

Acronis

Acronis Cyber Infrastructure 3.0

Self Service Guide

July 26, 2019

Copyright Statement

Copyright ©Acronis International GmbH, 2002-2019. All rights reserved.

"Acronis" and "Acronis Secure Zone" are registered trademarks of Acronis International GmbH.

"Acronis Compute with Confidence", "Acronis Startup Recovery Manager", "Acronis Instant Restore", and the Acronis logo are trademarks of Acronis International GmbH.

Linux is a registered trademark of Linus Torvalds.

VMware and VMware Ready are trademarks and/or registered trademarks of VMware, Inc. in the United States and/or other jurisdictions.

Windows and MS-DOS are registered trademarks of Microsoft Corporation.

All other trademarks and copyrights referred to are the property of their respective owners.

Distribution of substantively modified versions of this document is prohibited without the explicit permission of the copyright holder.

Distribution of this work or derivative work in any standard (paper) book form for commercial purposes is prohibited unless prior permission is obtained from the copyright holder.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Third party code may be provided with the Software and/or Service. The license terms for such third-parties are detailed in the license.txt file located in the root installation directory. You can always find the latest up-to-date list of the third party code and the associated license terms used with the Software and/or Service at <http://kb.acronis.com/content/7696>

Acronis patented technologies

Technologies, used in this product, are covered and protected by one or more U.S. Patent Numbers: 7,047,380; 7,246,211; 7,275,139; 7,281,104; 7,318,135; 7,353,355; 7,366,859; 7,383,327; 7,475,282; 7,603,533; 7,636,824; 7,650,473; 7,721,138; 7,779,221; 7,831,789; 7,836,053; 7,886,120; 7,895,403; 7,934,064; 7,937,612; 7,941,510; 7,949,635; 7,953,948; 7,979,690; 8,005,797; 8,051,044; 8,069,320; 8,073,815; 8,074,035; 8,074,276; 8,145,607; 8,180,984; 8,225,133; 8,261,035; 8,296,264; 8,312,259; 8,347,137; 8,484,427; 8,645,748; 8,732,121; 8,850,060; 8,856,927; 8,996,830; 9,213,697; 9,400,886; 9,424,678; 9,436,558; 9,471,441; 9,501,234; and patent pending applications.

Contents

1. About This Guide	1
2. Logging in to the Self Service Panel	2
3. Managing Users and Projects	3
3.1 Creating Users	3
3.2 Assigning Users to Projects	4
3.3 Viewing Project Quotas	7
4. Managing Compute Resources	8
4.1 Managing Virtual Machines	8
4.1.1 Supported Guest Operating Systems	9
4.1.2 Creating Virtual Machines	9
4.1.3 Virtual Machine Actions Overview	17
4.1.4 Enabling Logging inside Virtual Machines	18
4.1.5 Reconfiguring and Monitoring Virtual Machines	19
4.2 Managing Images	19
4.3 Managing Volumes	21
4.3.1 Creating, Editing, and Removing Volumes	21
4.3.2 Cloning Volumes	22
4.3.3 Attaching and Detaching Volumes	23
4.3.4 Creating Images from Volumes	24
4.3.5 Managing Volume Snapshots	25
4.4 Managing Private Virtual Networks	27
4.5 Managing Virtual Routers	30
4.5.1 Managing Router Interfaces	32
4.5.2 Managing Static Routes	34
4.6 Managing Floating IP Addresses	36

4.7 Managing SSH Keys 37

CHAPTER 1

About This Guide


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

CHAPTER 2

Logging in to the Self Service Panel

To log in to the self-service panel, do the following:

1. Visit the panel's IP address on port 8800.
2. Enter your domain name (case sensitive) as well as user name and password. Alternatively, if you are given the link to the self-service panel for a specific domain, you will only need to provide the user name and password.



The screenshot shows a login interface with the following elements:

- Sign in** (Title)
- Domain** input field containing `Domain1`
- Login** input field containing `domadmin1`
- Password** input field containing `*****` and a toggle icon for visibility.
- Sign in** button (Submit)

CHAPTER 3

Managing Users and Projects

A user can be assigned one of the following roles:

- A domain administrator can manage virtual objects in all projects within the assigned domain as well as project and user assignment in the self-service panel.
- A project member acts as a project administrator in a specific domain in the self-service panel. A project member can be assigned to different projects and can manage virtual objects in them.

You can create, view, and edit users on the **All users** tab. Creating a user account differs slightly depending on the user role and is described in the following sections.

