	Retention period in the notification center
Info		Notifications about a task launch	3 seconds	10 minutes
Success		Notifications about successfully completed tasks	3 seconds	10 minutes
Error		Notifications about failed tasks	10 seconds	50 minutes
In progress		Long-running tasks, such as image upload or problem report creation	Task time	Task time

To view notifications

1. On any screen, click the bell icon in the top right corner. Next to the bell icon, you can see the notification counter, or the loading sign if you have a running task.
2. To view notifications of a particular type, click **All types**, and then select the notification type you wish to be displayed in the notification center.



To clear notifications

1. On any screen, click the bell icon in the top right corner.
2. To clear only one notification, click the cross icon next to it.
3. To clear all of the notifications, click **Clear all** above the notification list.

To configure pop-up notifications

1. On any screen, click the bell icon in the top right corner.
2. Click the cogwheel icon, and then clear notification types that you do not wish to be shown in a

pop-up window. Only the selected notification types will appear as pop-up messages.

Allow pop-up notifications

Select which notification types will be shown as pop-up messages.

- ☒ Error
- ☐ Info
- ☐ Success

To mute pop-up notifications

1. On any screen, click the bell icon in the top right corner.
2. Click the cogwheel icon, and then turn on the **Do not disturb** mode.

Notification settings


☒ Do not disturb

Mutes all pop-up notifications. Notifications will still be available in the notification center.

The bell icon will be greyed out, and the notification counter will disappear. While this mode is on, pop-up notifications are disabled. However, all notifications are still available in the notification center.

To unmute pop-up notifications

1. On any screen, click the greyed out bell icon in the top right corner.
2. Click **Turn off**, to turn off the **Do not disturb** mode.



Managing API credentials

API credentials allow applications, automation tools, and scripts to authenticate to the public compute API without using a user name and password.

Each API credential is associated with a user account and a project. It consists of a credential ID and an API secret key. After successful authentication, the application receives a project-scoped token and can access cloud resources according to the user's permissions in that project. An API credential cannot grant more permissions than the user already has. To access multiple projects, create a separate API credential for each project.

After creating an API credential, you can use it to authenticate OpenStack-compatible tools, scripts, and automation systems. For more information about OpenStack application credentials, refer to the [OpenStack documentation](#).

Limitations

- You can view, create, and revoke only your own API credentials.

To create an API credential

1. Go to **Settings > API credentials** and click **Create**.
2. In the **Create API credential** window, specify a unique credential name, select the project that the credential will access, and set the date when the credential expires. The expiration date cannot be later than one year from the credential creation date.
3. Click **Create**.
4. In the **Credentials created** window, copy and save the credential ID and API secret key. To copy both values at once, click **Copy token information**.

Important

Save the credential ID and API secret key immediately. The API secret key is shown only once and cannot be viewed again.

The new credential appears in the API credentials list. The API secret key is not available after credential creation.

A credential can have one of the following statuses:

- **Active.** The credential can be used for authentication.
- **Expiring soon.** The credential expires in less than 30 days.
- **Expired.** The credential can no longer be used for authentication.

To use an API credential

Set the following environment variables:

```
# export OS_AUTH_TYPE=v3applicationcredential
# export OS_AUTH_URL=<auth_url>
# export OS_APPLICATION_CREDENTIAL_ID=<credential_ID>
```

```
# export OS_APPLICATION_CREDENTIAL_SECRET=<api_secret_key>
# export OS_IDENTITY_API_VERSION=3
```

Then run OpenStack commands. For example:

```
# openstack --insecure server list
```

If the API endpoint uses a certificate trusted by your client, omit the `--insecure` option.

To revoke an API credential

1. On the **Settings > API credentials** screen, select the required credential.
2. On the credential right pane, click **Delete**.
3. Click **Delete** in the confirmation window.

Applications using the revoked credential can no longer authenticate. Expired credentials are also rejected on API calls.

Managing compute resources

This section explains how to manage compute resources in the self-service panel.

Note

If you have read-only access to a project, a **Read-only mode** label is displayed in the upper-right corner of the panel. This means you can view resources but cannot perform any management actions.

Managing virtual machines

Each virtual machine (VM) is an independent system with an independent set of virtual hardware. Its main features are the following:

- A virtual machine resembles and works like a regular computer. It has its own virtual hardware. Software applications can run in virtual machines without any modifications or adjustment.
- Virtual machine configuration can be changed easily, for example, by adding new virtual disks or memory.
- Although virtual machines share physical hardware resources, they are fully isolated from each other (file system, processes, sysctl variables) and the compute node.
- A virtual machine can run any supported guest operating system.

The following table lists the current virtual machine configuration limits:

Resource	Limit
RAM	1 TiB
CPU	64 virtual CPUs
Storage	15 volumes, 512 TiB each
Network	15 NICs

Supported guest operating systems

The guest operating systems listed below have been tested and are supported in virtual machines.

Note

Only the x64 architecture is supported.

Windows

Version	Edition	CPU hot plug support	RAM hot plug support
Windows Server 2025	Essentials	No	No
	Standard, Datacenter	Yes	Yes
Windows Server 2022	Essentials	No	No
	Standard, Datacenter	Yes	Yes
Windows Server 2019	Essentials	No	No
	Standard, Datacenter	Yes	Yes
Windows Server 2016	Essentials	No	No
	Standard, Datacenter	Yes*	Yes
Windows Server 2012 R2	Essentials, Standard, Datacenter	Yes	Yes
Windows Server 2012	Standard, Datacenter	Yes	Yes
Windows Server 2008 R2	Standard, Datacenter	No	No
Windows 11	Home, Professional, Enterprise	No	No
Windows 10	Home, Professional, Enterprise, Enterprise 2016 LTSB	No	No
Windows 8.1	Home, Professional, Enterprise	No	No

* CPU hot plug does not work properly due to a Windows bug with a wrongly installed driver. To fix the issue, refer to [this solution](#).

Note

For a Windows in-place upgrade to work, install the guest tools inside a virtual machine and restart it, as described in "Installing guest tools" (p. 30).

Linux

Distribution	Version	CPU hot plug support	RAM hot plug support
Rocky Linux	10.x, 9.x, 8.x	Yes	Yes
AlmaLinux	10.x, 9.x, 8.x	Yes	Yes

Distribution	Version	CPU hot plug support	RAM hot plug support
CentOS	Stream 10.x, 9.x, 8.x, 7.x	Yes	Yes
	6.x	No	No
Red Hat Enterprise Linux	10.x, 9.x, 8.x, 7.x	Yes	Yes
Debian	13.x, 12.x, 11.x, 10.x, 9.x	Yes	Yes
Ubuntu	24.04.x LTS, 22.04.x, 20.04.x, 18.04.x	Yes	Yes
Oracle Linux	10.x, 9.x, 8.x, 7.3, 7.9	Yes	Yes
SUSE Linux Enterprise	15.x (SP3–SP7)	Yes	Yes

Creating virtual machines

Limitations


- UEFI boot is not supported for CentOS 7.x virtual machines with less than 1 GiB of RAM.

Prerequisites


- You have a guest OS source prepared, as described in "Managing images" (p. 60).
- One or more compute networks are created by using the instructions in "Managing virtual networks" (p. 76).
- [Optional] Custom security groups are configured, as instructed in "Managing security groups" (p. 92).
- [Optional] An SSH key is added, as outlined in "Managing SSH keys" (p. 110). You can specify an SSH key only when creating VMs from a template or boot volume.

To create a virtual machine

1. On the **Virtual machines** screen, click **Create virtual machine**. A window will open where you will need to specify the VM parameters.
2. Specify a name for the new VM.
3. Select the VM boot media:
 - If you have an ISO image or a template
 - a. Select **Image** in the **Deploy from** section, and then click **Specify** in the **Image** section.
 - b. In the **Images** window, select the ISO image or template, and then click **Done**.



- If you have a compute boot volume
 - a. Select **Volume** in the **Deploy from** section, and then click **Specify** in the **Volumes** section.
 - b. In the **Volumes** window, click **Attach**.
 - c. In the **Attach volume** window, find and select the volume, and then click **Attach**.



If you attach more than one volume, the first attached volume becomes the boot volume, by default. To select another volume as bootable, place it first in the list by clicking the up arrow button next to it.

Note

If you select an image or volume with an assigned placement, the created VM will also inherit this placement.

After selecting the boot media, volumes required for this media to boot will be automatically added to the **Volumes** section.

4. Configure the VM disks:
 - a. In the **Volumes** window, make sure the default boot volume is large enough to accommodate the guest OS. Otherwise, click the ellipsis icon next to it, and then **Edit**. Change the volume size and click **Save**.
 - b. [Optional] Add more disks to the VM by creating or attaching volumes. To do this, click the pencil icon in the **Volumes** section, and then **Add** or **Attach** in the **Volumes** window.
 - c. Select volumes that will be removed during the VM deletion. To do this, click the pencil icon in the **Volumes** section, click the ellipsis icon next to the needed volume, and then **Edit**. Enable **Delete on termination** and click **Save**.
 - d. When you finish configuring the VM disks, click **Done**.
5. Choose the amount of RAM and CPU resources that will be allocated to the VM in the **Flavor** section. In the **Flavor** window, select a flavor, and then click **Done**.

Important

When choosing a flavor for a VM, ensure it satisfies the hardware requirements of the guest OS.

Note

To select a flavor with an assigned placement, you can filter flavors by placement. The VM created from such a flavor will also inherit this placement.

Flavor ×

Search Filter by placements:

Name ↓	vCPU ↓	Memory	Placement
 tiny	1	512 MiB	—
 small	1	2 GiB	placement1
 medium	2	4 GiB	placement1
 large	4	8 GiB	—
 xlarge	8	16 GiB	—

6. Add network interfaces to the VM in the **Networks** section:
 - a. In the **Network interfaces** window, click **Add** to attach a network interface.
 - b. In the **Add network interface** window, select a compute network to connect to, and then specify MAC address, IPv4 and/or IPv6 addresses, and security groups. By default, MAC and primary IP addresses are assigned automatically. To specify them manually, clear the **Assign automatically** check boxes, and enter the desired addresses. Optionally, assign additional IP addresses to the network interface in the **Secondary IP addresses** section. Note that a secondary IPv6 address is not available for an IPv6 subnet that works in the SLAAC or DHCPv6 stateless mode.

Note

Secondary IP addresses, unlike the primary one, will not be automatically assigned to the network interface inside the virtual machine guest OS. You should assign them manually.

- If you selected a virtual network with enabled IP address management
In this case, spoofing protection is enabled and the **default** security group is selected by default. This security group allows all incoming and outgoing traffic on all the VM ports. If required, you can select another security group or multiple security groups.
To disable spoofing protection, clear all of the check boxes and turn off the toggle switch. Security groups cannot be configured with disabled spoofing protection.
- If you selected a virtual network with disabled IP address management

In this case, spoofing protection is disabled by default and cannot be enabled. Security groups cannot be configured for such a network.

- If you selected a shared physical network

In this case, spoofing protection cannot be configured by a self-service user. If you want to enable or disable spoofing protection, contact your system administrator.

Add network interface

Network
net1: 10.136.16.0/22, 2001:bd8::/64

MAC address
Auto

☒ Assign automatically

Primary IP address ⓘ

IPv4: Assign automatically

☒ Assign automatically

ⓧ

Add

Secondary IP addresses ⓘ

IPv4 addresses

Add

Security groups
default

☐ Spoofing protection

Cannot configure spoofing protection if at least one security group is selected.

Cancel

Add

After specifying the network interface parameters, click **Add**. The network interface will appear in the **Network interfaces** list.

- [Optional] If required, edit IP addresses and security groups of newly added network interfaces. To do this, click the ellipsis icon, click **Edit**, and then set the parameters.
 - When you finish configuring the VM network interfaces, click **Done**.
7. [Optional] If you have chosen to boot from a template or volume, which has cloud-init and OpenSSH installed:

Important

Cloud images do not have a default password. By default, you can access VMs deployed from them using SSH key authentication. Alternatively, you can configure user credentials and enable password-based access using cloud-init.

- Add an SSH key to the VM, to be able to access it via SSH without a password.
In the **Select an SSH key** window, select an SSH key and then click **Done**.

Select an SSH key

×

Search

+

Add

Name ↑	Description ↑	Created on	
root_node001vstoragedom	My public key	June 11, 2019 11:34 AM	...

ⓘ

To be able to manage SSH keys, make sure the VM template has cloud-init installed.

Cancel

Done

- Add user data to customize the VM after launch, for example, change a user password.
Write a cloud-config or shell script in the **Customization script** field or browse a file on your local server to load the script from.

Provide a customization script

×

Provide user data to customize the VM after launch. User data can be in one of two formats: cloud-config or shell script. For the guest OS to be customizable, the template must have cloud-init installed.

Customization script

```
#cloud-config
user: myuser
password: password
chpasswd: {expire: False}
ssh_pwauth: True
```

Load from file

user-data

Browse

Cancel

Save

17

Copyright © 2016-2026 Virtuozzo International GmbH

To inject a script in a Windows VM, refer to the [Cloudbase-Init documentation](#). For example, you can set a new password for the account using the following script:

```
#ps1
net user <username> <new_password>
```

8. Specify the boot parameters:

- [For Windows 11 and Windows Server 2025] Select **Secure Boot** to ensure that only trusted software can run during VM startup.
- Select **UEFI boot** to configure the VM to boot in UEFI mode.
- Select **vTPM** to enable a virtual Trusted Platform Module (TPM) that provides enhanced security to the guest OS. Enabling vTPM also enables UEFI boot, since vTPM requires UEFI mode.

Note

UEFI boot and vTPM can be configured when creating a VM from an ISO, but are inherited when creating a VM from a template or volume. For Windows 11 and Windows Server 2025, these parameters are automatically enabled and cannot be disabled.

9. [Optional] Enable CPU and RAM hot plug for the VM in **Advanced options**, to be able to change its flavor when the VM is running. You can also enable hot plug after the VM is created.

Note

If you do not see this option, CPU and RAM hot plug is disabled in your project. To enable it, contact your system administrator.

10. After configuring all of the VM parameters, click **Deploy** to create and boot the VM.

If you are deploying the VM from an ISO image, you need to install the guest OS inside the VM by using the built-in VNC console. For VMs with UEFI boot enabled, open the VNC console, and then press any key to boot from the chosen ISO image. Virtual machines created from a template or a boot volume already have a preinstalled guest OS.

When you deploy a virtual machine with cloud-init installed, you can verify that cloud-init has completed successfully by running the following command inside the guest OS:

```
cloud-init status --long
```

If cloud-init has finished successfully, the status will be reported as done.

Connecting to virtual machines

Prerequisites

- Virtual machines are created, as described in "Creating virtual machines" (p. 13).
- To be able to connect via SSH, the virtual machine must have cloud-init and OpenSSH installed.

To connect to a virtual machine via the VNC console

Select a VM, and then click **Console** on its right pane. The console will open in a separate browser window. In the console, you can send a key combination to a VM, take a screenshot of the console window, and download the console log (refer to "Troubleshooting virtual machines" (p. 32)).

To connect to a virtual machine via SSH

Note

For Linux cloud images, access credentials are configured during deployment using cloud-init or SSH key injection. For details, refer to "Preparing Linux templates" (p. 63).

Specify the username and VM IP address in the SSH terminal:

```
# ssh <username>@<VM_IP_address>
```

Linux cloud images may define a default user name, depending on the operating system (for example, centos or ubuntu). However, access credentials must be configured during deployment using cloud-init or SSH keys.

To connect to a Windows VM, enter the username that you specified during Cloudbase-Init installation.

If you have deployed a VM without specifying an SSH key, you also need to enter a password to log in to the VM.

Setting a password inside virtual machines

Instead of an SSH key, you can use a password of the default administrator, to access a virtual machine created from a template.

Setting a password inside virtual machines is supported for both Linux and Windows guest operating systems.

Note

If you do not have this functionality available for a virtual machine, contact your system administrator.

To set a password inside a virtual machine

1. On the **Virtual machines** screen, click the required VM.
2. On the VM right pane, click **Set password**.
3. In the **Set password** window, specify a password for the default administrator login. The password must meet the following complexity requirements:
 - It must be at least 12 characters long.
 - It must contain characters from all of the following categories:
 - Uppercase Latin letters
 - Lowercase Latin letters

- Base 10 digits (0 through 9)
- Non-alphanumeric characters (special characters)

Alternatively, click **Generate** to automatically generate a random password and copy it to the clipboard.

Important

Save this password. After closing this window, the password will be hidden and unavailable for recovery.

Set password

Specify a password for the default login.

Password
GK11rQLePcHYVgE+

Generate

i The password must be at least 12 characters long. It must contain characters from all of these categories: uppercase and lowercase letters, digits, and special characters.
Save this password, as it will be hidden and unavailable for recovery.

Cancel Set

4. Click **Set** to set the specified password for the default administrator account inside the VM.

Once the password is injected inside the virtual machine, you can use it to log in to the guest operating system with the default admin login. The **Default admin login** is displayed on the VM right pane in the VM properties.

Managing virtual machine power state

Prerequisites

- Virtual machines are created, as described in "Creating virtual machines" (p. 13).

To manage the power state of a virtual machine

Click the virtual machine or the ellipsis button next to it to see the full list of actions available for the current state.

- To power up a VM, click **Run**.
- To gracefully shut down a running VM, click **Shut down**. The default shutdown timeout, after which a virtual machine will be powered off, is 10 minutes.
- To forcibly cut off power from a VM, click **Power off**.
- To softly reboot a running VM, click **Reboot**.
- To reboot a VM without the guest OS graceful shutdown, click **Hard reboot**.
- To save the current VM state to a file, click **Suspend**. This may prove useful, for example, if you need to restart the host but do not want to quit the applications currently running in the VM or restart its guest OS.
- To restore a VM from the suspended state, click **Resume**.

Attaching ISO images to virtual machines

You can attach ISO images to running or stopped virtual machines, for example, to install additional software inside them or to restore their operating system in the rescue mode. To attach an ISO image, you need to convert it to a volume, and then attach this volume to a VM.

When you finish installing software from an ISO volume, you can detach it without stopping the VM first.

To create a volume from an ISO image

1. On the **Images** screen, click the required ISO image.
2. On the image right pane, click **Create volume**.
3. In the **Create volume from image** window, specify a name for the volume, and then click **Create**.

To attach an ISO volume to a virtual machine

1. On the **Virtual machines** screen, click the required VM.
2. On the **Overview** tab, click the pencil icon in the **Volumes** field.
3. In the **Volumes** window, click **Attach**.
4. In the **Attach volume** window, select the created volume, and then click **Attach**. The attached volume will be marked as ISO.
5. In the **Volumes** window, click **Done** to save your changes.

The attached volume will appear inside the VM operating system.

To detach an ISO volume from a virtual machine

1. On the **Virtual machines** screen, click the required VM.
2. On the **Overview** tab, click the pencil icon in the **Volumes** field.
3. In the **Volumes** window, click the ellipsis icon next to the ISO volume, and then click **Force detach**.
4. Click **Done** to save your changes.

Reconfiguring virtual machines

Once you create a virtual machine, you can manage its CPU and RAM resources, as well as network interfaces and volumes.

Prerequisites

- Virtual machines are created, as described in "Creating virtual machines" (p. 13).

Changing virtual machine resources

You can change amount of CPU and RAM resources used by a virtual machine by applying another flavor to it. To be able to resize a running VM, you need to enable CPU and RAM hot plug for it first. You can change the hot plug settings for both new and existing VMs.

A running virtual machine has a resize limit, which defines the maximum number of vCPUs and the maximum amount of RAM you can allocate to the VM. The resize limit on vCPUs is static and equal to 64 for all VMs. The resize limit on RAM, on the contrary, is dynamic and depends on the amount of RAM a running VM is currently using. This limit is updated on a VM startup, and its values are listed in the table below.

Current RAM size, in GiB	RAM size limit, in GiB
1-4	16
5-8	32
9-16	64
17-32	128
33-64	256
65-128	512
129-256	1024

For example, you can resize a running VM with a flavor that has 16 GiB to a flavor with 256 GiB in two iterations:

1. Resize the VM to a flavor with 64 GiB.
2. Restart the VM to update the RAM size limit.
3. Resize the VM to a flavor with 256 GiB.

Limitations

- You cannot change the flavor for shelved VMs. To resize such a VM, unshelve it first.
- You cannot decrease the number of CPUs and the amount of RAM for running VMs.
- [For all Linux guests] If a VM has no guest tools installed, new cores may be offline after CPU hot plugging

You can verify which CPU cores are online by using the command:

```
# cat /sys/devices/system/cpu/online
```

To activate offline CPU cores, run:

```
# echo 1 > /sys/devices/system/cpu/cpu<cpu_number>/online
```

Prerequisites

- Resizing a VM requires the full set of resources defined by the target flavor to be available, not only the difference between the current and target flavors. Ensure that project quotas and available node resources can satisfy the target flavor, including vCPUs, RAM, PCI devices, and other extra specifications. For example, to resize a VM to the **large** flavor, the host must have at least 4 vCPUs and 8 GiB of RAM free.
- CPU and RAM hot plug is enabled by the system administrator.
- Before resizing a running VM, ensure that the guest operating system supports CPU and RAM hot plug (refer to "Supported guest operating systems" (p. 11)). Note that otherwise the guest operating system may become unstable after a resize. To increase CPU or RAM resources for such a guest operating system, you need to stop the virtual machine first.
- Before resizing a running VM, ensure that the guest operating system has the latest updates installed.

To enable or disable CPU and RAM hot plug for a virtual machine

1. On the **Virtual machines** screen, ensure that the required virtual machine in the "Shut down" state, and then click it.
2. On the **Overview** tab, click the pencil icon in the **CPU and RAM hot plug** field.

Note

If you do not see this field, CPU and RAM hot plug is disabled in your project. To enable it, contact your system administrator.

3. Select or clear the **Enable hot plug** check box, and then click the tick icon to save the changes.

With CPU and RAM hot plug enabled, you can change the flavor of a running VM.

To change the virtual machine flavor

1. On the **Virtual machines** screen, click the required virtual machine.
2. On the **Overview** tab, click the pencil icon in the **Flavor** field.
3. In the **Flavor** window, select a new flavor, and then click **Done**.

Configuring network interfaces of virtual machines

You can add new network interfaces to your virtual machines, edit IP addresses and security groups for the existing interfaces, and remove network interfaces by detaching them.

Limitations

- You cannot manage network interfaces of shelved VMs.
- A VM that is connected to a dual-stack network always receives an IPv6 address, if the IPv6 subnet is in the SLAAC or DHCPv6 stateless mode.

To attach a network interface to a virtual machine

1. On the **Virtual machines** screen, click the required virtual machine.
2. On the **Overview** tab, click **Edit** in the **Network interfaces** section.
3. In the **Network interfaces** window, click **Add** to attach a network interface.
4. In the **Add network interface** window, select a compute network to connect to, and then specify MAC address, IPv4 and/or IPv6 addresses, and security groups. By default, MAC and primary IP addresses are assigned automatically. To specify them manually, clear the **Assign automatically** check boxes, and enter the desired addresses. Optionally, assign additional IP addresses to the network interface in the **Secondary IP addresses** section. Note that a secondary IPv6 address is not available for an IPv6 subnet that works in the SLAAC or DHCPv6 stateless mode.

Note

Secondary IP addresses, unlike the primary one, will not be automatically assigned to the network interface inside the virtual machine guest OS. You should assign them manually.

- If you selected a virtual network with enabled IP address management
In this case, spoofing protection is enabled and the **default** security group is selected by default. This security group allows all incoming and outgoing traffic on all the VM ports. If required, you can select another security group or multiple security groups.
To disable spoofing protection, clear all of the check boxes and turn off the toggle switch. Security groups cannot be configured with disabled spoofing protection.
- If you selected a virtual network with disabled IP address management
In this case, spoofing protection is disabled by default and cannot be enabled. Security groups cannot be configured for such a network.
- If you selected a shared physical network
In this case, spoofing protection cannot be configured by a self-service user. If you want to enable or disable spoofing protection, contact your system administrator.

After specifying the network interface parameters, click **Add**.

5. Click **Done** to finish editing VM network interfaces and save your changes.

To edit a network interface of a virtual machine

1. On the **Virtual machines** screen, click the required virtual machine.
2. On the **Overview** tab, click **Edit** in the **Network interfaces** section.
3. In the **Network interfaces** window, click the ellipsis button next to the interface you want to edit, and then click **Edit**.
4. In the **Edit network interface** window, modify the network interface parameters as follows:

- Change the primary IP address. To update the address inside the VM guest OS, restart the network interface.
- Add or remove secondary IP addresses.
- Modify security groups assigned to the VM.

After updating the required parameters, click **Save**.

5. Click **Done** to finish editing VM network interfaces and save your changes.

To detach a network interface from a virtual machine

1. On the **Virtual machines** screen, click the required virtual machine.
2. On the **Overview** tab, click **Edit** in the **Network interfaces** section.
3. In the **Network interfaces** window, click the ellipsis button next to the interface you want to detach, and then click **Remove**.
4. Click **Done** to finish editing VM network interfaces and save your changes.

Configuring virtual machine volumes

You can add new volumes to your virtual machines, attach existing volumes, and detach unneeded volumes from virtual machines.

Limitations

- You cannot change, detach, or delete the boot volume.
- You can only attach and detach non-boot volumes.
- You cannot manage volumes of shelved VMs.

Prerequisites

- To be able to use volumes attached to VMs, they must be initialized inside the guest OS by standard means.

To attach a volume to a virtual machine

1. On the **Virtual machines** screen, click the required virtual machine.
2. On the **Overview** tab, click the pencil icon in the **Disks** field.
3. In the **Volumes** window:
 - Click **Attach** to attach an existing volume, and then select the volume in the **Attach volume** window.
 - Click **Add** to create a new volume, and then specify the volume name, size, and storage policy. The created volume will be automatically added to the VM disks.
4. Click **Done** to finish editing VM disks and save your changes.

To detach a volume from a virtual machine

1. On the **Virtual machines** screen, click the required virtual machine.
2. On the **Overview** tab, click the pencil icon in the **Disks** field.
3. In the **Volumes** window:

- Click **Detach** to detach a volume from a stopped virtual machine.
- Click **Force detach** to detach a volume from a running virtual machine.

Warning!

There is a risk of data loss.

4. Click **Done** to finish editing VM disks and save your changes.

Monitoring virtual machines

Prerequisites

- Virtual machines are created, as described in "Creating virtual machines" (p. 13).

To monitor virtual machine's CPU, storage, and network usage

Select a virtual machine and open the **Monitoring** tab. The following performance charts are available for virtual machines:

CPU

VM CPU usage from the hypervisor's perspective. The chart shows the CPU usage of the VM processes on the node, not the CPU usage within the VM.

RAM

VM memory usage from the hypervisor's perspective. The chart is based on the resident memory of the QEMU process on the node, not on the memory usage reported inside the guest OS. If the guest OS does not release unused memory to the hypervisor, the chart may show usage close to the amount of RAM allocated to the VM.

Network interface: <network_name> / MAC: <mac_address>

The VM interface parameters:

- **Speed:** the port transmit (TX) and receive (RX) speed, in bytes per second.
- **Packets:** the number of TX and RX packets per second on the port.
- **Drop rate:** the number of TX and RX packets dropped per second on the port.

To see a list of VM interfaces and hide all other charts, click **Only network interfaces** above the charts.

Volume: <volume_name> / <volume_id>

The VM disk parameters:

- **Storage read:** the amount of data being read by the VM, in bytes and operations per second.
- **Storage write:** the amount of data being written by the VM, in bytes and operations per second.

- **Read latency:** the disk latency while reading data. Hovering the mouse cursor over a point on the chart, you can also see the average and maximum latency for that moment, as well as the 95 and 99 percentiles.
- **Write latency:** the disk latency while writing data. Hovering the mouse cursor over a point on the chart, you can also see the average and maximum latency for that moment, as well as the 95 and 99 percentiles.
- **Flush latency:** the disk latency while flushing data. Hovering the mouse cursor over a point on the chart, you can also see the average and maximum latency for that moment, as well as the 95 and 99 percentiles.

To see a list of VM disks and hide all other charts, click **Only volumes** above the charts.

Note

Averaged values are calculated every five minutes.

The default time interval for the charts is twelve hours. To zoom into a particular time interval, select the interval with the mouse; to reset zoom, double-click any chart.

Shelving virtual machines

You can unbind a stopped VM from the node it is hosted on and release its reserved resources such as CPU and RAM. A shelved VM remains bootable and retains its configuration, including the IP addresses.

Prerequisites

- Virtual machines are created, as described in "Creating virtual machines" (p. 13).

To shelve a virtual machine

1. Click the desired virtual machine.
2. If the VM is stopped, click **Shelve** on its right pane.
3. If the VM is running or suspended, click **Shut down** or **Power off** on its right pane, and then select **Shelve virtual machine** in the confirmation window.

To spawn a shelved VM on a node with enough resources to host it

1. Click a shelved virtual machine.
2. On the VM right pane, click **Unshelve**.

Rescuing virtual machines

If a VM experiences boot problems, you can send it to rescue mode to access its boot volume. When a VM in the "Active" state is sent to rescue mode, it is shut down softly first. Once the VM is in rescue mode, you can connect to it via SSH or via the console. Its previous boot disk is now attached as a secondary one. You can mount the disk and repair it.

Limitations

- Rescue mode can use ISO images for booting both Linux and Windows virtual machines and QCOW2 images (templates) for booting Linux VMs. For instructions on making templates, refer to "Preparing templates" (p. 63).
- You can send a VM to rescue mode only if its current status is "Active" or "Shut down".
- There are only three actions available for the VM in rescue mode: **Console**, **Exit rescue mode**, and **Delete**.
- If a rescue image has cloud-init installed, then the VM booted from it can be accessed with the same SSH key that was used for its creation.
- Rescue mode is not supported for virtual machines with vTPM enabled.

Prerequisites

- Virtual machines are created, as described in "Creating virtual machines" (p. 13).

To put a virtual machine to rescue mode

1. On the **Virtual machines** screen, click the required VM on the list.
2. On the VM right pane, click the ellipsis button on the toolbar. Then, click **Enter rescue mode**.
3. In the **Enter rescue mode** window, select an image to rescue the VM with. By default, the initial image used for creating the VM is selected. Click **Enter**.

Enter rescue mode

Select an ISO image or template to rescue the virtual machine "centos7" with.

☒ ISO image
 ☐ Template

ISO image

centos-7-min.iso

▼

Cancel

Enter

The machine status changes to "Rescue".

To return a virtual machine to normal operation

1. On the **Virtual machines** screen, click the required VM on the list.
2. On the VM right pane, click **Exit rescue mode**.
3. In the **Exit rescue mode** window, click **Exit**. The VM will be automatically rebooted.

The VM status changes to "Active" and it boots from the original root disk.

Note

If the VM status changes to “Error” when exiting rescue mode, you can reset its status with the **Reset state** action. The VM should then return to the “Rescue” status again.

To exit rescue mode for a Windows VM

There might be an issue of exiting rescue mode for a Windows VM. If in rescue mode you set the original system disk online, its ID becomes the same as that of the rescue disk. Then, when you try to exit rescue mode, the boot loader cannot find the proper boot disk. To resolve the ID conflict, follow the steps:

1. With the VM in rescue mode, open the **Disk Management** window and note the numbers of the original system disk (offline) and the rescue disk (online). Set the original system disk to **Online**.
2. To edit the boot configuration, enter the following command in the **Command Prompt** window:

```
> bcdedit /store <the original system disk name>:\boot\bcd
```

3. Review the output and check that the rescue disk is the target for objects in the output (partition=<the rescue disk name>).

If the objects do not point to drive C, fix it with the following commands:

```
> bcdedit /store <the original system disk name>:\boot\bcd \
/set {default} osdevice partition=<the rescue disk name>:
> bcdedit /store <the original system disk name>:\boot\bcd \
/set {default} device partition=<the rescue disk name>:
> bcdedit /store <the original system disk name>:\boot\bcd \
/set {bootmgr} device partition=<the rescue disk name>:
> bcdedit /store <the original system disk name>:\boot\bcd \
/set {memdiag} device partition=<the rescue disk name>:
```

4. To view the available disks, enter the following commands in the command line:

```
> DISKPART
> LIST DISK
```

Match the disk number and name to those displayed in the **Disk Management** window.

5. To get the ID of the rescue disk, run the following commands:

```
> SELECT DISK <the rescue disk number>
> UNIQUEID DISK
```

Record the disk ID, you will need it later.

6. Change this ID by using the following command:

```
> UNIQUEID DISK id=<any hex value of 8 characters>
```

Make sure that the value has changed with the UNIQUEID DISK command.

7. Assign the ID that you recorded previously to the original system disk:

```
> SELECT DISK <the original system disk number>
> UNIQUEID DISK id=<the recorded disk ID>
```

Make sure that the value has changed with the `UNIQUEID DISK` command.

You should now be able to exit rescue mode.

Managing guest tools

This section explains how to install and uninstall the guest tools. This functionality is required for creating consistent snapshots of a running VM's disks.

Limitations

- Guest tools rely on the QEMU guest agent that is installed alongside the tools. The agent service must be running for the tools to work.

Prerequisites

- Virtual machines are created, as described in "Creating virtual machines" (p. 13).
- The virtual machine has a guest operating system installed.

Installing guest tools

1. Create a compute volume from the **vz-guest-tools-win** or **vz-guest-tools-lin** image, depending on the VM operating system:

Note

If you do not have these images in your project, contact your system administrator.

- a. On the **Images** screen, click the **vz-guest-tools-win** or **vz-guest-tools-lin** image.
 - b. On the image right pane, click **Create volume**.
 - c. In the **Create volume from image** window, specify a name for the volume, and then click **Create**.
2. Attach the volume with the guest tools to the virtual machine:
 - a. On the **Virtual machines** screen, click the required VM.
 - b. On the VM right pane, click the pencil icon in the **Volumes** field.
 - c. In the **Volumes** window, click **Attach**.
 - d. In the **Attach volume** window, select the created volume with the guest tools, and then click **Attach**. The attached volume will be marked as ISO.
 - e. In the **Volumes** window, click **Done**, to save your changes.
 3. Log in to the virtual machine.
 4. Inside the VM, do the following:
 - Inside a Windows VM, go to the mounted optical drive in Explorer and install the guest tools by running `setup.exe`. After the installation is complete, restart the VM.

- Inside a Linux VM, create a mount point for the optical drive with the guest tools image and run the installer:

```
# mkdir /mnt/cdrom
# mount <path_to_guest_tools_iso> /mnt/cdrom
# bash /mnt/cdrom/install
```

Uninstalling guest tools

If you find out that the guest tools are incompatible with some software inside a virtual machine, you can uninstall them by doing the following:

- Inside a Windows VM:
 1. Remove the QEMU device drivers from the device manager.

Important

Do not remove the VirtIO/SCSI hard disk driver and NetKVM network driver. Without the former, the VM will not boot; without the latter, the VM will lose network connectivity.

2. Uninstall the QEMU guest agent and guest tools from the list of installed applications.
3. Stop and delete **Guest Tools Monitor**:

```
> sc stop VzGuestToolsMonitor
> sc delete VzGuestToolsMonitor
```

4. Unregister **Guest Tools Monitor** from **Event Log**:

```
> reg delete HKLM\SYSTEM\CurrentControlSet\services\eventlog\Application\VzGuestToolsMonitor
```

5. Delete the autorun registry key for **RebootNotifier**:

```
> reg delete HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Run /v \
VzRebootNotifier
```

6. Delete the C:\Program Files\Qemu-ga\ directory.

If VzGuestToolsMonitor.exe is locked, close all the Event Viewer windows. If it remains locked, restart the eventlog service:

```
> sc stop eventlog
> sc start eventlog
```

After removing the guest tools, restart the virtual machine.

- Inside a Linux VM:
 1. Remove the packages:
 - a. On RPM-based systems (CentOS and other):

```
# yum remove dkms-vzvirtio_balloon prl_nettool qemu-guest-agent-vz \
vz-guest-udev
```

- b. On DEB-based systems (Debian and Ubuntu):

```
# apt-get remove vzvirtio-balloon-dkms prl-nettool qemu-guest-agent-vz \
vz-guest-udev
```

If any of the packages listed above are not installed on your system, the command will fail. In this case, exclude these packages from the command and run it again.

2. Remove the files:

```
# rm -f /usr/bin/prl_backup /usr/share/qemu-ga/VERSION \
/usr/bin/install-tools \
/etc/udev/rules.d/90-guest_iso.rules /usr/local/bin/fstrim-static \
/etc/cron.weekly/fstrim
```

3. Reload the udev rules:

```
# udevadm control --reload
```

After removing guest tools, restart the virtual machine.

Troubleshooting virtual machines

If a virtual machine fails to deploy

Review the error message on the VM right pane. One of the possible root causes is that compute nodes lack free RAM or CPU resources to host the VM.

If a virtual machine fails to boot

Examine the VM console log by clicking **Download console log** on the VM right pane. The log will contain log messages only if logging is enabled inside the VM (refer to "Enabling logging for virtual machines" (p. 67)).

If a virtual machine is in the error state

Examine the VM history in the **History** tab on the VM right pane. The event log lists all VM management operations performed in the user or command-line interface within the last 60 days. You can expand each log entry to view operation details by clicking the arrow icon next to it. The details include the operation name, date and time, status, initiator, and request ID.

If a virtual machine is stuck in the "Spawning" task state

Delete the VM, and then try creating a new one.

1. Click the stuck VM.
2. On the VM right pane, click **Delete**.

If a virtual machine is stuck in a failed or transitional state

Reset the VM to its last stable state: active, shut down or shelved.

Important

Perform this operation only as a last resort.

1. Click the stuck VM.
2. On the VM right pane, click **Reset state**.

Deleting virtual machines

Limitations

- A VM is removed along with its disks that have the **Delete on termination** option enabled during the VM deployment.

Prerequisites

- Virtual machines are created, as described in "Creating virtual machines" (p. 13).

To remove one virtual machine

1. Click the ellipsis button next to a VM you want to delete, and then click **Delete**.
2. Click **Delete** in the confirmation window.

To remove multiple virtual machines

1. Select the check boxes next to VMs you want to delete.
2. Over the VM list, click **Delete**.
3. Click **Delete** in the confirmation window.

Managing Kubernetes clusters

Self-service users can deploy ready-to-use Kubernetes clusters with persistent storage for managing containerized applications.

Limitations

- Kubernetes cluster certificates are issued for five years. To renew the certificates, use the `openstack coe ca rotate` command, as described in the [OpenStack documentation](#).
- Starting with version 1.31.x, Envoy proxy can be installed automatically with Cilium if SELinux is disabled on Kubernetes nodes.

Supported Kubernetes and Fedora CoreOS versions

This section lists the Kubernetes and Fedora CoreOS versions supported in VirtuoZZo Cloud.

Version support policy

VirtuoZZo Cloud supports multiple Kubernetes versions simultaneously. Each release includes the latest patch release for every supported Kubernetes version and may introduce support for newer Kubernetes and Fedora CoreOS versions.

Supported Kubernetes and Fedora CoreOS versions are updated in product releases to stay close to upstream releases while maintaining compatibility with existing deployments.

A Kubernetes version will be deprecated when:

- it reaches upstream end of life;
- it falls outside the range that Virtuozzo Cloud can actively maintain; or
- it is no longer required to support supported upgrade paths.

Supported Kubernetes versions

In Virtuozzo Cloud 7.4, the following Kubernetes versions are supported:

Kubernetes version	Patch version	Supported since
1.35	1.35.3	7.4
1.34	1.34.6	7.4
1.33	1.33.10	7.4
1.32	1.32.13	7.4
1.31	1.31.2	6.3 Update 1

Deprecated Kubernetes versions

Kubernetes clusters created with deprecated versions are marked with the **Deprecated** tag. Upgrade such clusters to a supported Kubernetes version.

Kubernetes version	Supported since	Deprecated since
1.30	6.3	7.4
1.29	6.2	7.4
1.28	6.1	7.4
1.27	6.1	6.3 Update 1
1.26	6.1	6.3
1.25	5.4 Update 3	6.2 Update 1
1.24	5.3	6.2 Update 1
1.23	5.1	6.2
1.22	5.0	6.1
1.21 and earlier	See older release notes	

Kubernetes and Fedora CoreOS versions by release

The following table lists the Kubernetes and Fedora CoreOS versions included in VirtuoZZo Cloud releases.

Release	Kubernetes versions	Fedora CoreOS version
7.4	1.31.2, 1.32.13, 1.33.10, 1.34.6, 1.35.3	43
7.3	1.28.4, 1.29.3, 1.30.3, 1.31.2	41
7.2	1.28.4, 1.29.3, 1.30.3, 1.31.2	41
7.1	1.28.4, 1.29.3, 1.30.3, 1.31.2	41
7.0	1.28.4, 1.29.3, 1.30.3, 1.31.2	41
6.3 Update 1	1.28.4, 1.29.3, 1.30.3, 1.31.2	39
6.3	1.27.8, 1.28.4, 1.29.3, 1.30.3	39
6.2 Update 1	1.26.11, 1.27.8, 1.28.4, 1.29.3	39
6.2	1.24.3, 1.25.7, 1.26.11, 1.27.8, 1.28.4, 1.29.3	39
6.1	1.23.5, 1.24.3, 1.25.7, 1.26.11, 1.27.8, 1.28.4	39
6.0	1.22.2, 1.23.5, 1.24.3, 1.25.7	37

Creating and deleting Kubernetes clusters

When creating a Kubernetes cluster, you can specify supplementary parameters by using the following labels:

selinux_mode

Choose the SELinux mode. Possible values: enforcing, permissive, and disabled. The default is permissive. Starting with version 1.31.x, you can disable SELinux to automatically install Envoy proxy with Cilium.

cilium_ebpf_enabled [starting with version 1.30.x]

Choose whether to use eBPF-based host-routing to optimize the host-internal packet routing. The default is false.

cilium_ipv4pool [starting with version 1.30.x]

Configure the IP pool for assigning pod IPs. The default is 10.100.0.0/16.

cilium_ipv4pool_mask_size [starting with version 1.30.x]

Set the size of a subnet assigned to each minion. The default is 24.

cilium_routing_mode [starting with version 1.30.x]

Enable native-routing mode or tunneling mode. Possible values: tunnel and native. The default is tunnel.

cilium_hubble_enabled [starting with version 1.30.x]

Choose whether to enable Hubble, a fully distributed networking and security observability platform. Disabled by default to maximize performance. The default is false.

cilium_tag [starting with version 1.30.x]

Specify a tag for the operator and agent used to provision the Cilium node.

To see the full list of supported labels, refer to the [OpenStack documentation](#).

Limitations

- Only users that have access to the corresponding project can perform operations with Kubernetes clusters.

Prerequisites

- The Kubernetes-as-a-service component is installed by a system administrator. It can be deployed along with the compute cluster or later.
- You have a network that will interconnect the Kubernetes master and worker nodes. It can be either a shared physical network or a virtual network linked to a physical one via a virtual router. The virtual network needs to have a gateway and a DNS server specified.
- An SSH key is added. It will be installed on both the master and worker nodes.
- You have enough resources for all of the Kubernetes nodes, taking their flavors into account.
- It is also required that the network where you create a Kubernetes cluster does not overlap with these default networks:
 - 10.100.0.0/16—Used for pod-level networking
 - 10.254.0.0/16—Used for allocating Kubernetes cluster IP addresses

To create a Kubernetes cluster

1. Go to the **Kubernetes clusters** screen, and then click **Create** on the right. A window will open where you can set your cluster parameters.
2. Enter the cluster name, and then select a Kubernetes version and an SSH key.
3. In the **Network** section:
 - a. Select a network that will interconnect the Kubernetes nodes in the cluster.
 - b. When selecting a virtual network, decide whether you need access to your Kubernetes cluster via a floating IP address:
 - If you select **None**, you will not have access to the Kubernetes API.
 - If you select **For Kubernetes API**, a floating IP address will be assigned to the master node or to the load balancer if the master node is highly available.
 - If you select **For Kubernetes API and nodes**, floating IP addresses will be additionally assigned to all of the Kubernetes nodes (masters and workers).

- c. If you require access to the Kubernetes cluster and your virtual network is linked to multiple physical networks via routers, select the network to pick up a floating IP address from.
- d. Then, choose whether or not to enable **High availability** for the master node. If you enable high availability, three master node instances will be created. They will work in the Active/Active mode.

The screenshot shows a configuration panel with the following elements:

- A dropdown menu labeled "Network" with the selected value "private1 (192.168.128.0/24)".
- A dropdown menu labeled "Floating IP address" with the selected value "For Kubernetes API".
- A dropdown menu labeled "Floating IP network" with the selected value "public (10.136.16.0/20)".
- A checkbox labeled "High availability" which is checked.
- Below the checkbox, a text description: "Enable this feature to create three master nodes working in the Active/Active mode. Leave it disabled to create a single master node."

4. In the **Master node** section, select a flavor for the master node. For production clusters, it is strongly recommended to use a flavor with at least 2 vCPUs and 8 GiB of RAM.
5. Optionally, enable **Integrated monitoring** to automatically deploy the cluster-wide monitoring solution, which includes the following components: Prometheus, Alertmanager, and Grafana.

Warning!

This feature is experimental and not intended for use in production environments.

6. In the **Container volume** section, select a storage policy, and then enter the size for volumes on both master and worker nodes.

Important

After the cluster is created, you cannot change the volume size or storage policy for the master or worker nodes.

7. In the **Default worker group** section, select a flavor for each worker, and then decide whether you want to allow automatic scaling of the worker group:
 - With **Autoscaling** enabled, the number of workers will be automatically increased if there are pods stuck in the pending state due to insufficient resources, and reduced if there are workers with no pods running on them. For scaling of the worker group, set its minimum and maximum size.

Important

Some types of pods can prevent the autoscaler from removing a worker. To see a list of such pod types, refer to the official [Kubernetes Autoscaler documentation](#).

- With **Autoscaling** disabled, the number of worker nodes that you set will be permanent.

Default worker group

Flavor
small — 1 vCPU, 2 GiB RAM

☒ Autoscaling ⓘ

Minimum Maximum

— 1 + — 3 + Number of workers

8. In the **Labels** section, enter labels that will be used to specify supplementary parameters for this Kubernetes cluster as key/value pairs. For example: `selinux_mode=permissive`.
9. Click **Create**.

Creation of the Kubernetes cluster will start. The master and worker nodes will appear on the **Virtual machines** screen, while their volumes will show up on the **Volumes** screen.

After the cluster is ready, click **Kubernetes access** for instructions on how you can access the dashboard. You can also access the Kubernetes master and worker nodes via SSH, by using the assigned SSH key and the user name **core**.

To delete a Kubernetes cluster

Click the required Kubernetes cluster on the **Kubernetes clusters** screen and click **Delete**. The master and worker VMs will be deleted along with their volumes.

Managing Kubernetes worker groups

To meet system requirements of applications running in Kubernetes clusters, you can have worker nodes with different number of CPUs and amount of RAM. Creating workers with different flavors is possible by using worker groups.

When creating a Kubernetes cluster, you can specify the configuration of only one worker group, the default worker group. After the cluster is created, add as many worker groups as you need. If required, you can also edit the number of workers in a group later.

Limitations

- The default worker group in a Kubernetes cluster cannot be removed or replaced because it is part of the initial cluster stack. However, you can stop using it by cordoning and draining its node and letting the autoscaler scale it down to zero.

Prerequisites

- A Kubernetes cluster is created, as described in "Creating and deleting Kubernetes clusters" (p. 35).

To add a worker group

1. On the **Kubernetes clusters** screen, click a Kubernetes cluster.
2. On the cluster right pane, navigate to the **Groups** tab.
3. In the **Workers** section, click **Add**.
4. In the **Add worker group** window, specify a name for the group.
5. In the **Worker group** section, select a flavor for each worker, and then decide whether you want to allow automatic scaling of the worker group:
 - With **Autoscaling** enabled, the number of workers will be automatically increased if there are pods stuck in the pending state due to insufficient resources, and reduced if there are workers with no pods running on them. For scaling of the worker group, set its minimum and maximum size.

Important

Some types of pods can prevent the autoscaler from removing a worker. To see a list of such pod types, refer to the official [Kubernetes Autoscaler documentation](#).

- With **Autoscaling** disabled, the number of worker nodes that you set will be permanent.
6. In the **Labels** section, enter labels that will be used to specify supplementary parameters for this Kubernetes cluster as key/value pairs. For example: `selinux_mode=permissive`.
 7. Click **Add**.

Add workers
×

Name
mygroup

Worker group

Flavor
small — 1 vCPU, 2 GiB RAM

☒ Autoscaling ⓘ

Minimum

— 1 +

Maximum

— 3 +

Number of workers

Labels

With labels, you can specify supplementary parameters specific to certain Kubernetes clusters or associated with certain options. Labels are key/value pairs that are interpreted and validated by the drivers that use them.

Label
selinux_mode=permissive,cert_manager_api=true

Specify labels in the format: example1=true, example2=false

Cancel

Add

When the worker group is created, you can assign pods to these worker nodes, as explained in "Assigning Kubernetes pods to specific nodes" (p. 56).

To edit the number of workers in a group

1. On the Kubernetes cluster right pane, navigate to the **Groups** tab.
2. In the **Workers** section, click the pencil icon for the default worker group or the ellipsis icon for all other groups, and then select **Edit**.
3. In the **Edit workers** window, enable or disable **Autoscaling**, or change the number of workers in the group.
4. Click **Save**.

To delete a worker group

Click the ellipsis icon next to the required worker group, and then select **Delete**. The worker group will be deleted along with all of its workers. After the deletion, the worker group data will be lost.

To remove the default worker group

1. Enable the Kubernetes autoscaler and set the minimum node count of the default worker group to 0.

- a. On the Kubernetes cluster right pane, open the **Groups** tab.
- b. In the **Workers** section, click the pencil icon next to the default worker group and select **Edit**.
- c. In the **Edit workers** window, enable **Autoscaling** and set the number of workers in the group to 0.
- d. Click **Save**.

This allows the autoscaler to scale the group down automatically once the node is empty.