To edit the user credentials or permissions, click the ellipsis button next to the user and then click **Edit**.

Enabling and disabling a user account means allowing and prohibiting user login, respectively.

To enable/disable or remove a user, click the corresponding ellipsis button and select the desired action.

3.1 Creating Users

To create a user, do as follows:

1. Select the domain in the drop-down list in the top right corner.
2. Switch to **All users** and click **Create user**.
3. In the **Create user** window, specify the user name, password, and, if required, a user e-mail address and description in the corresponding fields. The user name must be unique within a domain.
4. Select the desired role from the **Role** drop-down menu.

5. Click **Create**.

Create user ✕

Login
user1

Email (optional)
user1@example.com

Password
••••••••

Description (optional)

Role
Domain administrator

Can create and manage projects and services in the assigned domain.

Cancel Create

3.2 Assigning Users to Projects

Domain administrators can manage project members assignment on the **All projects** and **All users** screens.


To assign a user to a project, do one of the following:

- On the **All projects** screen:
 1. Click the project to which you want to assign users (not the project name).
 2. On the project panel, click **Assign members**.
 3. In the **Assign members** window, choose one or multiple users to assign to the project. Only user accounts with the **Project member** role are displayed. Optionally, click **Create project member** to create a new project member in a new window.
 4. Click **Assign**.

Assign members ✕

Select users to assign as members to the project "dom1project1".

Search + Create project member


<input checked="" type="checkbox"/>	Login ↑	Email
<input checked="" type="checkbox"/>	 projectmember1	—

- On the **All users** screen:
 1. Click the user account with the **Project member** role whom you want to assign to the project.
 2. On the user panel, click **Assign to project**.
 3. On the **Assign user to projects** window, select one or multiple projects and click **Assign**.

Assign user to projects ✕

Select projects to assign to the user "user1".


Search Q

<input checked="" type="checkbox"/>	Name ↑	Description
<input checked="" type="checkbox"/>	 project1	A custom project






To unassign a user from a project, do one of the following:

- On the **All projects** screen:
 1. Click the project to unassign users from.
 2. On the project panel, open the **Members** tab.


3. Click the cross icon next to a user you want to unassign.





project1 ✕


 Edit  Assign members  Edit quotas  Disable  Delete

Properties Members (1) Quotas





Search 

Login ↑	Email	
 user1	user1@example.com	


- On the **All users** tab:
 1. Click the user to unassign from the project.
 2. On the user panel, open the **Projects** tab.
 3. Click the cross icon next to the project from which you want to unassign the user.





user1 ✕

 Edit  Assign to project  Disable  Delete

Properties Projects (1)





Search 

Name ↑	Description	
 project1	A custom project	

3.3 Viewing Project Quotas

Each project is allocated a certain amount of compute resources by means of quotas. To view quotas of a project, open **PROJECTS**, click the desired project in the list, and switch to the **Quotas** tab.

The screenshot shows a user interface for viewing project quotas. At the top, there are three tabs: 'Properties', 'Members', and 'Quotas', with 'Quotas' being the active tab. Below the tabs is a table of resource quotas. The table has five rows, each representing a different resource. Each row includes an icon, the resource name, a progress bar, and the current usage relative to the limit.

Compute			
	vCPUs	<div style="width: 0%;"></div>	0 / 8 cores
	RAM	<div style="width: 0%;"></div>	0 / 16 GiB
	Storage policy		
	default	<div style="width: 0%;"></div>	0 / 1000 GiB
	Floating IPs	<div style="width: 0%;"></div>	0 / 32

CHAPTER 4

Managing Compute Resources

4.1 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, e.g., 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:

Table 4.1.1: Virtual machine hardware

Resource	Limit
RAM	1 TiB
CPU	48 logical CPUs
Storage	15 volumes, 512 TiB each
Network	15 NICs

A logical CPU is a core (thread) in a multicore (multithreading) processor.

4.1.1 Supported Guest Operating Systems

The following guest operating systems have been tested and are supported in virtual machines:

Table 4.1.1.1: Windows guest operating systems

Operating System	Edition	Architecture
Windows Server 2019	Essentials, Standard, Datacenter	x64
Windows Server 2016	Essentials, Standard, Datacenter	x64
Windows Server 2012 R2	Essentials, Standard, Datacenter	x64
Windows Server 2012	Standard, Datacenter	x64
Windows Server 2008 R2	Standard, Datacenter	x64
Windows Server 2008	Standard, Datacenter	x64
Windows 10	Home, Professional, Enterprise, Enterprise 2016 LTSB	x64
Windows 8.1	Home, Professional, Enterprise	x64
Windows 7	Home, Professional, Enterprise	x64

Table 4.1.1.2: Linux guest operating systems

Operating System	Architecture
CentOS 7.x	x64
CentOS 6.x	x64
RHEL 8.x	x64
RHEL 7.x	x64
Debian 9.x	x64
Ubuntu 18.04.x	x64
Ubuntu 16.04.x	x64

4.1.2 Creating Virtual Machines

Before you proceed to creating VMs, check that you have these:

- A guest OS source (see *Managing Images* (page 19)):
 - a distribution ISO image of a guest OS to install in the VM, or

- a boot volume template, or
- a boot volume.

Note: To obtain a boot volume, create a volume as described in *Managing Volumes* (page 21), attach it to a VM, install an operating system in it, then delete the VM.

- One or more virtual networks (see *Managing Private Virtual Networks* (page 27)).
- An SSH key (see *Managing SSH Keys* (page 37)).

Note: You can specify an SSH key only when creating VMs from a template or boot volume.

Note: Virtual machines are created with the host CPU model by default. Having compute nodes with different CPUs may lead to live migration issues. To avoid them, you can manually set CPU model for all new VMs as described in the *Administrator's Command Line Guide*.

To create a VM, do the following:









1. On the **Virtual machines** screen, click **Create virtual machine**. A window will open where you will need to specify VM parameters.

Create virtual machine ✕

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

Name


Deploy from: Image Volume


	Image	Specify	
	Volumes	Specify	
	Flavor	Specify	
	Networks	Specify	

[Deploy](#)

- Specify a name for the new VM.
- In **Deploy from**, choose **Volume** if you have a boot volume or want to create one. Otherwise, choose **Image**.
- Depending on your choice, click the pencil icon in the **Volumes** or **Image** section and do one of the following:
 - In the **Images** window, select the ISO image or template and click **Done**.

Images ✕



	Name ↑	Type	Min. volume size	OS Type	Size
<input checked="" type="radio"/>	 cirros	Template	1 GB	linux	13 MB

You can add images to this list on the [Images tab](#). Then [reload](#) the page.






[Cancel](#) [Done](#)

- In the **Volumes** window, do one of the following:
 - If you have prepared a volume with an installed guest OS, click **Attach**, find and select the volume, and click **Done**.

The image displays two screenshots of the OpenStack Volumes interface. The top screenshot shows the 'Attach volume' dialog box. It has a title bar with 'Attach volume' and a close button (X). Below the title bar is a dropdown menu labeled 'Volume' with the selected option 'vol1 (f71f6053-5b9b-4e33-8046-80b11139ab07), 1 ...'. At the bottom of the dialog are two buttons: 'Cancel' and 'Attach'. The bottom screenshot shows the 'Add volume' dialog box. It has a title bar with 'Add volume' and a close button (X). Below the title bar are several input fields: 'Name' with the value 'vol1', 'Size (GiB)' with the value '1', and 'Storage policy' with the value 'default'. To the right of the 'Size (GiB)' field, there is text indicating 'Min. 1 GiB, Max. 512 TiB'. Below these fields is a toggle switch for 'Delete on termination', which is currently turned off. At the bottom of the dialog are two buttons: 'Cancel' and 'Add'.

5. Optionally, in the **Volumes** window, click **Add** or **Attach** to create or attach any other volumes you need. To select a volume as bootable, place it first in the list by clicking the up arrow button next to it.
6. In the **Flavor** window, choose a flavor and click **Done**.

Flavor ✕

	Name ↑	vCPU ↑	Memory
<input checked="" type="radio"/>	 tiny	1	512 MiB
<input type="radio"/>	 small	1	2 GiB
<input type="radio"/>	 medium	2	4 GiB
<input type="radio"/>	 large	4	8 GiB
<input type="radio"/>	 xlarge	8	16 GiB


You can add flavors to this list on the [Flavors tab](#). Then [reload](#) the page.

7. In the network window, click **Add**, select a virtual network interface and click **Add**. It will appear in the **Network interfaces** list.