2. Prevent new workloads from being scheduled on the default worker node:

```
kubect1 cordon <node-name>
```

3. Drain existing workloads from the default node:

```
kubect1 drain <node-name> --ignore-daemonsets --delete-local-data --force
```

This evicts all pods (except DaemonSets) and prepares the node for removal.

4. Wait for the autoscaler to remove the node. After the last pod is evicted, the node will be automatically deleted after the standard autoscaler timeout (approximately 10 minutes).

Once the default worker node is removed, your new worker group will take over and run all workloads going forward.

Updating Kubernetes clusters

When a new Kubernetes version becomes available, you can update your Kubernetes cluster to it. An update is non-disruptive for Kubernetes worker nodes, which means that these nodes are updated one by one, with the data availability unaffected. The Kubernetes API will be unavailable during an update, unless high availability is enabled for the master node.

During the update procedure, Kubernetes virtual machines are re-created based on a newer Fedora CoreOS image. Such a rolling update is used to preserve the cluster data. Before starting the update, you need to make sure that the compute cluster has enough resources and quotas for at least one extra VM of the largest flavor used by your Kubernetes cluster. If the master and worker node flavors differ, then you should take into account the largest one of them.

Limitations

- You can update Kubernetes clusters only to the next minor version in one iteration. For example, to update a cluster from version 1.31 to 1.33, you need to update it to version 1.32 first.
- Kubernetes clusters can have only one minor version difference between node groups (for example, 1.31 and 1.32).
- You cannot manage Kubernetes clusters in the self-service panel during an update.

Prerequisites

- A Kubernetes cluster is created, as described in "Creating and deleting Kubernetes clusters" (p. 35).

- Knowledge of the supported Kubernetes and Fedora CoreOS versions listed in "Supported Kubernetes and Fedora CoreOS versions" (p. 33).

To update a Kubernetes cluster

1. Click a Kubernetes cluster that is marked with the **Update available** tag.
2. On the Kubernetes cluster pane, click **Update** in the **Kubernetes version** field.
3. In the **Update** window, do the following:
 - a. Select a Kubernetes version to update to and follow the provided link to read about API resources that are deprecated or obsoleted in the selected version. Then, click **Next**.
 - b. Choose how to proceed:
 - Select **Update all** to update all of the node groups in the Kubernetes cluster.
 - Select **Custom update** to update only specific node groups. The master node group is selected automatically and is mandatory for an update.

Update

×

Select node groups in the Kubernetes cluster "kube1" that you want to update to version "v1.28.4".

☐ Update all
 ☒ Custom update

☐	Group ↓	Nodes ↓	Version ↓
☒	default-master Management	1	v1.27.8
☒	default-worker	1	v1.27.8
☐	group1	1	v1.27.8
☐	group2	1	v1.27.8

Back

Update

4. In the confirmation window, click **Confirm**. The update process will start.

Warning!

Do not manage Kubernetes virtual machines during the update as it may lead to disruption of the update process and cluster inoperability.

When node groups in a Kubernetes cluster have different versions, the cluster tag changes to **Partially updated**. In this case, new worker groups will be created with the version of the master node group. To finish the Kubernetes cluster upgrade, you need to repeat the update procedure for worker groups with an older version.

Using persistent volumes for Kubernetes pods

Kubernetes allows using compute volumes as persistent storage for pods. Persistent volumes (PV) exist independently of pods, meaning that such a volume persists after the pod it is mounted to is deleted. This PV can be mounted to other pods for accessing data stored on it. You can provision PVs dynamically, without having to create them manually, or statically, using volumes that exist in the compute cluster.

Creating storage classes

In Virtuozzo Cloud, storage classes map to compute storage policies defined in the admin panel. Creating a storage class is required for all storage operations in a Kubernetes cluster.

To create a storage class

Click **+ Create** on the Kubernetes dashboard and specify a YAML file that defines this object. For example:

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: default
  annotations:
    storageclass.kubernetes.io/is-default-class: "true"
provisioner: cinder.csi.openstack.org
parameters:
  type: default
```

This manifest describes the storage class `default` with the storage policy `default`. It also marks this storage policy as default for the Kubernetes cluster. The storage policy must exist in the compute cluster and be specified in the storage quotas to the current project.

Dynamically provisioning persistent volumes

Persistent volumes can be dynamically provisioned via persistent volume claims (PVC). A PVC requests for a PV of a specific storage class, access mode, and size. If a suitable PV exists in the cluster, it is bound to the claim. If suitable PVs do not exist but can be provisioned, a new volume is created and bound to the claim. Kubernetes uses a PVC to obtain the PV backing it and mounts it to the pod.

Prerequisites

- A pod and the persistent volume claim it uses must exist in the same namespace.

To dynamically provision a PV to a pod

1. Access the Kubernetes cluster via the dashboard. Click **Kubernetes access** for instructions.
2. On the Kubernetes dashboard, create a storage class, as described in "Creating storage classes" (p. 43).

3. Create a persistent volume claim. To do it, click **+ Create** and specify the following YAML file:

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: mypvc
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 10Gi
  storageClassName: default
```

This manifest specifies the persistent volume claim `mypvc` that requests from the storage class `default` a volume of at least 10 GiB that can be mounted in the read/write mode by a single node.

Creation of the PVC triggers dynamic provisioning of a persistent volume that satisfies the claim's requirements. Kubernetes then binds it to the claim.

Metadata

Name	Namespace	Created	Age	UID
mypvc	default	Oct 31, 2023	5 minutes ago	12e3ae0c-fe23-43a2-93db-44cfd00fb40

Annotations

pv.kubernetes.io/bind-completed: yes

pv.kubernetes.io/bound-by-controller: yes

volume.beta.kubernetes.io/storage-provisioner: cinder.csi.openstack.org

volume.kubernetes.io/storage-provisioner: cinder.csi.openstack.org

Resource information

Status	Storage Class	Volume Name
Bound	default	pvc-12e3ae0c-fe23-43a2-93db-44cfd00fb40

Capacity

storage: 10Gi

Access Modes

ReadWriteOnce

4. Create a pod and specify the PVC as its volume. To do it, click **+ Create** and enter the following YAML file:

```
apiVersion: v1
kind: Pod
metadata:
  name: nginx
spec:
  containers:
    - image: nginx
      imagePullPolicy: IfNotPresent
      name: nginx
      ports:
        - containerPort: 80
          protocol: TCP
      volumeMounts:
        - mountPath: /var/lib/www/html
          name: mydisk
```

```
volumes:
- name: mydisk
  persistentVolumeClaim:
    claimName: mypvc
    readOnly: false
```


This configuration file describes the pod `nginx` that uses the persistent volume claim `mypvc`. The persistent volume bound to the claim will be accessible at `/var/lib/www/html` inside the `nginx` container.

Statically provisioning persistent volumes

You can mount existing compute volumes to pods using static provisioning of persistent volumes.

To mount a compute volume

1. In the self-service panel, obtain the ID of the desired volume.



2. Access the Kubernetes cluster via the dashboard. Click **Kubernetes access** for instructions.
3. On the Kubernetes dashboard, create a storage class, as described in "Creating storage classes" (p. 43).
4. Create a persistent volume. To do it, click **+ Create** and specify the following YAML file:

```
apiVersion: v1
kind: PersistentVolume
metadata:
  annotations:
    pv.kubernetes.io/provisioned-by: cinder.csi.openstack.org
  name: mypv
spec:
  accessModes:
    - ReadWriteOnce
  capacity:
    storage: 10Gi
  csi:
```

```

driver: cinder.csi.openstack.org
fsType: ext4
volumeHandle: c5850e42-4f9d-42b5-9bee-8809dedae424
persistentVolumeReclaimPolicy: Delete
storageClassName: default

```

This manifest specifies the persistent volume `mypv` from the storage class `default` that has 10 GiB of storage and access mode that allows it to be mounted in the read/write mode by a single node. The PV `mypv` uses the compute volume with the ID `c5850e42-4f9d-42b5-9bee-8809dedae424` as backing storage.

5. Create a persistent volume claim. Before you define the PVC, make sure the PV is created and has the status "Available". The existing PV must meet the claim's requirements to storage size, access mode and storage class. Click **+ Create** and specify the following YAML file:

```

apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: mypvc
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 10Gi
  storageClassName: default

```

Once the persistent volume claim `mypvc` is created, the volume `mypv` is bound to it.

Metadata

Name	Namespace	Created	Age	UID
mypvc	default	Oct 31, 2023	59 seconds ago	b27fac2a-fd97-41c0-9f06-fcdd1a1d4fe6

Annotations

pv.kubernetes.io/bind-completed: yes pv.kubernetes.io/bound-by-controller: yes

Resource information

Status	Storage Class	Volume Name
Bound	default	mypv

Capacity


storage: 10Gi

Access Modes

ReadWriteOnce

6. Create a pod and specify the PVC as its volume. Use the example from Step 4 in "Dynamically provisioning persistent volumes" (p. 43).

In the self-service panel, the compute volume will be mounted to the virtual machine running the Kubernetes pod.



Making Kubernetes deployments highly available

If a node that hosts a Kubernetes pod fails or becomes unreachable over the network, the pod is stuck in a transitional state. In this case, the pod's persistent volumes are not automatically detached, and it prevents the pod redeployment on another worker node. To make your Kubernetes applications highly available, you need to enforce the pod termination in the event of node failure by adding rules to the pod deployment.

To terminate a stuck pod

Add the following lines to the spec section of the deployment configuration file:

```
terminationGracePeriodSeconds: 0
tolerations:
- effect: NoExecute
  key: node.kubernetes.io/unreachable
  operator: Exists
  tolerationSeconds: 2
- effect: NoExecute
  key: node.kubernetes.io/not-ready
  operator: Exists
  tolerationSeconds: 2
```

If the node's state changes to "NotReady" or "Unreachable", the pod will be automatically terminated in 2 seconds.

The entire YAML file of a deployment may look as follows:

```
apiVersion: apps/v1
kind: Deployment
metadata:
```

```

name: nginx
spec:
  replicas: 1
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      terminationGracePeriodSeconds: 0
      tolerations:
        - effect: NoExecute
          key: node.kubernetes.io/unreachable
          operator: Exists
          tolerationSeconds: 2
        - effect: NoExecute
          key: node.kubernetes.io/not-ready
          operator: Exists
          tolerationSeconds: 2
      containers:
        - image: nginx
          imagePullPolicy: IfNotPresent
          name: nginx
          ports:
            - containerPort: 80
              protocol: TCP
          volumeMounts:
            - mountPath: /var/lib/www/html
              name: mydisk
      volumes:
        - name: mydisk
          persistentVolumeClaim:
            claimName: mypvc

```

The manifest above describes the deployment `nginx` with one pod that uses the persistent volume claim `mypvc` and will be automatically terminated in 2 seconds in the event of node failure.

Managing volume snapshots in Kubernetes

You can create a snapshot of a Kubernetes volume to copy its contents at a specific moment in time. This can be useful for restoring volume data in case of data loss.

Volume snapshots are based on [custom resource definitions](#) (CRD), so you need to add them to your Kubernetes cluster first.

Prerequisites

- A persistent volume claim is created, as described in "Dynamically provisioning persistent volumes" (p. 43).

To add custom resource definitions

Run the following commands:

```
# git clone https://github.com/kubernetes-csi/external-snapshotter/
cd ./external-snapshotter
# git checkout release-5.0
# kubectl apply -f client/config/crd/snapshot.storage.k8s.io_volumesnapshotclasses.yaml
# kubectl apply -f client/config/crd/snapshot.storage.k8s.io_volumesnapshotcontents.yaml
# kubectl apply -f client/config/crd/snapshot.storage.k8s.io_volumesnapshots.yaml
# kubectl apply -f deploy/kubernetes/snapshot-controller/rbac-snapshot-controller.yaml -n kube-system
# kubectl apply -f deploy/kubernetes/snapshot-controller/setup-snapshot-controller.yaml -n kube-system
```

Important

The CSI snapshotter version must not be higher than release-5.0.

These commands will create CRDs for VolumeSnapshotClass, VolumeSnapshotContent, VolumeSnapshot, as well as required ClusterRole, ServiceAccount, ClusterRoleBinding, Role, RoleBinding, and finally the snapshot-controller deployment.

You can check that the required resources are successfully created by using [Lens](#), an easy tool for managing Kubernetes clusters and resources.

To create a volume snapshot

1. Create the snapshot-class.yaml file that defines the VolumeSnapshotClass object:

```
cat > snapshot-class.yaml <<\EOT
apiVersion: snapshot.storage.k8s.io/v1
kind: VolumeSnapshotClass
metadata:
  name: mysnapclass
driver: cinder.csi.openstack.org
deletionPolicy: Delete
parameters:
  force-create: "true"
EOT
```

This manifest describes the volume snapshot class `mysnapclass` with the deletion policy `Delete`. This policy allows deleting the underlying storage snapshot along with the VolumeSnapshotContent object. To keep both the underlying snapshot and VolumeSnapshotContent when deleting the VolumeSnapshot object, set the deletion policy to `Retain`.

2. Create a snapshot class:

```
$ kubectl apply -f snapshot-class.yaml
volumesnapshotclass.snapshot.storage.k8s.io/mysnapclass created
```

3. Create the snapshot.yaml file that defines the VolumeSnapshot object:

```
cat > snapshot.yaml <<\EOT
apiVersion: snapshot.storage.k8s.io/v1
kind: VolumeSnapshot
metadata:
  name: mysnapshot
spec:
  volumeSnapshotClassName: mysnapclass
  source:
    persistentVolumeClaimName: mypvc
EOT
```

This manifest specifies the volume snapshot `mysnapshot` that uses the volume snapshot class `mysnapclass` and creates a snapshot of the previously created volume bound to the persistent volume claim `mypvc`.

4. Create a volume snapshot:

```
$ kubectl create -f snapshot.yaml
volumesnapshot.snapshot.storage.k8s.io/mysnapshot created
```

In the self-service panel, you can find the newly created snapshot `mysnapshot` on the **Compute** -> **Volumes** -> **<pv_name>** -> **Snapshots** tab.

To delete a volume snapshot

Run the following command:

```
$ kubectl delete -f snapshot.yaml
volumesnapshot.snapshot.storage.k8s.io "mysnapshot" deleted
```

Creating external load balancers in Kubernetes

In Kubernetes, you can create a service with an external load balancer that provides access to it from public networks. The load balancer will receive a publicly accessible IP address and route incoming requests to the correct port on the Kubernetes cluster nodes.

Prerequisites

- To be able to assign a specific floating IP address to an external load balancer during its deployment, this floating IP address must be created in advance, as described in "Managing floating IP addresses" (p. 91).

To create a service with an external load balancer

1. Access the Kubernetes cluster via the dashboard. Click **Kubernetes access** for instructions.
2. On the Kubernetes dashboard, create a deployment and service of the **LoadBalancer** type. To do it, click **+ Create** and specify a YAML file that defines these objects. For example:
 - If you have deployed the Kubernetes cluster in a shared physical network, specify the following manifest:

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx
spec:
  replicas: 2
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
        - name: nginx
          image: nginx
          ports:
            - containerPort: 80
---
kind: Service
apiVersion: v1
metadata:
  name: load-balancer
  annotations:
    service.beta.kubernetes.io/openstack-internal-load-balancer: "true"
spec:
  selector:
    app: nginx
  type: LoadBalancer
  ports:
    - port: 80
      targetPort: 80
      protocol: TCP

```

The manifest above describes the deployment nginx with a replica set of two pods and the service load-balancer with the LoadBalancer type. The annotation used for the service indicates that the load balancer will be internal.

Once the load balancer is created, it will be allocated an IP address from the shared physical network and can be accessed at this external endpoint.

Details

Name: load-balancer	Connection
Namespace: default	Cluster IP: 10.254.147.243
Annotations: service.beta.kubernetes.io/openstack-internal-load-balancer: true	Internal endpoints: load-balancer:80 TCP load-balancer:32069 TCP
Creation Time: 2020-05-26T14:37 UTC	External endpoints: 10.94.156.196:80 
Label selector: app: nginx	
Type: LoadBalancer	
Session Affinity: None	

- If you have deployed the Kubernetes cluster in a virtual network linked to a physical one via a virtual router, you can use the YAML file above without the annotations section for the load-

balancer service. The created load balancer will receive a floating IP address from the physical network and can be accessed at this external endpoint. To use a specific floating IP address, create it in the self-service panel in advance, and then specify it with the `loadBalancerIP` parameter:

```
<...>
---
kind: Service
apiVersion: v1
metadata:
  name: load-balancer
spec:
  selector:
    app: nginx
  type: LoadBalancer
  loadBalancerIP: 10.10.10.100
  ports:
  - port: 80
    targetPort: 80
    protocol: TCP
```

- If you want to choose whether to create highly available load balancers for your service or not, you can make use of load balancer flavors. To specify a flavor for a load balancer add `loadbalancer.openstack.org/flavor-id: <flavor-id>` to the annotations section. The flavor ID can be obtained from your system administrator.

The load balancer will also appear in the self-service panel, where you can monitor its performance and health. For example:

Load balancers

Filters		Search		+ Create load balancer			
☐	Name ↑	Status ↓	IP address ↓	Floating IP ↓	Members state	Members ... ↓	⚙
☐	 kube_service_d66...	 Active	192.168.10.201	10.94.129.73		2	...

Using network policies in Kubernetes

[Network policies](#) are used to control network traffic in a Kubernetes cluster. By default, all outbound and inbound connections are allowed for a pod. You can provide network isolation for your pods by using network policies.

A network policy applies to a pod or a group of pods and sets a list of allowed targets for its ingress and egress rules. These targets can be one of the following:

- Specific pods (pods matching a label are allowed)
- Specific namespaces (all pods in the namespace are allowed)
- IP address blocks (endpoints with an IP address in the block are allowed)

If you want to isolate a pod for outbound connections, you need to create a network policy that selects this pod and has Egress in its `policyTypes`. Similarly, to isolate a pod for inbound

connections, a network policy should select this pod and have Ingress in its policyTypes. In this case, the only allowed connections to and from the pod will be those allowed by the ingress and egress lists. Note that traffic to and from the node where a pod is running is always allowed, as well as reply traffic for allowed connections.

Network policies are additive, so you can have multiple policies for a pod. Allowed connections to and from the pod will be the sum of the allow rules in the applicable policies. If the allow rules are not specified, all of the connections will be blocked.

To allow a connection from one pod to another, both of these pods must have respective egress and ingress policies allowing each other. If either side does not allow the connection, it will not happen.

To ensure that your Kubernetes cluster is protected from accidental network exposure, you can create a default deny all policy, and then add specific allow policies for the required traffic flows.

Limitations

- The Flannel network plugin, default for Kubernetes version 1.29.x and earlier, does not allow using network policies. Starting with version 1.30.x, Kubernetes clusters are created with the [Cilium](#) network plugin, which supports network policies.

To create a default deny all network policy

Use the following default-deny-all.yaml file:

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: default-deny-all
spec:
  podSelector: {}
  policyTypes:
    - Ingress
    - Egress
```

This manifest specifies the network policy default-deny-all that applies to all pods in the namespace and prevents all ingress and egress traffic.

To create a default allow all network policy for a pod

Use the following allow-all-demo.yaml file:

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: allow-all-demo
spec:
  podSelector:
    matchLabels:
      app: demo
```

```

policyTypes:
  - Ingress
  - Egress
ingress:
  - {}
egress:
  - {}

```

This manifest specifies the network policy `allow-all-demo` that applies to all pods labeled `app=demo` and allows all ingress and egress traffic.

To create a pod-based network policy

Use the following `pod-based-policy.yaml` file that defines the `NetworkPolicy` object:

```

apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: pod-based-policy
  namespace: default
spec:
  podSelector:
    matchLabels:
      app: pod1
  policyTypes:
    - Ingress
    - Egress
  ingress:
    - from:
        - podSelector:
            matchLabels:
              app: pod2
  egress:
    - to:
        - podSelector:
            matchLabels:
              app: pod2

```

This manifest specifies the network policy `pod-based-policy` that applies to all pods labeled `app=pod1` and allows ingress and egress traffic from all pods with the label `app=pod2`.

To create a namespace-based network policy

Use the following `namespace-based-policy.yaml` file that defines the `NetworkPolicy` object:

```

apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: namespace-based-policy
  namespace: default
spec:

```

```

podSelector:
  matchLabels:
    app: pod1
policyTypes:
  - Ingress
  - Egress
ingress:
  - from:
    - namespaceSelector:
        matchLabels:
          kubernetes.io/metadata.name: demo-namespace
egress:
  - to:
    - namespaceSelector:
        matchLabels:
          kubernetes.io/metadata.name: demo-namespace

```

This manifest specifies the network policy namespace-based-policy that applies to all pods labeled app=pod1 and allows ingress and egress traffic from all pods running in the namespace demo-namespace.

To create an IP-based network policy

Use the following ip-based-policy.yaml file that defines the NetworkPolicy object:

```

apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: ip-based-policy
  namespace: default
spec:
  podSelector:
    matchLabels:
      app: pod1
  policyTypes:
    - Ingress
    - Egress
  ingress:
    - from:
      - ipBlock:
          cidr: 10.10.10.0/24
  egress:
    - to:
      - ipBlock:
          cidr: 10.10.10.0/24

```

This manifest specifies the network policy ip-based-policy that applies to all pods labeled app=pod1 and allows ingress and egress traffic from the subnet 10.10.10.0/24.

Assigning Kubernetes pods to specific nodes

By using worker groups, you can assign a pod in Kubernetes to specific nodes. When you create a custom worker group, its nodes are added a label with the group name. If you want your pod to be scheduled on a node from a specific worker group, add the node selector section with the node label to the pod's configuration file.

To create a pod that will be scheduled on a specific node

Click **+ Create** on the Kubernetes dashboard and specify a YAML file that defines this object. For example:

```
apiVersion: v1
kind: Pod
metadata:
  name: nginx
  labels:
    env: test
spec:
  containers:
  - name: nginx
    image: nginx
    imagePullPolicy: IfNotPresent
  nodeSelector:
    magnum.openstack.org/nodegroup: mygroup
```

This manifest describes the pod nginx that will be assigned to a node from the node group mygroup.

When the pod is created, check that the hosting node belongs to the specified worker group.

Pods								
Name	Namespace Labels		Node	Status	Restarts	CPU Usage (cores)	Memory Usage (bytes)	Created
 nginx	default	env: test	kube1-mygroup-vogevh53o-node-1	Running	0	-	-	a minute ago

Enabling GPU support for Kubernetes nodes

To enable GPU support for your Kubernetes cluster, you need to deploy the NVIDIA device plugin for Kubernetes.

Note

This guide allows to deploy the latest version of the driver. If you need a specific version, you can build your own nvidia-driver-installer container image by following the instructions on this [GitHub page](#).

To deploy the NVIDIA device plugin for Kubernetes

1. Disable SELinux on Kubernetes worker nodes with GPU by using the `selinux_mode=disabled` label during the worker group creation.
2. Deploy [Node Feature Discovery](#) (NFD), a Kubernetes add-on for detecting hardware features and system configuration, to automatically discover GPU devices on your Kubernetes nodes and add the required labels:

```
# kubectl apply -f https://raw.githubusercontent.com/virtuozzo/nvidia-driver-installer/main/daemonsets/node-feature-discovery.yaml
```

3. Deploy the NVIDIA device plugin for Kubernetes:

```
# kubectl apply -f https://raw.githubusercontent.com/virtuozzo/nvidia-driver-installer/main/daemonsets/nvidia-gpu-driver.yaml
```

This daemon set will automatically distribute pods to all of your worker nodes with the required labels. For more details, refer to the [official guide](#).

You can check that the plugin is installed correctly by doing as follows:

1. Run a test pod:

```
# kubectl apply -f https://raw.githubusercontent.com/virtuozzo/nvidia-driver-installer/main/tests/gpupod.yaml
```

2. Check the pod logs:

```
# kubectl logs gpu-pod
[Vector addition of 50000 elements]
Copy input data from the host memory to the CUDA device
CUDA kernel launch with 196 blocks of 256 threads
Copy output data from the CUDA device to the host memory
Test PASSED
Done
```

Monitoring Kubernetes clusters

Warning!

This feature is experimental and not intended for use in production environments.

If you have enabled integrated monitoring during your Kubernetes cluster deployment, that means that the cluster has the `monitoring_enabled=true` label and the following components installed:

- Prometheus for data collection, storage, and search:
 - `node-exporter` exposes various server-level and OS-level metrics.
 - `kube-state-metrics` generates metrics on the state of Kubernetes objects.
- Alertmanager for alarm aggregation, processing, and dispatch.
- Grafana server for metrics visualization.

For instructions on how to create and configure Alertmanager and Prometheus instances, refer to the [kube-prometheus documentation](#).

The Grafana server is accessible from within a Kubernetes cluster at the **magnum-grafana.kube-system.svc.cluster.local** DNS name and TCP port 80.

The metrics on the state of Kubernetes objects are exported at the **/metrics** HTTP endpoint on the listening port: **magnum-kube-state-metrics.kube-system.svc.cluster.local:8080/metrics**. The metrics can be consumed either by Prometheus itself or by a scraper that is able to scrape a Prometheus client endpoint. For the list of exposed metrics, refer to [kube-state-metrics documentation](#).

Prerequisites

- A Kubernetes cluster with enabled integrated monitoring is created, as described in "Creating and deleting Kubernetes clusters" (p. 35).

To access the Kubernetes Grafana dashboards

1. On the **Kubernetes clusters** screen, click a Kubernetes cluster.
2. On the cluster right pane, click **Download kubeconfig**. The `.kubeconfig` file will be downloaded to your client machine.
3. On your client machine, install and set up the `kubectl` tool, to be able to run commands against Kubernetes clusters, as described in the [official documentation](#).
4. Specify the path to your Kubernetes configuration file in the `KUBECONFIG` environment variable:

```
# export KUBECONFIG=<path_to_kubeconfig>
```

5. Check that the kube-prometheus stack is installed:

```
# kubectl --namespace kube-system get pods -l "release=magnum"
NAME                                READY   STATUS    RESTARTS   AGE
magnum-kube-prometheus-sta-operator-85f757c5dc-ck1lb  1/1     Running   0           3d17h
magnum-kube-state-metrics-5cc46cbc5f-tclcv           1/1     Running   0           3d17h
magnum-prometheus-node-exporter-99kfc               1/1     Running   0           3d3h
magnum-prometheus-node-exporter-gwgzr              1/1     Running   0           3d17h
magnum-prometheus-node-exporter-q2pm2              1/1     Running   0           3d17h
magnum-prometheus-node-exporter-sqsl7              1/1     Running   0           2d22h
```

6. Obtain the password of the admin user:


```
# kubectl get secret --namespace kube-system magnum-grafana \
-o jsonpath="{.data.admin-password}" | base64 --decode ; echo
```

7. Configure the port forwarding for the Grafana pod:

```
# kubectl --namespace kube-system port-forward service/magnum-grafana 3000:80
```

8. Log in to `http://localhost:3000` under the admin user by specifying its username and password obtained in step 6.

- In the left menu, click **Dashboards > Browse**, and then select the dashboard you want to view.



To access the Prometheus user interface

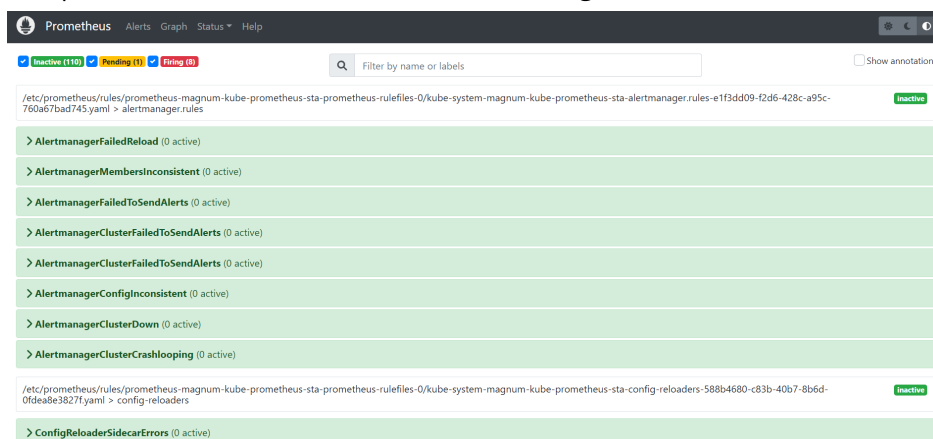
- On the **Kubernetes clusters** screen, click a Kubernetes cluster.
- On the cluster right pane, click **Download kubeconfig**. The .kubeconfig file will be downloaded to your client machine.
- On your client machine, install and set up the kubectl tool, to be able to run commands against Kubernetes clusters, as described in the [official documentation](#).
- Specify the path to your Kubernetes configuration file in the KUBECONFIG environment variable:

```
# export KUBECONFIG=<path_to_kubeconfig>
```

- Configure the port forwarding for the Prometheus pod:

```
# kubectl --namespace kube-system port-forward service/magnum-kube-prometheus-sta-prometheus 9090
```

- Visit <http://localhost:9090/graph> to use the Prometheus expression browser and to graph expressions. You can also navigate to <http://localhost:9090/metrics> to view the list of exported metrics, or <http://localhost:9090/alerts> to view the alerting rules.



To access the Alertmanager user interface


1. On the **Kubernetes clusters** screen, click a Kubernetes cluster.
2. On the cluster right pane, click **Download kubeconfig**. The `.kubeconfig` file will be downloaded to your client machine.
3. On your client machine, install and set up the `kubectl` tool, to be able to run commands against Kubernetes clusters, as described in the [official documentation](#).
4. Specify the path to your Kubernetes configuration file in the `KUBECONFIG` environment variable:

```
# export KUBECONFIG=<path_to_kubeconfig>
```

5. Configure the port forwarding for the Alertmanager pod:

```
# kubectl --namespace kube-system port-forward service/magnum-kube-prometheus-sta-alertmanager 9093
```

6. Visit <http://localhost:9093> to access the Alertmanager user interface.



Managing images

Virtuozzo Cloud allows you to upload ISO images and templates that can be used to create VM volumes:

- An ISO image is a typical OS distribution that needs to be installed on disk. You can upload an ISO image to the compute cluster.
- A template is a ready boot volume in the QCOW2 (rarely RAW or IMG) format with an installed operating system and applications. Many OS vendors offer templates of their operating systems under the name "cloud images". You can upload a cloud image from the [OS official repository](#) or prepare your own template in the compute cluster.

Prerequisites

- Knowledge of the supported guest operating systems listed in "Supported guest operating systems" (p. 11).

Uploading images

Limitations

- The image file size must not exceed 500 GiB.

To upload an image

1. On the **Images** screen, click **Add image**.
2. In the **Add image** window, click **Browse** and select a file in one of the supported formats: .iso, .img, .qcow2, .raw.
3. Specify an image name and select the correct OS type from the drop-down list.

Important

The OS type affects VM parameters such as hypervisor settings. VMs created from an image with an incorrect OS type may not work correctly, for example, they may crash.

4. [Optional] For an image in the QCOW2, RAW, or IMG format:
 - Select **UEFI boot** to mark the image as UEFI bootable.
 - Select **vTPM** to enable a virtual Trusted Platform Module (TPM) that provides enhanced security to the guest OS. Enabling vTPM also enables UEFI boot, since vTPM requires UEFI mode.

Note

These parameters cannot be changed after the image is uploaded. They will be inherited by all VMs and volumes provisioned from the image.

Important

For Windows 11 and Windows Server 2025, these parameters are mandatory and enabled automatically.

Add image

×

Image file

centos7-minimal.qcow2

×

Browse

Name

centos7-minimal.qcow2

Select OS distribution

CentOS 7

▼

☒ UEFI boot

Cancel

Add

5. Click **Add** to start uploading the image. The upload progress will be shown in the bottom right corner.

You can hide the pop-up window without interrupting the upload process. The upload progress will be available in the notification center.

Creating volumes from images

You can create volumes from both ISO images and templates.

To make a volume from an image

1. Go to the **Images** screen, and then click the required image.
2. On the image panel, click **Create volume**.
3. In the **Create volume** window, specify the volume name, size, and select a storage policy.

Add image

Image file
windows11.qcow2

X

Browse

Name
win11

Select OS distribution
Windows 11

▼

Boot parameters

☒ UEFI boot

☒ vTPM ?

Cancel

Add

4. Click **Create**.

The new volume will appear on the **Volumes** screen.

Preparing templates

You may need to create a template in these cases:

- To rescue a virtual machine
- To create a VM accessible via SSH
- To create a VM customizable with user data

Preparation overview

1. Install cloud-init and OpenSSH Server in the virtual machine.
2. [Optional] Enable logging for virtual machines that will be created from the template.
3. Convert the VM boot volume to the template, as described in "Creating images from volumes" (p. 71).

Preparing Linux templates

As all Linux guests have OpenSSH Server preinstalled by default, you only need to ensure that a Linux template has cloud-init installed.

Cloud-init is required to configure user credentials and SSH access for virtual machines created from cloud images.

The easiest way to get a Linux template with cloud-init installed is to obtain it from [its official repository](#). You can also create a Linux template from an existing boot volume.

Using cloud-init with Linux cloud images

When deploying virtual machines from Linux cloud images (for example, Debian or Ubuntu), cloud-init applies the initial system configuration.

Cloud images do not contain predefined credentials. User accounts, passwords, and SSH settings are defined through cloud-init.

Access to such virtual machines is typically provided in one of the following ways:

- by injecting an SSH key
- by specifying user credentials in a cloud-init configuration

SSH keys selected at deployment are automatically injected into the guest operating system via cloud-init.

If neither a cloud-init configuration nor an SSH key is provided, you may not be able to log in to the virtual machine.

The following example creates a user account, sets a password, and enables SSH password authentication:

```
#cloud-config
users:
  - name: clouduser
    groups: sudo
    shell: /bin/bash
    sudo: ALL=(ALL) NOPASSWD:ALL
    lock_passwd: false

ssh_pwauth: true

chpasswd:
  expire: false
  list: |
    clouduser:<password>

runcmd:
  - systemctl restart ssh
```

Replace <password> with a secure password.

Cloud-init configuration is provided when creating the virtual machine. For details, refer to "Creating virtual machines" (p. 13).

Preparing Windows templates

Windows guests have neither Cloudbase-Init nor OpenSSH Server preinstalled by default. You need to install and configure them manually.

To install Cloudbase-Init and OpenSSH Server inside a Windows virtual machine

1. Log in to a Windows VM.
2. Create a new administrator account that will be used for SSH connections and log in with it.
3. To install and configure OpenSSH Server:

- a. Run Windows PowerShell with administrator privileges and set the execution policy to unrestricted to be able to run scripts:

```
> Set-ExecutionPolicy Unrestricted
```

- b. Download OpenSSH Server (for example, from the [GitHub repository](#)), extract the archive into the C:\Program Files directory, and then install it by running:

```
> & 'C:\Program Files\OpenSSH-Win64\install-sshd.ps1'
```

- c. Start the sshd service and set its startup type to "Automatic":

```
> net start sshd  
> Set-Service sshd -StartupType Automatic
```

- d. Open TCP port 22 for the OpenSSH service in the Windows Firewall:

- On Windows 8.1, Windows Server 2012, and newer versions, run:

```
> New-NetFirewallRule -Protocol TCP -LocalPort 22 -Direction Inbound -Action Allow -DisplayName OpenSSH
```

- On Windows Server 2008/2008 R2, run:

```
> netsh advfirewall firewall add rule name=sshd dir=in action=allow  
protocol=TCP localport=22
```

- e. Open the C:\ProgramData\ssh\sshd_config file:

```
> notepad 'C:\ProgramData\ssh\sshd_config'
```

Comment out the following lines at the end of the file:

```
#Match Group administrators  
#AuthorizedKeysFile __PROGRAMDATA__/ssh/administrators_authorized_keys
```

Save the changes.

- f. Create the .ssh directory in C:\Users\<current_user> and an empty authorized_keys file inside it:

```
> cd C:\Users\<current_user>
> mkdir .ssh
> notepad .\.ssh\authorized_keys
```

Remove the .txt extension from the created file:

```
> move .\.ssh\authorized_keys.txt .\.ssh\authorized_keys
```

- g. Modify the permissions for the created file to disable inheritance:

```
> icacls .\.ssh\authorized_keys /inheritance:r
```

4. Download Cloudbase-Init from <https://cloudbase.it/cloudbase-init/#download>, and then install it by following the procedure from the **Installation** section at <https://cloudbase.it/cloudbase-init/>.

Important

- a. The password for the user specified during the Cloudbase-Init installation will be reset on the next VM startup. If this user does not exist, a new user account will be created. You will be able to log in with this account by using the key authentication method or you can set a new password with a customization script. If there are multiple Windows users at the image preparation time, the passwords for other users will not be changed.
 - b. When the Cloudbase-Init installation is complete, do not select the option to run Sysprep before clicking **Finish**. Otherwise, you will not be able to modify `cloudbase-init.conf`.
-
5. Run Windows PowerShell with administrator privileges and open the file `C:\Program Files\Cloudbase Solutions\Cloudbase-Init\conf\cloudbase-init.conf`:

```
> notepad 'C:\Program Files\Cloudbase Solutions\Cloudbase-Init\conf\cloudbase-init.conf'
```

Add `metadata_services` and `plugins` on two lines:

```
metadata_services=\
cloudbaseinit.metadata.services.configdrive.ConfigDriveService,\
cloudbaseinit.metadata.services.httpservice.HttpService
plugins=cloudbaseinit.plugins.common.mtu.MTUPlugin,\
cloudbaseinit.plugins.windows.ntpclient.NTPClientPlugin,\
cloudbaseinit.plugins.common.sethostname.SetHostNamePlugin,\
cloudbaseinit.plugins.windows.createuser.CreateUserPlugin,\
cloudbaseinit.plugins.common.networkconfig.NetworkConfigPlugin,\
cloudbaseinit.plugins.windows.licensing.WindowsLicensingPlugin,\
cloudbaseinit.plugins.common.sshpublickeys.SetUserSSHPublicKeysPlugin,\
cloudbaseinit.plugins.windows.extendvolumes.ExtendVolumesPlugin,\
cloudbaseinit.plugins.common.setuserpassword.SetUserPasswordPlugin,\
cloudbaseinit.plugins.common.userdata.UserDataPlugin,\
cloudbaseinit.plugins.windows.winrmlistener.ConfigWinRMListenerPlugin,\
cloudbaseinit.plugins.windows.winrmcertificateauth.\
ConfigWinRMCertificateAuthPlugin,\
cloudbaseinit.plugins.common.localscripts.LocalScriptsPlugin
```

Note

Make sure to remove all backslashes in the lines above.

Save the changes.

6. Run the built-in Sysprep tool:

```
> sysprep /generalize /oobe /shutdown
```

Enabling logging for virtual machines

The console log of a virtual machine can be used for troubleshooting boot issues. The log contains messages only if logging is enabled inside the VM, otherwise the log is empty.

The logging can be turned on by enabling the TTY1 and TTYS0 logging levels in Linux VMs and Emergency Management Services (EMS) console redirection in Windows VMs. You may also enable driver status logging in Windows VMs, to see the list of loaded drivers. This can be useful for troubleshooting a faulty driver or long boot process.

To enable TTY1 and TTYS0 logging in Linux virtual machines

1. Add the line `GRUB_CMDLINE_LINUX_DEFAULT="console=tty1 console=ttys0"` to the file `/etc/default/grub`.
2. Depending on the boot loader, run either

```
# grub-mkconfig -o /boot/grub/grub.cfg
```

or

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

3. Reboot the VM.

To enable EMS console redirection in Windows virtual machines

1. Start **Windows PowerShell** by using administrator privileges.
2. In the PowerShell console, set the COM port and baud rate for EMS console redirection. As Windows VMs have only the COM1 port with the transmission rate of 9600 bps, run:

```
bcdedit /emssettings EMSPORT:1
```

3. Enable EMS for the current boot entry:

```
bcdedit /ems on
```

To enable driver status logging in Windows virtual machines

1. Start **System Configuration** by using administrator privileges.
2. In the **System Configuration** windows, open the **Boot** tab, and select the check boxes **OS boot information** and **Make all boot settings permanent**.
3. Confirm the changes and restart the system.

Managing volumes

A volume in Virtuozzo Cloud is a virtual disk drive that can be attached to a virtual machine. A virtual disk drive can emulate either a hard disk (HDD) or a CD-ROM. This refers to the type of device that the guest operating system inside the virtual machine recognizes and interacts with. The integrity of data in volumes is protected by the redundancy mode specified in the storage policy.


Creating and deleting volumes

Limitations

- A volume is removed along with all of its snapshots.

To create a volume

1. On the **Volumes** screen, click **Create volume**.
2. In the **Create volume** window, specify a volume name and size in gigabytes, select a storage policy, and then click **Create**.



Create volume [X]

Name
vol1

Size (GiB)
1

Min. 1 GiB,
Max. 512 TiB

Storage policy
default [v]

Cancel Create

To remove a volume

1. On the **Volumes** tab, check the status of the volume you want to remove.
2. If the status is "In use", click the volume, and then click **Force detach** on the volume right pane.
3. If the status is "Available", click the volume, and then click **Delete** on the volume right pane.

4. If the volume has recovery points, you can delete them along with the volume:
 - a. Select **Delete recovery points**.
 - b. Enter "Delete" for confirmation.

Delete volume ✕

Are you sure that you want to delete the volume
"win10/windows_10_multi-edition_1709/Boot volume"?

 The volume has recovery points. You can delete them along with the volume.

☒ Delete recovery points

Enter "**Delete**" for confirmation:

Enter confirmation

Delete

CancelDelete

5. Click **Delete** in the confirmation window.

Attaching and detaching volumes

Limitations

- You can only attach and detach non-boot volumes.

Prerequisites

- A volume is created, as described in "Creating and deleting volumes" (p. 68).
- To be able to use volumes attached to VMs, they must be initialized inside the guest OS by standard means.

To attach a volume to a virtual machine

1. On the **Volumes** screen, click an unused volume.
2. On the volume right pane, click **Attach**.
3. In the **Attach volume** window, select the VM from the drop-down list, and then click **Done**.

Attach volume

Choose a volume to attach

Volume

vol1

Virtual machine

vm1

Cancel

Done

To detach a volume from a virtual machine

1. On the **Volumes** screen, click a volume that is in use.
2. If the VM is stopped, click **Detach** on the volume right pane.
3. If the VM is running, click **Force detach** on the volume right pane.

Warning!

There is a risk of data loss.

Resizing volumes

You can change volume size only by increasing it. Volumes can be extended for both running (online resizing) and stopped (offline resizing) virtual machines. Online volume resizing allows users to avoid downtime and enables scaling VM storage capacity on the fly without service interruption.

Limitations

- You cannot shrink volumes.
- During volume resizing, the file system inside the guest OS is not extended.
- If you revert a volume to a snapshot that was taken before the volume extension, the new volume size will be retained.

Prerequisites

- A volume is created, as described in "Creating and deleting volumes" (p. 68).

To extend a volume

1. On the **Volumes** screen, click a volume.
2. Click the pencil icon in the **Size** field.

3. Enter the desired volume capacity, and then click the tick icon.

After the volume is extended, you will need to re-partition the disk inside the guest OS to allocate the added disk space.

Changing the storage policy for volumes

If you use redundancy by replication for a compute volume, you can update the chosen redundancy scheme by changing the storage policy. With redundancy by erasure coding, however, changing the redundancy scheme applied to the volume is disabled.

The storage policy can be changed for detached volumes and volumes attached to running or stopped virtual machines.

Limitations

- Only storage policies enabled by project quotas will be available for selection.
- Changing the storage policy with the erasure coding redundancy type is disabled.

Prerequisites

- A volume is created, as described in "Creating and deleting volumes" (p. 68).

To change the storage policy of a volume

1. On the **Volumes** screen, click a volume.
2. Click the pencil icon in the **Storage policy** field.
3. Select a new storage policy, and then click the tick icon. You can choose only between storage policies with the replication redundancy type.

Creating images from volumes


To create multiple VMs with the same boot volume, you can create a template from an existing boot volume and deploy VMs from it.

Prerequisites

- Linux virtual machines have cloud-init installed, as described in "Preparing Linux templates" (p. 63).
- Windows virtual machines have Cloudbase-Init and OpenSSH Server installed, as described in "Preparing Windows templates" (p. 65).
- [Optional] Logging is enabled inside a virtual machine, as instructed in "Enabling logging for virtual machines" (p. 67).

To create a template from a boot volume

1. Power off the VM that the original volume is attached to.
2. Switch to the **Volumes** screen, click volume's ellipsis button and select **Create image**.
3. In the **Create image** window, enter an image name, and then click **Create**.

A dialog box titled "Create image" with a close button (X) in the top right corner. Inside the dialog, there is a text input field labeled "Name" containing the text "vol1-image". Below the input field, it says "Volume: vm1/cirros/Boot volume". At the bottom right of the dialog, there are two buttons: "Cancel" and "Create".

Create image

Name
vol1-image

Volume: vm1/cirros/Boot volume

Cancel Create

The new image will appear on the **Images** screen.

Cloning volumes

Limitations

- You can clone volumes that are not attached to VMs or attached to stopped VMs.

Prerequisites

- A volume is created, as described in "Creating and deleting volumes" (p. 68).

To clone a volume

1. On the **Volumes** screen, click a volume.
2. On the volume right pane, click **Clone**.
3. In the **Clone volume** window, specify a volume name, size, and storage policy. Click **Clone**.

Clone volume

Name

Clone_vol1

Size (GiB)

1

Min. 1 GiB,
Max. 512 TiB

Storage policy

default

Cancel

Clone

Managing volume snapshots

You can save the current state of a VM file system or user data by creating a snapshot of a volume. A snapshot of a boot volume may be useful, for example, before updating VM software. If anything goes wrong, you will be able to revert the VM to a working state at any time. A snapshot of a data volume can be used for backing up user data and testing purposes.

When a snapshot is created, the current read-write file of a volume becomes a read-only backing file (the snapshot), and a new overlay file is created. All changes to the volume are written to this overlay file, preserving the original data in the snapshot. The more changes made after the snapshot, the larger the overlay file becomes. When the snapshot is deleted, the data that was written after the snapshot is merged back with the original data, making the changes permanent, and then the overlay file is removed.

Limitations


- A volume supports a maximum of 32 snapshots.

Prerequisites

- To create a consistent snapshot of a running VM's volume, the guest tools must be installed in the VM, as described in "Installing guest tools" (p. 30). The QEMU guest agent included in the guest tools image automatically quiesces the filesystem during snapshotting.


To create a snapshot of a volume

1. On the **Volumes** screen, click a volume.
2. In the volume right pane, switch to **Snapshots**, and then click **Create snapshot**.



To manage a volume snapshot

Select a volume and open the **Snapshots** tab on its right pane.



You can do the following:

- Create a new volume from the snapshot.
- Create a template from the snapshot.
- Discard all changes that have been made to the volume since the snapshot was taken. This action is available only for VMs with the "Shut down" and "Shelved offloaded" statuses.

Warning!

As each volume has only one snapshot branch, all snapshots created after the snapshot you are reverting to will be deleted. If you want to save a subsequent snapshot before reverting, create a volume or an image from it first.

- Change the snapshot name and description.

Important

A description should not contain any personally identifiable information or sensitive business data.

- Reset the snapshot stuck in an "Error" state or transitional state to the "Available" state.
- Remove the snapshot.

To perform these actions, click the ellipsis button next to a snapshot, and then click the corresponding action.

Transferring volumes between projects

There is no direct way to migrate a virtual machine between different projects. However, you can transfer the VM boot volume, and then create a new VM from it. You can transfer both boot and non-boot volumes to projects within different domains.

Limitations

- You can only transfer volumes with the "Available" status.
- Transferring volumes that have snapshots breaks the snapshots.

Prerequisites

- Access to the compute API depends on your provider's settings. You need to obtain from your provider the instruction how to connect to the API.
- You have login credentials for the source and destination projects.
- If you want to transfer a boot volume that is attached to a VM, clone this volume first, as described in "Cloning volumes" (p. 72).
- If you want to transfer a non-boot volume that is attached to a VM, detach it first, as described in "Attaching and detaching volumes" (p. 69).

To transfer a volume between two projects

1. Log in to the source project by changing the environment variables to the project credentials. For example:

```
export OS_PROJECT_DOMAIN_NAME=domain1
export OS_USER_DOMAIN_NAME=domain1
export OS_PROJECT_NAME=project1
export OS_USERNAME=user1
export OS_PASSWORD=password
```

2. List all volumes within your project to find out the ID of the volume you want to transfer:

```
# openstack --insecure volume list
+-----+-----+-----+-----+
| ID                                     | Name                               | Status | Size |
+-----+-----+-----+-----+
```

```
+-----+-----+-----+-----+
| 2c8386fa-331b-4ba8-9e4c-de690969a4c8 | win10/Boot volume | available | 64 |
+-----+-----+-----+-----+
```

3. Create a transfer request by specifying the ID of the chosen volume. For example:

```
# openstack --insecure volume transfer request create c0d4cf0e-48e3-417d-b6fc-
f1fb36571c5f
+-----+-----+
| Field      | Value                               |
+-----+-----+
| auth_key   | 75fcf37d56f40182                   |
| created_at | 2022-04-27T09:00:11.776511         |
| id         | b9b835a3-ed41-489a-9552-483fae33c549 |
| name       | None                                |
| volume_id  | c0d4cf0e-48e3-417d-b6fc-f1fb36571c5f |
+-----+-----+
```

Save the request id and auth-key from the command output, to accept the transfer in the other project.

4. Log in to the destination project by changing the environment variables to the project credentials. For example:

```
export OS_PROJECT_DOMAIN_NAME=domain1
export OS_USER_DOMAIN_NAME=domain1
export OS_PROJECT_NAME=project2
export OS_USERNAME=user2
export OS_PASSWORD=password
```

5. Accept the transfer request by specifying the request ID and authorization key. For example:

```
# openstack --insecure volume transfer request accept --auth-key 75fcf37d56f40182 \
b9b835a3-ed41-489a-9552-483fae33c549
```

Once the volume is moved to the other project, you can create a virtual machine from it, as described in "Creating virtual machines" (p. 13).

Managing virtual networks

Limitations

- You can delete a compute network only if no VMs are connected to it.

To add a new virtual network

1. On the **Networks** screen, click **Create virtual network**.
2. On the **Network configuration** step, do the following:
 - a. Enable or disable IP address management:
 - With IP address management enabled, VMs connected to the network will automatically be assigned IP addresses from allocation pools by the built-in DHCP server and use custom DNS servers. Additionally, spoofing protection will be enabled for all VM network ports by default. Each VM network interface will be able to accept and send IP packets only if it has IP and MAC addresses assigned. You can disable spoofing protection manually for a VM interface, if required.
 - With IP address management disabled, VMs connected to the network will obtain IP addresses from the DHCP servers in that network, if any. Also, spoofing protection will be disabled for all VM network ports, and you cannot enable it manually. This means that each VM network interface, with or without assigned IP and MAC addresses, will be able to accept and send IP packets.

In any case, you will be able to manually assign static IP addresses from inside the VMs.

- b. Specify a name, and then click **Next**.

The screenshot shows a 'Create virtual network' dialog box with a close button (X) in the top right corner. On the left is a sidebar with three steps: 'Network configuration' (selected), 'IP address management', and 'Summary'. The main area shows a toggle switch for 'IP address management' which is turned on. Below the toggle is a text input field labeled 'Name' with the value 'net1' entered. At the bottom right are 'Cancel' and 'Next' buttons.

3. If you enabled IP address management, you will move on to the **IP address management** step, where you can add an IPv4 subnet:
 - a. In the **Subnets** section, click **Add** and select **IPv4 subnet**.
 - b. In the **Add IPv4 subnet** window, specify the network's IPv4 address range and, optionally, specify a gateway. If you leave the **Gateway** field blank, the gateway will be omitted from network settings.
 - c. Enable or disable the built-in DHCP server:
 - With the DHCP server enabled, VM network interfaces will automatically be assigned IP addresses: either from allocation pools or, if there are no pools, from the network's entire IP range. The DHCP server will receive the first two IP addresses from the IP pool. For example:

- In a subnet with CIDR 192.168.128.0/24 and without a gateway, the DHCP server will be assigned the IP addresses 192.168.128.1 and 192.168.128.2.
- In a subnet with CIDR 192.168.128.0/24 and the gateway IP address set to 192.168.128.1, the DHCP server will be assigned the IP addresses 192.168.128.2 and 192.168.128.3.
- With the DHCP server disabled, VM network interfaces will still get IP addresses, but you will have to manually assign them inside VMs.

The virtual DHCP service will work only within the current network and will not be exposed to other networks.

- Specify one or more allocation pools (ranges of IP addresses that will be automatically assigned to VMs).
- Specify DNS servers that will be used by virtual machines. These servers can be delivered to VMs via the built-in DHCP server or by using the cloud-init network configuration (if cloud-init is installed in the VM).
- Click **Add**.

Add IPv4 subnet

×

CIDR

10.10.10.0/24

Gateway (optional)

10.10.10.1

☒ Built-in DHCP server ⓘ

Allocation pools

+ Add

10.10.10.100 — 10.10.10.200

101 addresses available

✎

🗑

DNS servers

+ Add

8.8.8.8

✎

🗑

Cancel

Add

- On the **Summary** step, review the configuration, and then click **Create virtual network**.

To edit parameters of a virtual network

- On the **Networks** screen, click the required network.
- On the network right pane, click the pencil icon next to the network name or IPv4 subnet.
- Make changes and save them.

To delete a compute network

Click the ellipsis icon next to the required network, and then click **Delete**. To remove multiple compute networks at once, select them, and then click **Delete**.


Managing VPN connections

With Virtual Private Network (VPN) as a service, self-service users can extend virtual networks across public networks, such as the Internet. To connect two or more remote endpoints, VPNs use virtual connections tunneled through physical networks. To secure VPN communication, the traffic that flows between remote endpoints is encrypted. The VPN implementation uses the Internet Key Exchange (IKE) and IP Security (IPsec) protocols to establish secure VPN connections and is based on the strongSwan IPsec solution.

VPN as a service can be used to establish a Site-to-Site VPN connection between a virtual network configured in Virtuozzo Cloud and any other network with a VPN gateway that uses the IPsec and IKE protocols. With VPN as a service, you can connect the following workloads:


- On-premises workloads with workloads hosted in Virtuozzo Cloud
- Workloads hosted in other clouds with workloads hosted in Virtuozzo Cloud
- Workloads hosted in different Virtuozzo Cloud clusters

To better understand how a VPN works, consider the following example:



- In the **cluster 1**, the virtual machine **VM1** is connected to the virtual network **privnet1** (192.168.10.0/24) via the network interface with IP address 192.168.10.10. The network **privnet1** is exposed to public networks via the router **router1** with the external port 10.10.10.5.
- In the **cluster 2**, the virtual machine **VM2** is connected to the virtual network **privnet2** (192.168.20.0/24) via the network interface with IP address 192.168.20.20. The network **privnet2** is exposed to public networks via the router **router2** with the external port 10.10.10.4.
- The VPN tunnel is created between the routers **router1** and **router2** that serve as VPN gateways, thus allowing mutual connectivity between the networks **privnet1** and **privnet2**.
- The virtual machines **VM1** and **VM2** are visible to each other at their private IP addresses. That is, **VM1** can access **VM2** at 192.168.20.20, and **VM2** can access **VM1** at 192.168.10.10.

For key exchange between communicating parties, two IKE versions are available: IKE version 1 (IKEv1) and IKE version 2 (IKEv2). IKEv2 is the latest version of the IKE protocol and it supports connecting multiple remote subnets.



In the example above:

- **VPN1** uses the IKEv1 and connects the network **network1** with the **network3**.
- **VPN2** uses the IKEv2 and connects the network **network2** with the two networks **network4** and **network5**.

Limitations

- Currently, we support only Site-to-Site VPN connections. Point-to-Site VPN connections are not supported.

Creating VPN connections

Prerequisites

- You have a virtual router created, as described in "Managing virtual routers" (p. 85).
- The virtual router connects the physical network with virtual networks that you want to be exposed.
- Networks that will be connected via a VPN tunnel must have non-overlapping IP ranges.
- [For Virtuoizzo Cloud 5.4 Update 1 and earlier versions] If a virtual machine has a floating IP address assigned to its private network interface, configure static routes of a virtual router, for the VM traffic to be routed through a VPN tunnel.

In this case, you need to add static routes to your virtual router for remote subnets that you want to access via a VPN tunnel. The next hop IP address will be the IP address of the internal SNAT router interface. To find out this IP address, run:

```
# openstack --insecure port list --device-id <router_id> --device-owner
network:router_centralized_snat -c fixed_ips
+-----+
| Fixed IP Addresses |
+-----+
| ip_address='192.168.128.69', subnet_id='c33e75f3-8ede-4899-a6cb-6f9d87a61714' |
+-----+
```

In this example, 192.168.128.69 is the IP address of the internal SNAT router interface. A router, however, may have multiple internal SNAT router interfaces. You can specify any of them as the next hop IP address. For more details on adding static routes, refer to "Managing static routes" (p. 89).

To create a VPN connection

1. On the **VPN** screen, click **Create VPN**.
2. On the **Configure IKE** step, specify parameters for the IKE policy that will be used to establish a VPN connection. You can choose to use an existing IKE policy or create a new one. For the new IKE policy, do the following:
 - a. Specify a custom name for the IKE policy.
 - b. Specify the key lifetime, in seconds, that will define the rekeying interval. The IKE key lifetime must be greater than that of the IPsec key.
 - c. Select the authentication algorithm that will be used to verify the data integrity and authenticity.
 - d. Select the encryption algorithm that will be used to ensure that data is not viewable while in transit.
 - e. Select the IKE version 1 or 2. Version 1 has limitations, for example, it does not support multiple subnets.
 - f. Select the Diffie-Hellman (DH) group that will be used to build the encryption key for the key exchange process. Higher group numbers are more secure but require additional time for the key to compute.
 - g. Click **Next**.

Create VPN

×

● Configure IKE

● Configure IPsec

● Create endpoint groups

● Configure VPN

● Summary

Key lifetime (in seconds)

− 3600 + ⓘ

Authentication algorithm

☐ SHA-1 ☒ SHA-256 ☐ SHA-384 ☐ SHA-512

Encryption algorithm

☐ 3DES ☒ AES-128 ☐ AES-192 ☐ AES-256

IKE version ⓘ

☐ v1 ☒ v2

Diffie-Hellman group ⓘ

☐ group2 ☒ group5 ☐ group14

Cancel

Next

3. On the **Configure IPsec** step, specify parameters for the IPsec policy that will be used to encrypt the VPN traffic. You can choose to use an existing IPsec policy or create a new one. For the new IPsec policy, do the following:
 - a. Specify a custom name for the IPsec policy.
 - b. Specify the key lifetime, in seconds, that will define the rekeying interval. The IPsec key lifetime must not be greater than that of the IKE key.
 - c. Select the authentication algorithm that will be used to verify the data integrity and authenticity.
 - d. Select the encryption algorithm that will be used to ensure that data is not viewable while in transit.
 - e. Select the Diffie-Hellman (DH) group that will be used to build the encryption key for the key exchange process. Higher group numbers are more secure but require additional time for the key to compute.
 - f. Click **Next**.

Create VPN

- Configure IKE
- **Configure IPsec**
- Create endpoint groups
- Configure VPN
- Summary

IPsec policy
New IPsec policy

Policy name
ipsec1

Key lifetime (in seconds)
3600

Authentication algorithm
☐ SHA-1 ☒ SHA-256 ☐ SHA-384 ☐ SHA-512

Encryption algorithm
☐ 3DES ☒ AES-128 ☐ AES-192 ☐ AES-256

Diffie-Hellman group
☐ group2 ☒ group5 ☐ group14

Back Next

4. On the **Create endpoint groups** step, select a virtual router and specify local and remote subnets that will be connected by the VPN tunnel. You can choose to use existing local and remote endpoints, or create new ones. For the new endpoints, do the following:
 - a. Specify a custom name for the local endpoint, and then select local subnets.
 - b. Specify a custom name for the remote endpoint, and then add remote subnets in the CIDR format.
 - c. Click **Next**.

Create VPN

×

- Configure IKE
- Configure IPsec
- **Create endpoint groups**
- Configure VPN
- Summary

Subnets

private1: 10.10.10.0/24

▼

Remote endpoint

Remote endpoint

Create endpoint group

▼

Group name

remote-endpoint1

Subnets

+

Add

10.10.20.0/24

✕

10.10.30.0/24

✕

Back

Next

5. On the **Configure VPN** step, specify parameters to establish the VPN connection with a remote gateway:
 - a. Specify a custom name for the VPN connection.
 - b. Specify the public IPv4 address of the remote gateway, that is, peer IP address.
 - c. Generate the pre-shared key that will be used for the peer authentication.
 - d. [Optional] If necessary, you can also configure additional settings by selecting **Advanced settings** and specifying the following parameters:
 - The peer ID for authentication and the mode for establishing a connection.
 - The Dead Peer Detection (DPD) policy, interval, and timeout, in seconds.
 - e. Click **Next**.

Create VPN

×

- Configure IKE
- Configure IPsec
- Create endpoint groups
- **Configure VPN**
- Summary

Specify parameters to establish the VPN connection with a remote gateway.

☒ Basic settings
 ☐ Advanced settings

VPN name

vpn1

Public IPv4 address (Peer IP)

10.136.18.134

?

Pre-shared key (PSK)

psk

📋

🔄 Generate

Back

Next

6. On the **Summary** step, review the configuration, and then click **Create**.

When the VPN connection is created, its status will change from "Pending creation" to "Down". The connection will become active once the VPN tunnel is configured by the other VPN party and the IKE authorization is successful.

Important

The IKE and IPsec configuration must match for both communicating parties. Otherwise, the VPN connection between them will not be established.

Editing VPN connections

After a VPN connection is created, you can change its endpoint groups and VPN settings at any time.

Limitations

- You cannot change the virtual router and security policies used to establish a VPN connection.

Prerequisites

- A VPN connection is created, as described in "Creating VPN connections" (p. 80).

To edit a VPN connection

1. On the **VPN** screen, click a VPN connection to modify.
2. On the connection right pane, click **Edit**.
3. In the **Edit VPN** window, configure local and remote endpoints, if required, and then click **Next**.
4. On the next step, change VPN parameters such as the VPN connection name, peer IP address, and PSK key. If necessary, you can also configure additional settings by selecting **Advanced**

settings and editing the required parameters.

5. Click **Save** to apply your changes.

After you update the connection parameters, its status will change to "Down". The connection will re-initiate once the parameters are similarly updated by the other VPN party.

Important

The IKE and IPsec configuration must match for both communicating parties. Otherwise, the VPN connection between them will not be established.

Restarting and deleting VPN connections

You can forcefully re-initiate a VPN connection by manually restarting it. When you delete a VPN connection, you also delete the IKE and IPsec policies and endpoint groups that were created during the VPN creation.

Prerequisites

- A VPN connection is created, as described in "Creating VPN connections" (p. 80).

To restart a VPN connection

1. On the **VPN** screen, click a VPN connection to restart.
2. On the connection right pane, click **Restart**.
3. Click **Restart VPN** in the confirmation window.

To delete a VPN connection

1. On the **VPN** screen, click a VPN connection to delete.
2. On the connection right pane, click **Delete**.
3. Click **Delete** in the confirmation window.

Managing virtual routers

Virtual routers provide Layer 3 (L3) networking services such as routing and Source Network Address Translation (SNAT) between virtual and physical networks, as well as between different virtual networks.

A virtual router connecting a virtual network to a physical network enables virtual machines to access external networks, such as the Internet. When a router connects multiple virtual networks, it enables communication between VMs on those networks.

A virtual router has two types of ports:

- An external gateway is connected to a physical network and used for outbound traffic and floating IP access.
- An internal interface is connected to a virtual network and used for communication with VMs.

Traffic flow and address translation

Virtual routers apply different types of network address translation depending on traffic direction.

- For outbound traffic, SNAT is used: the VM's private IP address is translated to the router's external IP.
- For inbound traffic, floating IPs are used: traffic sent to a floating IP is translated (DNAT) to the VM's internal IP address.

When a floating IP is assigned to a VM, inbound traffic is forwarded through the virtual router using DNAT. The original source IP address of the external client is preserved, and no source NAT is applied. As a result, the VM sees the real client IP address.

This allows source IP-based access control inside the VM (for example, using firewall rules).

In some cases, the source IP may not be preserved, for example, when traffic passes through a load balancer or proxy. In such scenarios, the VM may see the IP address of the intermediary instead of the original client.

Routing architecture

Virtuozzo Cloud uses a distributed routing architecture. Routing and floating IP processing are performed directly on compute nodes where VMs run. This allows traffic between VMs and inbound traffic from external networks to be handled locally, reducing latency and improving performance. At the same time, outbound traffic that requires SNAT is processed on management nodes.

Limitations

- A router can only connect networks that have IP management enabled.
- You can delete a virtual router if no floating IP addresses are associated with any network it is connected to.

Prerequisites

- Compute networks are created, as described in "Managing virtual networks" (p. 76).
- The compute networks that are to be connected to a router have a gateway specified.

To create a virtual router

1. Navigate to the **Routers** screen, and then click **Add router**.
2. In the **Add router** window:
 - a. Specify a router name.
 - b. From the **Network** drop-down menu, select a physical network through which external access will be provided via an external gateway. The new external gateway will pick an unused IP address from the selected physical network.
 - c. In the **Add internal interfaces** section, select one or more virtual networks to connect to a router via internal interfaces. The new internal interfaces will attempt to use the gateway IP address of the selected virtual networks by default.

- d. [Optional] Select or deselect the **SNAT** check box to enable or disable SNAT on the external gateway of the router. With SNAT enabled, the router replaces VM private IP addresses with the public IP address of its external gateway.

Add virtual router ✕

Name

router1 *

Specify a network through which physical networks will be accessed.

Network

public: 10.136.16.0/20 ▼

☒ SNAT ?

Add internal interfaces + Add

private: 192.168.128.0/24 ▼

🗑

Cancel

Create

3. Click **Create**.

Managing router interfaces

Prerequisites

- You have a virtual router created, as described in "Managing virtual routers" (p. 85).

To add an external router interface

1. On the **Routers** screen, click the router name. Open the **Interfaces** tab to view the list of its interfaces.
2. Click **Add interface**.
3. In the **Add interface** window, do the following:

87

Copyright © 2016-2026 Virtuozzo International GmbH

- a. Select **External gateway**.
- b. From the **Network** drop-down menu, select a physical network to connect to the router. The new interface will pick an unused IP address from the selected physical network. You can also provide a specific IP address from the selected physical network to assign to the interface in the **IP address** field.
- c. [Optional] Select or deselect the **SNAT** check box to enable or disable SNAT on the external gateway of the router. With SNAT enabled, the router replaces VM private IP addresses with the public IP address of its external gateway.

Add interface

×

☒ External gateway
 ☐ Internal interface

Specify new interface parameters

Network

public: 10.136.16.0/20

▼

IP address (optional)

By adding a router interface you connect the selected network to the router. The new interface will pick an unused IP address from the selected physical network. You can also provide a specific IP address from the selected physical network to assign to the interface.

☒ SNAT
 i

Cancel

Add

4. Click **Add**.

To add an internal router interface

1. On the **Routers** screen, click the router name. Open the **Interfaces** tab to view the list of its interfaces.
2. Click **Add interface**.
3. In the **Add interface** window, select a network to connect to the router from the **Network** drop-down menu. The new interface will attempt to use the gateway IP address of the selected virtual network by default. If it is in use, specify an unused IP address from the selected virtual network to assign to the interface in the **IP address** field.

Add interface



Specify new interface parameters

Network
private2: 192.168.30.0/24



IP address (optional)

By adding a router interface you connect the selected network to the router. The new interface will attempt to use the gateway IP address of the selected virtual network by default. If it is in use, specify an unused IP address from the selected virtual network to assign to the interface.

Cancel

Add

4. Click **Add**.

To edit external interface parameters

1. On the **Routers** screen, click the line with the required router, and then **Edit gateway** on the router right pane. Alternatively, on the **Interfaces** tab, click the ellipsis icon next to the external interface, and then click **Edit**.
2. In the **Edit interface** window, change the IP address or configure SNAT.
3. Click **Save** to save your changes.

To remove a router interface

1. Select the interface you want to remove.
2. Click the ellipsis icon next to it, and then click **Delete**.
3. In the confirmation window, click **Delete**.

Managing static routes

You can also configure static routes of a router by manually adding entries into its routing table. This can be useful, for example, if you do not need a mutual connection between two virtual networks and want only one virtual network to be accessible from the other.

Consider the following example:

- The virtual machine **VM1** is connected to the virtual network **private1** (192.168.128.0/24) via the network interface with IP address 192.168.128.10.

- The virtual machine **VM2** is connected to the virtual network **private2** (192.168.30.0/24) via the network interface with IP address 192.168.30.10.
- The router **router1** connects the network **private1** to the physical network via the external gateway with the IP address 10.94.129.73.
- The router **router2** connects the network **private2** to the physical network via the external gateway with the IP address 10.94.129.74.

To be able to access **VM2** from **VM1**, you need to add a static route for **router1**, specifying the CIDR of **private2**, that is 192.168.30.0/24, as the destination subnet and the external gateway IP address of **router2**, that is 10.94.129.74, as the next hop IP address. In this case, when an IP packet for 192.168.30.10 reaches **router1**, it will be forwarded to **router2** and then to **VM2**.

Prerequisites

- You have a virtual router created, as described in "Managing virtual routers" (p. 85).

To create a static route for a router

1. On the **Routers** screen, click the router name. Open the **Static routes** tab, and then click **Add** on the right pane. If there are no routes to show, click **Add static route**.
2. In the **Add static route** window, specify the destination subnet range and mask in CIDR notation and the next hop's IP address. The next hop's IP address must belong to one of the networks that the router is connected to.

Add static route

×

Specify static route parameters

Destination subnet and mask

192.168.30.0/24

Next hop

10.94.129.74

The next hop's IP address must belong to one of the networks that the router is connected to.

Cancel

Add

3. Click **Add**.

To edit a static route

1. Click the ellipsis icon next to the required static route, and then click **Edit**.
2. In the **Edit static route** window, change the desired parameters, and then click **Save**.

To remove a static route

Click the ellipsis icon next to the static route you want to remove, and then click **Delete**.

Managing floating IP addresses


A virtual machine connected to a virtual network can be accessed from public networks, such as the Internet, by means of a floating IP address. Such an address is picked from a physical network and mapped to the VM's private IP address. The floating and private IP addresses are used at the same time on the VM's network interface. The private IP address is used to communicate with other VMs on the virtual network. The floating IP address is used to access the VM from public networks. The VM guest operating system is unaware of the assigned floating IP address.

Prerequisites

- You have a virtual router created, as described in "Managing virtual routers" (p. 85).
- The virtual machine to assign a floating IP to has a fixed private IP address.
- The virtual router connects the physical network, from which a floating IP will be picked, with the VM's virtual network.

To create a floating IP address and assign it to a virtual machine

1. On the **Floating IPs** screen, click **Add floating IP**.
2. In the **Add floating IP address**, select a physical network, from which a floating IP will be picked, and a VM network interface with a fixed private IP address.



The screenshot shows a modal dialog titled "Add floating IP address" with a close button (X) in the top right corner. The dialog contains two dropdown menus. The first dropdown is labeled "Network" and displays "public: 10.94.0.0/16". The second dropdown is labeled "Virtual machine" and displays "myvm — private: 192.168.128.6". At the bottom right of the dialog are two buttons: "Cancel" and "Add".

3. Click **Add**.

To re-assign a floating IP address to another virtual machine

1. Click the ellipsis icon next to the floating IP address, and then click **Unassign**.
2. Once the VM name disappears in the **Assigned to** column, click the ellipsis icon again, and then select **Assign**.

3. In the **Assign floating IP address** window, select a VM network interface with a fixed private IP address.
4. Click **Assign**.

To remove a floating IP address

1. Unassign it from a virtual machine. Click the ellipsis icon next to the floating IP address, and then click **Unassign**.
2. Click the ellipsis icon again, and then select **Delete**.

Managing security groups

A security group is a set of network access rules that control incoming and outgoing traffic to virtual machines assigned to this group. With security group rules, you can specify the type and direction of traffic that is allowed access to a virtual interface port.

All security groups applied to a VM are processed independently, with rule order—whether within a single security group or across multiple groups—having no impact on evaluation. The system checks each packet against all applicable rules and allows it if at least one rule permits it; otherwise, the packet is denied by default. Since security group rules are cumulative, overlapping or redundant rules do not create conflicts or affect processing. The system prioritizes allowing traffic when a matching rule exists and continues evaluating all rules across all assigned security groups, rather than stopping at the first match.

For each project, the **default** security group is automatically created in the compute cluster. This group allows all traffic on all ports for all protocols and cannot be deleted. When you attach a network interface to a VM, the interface is associated with the **default** security group, unless you explicitly select a custom security group.

You can assign one or more security groups to both new and existing virtual machines. When you add rules to security groups or remove them, the changes are enforced at runtime.

Creating and deleting security groups

Limitations

- You cannot delete a security group if it is assigned to a VM.

To create a security group

1. On the **Security groups** screen, click **Add security group**.
2. In the **Add security group** window, specify a name and description for the group, and then click **Add**.

Important

A description should not contain any personally identifiable information or sensitive business data.

Add security group



Name
mygroup

Description (optional)
A custom security group

Cancel

Add

By default, the new security group will deny all incoming traffic and allow only outgoing traffic to assigned virtual machines.

To delete a security group

1. On the **Security groups** screen, click the required security group.
2. On the group right pane, click **Delete**.
3. Click **Delete** in the confirmation window.

Managing security group rules

You can modify security groups by adding and removing rules. Rules can be created for both IPv4 and IPv6 traffic. However, existing rules cannot be edited. To change a rule, you must first delete it and then recreate it with the required parameters.

Prerequisites

- You have a security group created, as described in "Creating and deleting security groups" (p. 92).

To add a rule to a security group

1. On the **Security groups** screen, click the security group to add a rule to.
2. On the group right pane, click **Add** in the **Inbound** or **Outbound** section to create a rule for incoming or outgoing traffic.
3. Specify the rule parameters:
 - a. Select a protocol from the list or enter a number from 0 to 255.
 - b. Enter a single port or a port range. Some protocols already have a predefined port range. For example, the port for SSH is 22.
 - c. Select a predefined subnet CIDR or an existing security group.

Protocol ⓘ	Port range	Source ⓘ	
SSH ▾	22	0.0.0.0/0 ▾	✓ ✕

4. Click the check mark to save the changes.

As soon as the rule is created, it is applied to all of the virtual machines assigned to the security group.

To remove a rule from a security group

1. On the **Security groups** screen, click the required security group.
2. On the group right pane, click the bin icon next to a rule you want to remove.

As soon as the rule is removed, this change is applied to all of the virtual machines assigned to the security group.

Changing security group assignment

When you create a VM, you select security groups for the VM network interfaces. You can also change assigned security groups later.

Limitations

- You cannot configure security groups if spoofing protection is disabled or IP address management is disabled for the selected network.

To view virtual machines assigned to a security group

1. On the **Security groups** screen, click the required security group.
2. On the group right pane, navigate to the **Assigned VMs** tab. All the assigned virtual machines will be shown along with their status.

You can click the VM name to go to the VM **Overview** pane and change the security group assignment for its network interfaces.

To assign a security group to a virtual machine

1. On the **Virtual machines** screen, click the required virtual machine.
2. On the **Overview** tab, click the pencil icon in the **Networks** section.
3. Click the ellipsis icon next to the network interface to assign a security group to, and then click **Edit**.
4. In the **Edit network interface** window, go to the **Security groups** tab.
5. Select one or more security groups from the drop-down list, and then click **Save**.

The rules from chosen security groups will be applied at runtime.

Managing load balancers

Virtuozzo Cloud offers load balancing as a service for the compute infrastructure. Load balancing ensures fault tolerance and improves performance of web applications by distributing incoming network traffic across virtual machines from a balancing pool. A load balancer receives and then routes incoming requests to a suitable VM based on a configured balancing algorithm and VM health.

Creating load balancers

Limitations

- The forwarding rule and protocol cannot be changed after the load balancer pool is added.
- If an IPv6 subnet where a load balancer will operate works in the SLAAC or DHCPv6 stateless mode, the load balancer will receive an IPv6 address automatically.

Prerequisites

- A network where a load balancer will operate has IP management enabled.
- All VMs that will be added in balancing pools have fixed IP addresses.

To create a load balancer with balancing pools

1. On the **Load balancers** screen, click **Create load balancer**.
2. In the **Create load balancer** window, do the following:
 - a. Specify a name and, optionally, description.

Important

A description should not contain any personally identifiable information or sensitive business data.

- b. Enable or disable high availability:
 - With high availability enabled, two load balancer instances will be created. They will work in the Active/Standby mode according to the Virtual Router Redundancy Protocol (VRRP).
 - With high availability disabled, a single load balancer instance will be created.
- c. Select a flavor for the load balancer:
 - If high availability is enabled, you can only choose between load balancer flavors that will create two instances, one active and one standby. If the active instance becomes unhealthy, the instance automatically fails over to the standby instance, making it active.
 - If high availability is disabled, you can only choose between load balancer flavors that will create a standalone instance.

The screenshot shows a form for creating a load balancer. It includes a 'Name' field with the value 'lbaas1', a 'Description (optional)' field with the value 'Custom load balancer', a 'High availability' toggle switch that is turned on, and a 'Flavor' dropdown menu with the selected value 'ACTIVE_STANDBY'.

3. In the **Network settings** section, select the network that the load balancer will operate in and, optionally, specify an IP address that will be allocated to the load balancer.

- If you selected a virtual network that is connected to a physical network via a router
In this case, you can assign a floating IP address to the load balancer. To do it, select **Use a floating IP address**, and then choose either to use an available floating IP address or to create a new one.

The screenshot shows a configuration panel for a floating IP address. It includes a 'Network' dropdown menu with the value 'private1: 192.168.128.0/24'. Below it is a 'Load balancer IP version' dropdown menu set to 'IPv4'. There is an 'IP address (optional)' text input field. A checkbox labeled 'Use a floating IP address' is checked. At the bottom, there is a 'Floating IP address' dropdown menu with the option 'Create new' selected.

- If you selected a shared physical network with both IPv4 and IPv6 subnets
In this case, you need to choose the IP version that will be used for the load balancer.

The screenshot shows a configuration panel for a shared physical network. It includes a 'Network' dropdown menu with the value 'public: 10.136.16.0/22, 2001:bd8::/64'. Below it is a 'Load balancer IP version' dropdown menu set to 'IPv4'. There is an 'IP address (optional)' text input field.

4. In the **Balancing pools** section, create a balancing pool to forward traffic from the load balancer to virtual machines by clicking **Add**. In the **Create balancing pool** window that opens, do the following:
 - a. In the **Forwarding rule** section, select a forwarding rule from the load balancer to the backend protocol:
 - With the **HTTPS -> HTTPS** rule
 - i. Specify ports for incoming and destination connections.
 - ii. Ensure that all virtual machines have the same SSL certificate (or a certificate chain).
 - iii. [Optional] Enable the PROXY protocol version 1 to add a human-readable header with connection information (the source IP address, destination IP address, and port numbers) as a part of the request header.
 - With the **HTTPS -> HTTP** rule
 - i. Specify ports for incoming and destination connections.
 - ii. Upload an SSL certificate (or a certificate chain) in the PEM format and a private key in the PEM format.
 - iii. [Optional] Choose HTTP headers to insert into the request.
 - iv. [Optional] Enable the TLS encryption to re-encrypt traffic from the load balancer to its members.

- v. [Optional] Enable the PROXY protocol version 1 to add a human-readable header with connection information (the source IP address, destination IP address, and port numbers) as a part of the request header.
- With the **HTTP -> HTTP** rule
 - i. Specify ports for incoming and destination connections.
 - ii. [Optional] Choose HTTP headers to insert into the request.
 - iii. [Optional] Enable the TLS encryption to re-encrypt traffic from the load balancer to its members.
 - iv. [Optional] Enable the PROXY protocol version 1 to add a human-readable header with connection information (source IP address, destination IP address, and port numbers) as a part of the request header.
- With the **TCP -> TCP** rule
 - i. Specify ports for incoming and destination connections.
 - ii. [Optional] Enable the TLS encryption to re-encrypt traffic from the load balancer to its members.
- With the **UDP -> UDP** rule

Specify ports for incoming and destination connections.

The screenshot shows the configuration for the **HTTP -> HTTP** rule. At the top, there are three input fields: the rule name 'HTTP -> HTTP', the 'LB port' set to '80', and the 'Backend port' set to '80'. Below these is a dropdown menu for 'X-Forwarded-For'. Underneath the dropdown are two checked checkboxes: 'Enable the TLS encryption' with the description 'Traffic from the load balancer to members will be re-encrypted with the TLS protocol.', and 'Enable the PROXY protocol' with the description 'The load balancer will use PROXY protocol version 1 with a human-readable header format.'

- b. In the **Balancing settings** section, do the following:
 - i. Select the balancing algorithm:
 - **Least connections.** New connections will be forwarded to the VM with the fewest active connections.
 - **Round robin.** New connections will be forwarded to VMs in sequence. Individual HTTP requests may reuse existing connections and may not appear sequentially distributed.
 - **Source IP.** Connections from a unique source IP address will be directed to the same VM.
 - ii. [Optional] Select **Sticky session** to enable session persistence. The load balancer will generate a cookie that will be inserted into each response. The cookie will be used to send future requests to the same VM.

Note

This option is not available in the SSL passthrough mode.

- c. In the **Members** section, add members, that is, virtual machines, to the balancing pool by clicking **Add**. Each VM can be included to multiple balancing pools. In the **Add members** window that opens, select the desired VMs, and then click **Add**.

Note

You can select only between VMs that are connected to the chosen network.

- d. [Optional] In the **Allowed CIDRs** section, specify IP address ranges in the CIDR format that will be allowed to interact with the balancing pool. This will limit incoming traffic to the specified IP addresses, any other incoming traffic will be rejected. For example:
- To limit traffic from the IP address 10.10.10.10, add the /32 suffix: 10.10.10.10/32.
 - To limit traffic from the subnet range 10.10.10.0–10.10.10.255, add the /24 suffix: 10.10.10.10/24.
 - To limit traffic from the subnet range 10.10.0.0 - 10.10.255.255, add the /16 suffix: 10.10.10.10/16.
- e. In the **Health monitor** section, select the protocol that will be used for monitoring members availability:
- **HTTP/HTTPS**. The HTTP/HTTPS method GET will be used to check for the response status code 200. Additionally, specify the URL path to the health monitor.
 - **TCP/UDP**. The health monitor will check the TCP/UDP connection on the backend port.
 - **PING**. The health monitor will check members' IP addresses.

By default, the health monitor removes a member from a balancing pool if it fails three consecutive health checks of five-second intervals. When a member returns to operation and responds successfully to three consecutive health checks, it is added to the pool again. You can manually set the health monitor parameters, such as the interval after which VM health is checked, the time after which the monitor times out, healthy and unhealthy thresholds. To change the default parameters, click **Edit parameters**, enter the desired values, and then click **Save**.

Protocol
HTTP

URL path
/

The HTTP method GET will be used to check for the response status code 200.

Edit parameters

- f. Click **Create**.
5. [Optional] Add more balancing pools, as described above.
6. Click **Create**.

Managing balancing pools

To see a list of balancing pools in a load balancer, click its name.

Load balancers > LBaaS1

Search		+ Create balancing pool			
☐	Balancing pool	Status	Members state	Members total	
☐	 HTTP on port 80 → HTTP on port 80	Active		3	...
☐	 HTTPS on port 443 → HTTPS on port 443	Active		3	...

You can open the pool right pane to monitor its performance and health on the **Overview** tab, see its parameters on the **Properties** tab, manage its members on the **Members** tab, and configure its allowed IP ranges on the **CIDRs** tab.

Creating balancing pools

Limitations

- The forwarding rule and protocol cannot be changed after the load balancer pool is added.

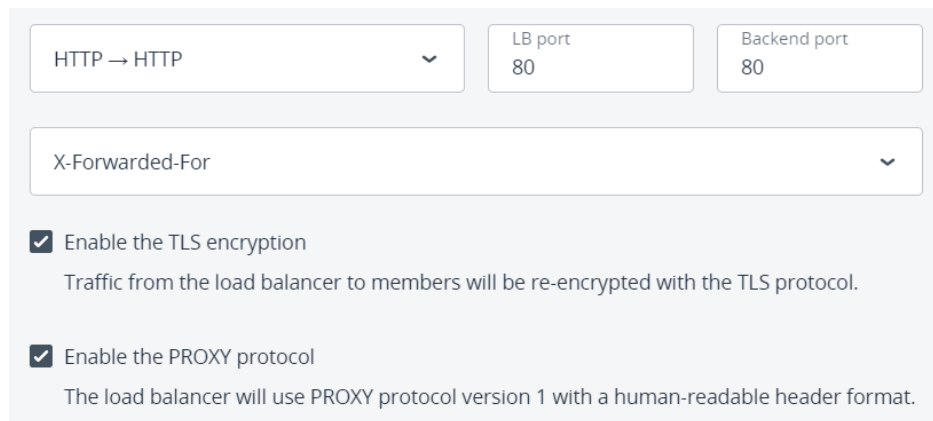
Prerequisites

- All VMs that will be added in balancing pools have fixed IP addresses.

To add a balancing pool to a load balancer

1. On the screen with balancing pools, click **Create balancing pool**.
2. In the **Forwarding rule** section, select a forwarding rule from the load balancer to the backend protocol:
 - With the **HTTPS -> HTTPS** rule
 - a. Specify ports for incoming and destination connections.
 - b. Ensure that all virtual machines have the same SSL certificate (or a certificate chain).
 - c. [Optional] Enable the PROXY protocol version 1 to add a human-readable header with connection information (the source IP address, destination IP address, and port numbers) as a part of the request header.
 - With the **HTTPS -> HTTP** rule
 - a. Specify ports for incoming and destination connections.
 - b. Upload an SSL certificate (or a certificate chain) in the PEM format and a private key in the PEM format.
 - c. [Optional] Choose HTTP headers to insert into the request.
 - d. [Optional] Enable the TLS encryption to re-encrypt traffic from the load balancer to its members.
 - e. [Optional] Enable the PROXY protocol version 1 to add a human-readable header with connection information (the source IP address, destination IP address, and port numbers) as a part of the request header.
 - With the **HTTP -> HTTP** rule
 - a. Specify ports for incoming and destination connections.
 - b. [Optional] Choose HTTP headers to insert into the request.

- c. [Optional] Enable the TLS encryption to re-encrypt traffic from the load balancer to its members.
- d. [Optional] Enable the PROXY protocol version 1 to add a human-readable header with connection information (source IP address, destination IP address, and port numbers) as a part of the request header.
- With the **TCP -> TCP** rule
 - a. Specify ports for incoming and destination connections.
 - b. [Optional] Enable the TLS encryption to re-encrypt traffic from the load balancer to its members.
- With the **UDP -> UDP** rule
Specify ports for incoming and destination connections.



HTTP → HTTP

LB port: 80

Backend port: 80

X-Forwarded-For

☒ Enable the TLS encryption
Traffic from the load balancer to members will be re-encrypted with the TLS protocol.

☒ Enable the PROXY protocol
The load balancer will use PROXY protocol version 1 with a human-readable header format.

3. In the **Balancing settings** section, do the following:
 - a. Select the balancing algorithm:
 - **Least connections.** New connections will be forwarded to the VM with the fewest active connections.
 - **Round robin.** New connections will be forwarded to VMs in sequence. Individual HTTP requests may reuse existing connections and may not appear sequentially distributed.
 - **Source IP.** Connections from a unique source IP address will be directed to the same VM.
 - b. [Optional] Select **Sticky session** to enable session persistence. The load balancer will generate a cookie that will be inserted into each response. The cookie will be used to send future requests to the same VM.

Note

This option is not available in the SSL passthrough mode.

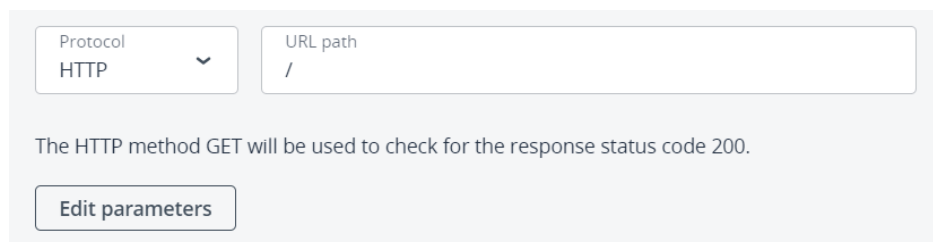
4. In the **Members** section, add members, that is, virtual machines, to the balancing pool by clicking **Add**. Each VM can be included to multiple balancing pools. In the **Add members** window that opens, select the desired VMs, and then click **Add**.

Note

You can select only between VMs that are connected to the chosen network.

5. [Optional] In the **Allowed CIDRs** section, specify IP address ranges in the CIDR format that will be allowed to interact with the balancing pool. This will limit incoming traffic to the specified IP addresses, any other incoming traffic will be rejected. For example:
 - To limit traffic from the IP address 10.10.10.10, add the /32 suffix: 10.10.10.10/32.
 - To limit traffic from the subnet range 10.10.10.0–10.10.10.255, add the /24 suffix: 10.10.10.10/24.
 - To limit traffic from the subnet range 10.10.0.0 - 10.10.255.255, add the /16 suffix: 10.10.10.10/16.
6. In the **Health monitor** section, select the protocol that will be used for monitoring members availability:
 - **HTTP/HTTPS**. The HTTP/HTTPS method GET will be used to check for the response status code 200. Additionally, specify the URL path to the health monitor.
 - **TCP/UDP**. The health monitor will check the TCP/UDP connection on the backend port.
 - **PING**. The health monitor will check members' IP addresses.

By default, the health monitor removes a member from a balancing pool if it fails three consecutive health checks of five-second intervals. When a member returns to operation and responds successfully to three consecutive health checks, it is added to the pool again. You can manually set the health monitor parameters, such as the interval after which VM health is checked, the time after which the monitor times out, healthy and unhealthy thresholds. To change the default parameters, click **Edit parameters**, enter the desired values, and then click **Save**.



7. Click **Create**.

The newly added pool will appear in the list of balancing pools.

Editing and deleting balancing pools

For a balancing pool, you can edit the balancing settings such as the HTTP headers, TLS encryption, balancing algorithm, and session persistence, as well as IP ranges that are allowed to interact with it and the health monitor parameters.

To edit the balancing settings

1. On the screen with balancing pools, click the required balancing pool.
2. On the right pane, click **Edit**.
3. Make the necessary changes and click **Save**.

To edit allowed CIDRs

1. On the screen with balancing pools, click the required balancing pool.
2. On the right pane, click **Edit allowed CIDRs**.
3. In the **Edit allowed CIDRs** window, do the following:
 - To add a CIDR, click **Add** and specify an IP address range in the CIDR format that will be allowed to interact with the balancing pool.
 - To change a CIDR, click the pencil icon next to the desired CIDR and make your edits.
 - To remove a CIDR, click the bin icon next to the desired CIDR.
4. Click **Save** to apply your changes.

To edit the health monitor parameters

1. On the screen with balancing pools, click the required balancing pool.
2. On the right pane, click **Edit health monitor**.
3. Make the necessary changes and click **Save**.

To delete a balancing pool

1. On the screen with balancing pools, click the required balancing pool.
2. On the right pane, click **Delete**.
3. Click **Delete** in the confirmation window.

Managing balancing pool members

Prerequisites

- All VMs that will be added in balancing pools have fixed IP addresses.

To add members to a balancing pool

1. On the screen with balancing pools, click the required balancing pool.
2. On the right pane, click **Add members**.
3. In the **Add members** window, select virtual machines to be added to the balancing pool, and then click **Add**.

To disable or enable members of a balancing pool

1. On the screen with balancing pools, click the required balancing pool.
2. On the right pane, go to the **Members** tab.
3. Click the ellipsis icon next to the required member, and then click **Disable** or **Enable**, depending on the member's current state.

To remove members from a balancing pool

1. On the screen with balancing pools, click the required balancing pool.
2. On the right pane, go to the **Members** tab.
3. Click the ellipsis icon next to the required member, and then click **Remove**.

Editing and deleting load balancers

To edit the name or description of a load balancer

1. On the **Load balancers** screen, click a load balancer you want to edit.
2. On the load balancer right pane, click **Edit**.
3. In the **Edit load balancer** window, modify the name or description, and then click **Save**.

Important

A description should not contain any personally identifiable information or sensitive business data.

To disable or enable a load balancer

1. On the **Load balancers** screen, click a load balancer you want to change.
2. On the load balancer right pane, click **Disable** or **Enable**, depending on the load balancer's current state.

To remove a load balancer

1. On the **Load balancers** screen, click a load balancer to delete.
2. On the load balancer right pane, click **Delete**.
3. Click **Delete** in the confirmation window.

Monitoring load balancers

To monitor performance and health of a load balancer

Open the **Overview** tab on the load balancer right pane.

The following charts are available:

Members state

The total number of members in the balancing pools grouped by status: "Healthy," "Unhealthy," "Error," and "Disabled".

Network

Incoming and outgoing network traffic.

Active connections

The number of active connections.

Error requests

The number of error requests.

Managing backups

A backup, or recovery point (these terms are used interchangeably), can be a copy either of a compute volume or only of data changes that is made at a specified time. With the backup service, you can create backups automatically by using backup plans or initiate backups manually. Backup plans define what data to back up, how frequently to create backups, and how long to keep them.

The backup service also allows you to restore virtual machines and volumes by creating new instances from backups.

The following backup types are supported:

- **Full.** A full backup contains a copy of an entire compute volume. It is self-sufficient, meaning that you can restore data from just one such a backup. Full backups are time-consuming and take up a large amount of storage space.
- **Incremental.** An incremental backup only copies changes to the data since the latest backup, regardless of its type. Incremental backups are usually faster and require less storage space. However, restoring data from incremental backups is more complex, as it requires the full backup chain consisting of the first full backup and all subsequent incremental backups.

Creating backup plans

You can schedule an automatic backup job for multiple volumes by creating a backup plan. A backup plan defines the following parameters:

- **Backup selection:** Select one or multiple volumes per backup plan. However, note that one volume cannot be added to multiple backup plans.
- **Backup frequency:** Choose when and how often backups will be created.
- **Backup retention:** Set the maximum number of full backups to retain at all times.

Backup chains

If incremental backups are enabled in your compute cluster, the total number of stored backups can be calculated as:

max. number of full backups * chain length

A backup chain consists of:

- one full backup (F)
- all incremental backups (i) created after it until the next full backup

By default, the chain length is set to 7 (one full backup and six incremental backups), but this value can be modified by the system administrator.

Backup rotation

When a new full backup is created and the retention limit is exceeded, the oldest full backup and its associated incremental backups are automatically deleted.

Consider the following examples:

- Backup retention is set to 1

Only one full backup and its incremental backups are retained. When a new full backup is created, the previous full backup and all its increments are deleted.

1. The backup chain reaches its maximum length:

```
F1 → i1 → i1 → i1 → i1 → i1 → i1
```

2. A new full backup is created:

```
F1 → i1 → i1 → i1 → i1 → i1 → i1 → F2
```

3. The previous chain is deleted:

```
F2
```

4. New incremental backups are created:

```
F2 → i2 → ...
```

- Backup retention is set to 2

Up to two full backup chains are preserved. When a third full backup is created, the oldest chain is deleted.

1. The second backup chain reaches its maximum length:

```
F1 → i1 → i1 → i1 → i1 → i1 → i1 → F2 → i2 → i2 → i2 → i2 → i2 → i2
```

2. A new full backup is created:

```
F1 → i1 → i1 → i1 → i1 → i1 → i1 → F2 → i2 → i2 → i2 → i2 → i2 → i2 → F3
```

3. The oldest chain is deleted:

```
F2 → i2 → i2 → i2 → i2 → i2 → i2 → F3
```

4. New incremental backups are created:

```
F2 → i2 → i2 → i2 → i2 → i2 → i2 → F3 → i3 → ...
```

At any given time, the system retains only the configured number of full backup chains. Older chains are deleted automatically once the retention threshold is exceeded.

To create a backup plan

1. On the **Backup plans** screen, click **Create backup plan**.
2. In the **Create backup plan** window, specify a name for the backup plan and, optionally, a description.

Important

A description should not contain any personally identifiable information or sensitive business data.

3. In **What to back up**, click **Manage**. In the **Manage volumes** window, select compute volumes that will be included in the backup plan, and then click **Save**.
4. In **Schedule**, select the schedule for the backup plan:
 - Select **Retention 7 days** to create a backup of the selected volumes every day at 2 AM (UTC time zone) and keep 7 full recovery points at most.
 - Select **Retention 14 days** to create a backup of the selected volumes every day at 2 AM (UTC time zone) and keep 14 full recovery points at most.
 - Select **Custom** to configure a custom schedule for the automatic backup. You can choose months, days of the week, days of the month, time, and the maximum number of full recovery points to keep.

Configure the schedule for the automatic backup (UTC time zone).

Backup schedule Custom	
Month Every month	Days of week Every day
Days of month All days	Start at 06:00
Max. number of recovery points 5	

5. Click **Create**.

Managing volumes in backup plans

You can manage compute volumes that you want to back up by adding them to or removing them from your backup plans. After removing a volume from a backup plan, all backups that have been already created remain intact.

Limitations

- A volume cannot be added to multiple backup plans.

Prerequisites

- A backup plan is created, as described in "Creating backup plans" (p. 104).

To add volumes to a backup plan

1. On the **Backup plans** screen, click the required backup plan.
2. On the plan right pane, navigate to the **Volumes to back up** tab. All the compute volumes included in the backup plan will be shown here.
3. Click **Manage** above the list of volumes.
4. In the **Manage volumes** window, select volumes that you want to back up, and then click **Save**.

To remove volumes from a backup plan

1. On the **Backup plans** screen, click the required backup plan.
2. On the plan right pane, navigate to the **Volumes to back up** tab. All the compute volumes included in the backup plan will be shown here.
3. Click **Manage** above the list of volumes.
4. In the **Manage volumes** window, remove the selection from volumes that you do not want to back up. To see only the volumes assigned to the backup plan, select **Show only selected items** next to the **Search** field. Then, click **Save**.

Editing and deleting backup plans

You can edit a backup plan's name and description, change its schedule, and delete it when it is no longer needed.

Prerequisites

- A backup plan is created, as described in "Creating backup plans" (p. 104).

To edit a backup plan

1. On the **Backup plans** screen, click the required backup plan.
2. On the plan right pane, click **Edit**.
3. In the **Edit backup plan** window, make the required changes, and then click **Save**.

Important

A description should not contain any personally identifiable information or sensitive business data.

To delete a backup plan

1. On the **Backup plans** tab, click the required backup plan.
2. On the plan right pane, click **Delete**.
3. If the backup plan has recovery points, you can delete them along with the backup plan:
 - a. Select **Delete recovery points**.
 - b. Enter "Delete" for confirmation.
4. Click **Delete** in the confirmation window.

Creating and deleting backups manually

You can initiate an instant backup job for a single volume by creating a backup manually. Such a backup does not have a retention policy and can only be deleted manually. Note that when you delete a backup, all dependent backups in the backup chain are also deleted.

To manually create a volume backup

1. On the **Volumes** screen, click a volume that you want to back up.
2. On the volume right pane, click **Create backup now**.

Once the backup is created, it will appear on the **Recovery points** screen.

To manually delete a volume backup

1. On the **Recovery points** screen, click the recovery point that you want to delete.
2. Click **Delete** in the confirmation window.

After deleting a recovery point, all its data will be lost.

Restoring volumes from backups

During the restore process, a new volume is created and the existing volume is not overwritten. You can restore a volume from a backup of a data or boot volume.

Prerequisites

- A volume backup is created automatically, as described in "Creating backup plans" (p. 104), or manually, as described in "Creating and deleting backups manually" (p. 108).

To restore a volume

1. On the **Recovery points** screen, click the recovery point from which you want to restore a volume.
2. On the right pane, click **Restore volume**.
3. In the **Restore volume** window, specify a volume name and select a storage policy, and then click **Restore**.

Restore volume

×

Create a new volume from the recovery point "mybackup".

Name

myvolume

Volume size

1 GiB

Storage policy

default

▼

Cancel

Restore

The new volume will appear on the **Volumes** screen.

Restoring virtual machines from backups

During the restore process, a new virtual machine is deployed and the existing VM is not overwritten. You can restore a VM from a backup of a boot volume only.

Prerequisites

- A volume backup is created automatically, as described in "Creating backup plans" (p. 104), or manually, as described in "Creating and deleting backups manually" (p. 108).

To restore a virtual machine

1. On the **Recovery points** screen, click the recovery point from which you want to restore a VM.
2. On the right pane, click **Restore virtual machine**.
3. In the **Restore virtual machine** window, the boot volume will be defined automatically and will be restored from the selected recovery point. Specify all other VM parameters, as described in "Creating virtual machines" (p. 13).

Restore virtual machine
✕

Review the virtual machine details and go back to change them if necessary.

Name
myvm

Deploy from:
☐ Image
☒ Volume

Volumes	Volume from recovery point — 1 GiB, default Boot	✎
Flavor	tiny — 1 vCPU, 512 MiB RAM	✎
Network interfaces	private — Auto Primary IP: Auto Security groups: 1	✎
SSH key (optional)	Specify	✎
Customization script (optional)	Specify	✎

Advanced options
>

Cancel

Deploy

4. Click **Deploy**.

The new virtual machine will appear on the **Virtual machines** screen.

Managing SSH keys

Use of SSH keys allows you to secure SSH access to virtual machines. You can generate a key pair on a client from which you will connect to VMs via SSH. The private key will be stored on the client and you will be able to copy it to other nodes. The public key will need to be uploaded to Virtuozzo Cloud and specified during VM creation. It will be injected into the VM by cloud-init and used for OpenSSH authentication. Keys injection is supported for both Linux and Windows virtual machines.

Limitations

- You can specify an SSH key only if you deploy a VM from a template or boot volume (not an ISO image).
- If a key has been injected into one or more VMs, it will remain inside those VMs even if you delete it from the panel.

Prerequisites

- The `cloud-init` utility and OpenSSH Server are installed in a VM template or boot volume, as instructed in "Preparing templates" (p. 63).

To add a public key

- Generate an SSH key pair on a client by using the `ssh-keygen` utility:

```
# ssh-keygen -t rsa
```

- On the **SSH keys** screen, click **Add key**.
- In the **Add SSH key** window, specify a key name and copy the key value from the generated public key located in `/root/.ssh/id_rsa.pub`. Optionally, you can add a key description.

Important

A description should not contain any personally identifiable information or sensitive business data.

Add SSH key ✕

For the key to be successfully injected into the VM, the template must contain the cloud-init package.

Name

root_node001vstoragedomain

Description (optional)

My public key

Key value

```
n1h0culzlqbj2AHYqglUWX7W3bE3nCCUxEX9DuHH2GJPy8Kz7Hka
RY0GULMIOJz7QRyzwbThgQ3T11YX+OJ5i7kbUek9hygy+RR/kjnMMI
rg6gyP2b4BrDflpZUNx4Nx1L9IGCGUoTWPieIc0n2LQMh2fAfxBBh
mSDVUPBLpowxuAlbOOkemW5IDJsKxuDuIqt35X27anWPcjFKTZN
47RnyCDT/X6tBYdxQJ6ARIQsp1JDWkjN7B65h9rwNZJ/PpyXI5wEVh
SLXrlMam93bh3YwMzQYhVILXGuvgbP+dF5Cq6Bg8FthXEfktpt121
5P/FD root@node001.vstoragedomain
```

Cancel

Add

To delete a public key

1. On the **SSH keys** screen, select the SSH key you want to delete, and then click **Delete**.
2. Click **Delete** in the confirmation window.

If this key has been injected into one or more virtual machines, it will remain inside those virtual machines.

Managing migrations from external platforms

This section describes migration of virtual machines from external platforms (such as VMware vSphere) into Virtuozzo Cloud.

Before starting a migration, ensure that your environment and virtual machines meet the supported configurations described in "Supported migration configurations" (p. 112) and the requirements described in "VMware migration requirements" (p. 114).

Migration is performed in multiple stages. First, you create a migration that defines the source environment, target configuration, and mapping settings. Then, you replicate data from the source platform to the target environment. Finally, you deploy the migrated virtual machines.

Only VM disks are replicated. VM memory and runtime state are not transferred.

Before creating migrations, you must add and configure source clouds.

Limitations

- This feature is currently in early access and intended for evaluation purposes only. Use in production environments is not recommended.
- Software RAID configurations are not supported for migration.
- Concurrent replication of multiple disks per VM is not supported. Currently, disks are replicated consecutively.

Prerequisites

- The migration service is installed by a system administrator. It can be deployed along with the compute cluster or later.

Supported migration configurations

This section describes supported configurations for migrating virtual machines from VMware vSphere to Virtuozzo Cloud. The table below summarizes compatibility across operating systems, firmware types, and migration options based on internal testing results.

Note

The migration service supports modern operating systems that are officially supported by their vendors, such as Windows 11. Legacy operating systems are supported by Virtuozzo on a best-effort basis with rollback safety. For example, Windows 10 migration has been tested even though it has reached end of life..

The table uses the following notation:

- **Yes** – supported configuration
- **No** – not supported configuration
- **N/A** – not tested

Windows

Guest OS	UEFI	Secure Boot	BIOS	CBT		DHCP		Fixed IPs	
				On	Off	On	Off	On	Off
Windows 11	Yes	N/A	N/A	Yes	N/A	Yes	N/A	N/A	Yes
Windows 10	Yes	N/A	Yes	Yes	N/A	Yes	N/A	N/A	Yes
Windows Server 2025	Yes	Yes	N/A	Yes	N/A	Yes	N/A	N/A	Yes
Windows Server 2022	Yes	Yes	Yes	Yes	N/A	Yes	N/A	N/A	Yes
Windows Server 2019	Yes	Yes	N/A	Yes	N/A	Yes	N/A	N/A	Yes
Windows Server 2016	Yes	Yes	Yes	Yes	N/A	Yes	N/A	N/A	Yes
Windows Server 2012 R2	N/A	N/A	Yes	Yes	N/A	Yes	N/A	N/A	Yes
Windows Server 2008 R2	N/A	N/A	Yes	Yes	N/A	Yes	N/A	N/A	Yes

Linux

Guest OS	Version	UEFI	Secure Boot	BIOS	CBT		DHCP		Fixed IPs	
					On	Off	On	Off	On	Off
Ubuntu	24.04.x LTS, 22.04.x LTS,	Yes	Yes	N/A	Yes	N/A	Yes	Yes	N/A	Yes
	20.04.x LTS	Yes	Yes	N/A	Yes	N/A	Yes	Yes	Yes	Yes
	18.04 LTS	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SUSE Linux Enterprise	15.x	N/A								
Red Hat Enterprise Linux	10.x, 9.x, 8.x, 7.x	Yes	Yes	Yes	Yes	N/A	Yes	Yes	Yes	Yes
Debian	13.x, 12.x, 11.x, 10.x	Yes	Yes	Yes	Yes	N/A	Yes	Yes	Yes	Yes
	9.x	N/A								
AlmaLinux	10.x, 9.x, 8.x	Yes	Yes	Yes	Yes	N/A	Yes	Yes	Yes	Yes
Oracle Linux	10.x, 9.x, 8.x	Yes	Yes	Yes	Yes	N/A	Yes	Yes	Yes	Yes

Guest OS	Version	UEFI	Secure Boot	BIOS	CBT		DHCP		Fixed IPs	
					On	Off	On	Off	On	Off
Rocky Linux	10.x, 9.x, 8.x	Yes	Yes	Yes	Yes	N/A	Yes	Yes	Yes	Yes
CentOS	Stream 10.x	Yes	Yes	Yes	Yes	N/A	Yes	Yes	Yes	Yes
	Stream 9.x	Yes	Yes	Yes	Yes	N/A	Yes	Yes	Yes	Yes
	8.x, 7.x	Yes	No	Yes	Yes	N/A	Yes	Yes	Yes	Yes

VMware migration requirements

Before migrating virtual machines from VMware vSphere, make sure that the VMware environment and source virtual machines meet the following requirements.

VMware connectivity requirements

When VMware vCenter is used as the migration source, the migration service connects directly to the selected ESXi hosts for disk transfer operations. Connectivity to vCenter alone is not sufficient.

- All ESXi hosts selected as migration sources and the vCenter server, if used, are reachable from the Virtuozzo Cloud cluster.
- TCP ports 443 and 902 are open between the Virtuozzo Cloud cluster and all ESXi hosts selected as migration sources. If vCenter is used, TCP port 443 must also be open to the vCenter server.
- DNS resolution and routing are configured correctly for all ESXi hosts. If ESXi hosts cannot be resolved through DNS, add the required host entries to `/etc/hosts` on the cluster nodes.

Source VM requirements

- The guest operating system and migration options are listed as supported in "Supported migration configurations" (p. 112).
- VMware Tools is installed and up to date on all source VMs. VMware Tools is required for collecting guest network information and creating quiesced snapshots of running VMs.
- Source VM disks use Changed Block Tracking (CBT)-compatible storage, such as VMFS or NFS.
- VM disks are not in Independent mode.
- VM disks are not encrypted with BitLocker, LUKS, or similar disk encryption software.
- Security software does not block access to the boot loader or disks. This includes VMware AppDefense Guest Integrity and antivirus or endpoint protection products such as Acronis, Kaspersky, vShield, Symantec, McAfee, and CrowdStrike.
- Linux VMs have a working package manager configuration. To verify that package repositories are accessible, run `apt update` on Debian-based distributions or `yum repolist` or `dnf repolist` on RHEL-based distributions.
- Linux VMs with static IP addresses store their network configuration in supported configuration files, such as Netplan or NetworkManager profiles.

- Linux VMs based on Oracle Linux, Red Hat Enterprise Linux, or CentOS have at least 500 MB of free space on the /boot partition.

If a Linux VM uses LVM volumes spanning multiple physical volumes (PVs), verify the disk layout before migration and make sure that the selected target flavor can boot from this disk layout:

```
lsblk -f
```

Managing source clouds

Before creating migrations, you must add a source cloud. A source cloud represents an external platform (for example, a VMware vSphere deployment) from which virtual machines are migrated.

After adding a cloud, you can validate the connection to ensure that the credentials and endpoint are correct. Once a cloud is validated, its virtual machines become available for selection when creating migrations. A source cloud can be reused for multiple migrations.

Limitations

- Only VMware vSphere (6.x–8.x) is supported as a source platform.

To add a source cloud

1. Go to the **Clouds** screen and click **Add cloud**.
2. In the **Add cloud** window, select the cloud type (currently, only **VMware** is available).
3. Specify the following parameters:
 - A name for the cloud and optionally its description
 - User account credentials for accessing the source platform
 - IP address or hostname of the VMware endpoint and the port number (for example, 443)

Add cloud

×

vmware®

Name

mycloud

Description (optional)

VMware vSphere hostname

10.10.10.2

Port

443

i

Username

admin

Password

.....

👁

☐

Disable SSL checks

i

Back

Add

4. [Optional] Select **Disable SSL checks** to allow connections to endpoints with self-signed or untrusted certificates.
5. Click **Add**.

The cloud will appear in the list.

To validate a source cloud

1. On the **Clouds** screen, select the required cloud.
2. On the cloud right pane, click **Validate**.

To edit a source cloud

1. On the **Clouds** screen, select the required cloud.
2. On the cloud right pane, click **Edit**.

3. In the **Edit cloud** window, make the required changes and click **Save**.

To remove a source cloud

1. On the **Clouds** screen, select the required cloud.
2. On the cloud right pane, click **Delete**.
3. In the confirmation window, click **Delete**.

Preparing the Windows migration worker template

To migrate Windows virtual machines, prepare a Windows migration worker template. During migration, the system creates a temporary worker VM from this template. After migration is complete, the temporary worker VM is deleted.

Windows migrations require both Linux and Windows worker templates: disk cloning uses a Linux worker, and OS morphing uses a Windows worker. In Virtuozzo Cloud, Linux worker templates are already available, so you only need to prepare the Windows template.

Important

The Windows worker template must be based on the same Windows version as the VMs you plan to migrate, or on a newer version. Also, the system language in the Windows template must be set to English.

To prepare the Windows migration worker template

1. Prepare a compute image with an official Windows ISO, as described in "Uploading images" (p. 61).
2. Create a Windows VM to use as the migration worker template, as described in "Creating virtual machines" (p. 13). Use a boot volume of no larger than 15 GB and ensure that the VM is connected to a network with Internet access.
3. Install Windows in the VM. If you are installing Windows Server, select an edition with the Desktop Experience option.
4. Download Cloudbase-Init from <https://cloudbase.it/cloudbase-init/#download>, and then install it by following the procedure from the **Installation** section at <https://cloudbase.it/cloudbase-init/>. During the installation, do the following:
 - a. Specify the default user name **Administrator**. The migration service relies on this user, so the **Administrator** account must exist and must not be disabled in the template.
 - b. Set a password for the **Administrator** account. You will need to set the same password for the migration service after uploading the temporary Windows worker image.
 - c. When closing the Cloudbase-Init setup wizard, do not run Sysprep.
5. Allow ICMPv4 echo requests in Windows Firewall. Start Windows PowerShell as administrator and run:

```
> New-NetFirewallRule -DisplayName "Allow ICMPv4 Echo Request" -Direction Inbound -  
Protocol ICMPv4 -IcmpType 8 -Action Allow -Profile Any
```

6. Configure WinRM over HTTPS, as described in [Microsoft documentation](#). The VM must have a valid local computer certificate for Server Authentication. The certificate must match the VM hostname and must not be expired, revoked, or self-signed.

Start Windows PowerShell as administrator and run:

```
> winrm quickconfig -transport:https
> New-NetFirewallRule -DisplayName "Allow WinRM HTTPS" -Direction Inbound -Protocol
TCP -LocalPort 5986 -Action Allow -Profile Any
> winrm enumerate winrm/config/listener
```

7. Obtain the VM ID and hosting node. On a cluster node, run:

```
# vinfra service compute server list
```

8. Download and save the guest tools ISO and Cloudbase-Init ZIP inside the VM. Log in to the node that hosts the VM and run:

```
# virsh x-exec <VM_ID> -- powershell -NoProfile -Command "$WebClient = New-Object
System.Net.WebClient; $WebClient.DownloadFile('https://repo.virtuozzo.com/vz-
platform/guest-tools/vz-guest-tools-win.iso', 'C:\vz-guest-tools-win.iso'); Write-
Host 'Downloaded vz-guest-tools-win.iso'"
# virsh x-exec <VM_ID> -- powershell -NoProfile -Command "$WebClient = New-Object
System.Net.WebClient; $WebClient.DownloadFile
('https://www.cloudbase.it/downloads/CloudbaseInitSetup_x64.zip',
'C:\cloudbaseinit.zip'); Write-Host 'Downloaded cloudbaseinit.zip'"
```

The files are only downloaded and saved to disk. Do not unpack or install them.

9. Mount C:\vz-guest-tools-win.iso inside the VM and copy all files from the ISO to C:\vztools.
10. Install Windows ADK Deployment Tools to provide the legacy DISM tool required by the migration worker:

```
> Invoke-WebRequest -Uri "https://download.microsoft.com/download/d5be3ad2-c9b2-4034-
b02d-5e6a3a0b9780/adk/adksetup.exe" -OutFile "C:\adksetup.exe"
> C:\adksetup.exe /quiet /features OptionId.DeploymentTools
> & 'C:\Program Files (x86)\Windows Kits\10\Assessment and Deployment Kit\Deployment
Tools\amd64\DISM\dism.exe'
```

11. Shut down the VM. Then obtain the VM boot volume ID from the volumes field and upload this volume to an image with the predefined name **vporter-minion-windows**:

```
# vinfra service compute server show <VM_ID>
# vinfra service compute volume upload-to-image --name vporter-minion-windows
<volume_id>
```

12. Set the same password for the temporary Windows worker that you specified for the **Administrator** account in step 4b:

```
# vinfra service compute maas set --minion-windows-password <password>
```

The migration service uses this image name to create the temporary Windows migration worker during migration.

Creating migrations

A migration defines how virtual machines are transferred from a source cloud to the target environment. It includes the selection of source VMs, target configuration, and mapping of networks and storage.

Changed Block Tracking (CBT) is enabled automatically on source VMs. It allows disk replication to run while source VMs remain online. Source VM downtime is expected only during the final switch to the target environment and OS morphing, when the source VM is stopped and the target VM is started.

Limitations

- Fixed IP addresses can be retained only if the source VM is powered on.

Prerequisites

- A source cloud is added, as described in "Managing source clouds" (p. 115).
- For Windows migrations, the migration worker template is prepared and uploaded into the compute cluster, as described in "Preparing the Windows migration worker template" (p. 117). (Linux migration worker templates are provided by default.)
- You have a network for target virtual machines. It can be either a shared physical network or a virtual network linked to a physical one via a virtual router. The virtual network needs to have a gateway and a DNS server specified.

To create a migration

1. Go to the **Migrations** screen and click **Migrate**. The migration wizard will open where you will need to specify the migration parameters.
2. On the **Source cloud** step, select a source cloud from which to migrate virtual machines.
3. On the **Source VMs** step, select one or more virtual machines to migrate.
4. On the **Source options** step, configure source-side options for the selected VMs. Select the **vixDiskLib compatibility mode** for your VMware environment. The default version 8.0 supports vSphere 6.7–8.x.
5. On the **Network mapping** step, select target networks for each source network.
6. On the **Target options** step, configure the target environment:
 - **Enable DHCP client** to automatically configure networking for the target VM.
 - Enable **Preserve fixed IPs** to retain source IP addresses on the target VM.
 - Select the correct **OS distribution** for the target VM. This parameter should match the source VM to ensure correct driver and configuration handling.
 - [Optional] Specify advanced settings (if not specified, default values are applied):
 - Select a **Flavor** (CPU and RAM configuration) for the target VM. If not specified, a suitable flavor is created automatically.

Important

The selected flavor must not be smaller than the source VM, as this may cause issues or result in an unusable VM due to insufficient resources.

- Select the **Disk bus** type: SCSI, IDE, or VirtIO. This parameter should match the source VM configuration to ensure compatibility and performance.
- Select the **SSH key pair** to inject into the target VM for access.
- Select the VM **Machine type** on the target platform: pc-i440fx or pc-q35. This parameter should match the source VM to ensure successful boot and compatibility.
- Enable **Retain user credentials** to preserve existing OS user credentials during migration.
- Enable **Delete disks on VM deletion** to automatically remove attached disks when the VM is deleted.
- If you selected a virtual compute network as a target network, you can enable **Use floating IPs** and select the **Floating IP pool** on the target platform to allocate floating IPs to the target VM.

The screenshot displays a configuration panel for a virtual machine. At the top, there are two toggle switches: 'Enable DHCP client' (turned on) and 'Preserve fixed IPs' (turned off). Below these is a text field for 'OS distribution' set to 'Ubuntu 24.04'. A section titled 'Advanced settings' with a dropdown arrow contains several more fields: 'Flavor' (set to 'Select'), 'Disk bus' (set to 'VirtIO'), 'SSH key pair' (set to 'Select'), and 'Machine type' (set to 'pc-q35'). Below the advanced settings, there are three more toggle switches: 'Retain user credentials' (turned on), 'Delete disks on VM deletion' (turned off), and 'Use floating IPs' (turned on). At the bottom, there is a text field for 'Floating IP pool' set to 'Public'.

7. On the **Storage mapping** step, select a storage policy for target VM disks.
8. On the **User scripts** step, upload custom scripts to run on the guest OS during migration. Scripts can be uploaded per OS type (Windows or Linux). They are executed before OS morphing and only when OS morphing is enabled.

Supported script types include:

- For Linux VMs: any scripting language available on a standard Linux system (bash, python, etc.)
- For Windows VMs: PowerShell and batch scripts are supported.

User scripts can be used for the following scenarios:

- Install or remove packages inside the VM
 - Make additional configuration changes
 - Download files on the VM
 - Install third-party drivers
9. On the **Migration options** step, configure how replication and deployment are performed:
- Enable **Start replication immediately** to begin replication after creating the migration. If enabled, you can configure the following options:
 - Enable **Shut down source VMs** to shut down source VMs during the final switch to the target environment.
 - Enable **Auto-deploy VMs** to automatically deploy VMs after replication.

Note

Auto-deploy is typically used for one-time migrations. For incremental synchronization workflows, keep this option disabled, run one or more replications, and then create the deployment manually after the final replication.

- Enable **Clone disks** to clone disks on the target platform instead of moving them. This allows continuous incremental synchronization to the same disks. If disabled, disks must be fully re-synced during the next replication.
- **Enable OS morphing** to prepare the guest OS to run on the target platform by applying required drivers and removing source-specific components.

Note

This step can be skipped when migrating between platforms using the same hypervisor (for example, KVM to KVM).

- Enable **Start target VMs after deployment** to start target VMs when they are deployed.
10. On the **Summary** step, review the configuration. You can return to previous steps to make changes without losing data. Click **Migrate** to create the migration.

After creation, the migration appears in the **Migrations** list.

Depending on the selected options, replication starts immediately, or the migration is created and can be replicated later.

You can monitor the migration progress and manage replications and deployments on the migration details page.

Viewing migration details

After migration creation, you can view its configuration, status, and progress, as well as monitor replications and deployments.

Migration, replication, and deployment statuses are updated automatically, but you can use **Fetch** to retrieve the latest state immediately.

Prerequisites

- A migration is created, as described in "Creating migrations" (p. 119).

To view migration details

On the **Migrations** screen, click the name of the required migration. The migration details page opens.

The **Overview** tab displays general information about the migration:

- Source platform used for migration
- Target domain and project for migrated VMs
- Current migration state
- Unique identifier of the migration
- Created and updated timestamps
- Configuration of source and target options
- Source VM details (you can use **Fetch** to retrieve the latest VM state)
- Mapping of source and target networks

The **Replications** tab lists all replications for the migration. For each replication, you can view its status, amount of replicated data, replication speed, and duration. Select a replication to view its details and associated tasks.

The **Deployments** tab lists all deployments for the migration. For each deployment, you can view its status and date. Select a deployment to view its details and associated tasks.

Replicating migrations

Replication transfers data from source virtual machines to the target environment. It creates or updates target disks without deploying the virtual machines.

You can run multiple replications for the same migration to keep data in sync before deployment. This reduces the amount of data transferred during the final synchronization and minimizes downtime when switching to the target environment.

You can start replication immediately when creating a migration or start it later on the **Replications** tab.

Prerequisites

- A migration must be created before replication, as described in "Creating migrations" (p. 119).

To replicate a migration

1. On the **Migrations** screen, click the name of the required migration.
2. On the **Replications** tab, click **Create replication**.
3. In the **Create replication** window, configure the replication options:
 - Enable **Shut down instances** to shut down source VMs during the final switch to the target environment.
 - Enable **Auto-deploy VMs** to automatically deploy VMs after replication completes.

Note

Auto-deploy is typically used for one-time migrations. For incremental synchronization workflows, keep this option disabled, run one or more replications, and then create the deployment manually after the final replication.

4. Click **Create.**

After replication starts, data is transferred from the source to the target environment. You can monitor the replication progress and status in the replication details.

If **Auto-deploy VMs** is enabled, deployment starts automatically after replication completes.

To cancel a replication

Important

Canceling a replication takes effect immediately and cannot be reversed.

1. On the **Migrations** screen, click the name of the required migration.
2. On the **Replications** tab, click the ongoing replication.
3. On the replication right pane, click **Cancel**.
4. In the confirmation window, click **Confirm**.

To delete a replication

1. On the **Migrations** screen, click the name of the required migration.
2. On the **Replications** tab, click the required replication.
3. On the replication right pane, click **Delete**.
4. In the confirmation window, click **Delete**.

Deploying migrations

Deployment creates virtual machines on the target platform from the replicated data. During deployment, the system prepares the guest OS and starts the VMs.

If no replication has been performed, deployment cannot be started. You can run additional replications before deployment to reduce downtime when switching to the target environment.

You can deploy a migration automatically after replication or start deployment manually on the **Deployments** tab.

Prerequisites

- Replication must be completed before deployment. It can run automatically after migration creation, as described in "Creating migrations" (p. 119), or be started manually, as described in "Replicating migrations" (p. 122).

To deploy a migration

1. On the **Migrations** screen, click the name of the required migration.
2. On the **Deployments** tab, click **Create deployment**.
3. In the **Create deployment** window, specify the VM name and configure the deployment options:
 - Enable **Clone disks** to reuse existing disks for incremental synchronization. If disabled, disks are fully re-synchronized.
 - **Enable OS morphing** to prepare the guest OS for the target platform by installing required drivers and removing source-specific components.

Note

This step can be skipped when migrating between platforms using the same hypervisor (for example, KVM to KVM).

- Enable **Start target VMs after deployment** to start target VMs when they are deployed.

Create deployment X

VM name
Ubuntu-22.04

☒ Clone disks ⓘ

☒ Enable OS morphing ⓘ

☒ Start target VMs after deployment ⓘ

User scripts ▾

Windows script file Browse

Linux script file Browse

Cancel Create

4. [Optional] Upload user scripts to run on the guest OS during deployment.
5. Click **Deploy**.

After deployment starts, virtual machines are created on the target platform. You can monitor the deployment progress and status in the deployment details.

To cancel a deployment

Important

Canceling a deployment takes effect immediately and cannot be reversed.

1. On the **Migrations** screen, click the name of the required migration.
2. On the **Deployments** tab, click the ellipsis icon next to the ongoing deployment, and then click **Cancel**.
3. In the confirmation window, click **Confirm**.

To recreate a deployment

1. On the **Migrations** screen, click the name of the required migration.
2. On the **Deployments** tab, click the ellipsis icon next to the required deployment, and then click **Recreate**.

The deployment will be recreated with the preconfigured options.

To delete a deployment

1. On the **Migrations** screen, click the name of the required migration.
2. On the **Deployments** tab, click the ellipsis icon next to the required deployment, and then click **Delete**.
3. In the confirmation window, click **Delete**.

Editing and deleting migrations

You can update migration settings or remove migrations that are no longer needed. Deleting a migration removes the migration record and associated metadata, it does not remove already deployed virtual machines.

You can also clean up temporary volumes created during replication. Source and target VM disks are not affected.

Prerequisites

- A migration is created, as described in "Creating migrations" (p. 119).

To edit a migration

1. On the **Migrations** screen, click the name of the required migration.
2. On the **Overview** tab, click **Edit** in the **Details** section.
3. Update the required settings and click **Save**.

To clean up disks

1. On the **Migrations** screen, click the name of the required migration.
2. On the **Overview** tab, click **Clean up disks** at the bottom.
3. In the confirmation window, click **Clean up**.

To delete a migration

1. On the **Migrations** screen, click the name of the required migration.
2. On the **Overview** tab, click **Delete** at the bottom.
3. In the confirmation window, click **Delete**.

Managing S3 resources

This section explains how to enable S3 access and manage S3 access keys and buckets in the self-service panel.

Note

If you have read-only access to a project, a **Read-only mode** label is displayed in the upper-right corner of the panel. This means you can view resources but cannot perform any management actions.

Enabling access to S3 storage

To be able to manage S3 resources, you need to enable access to the S3 storage in the self-service panel. This will automatically generate an access key pair, an access key ID and a secret access key, for the current user. You can think of the access key ID as the login and the secret access key as the password.

By default, the access to the S3 storage is disabled for all self-service users.

To enable access to S3 storage

Go to the **S3** screen and click **Enable S3 storage**.

Once the access is enabled, you can proceed to manage your buckets and access keys.

Managing access keys

After enabling access to the S3 storage, one access key pair is automatically generated for the current user. It is recommended to periodically delete old access key pairs and generate new ones. When you delete an access key, it cannot be retrieved.

Limitations

- You can have up to two access key pairs.
- If you have only one access key pair, it cannot be deleted.

Prerequisites

- Access to the S3 storage is enabled, as described in "Enabling access to S3 storage" (p. 127).

To create an S3 access key pair

1. Go to the **S3 > Access** screen.
2. In the **S3 access keys** section, click **Create**.

To copy an S3 access key pair

1. Go to the **S3 > Access** screen.
2. In the **S3 access keys** section, do the following:

- To copy an access key ID, click the copy icon next to the key.
- To copy a secret access key, click the ellipsis icon next to the key, and then click **Copy secret access key**.

To disable an S3 access key pair

1. Go to the **S3 > Access** screen.
2. In the **S3 access keys** section, click the ellipsis icon next to the required key, and then click **Disable**.

To delete an S3 access key pair

1. Go to the **S3 > Access** screen.
2. In the **S3 access keys** section, click the ellipsis icon next to the required key, and then click **Delete**.

Managing buckets

After enabling access to the S3 storage, you can start creating buckets and uploading data in them. In an Amazon S3-like storage, a bucket is a uniquely named container for files, known as objects. Buckets are used to group and isolate objects from those in other buckets.

Additionally, you can configure bucket policies to manage access to your S3 resources.

Prerequisites

- Access to the S3 storage is enabled, as described in "Enabling access to S3 storage" (p. 127).

Creating and deleting buckets

It is recommended to use bucket names that comply with DNS naming conventions:

- Must be from 3 to 63 characters long
- Can contain only lowercase letters, numbers, hyphens (-), and periods (.)
- Must start and end with a letter or number
- Can be a series of valid name parts separated by periods

Limitations

- You can only delete an empty bucket.

To create a bucket

1. Go to the **S3 > Buckets** screen, and click **Create bucket**.
2. In the **Create bucket** window, specify a name for the bucket, and then click **Create**.

To list the bucket contents

Go to the **S3 > Buckets** screen, and click the bucket name to open the list of its contents.

To delete a bucket

1. Go to the **S3 > Buckets** screen, and click the required bucket.
2. On the bucket right pane, click **Delete**.
3. In the conformation window, click **Delete**.

Managing bucket policies

Virtuozzo Cloud uses Policy-Based Access Control (PBAC) to manage user permissions by defining the actions and resources that authenticated users can access. Each policy specifies one or more actions and conditions that outline the permissions for a user or group of users. Only the bucket owner can attach a policy to a bucket, and the permissions specified in the policy apply to all objects in the bucket owned by the bucket owner.

The bucket owner retains ownership of all objects in the bucket and manages access exclusively through policies. Bucket policies are written using the JSON-based AWS Identity and Access Management (IAM) policy language, consisting of the following core elements:

Statement

The primary component of a policy, defining the permissions and containing other elements such as principals, resources, actions, and effects. Policies often include an array of statements.

Statement ID (Sid)

A unique identifier assigned to each policy statement.

Effect

Specifies whether the policy allows or denies an action. If no explicit permission is granted, the policy automatically denies access by default.

Action

Lists the specific S3 actions that the policy permits or denies.

Principal

Identifies the user, entity, or account granted permissions within the statement.

Resource

Specifies the S3 bucket or objects to which the policy applies.

Condition (optional)

Defines additional restrictions or requirements under which the policy applies.

Version (optional)

Indicates the policy language version in use.

Virtuozzo Cloud supports the following S3 actions, condition keys, and condition operators for bucket policies:

Supported S3 actions

Action	Access level	Resource	Description	Condition keys
s3:GetObject	Read	Object	Grants permission to retrieve objects from a bucket	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:x-amz-content-sha256 • aws:SourceIp
s3:GetObjectAcl	Read	Object	Grants permission to return the access control list (ACL) of an object	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:x-amz-content-sha256 • aws:SourceIp
s3:GetObjectVersion	Read	Object	Grants permission to retrieve a specific version of an object	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:versionid • s3:x-amz-content-sha256 • aws:SourceIp
s3:GetObjectVersionAcl	Read	Object	Grants permission to return the access control list (ACL) of a specific object version	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:versionid • s3:x-amz-content-sha256 • aws:SourceIp
s3:ListMultipartUploadParts	List	Object	Grants permission to list the parts that have been uploaded for a specific multipart upload	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:versionid • s3:x-amz-content-sha256 • aws:SourceIp

Action	Access level	Resource	Description	Condition keys
s3:ListBucket	List	Bucket	Grants permission to list some or all of the objects in a bucket (up to 1000).	• s3:authType • s3:max-keys • s3:prefix • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:x-amz-content-sha256 • aws:SourceIp
s3:ListBucketMultipartUploads	List	Bucket	Grants permission to list in-progress multipart uploads	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:x-amz-content-sha256 • aws:SourceIp
s3:ListBucketVersions	List	Bucket	Grants permission to list metadata about all the versions of objects in a bucket	• s3:authType • s3:max-keys • s3:prefix • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:x-amz-content-sha256 • aws:SourceIp
s3:GetBucketAcl	Read	Bucket	Grants permission to use the acl subresource to return the access control list (ACL) of a bucket	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:x-amz-content-sha256 • aws:SourceIp
s3:GetBucketCORS	Read	Bucket	Grants permission to return the CORS configuration information set for a bucket	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:x-amz-content-sha256 • aws:SourceIp

Action	Access level	Resource	Description	Condition keys
s3:GetBucketLocation	Read	Bucket	Grants permission to return the region that a bucket resides in	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:x-amz-content-sha256 • aws:SourceIp
s3:GetBucketLogging	Read	Bucket	Grants permission to return the logging status of a bucket and the permissions users have to view or modify that status	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:x-amz-content-sha256 • aws:SourceIp
s3:GetBucketNotification	Read	Bucket	Grants permission to get the notification configuration of a bucket	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:x-amz-content-sha256 • aws:SourceIp
s3:GetBucketPolicy	Read	Bucket	Grants permission to return the policy of the specified bucket	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:x-amz-content-sha256 • aws:SourceIp
s3:GetBucketVersioning	Read	Bucket	Grants permission to return the versioning state of a bucket	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:x-amz-content-sha256 • aws:SourceIp
s3:GetBucketWebsite	Read	Bucket	Grants permission to	• s3:authType • s3:signatureAge

Action	Access level	Resource	Description	Condition keys
			return the website configuration for a bucket	- s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:GetLifecycleConfiguration	Read	Bucket	Grants permission to return the lifecycle configuration information set on a bucket	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:GetReplicationConfiguration	Read	Bucket	Grants permission to get the replication configuration information set on a bucket	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:PutObject	Write	Object	Grants permission to add an object to a bucket	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-acl - s3:x-amz-content-sha256 - s3:x-amz-copy-source - s3:x-amz-grant-full-control - s3:x-amz-grant-read - s3:x-amz-grant-read-acp - s3:x-amz-grant-write - s3:x-amz-grant-write-acp - s3:x-amz-storage-class - s3:x-amz-website-

Action	Access level	Resource	Description	Condition keys
				redirect-location - s3:object-lock-mode - s3:object-lock-retain-until-date - s3:object-lock-remaining-retention-days - s3:object-lock-legal-hold - aws:SourceIp
s3:DeleteObject	Write	Object	Grants permission to remove the null version of an object and insert a delete marker, which becomes the current version of the object	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:DeleteObjectVersion	Write	Object	Grants permission to remove a specific version of an object	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:versionid - s3:x-amz-content-sha256 - aws:SourceIp
s3:AbortMultipartUpload	Write	Object	Grants permission to abort a multipart upload	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:DeleteBucket	Write	Bucket	Grants permission to delete the bucket named in the URI	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-

Action	Access level	Resource	Description	Condition keys
				sha256 aws:SourceIp
s3:PutBucketCORS	Write	Bucket	Grants permission to set the CORS configuration for a bucket	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:PutBucketLogging	Write	Bucket	Grants permission to set the logging parameters for a bucket	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:PutBucketNotification	Write	Bucket	Grants permission to receive notifications when certain events happen in a bucket	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:PutBucketRequestPayment	Write	Bucket	Grants permission to set the request payment configuration of a bucket	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:PutBucketVersioning	Write	Bucket	Grants permission to set the versioning state of an existing bucket	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:PutBucketWebsite	Write	Bucket	Grants	s3:authType

Action	Access level	Resource	Description	Condition keys
			permission to set the configuration of the website that is specified in the website subresource	- s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:PutLifecycleConfiguration	Write	Bucket	Grants permission to create a new lifecycle configuration for the bucket or replace an existing lifecycle configuration	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:PutReplicationConfiguration	Write	Bucket	Grants permission to create a new replication configuration or replace an existing one	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:PutBucketPolicy	Access management	Bucket	Grants permission to add or replace a bucket policy on a bucket	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3>DeleteBucketPolicy	Access management	Bucket	Grants permission to delete the policy on a specified bucket	- s3:authType - s3:signatureAge - s3:signatureversion - s3:TlsVersion - s3:x-amz-content-sha256 - aws:SourceIp
s3:PutObjectAcl	Access	Object	Grants	s3:authType

Action	Access level	Resource	Description	Condition keys
	management		permission to set the access control list (ACL) permissions for new or existing objects in a bucket	• s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:x-amz-acl • s3:x-amz-content-sha256 • s3:x-amz-grant-full-control • s3:x-amz-grant-read • s3:x-amz-grant-read-acp • s3:x-amz-grant-write • s3:x-amz-grant-write-acp • s3:x-amz-storage-class • aws:SourceIp
s3:PutObjectVersionAcl	Access management	Object	Grants permission to use the acl subresource to set the access control list (ACL) permissions for an object that already exists in a bucket	• s3:authType • s3:signatureAge • s3:signatureversion • s3:TlsVersion • s3:versionid • s3:x-amz-acl • s3:x-amz-content-sha256 • s3:x-amz-grant-full-control • s3:x-amz-grant-read • s3:x-amz-grant-read-acp • s3:x-amz-grant-write • s3:x-amz-grant-write-acp • s3:x-amz-storage-class • aws:SourceIp
s3:PutBucketAcl	Access management	Bucket	Grants permission to set the	• s3:authType • s3:signatureAge • s3:signatureversion

Action	Access level	Resource	Description	Condition keys
			permissions on an existing bucket using access control lists (ACLs)	• s3:TlsVersion • s3:x-amz-acl • s3:x-amz-content-sha256 • s3:x-amz-grant-full-control • s3:x-amz-grant-read • s3:x-amz-grant-read-acp • s3:x-amz-grant-write • s3:x-amz-grant-write-acp • aws:SourceIp

Supported condition keys

Condition keys can be used to restrict access based on request parameters, object metadata, authentication properties, and object tags.

Condition key	Description	Value
Request and header-based condition keys		
s3:x-amz-storage-class	Filters access by storage class	String
s3:x-amz-acl	Filters access by canned ACL in the request's x-amz-acl header	String
s3:x-amz-grant-full-control	Filters access by x-amz-grant-full-control (full control) header	String
s3:x-amz-grant-read	Filters access by x-amz-grant-read (read access) header	String
s3:x-amz-grant-read-acp	Filters access by the x-amz-grant-read-acp (read permissions for the ACL) header	String
s3:x-amz-grant-write	Filters access by the x-amz-grant-write (write access) header	String
s3:x-amz-grant-write-acp	Filters access by the x-amz-grant-write-acp (write permissions for the ACL) header	String
s3:x-amz-copy-source	Filters access by copy source bucket, prefix, or object in the copy object requests	String
s3:x-amz-content-	Filters access by unsigned content in your bucket	Valid value: UNSIGNED-

Condition key	Description	Value
sha256		PAYLOAD
s3:x-amz-website-redirect-location	Filters access by a specific website redirect location for buckets that are configured as static websites	String
Authentication and transport condition keys		
s3:TlsVersion	Filters access by the TLS version used by the client	Valid values: 1.2, 1.1, and 1.0
s3:signatureversion	Filters access by the version of AWS Signature used on the request	Valid values: - AWS identifies Signature Version 2 - AWS4-HMAC-SHA256 identifies Signature Version 4
s3:signatureAge	Filters access by the age in milliseconds of the request signature	Numeric
s3:authType	Filters access by authentication method	Valid values: REST-HEADER, REST-QUERY-STRING, and POST
aws:SourceIp	Filters access by IP range	String
Object and bucket attribute condition keys		
s3:object-lock-mode	Filters access by object retention mode	Valid values: COMPLIANCE and GOVERNANCE
s3:object-lock-retain-until-date	Filters access by object retain-until date	Date
s3:object-lock-legal-hold	Filters access by object legal hold status	String
s3:object-lock-remaining-retention-days	Filters access by remaining object retention days	Numeric
s3:prefix	Filters access by key name prefix	String
s3:versionid	Filters access by a specific object version	String
s3:max-keys	Filters access by maximum number of keys returned in a ListBucket request	Numeric

Supported condition operators

Condition operators define how policy conditions are evaluated against request context values, object attributes, and object tags.

Operator	Description
String operators	
StringEquals	Exact matching, case sensitive
StringNotEquals	Negated matching, case sensitive
StringEqualsIgnoreCase	Exact matching, ignoring case
StringNotEqualsIgnoreCase	Negated matching, ignoring case
StringLike	Case-sensitive matching. The values can include multi-character match wildcards (*) and single-character match wildcards (?) anywhere in the string. Specify wildcards to achieve partial string matches.
StringNotLike	Negated case-sensitive matching. The values can include multi-character match wildcards (*) or single-character match wildcards (?) anywhere in the string.
Numeric operators	
NumericEquals	Exact matching
NumericNotEquals	Negated matching
NumericLessThan	"Less than" matching
NumericLessThanEquals	"Less than or equals" matching
NumericGreaterThan	"Greater than" matching
NumericGreaterThanEquals	"Greater than or equals" matching
Date operators	
DateEquals	Matching a specific date
DateNotEquals	Negated matching
DateLessThan	Matching before a specific date and time
DateLessThanEquals	Matching at or before a specific date and time
DateGreaterThan	Matching after a specific a date and time
DateGreaterThanEquals	Matching at or after a specific date and time
Binary operator	
BinaryEquals	Matching in binary format. It compares the value of the specified key byte for byte against a base-64 encoded representation of the binary value. If the specified key is not present in the request context, the values do not

Operator	Description
	match.
IP address operators	
IpAddress	Matching the specified IP address or range
NotIpAddress	Matching all IP addresses except the specified IP address or range

To learn more about condition operators, refer to the [AWS Identity and Access Management User Guide](#).

Creating bucket policies

When configuring bucket policies, follow these best practices:

- Use Deny statements where possible to enforce strict security.
- Limit the use of "*" as Principal to avoid granting broad or unintended permissions.
- Apply Conditions to enforce additional restrictions or requirements.
- Regularly review and audit policies to ensure they align with security and compliance requirements.

Prerequisites

- A bucket is created, as described in "Creating and deleting buckets" (p. 128).

To add a policy to a bucket

1. Go to the **S3 > Buckets** screen, and click the line with the required bucket.
2. On the bucket right pane, click **Manage bucket policy**.
3. In the **Manage bucket policy** window, upload a JSON file with bucket policy statements or create them manually as follows:
 - a. In **Policy builder**, click **Add** to add a new statement.
 - b. In the **Add statement** window that opens, do the following:
 - i. Specify a statement ID (Sid). It can contain ASCII uppercase letters (A-Z), lowercase letters (a-z), and numbers (0-9), without spaces.
 - ii. Choose the statement type between **Allow** and **Deny** to determine whether the specified actions will be allowed or denied for the selected resources and principals.
 - iii. Select S3 actions from the drop-down menu.
 - iv. Choose resources that the statement will apply to:
 - Select **All objects in the bucket** if you want the statement to be applied to all objects in the bucket.
 - Select **Select prefix or object** if you want the statement to be applied to specific objects in the bucket. In the **Resources** section, click the folder icon and select the


required object. You can add as many objects as needed.



Resources		Add
arn:aws:s3::my-bucket/my-folder1/*		
arn:aws:s3::bucket1/my.file1		


v. Choose principals:

- Select **All S3 users** if you want to grant permissions defined by the statement to all S3 users.
- Select **Specify S3 users or domains** if you want to grant permissions only to specific users. In the **Principals** section, specify users or domains in the following format: `arn:aws:iam::<domain_id>:<user_id>`. You can add as many users as needed.



Principals		Add
arn:aws:iam::4de50c40e0254e8d987470bfb84980a1:91dd4459afbe09fe		
arn:aws:iam::4de50c40e0254e8d987470bfb84980a1:862871b203ca9580		

vi. [Optional] In the **Conditions** section, click **Add**, select a condition key and operator from the drop-down menus, and then specify the desired value.



Conditions (optional)			Add
Condition key	Operator	Value	
s3:prefix	StringEquals	my-folder	
aws:SourceIp	IpAddress	192.168.1.0/24	

vii. Click **Add** to create the statement.

c. Add as many statements as needed.

d. [Optional] In **JSON view**, check your policy in the JSON format.

e. [Optional] Click **Download file** to download the created policy as a JSON file to your machine.

4. Click **Save**.

You can view the applied bucket policy in the JSON format on the **Bucket policy** tab.

Bucket policy examples

Here are some practical S3 bucket policy examples that cover various use cases to help secure and manage access to your S3 buckets. To view more examples with bucket policies, refer to the [Amazon documentation](#).

Grant public read-only access to a bucket

Use this policy to allow public read-only access to objects in a bucket.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PublicReadOnlyAccess",
      "Effect": "Allow",
      "Principal": "*",
```

```

    "Action": "s3:GetObject",
    "Resource": "arn:aws:s3:::my-public-bucket/*"
  }
]
}

```

Deny public access to a bucket

Use this policy to explicitly deny public access to the entire bucket, even if other policies or ACLs allow it.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyPublicAccess",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:*",
      "Resource": ["arn:aws:s3:::my-private-bucket", "arn:aws:s3:::my-private-bucket/*"]
    }
  ]
}

```

Allow access from a specific IP range

Use this policy to grant read-only access to the bucket to users connecting from specific IP addresses.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowAccessFromIP",
      "Effect": "Allow",
      "Principal": "*",
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3:::my-secure-bucket/*",
      "Condition": {
        "IpAddress": {
          "aws:SourceIp": "192.168.1.0/24"
        }
      }
    }
  ]
}

```

Grant access to another S3 user

Use this policy to allow another S3 user to upload objects into your bucket.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowAccountUpload",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::4de50c40e0254e8d987470bfb84980a1:862871b203ca9580"
      },
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3::my-shared-bucket/*"
    }
  ]
}
```

Allow access to a specific folder

Use this policy to grant read access only to a specific folder inside the bucket.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowFolderAccess",
      "Effect": "Allow",
      "Principal": "*",
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3::my-bucket/folder1/*"
    }
  ]
}
```

Restrict access to objects within a specific path

Use this policy to allow users to list objects within a specific folder but not access other objects in the bucket.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowListSpecificPrefix",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::4de50c40e0254e8d987470bfb84980a1:862871b203ca9580"
      },
      "Action": "s3:ListBucket",
      "Resource": "arn:aws:s3::my-bucket",
      "Condition": {
        "StringEquals": {

```

```

        "s3:prefix": "my-folder/"
      }
    },
    {
      "Sid": "AllowGetSpecificFolder",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::4de50c40e0254e8d987470bfb84980a1:862871b203ca9580"
      },
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3::my-bucket/my-folder/*"
    }
  ]
}

```

Require a minimum TLS version

Use this policy to deny uploading objects in a bucket by clients that have a TLS version earlier than 1.2, for example, 1.1 or 1.0.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:PutObject",
      "Resource": [
        "arn:aws:s3::my-bucket",
        "arn:aws:s3::my-bucket/*"
      ],
      "Condition": {
        "NumericLessThan": {
          "s3:TlsVersion": 1.2
        }
      }
    }
  ]
}

```

Prevent deletion of objects

Use this policy to protect a bucket with important objects by denying DeleteObject actions.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyObjectDeletion",

```

```
    "Effect": "Deny",
    "Principal": "*",
    "Action": "s3:DeleteObject",
    "Resource": "arn:aws:s3::my-bucket/*"
  }
]
```

Editing and deleting bucket policies

Prerequisites

- A bucket policy is created, as described in "Creating bucket policies" (p. 141).

To edit a bucket policy

1. On the **S3 > Buckets** screen, and click the line with the required bucket.
2. On the bucket right pane, click **Manage bucket policy**.
3. In the **Manage bucket policy** window, you can do the following:
 - Add new statements by clicking **Add**.
 - Edit statements by clicking the pencil icon next to them.
 - Delete statements by clicking the bin icon next to them.
4. Click **Save** to apply your changes.

To remove a policy from a bucket

1. On the **S3 > Buckets** screen, and click the line with the required bucket.
2. On the bucket right pane, navigate to the **Bucket policy** tab.
3. Click the ellipsis button and select **Delete**.
4. In the confirmation window, click **Delete**.

Managing files and folders

Once you have a bucket, you can start populating it with data.

Limitations

- Folder size reporting is not available.

Prerequisites

- A bucket is created, as described in "Creating and deleting buckets" (p. 128).

To create folders in a bucket

1. Go to the **S3 > Buckets** screen, and click the bucket name.
2. On the bucket screen, click **Create folder**.
3. In the **Create folder** window, specify a name for the folder, and then click **Create**.

To upload files in a bucket

1. Go to the **S3 > Buckets** screen, click the bucket name and optionally the folder name.
2. On the bucket or folder screen, drag and drop files to upload, or click **Upload files**.
3. In the **Upload files** window, browse files to upload.

To download files from a bucket

1. Go to the **S3 > Buckets** screen, click the bucket name and optionally the folder name.
2. On the bucket or folder screen, click the required file.
3. In the file right pane, click **Download**.

To delete a folder or a file

1. Go to the **S3 > Buckets** screen, click the bucket name.
2. On the bucket, click the required file or folder to delete.
3. In the file or folder right pane, click **Delete**.