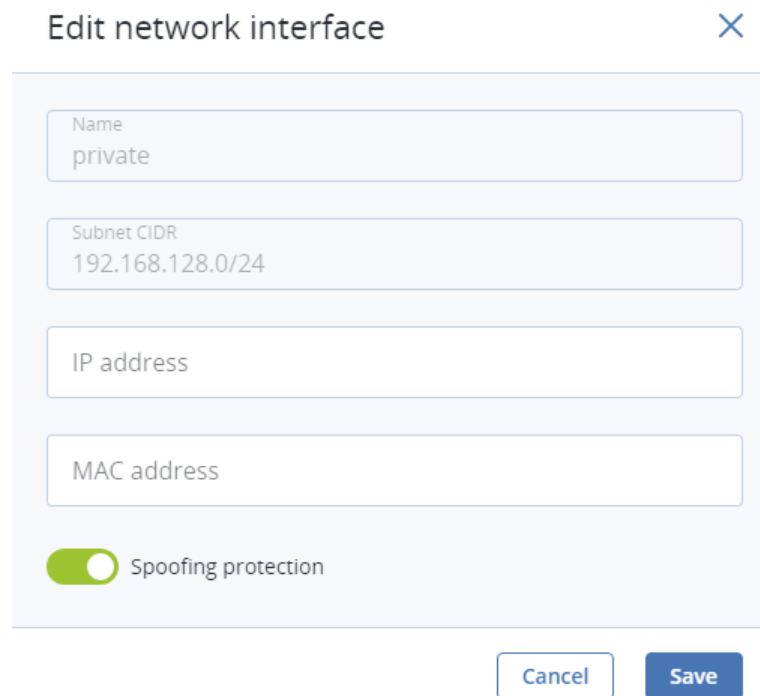
Add network interface ✕

Network
public ▼

To edit additional parameters of network interfaces that you have just added (IP and MAC addresses and spoofing protection), click interface's ellipsis icon, click **Edit**, and specify needed details in the **Edit network interface** window. Click **Save**.



You will not be able to edit these parameters later. Instead, you will be able to delete the old network interface and replace it with a new one.



Click **Done**.

- (Optional) If you are deploying the VM from a template or boot volume (not an ISO image), you can specify the following:

- An SSH key to be injected into the VM. To do it, select an SSH key in the **Select an SSH key** window, and click **Done**.

Select an SSH key ✕

🔍
+ Add

	Name ↑	Description ↑	Created on	
<input checked="" type="radio"/>	root_node001vstoragedom	My public key	June 10, 2019 4:23 PM	⋮

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

Cancel
Done

Note: To be able to connect to the VM via SSH, make sure the VM template or boot volume has cloud-init and OpenSSH installed (see the section “Creating SSH-Enabled Templates” in the *Administrator’s Command Line Guide*).

- User data to customize the VM after launch. To do it, write a script in the **Customization script** field or browse a file on your local server to load the script from.

You can specify user data in one of two formats: cloud-config or shell script. To inject a script in a Windows VM, refer to the [Cloudbase-Init documentation](#).

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 Browse

Cancel Save

Note: For the guest OS to be customizable, make sure the VM template or boot volume has cloud-init installed.

9. Back in the **Create virtual machine** window, click **Deploy** to create and boot the VM.
10. If you are deploying the VM from an ISO image (not a boot volume template or a volume with a pre-installed guest OS), select the VM, click **Console**, and install the guest OS using the built-in VNC console.
11. (Optional) If you are deploying the VM from a prepared template with an injected SSH key, you can connect to it via SSH using the username and the VM IP address:
 - for Linux templates, enter the username that is default for the cloud image OS (for example, for a CentOS cloud image, the default login is centos);

- for Windows templates, enter the username that you specified during Cloudbase-Init installation;
- for VMs customized with user data, enter the username specified in the script.

For example:

```
# ssh myuser@10.10.10.10
```

4.1.3 Virtual Machine Actions Overview

After you create a virtual machine, you can manage it using the actions available for its current state. To see the full list of available actions, click the ellipsis button next to a VM or on top of its panel. Actions include:

- **Run** powers up a VM.
- **Console** connects to running VMs via the built-in VNC console. In the console browser window, you can send a key combination to a VM, take a screenshot of the console window, and download the console log.
- **Reboot** soft-reboots a running VM.
- **Shut down** gracefully shuts down a running VM.
- **Hard reboot** cuts off and restores power, then starts a VM.
- **Power off** forcibly cuts off power from a VM.
- **Shelve** unbinds a stopped VM from the node it is hosted on and releases its reserved resources such as CPU and RAM. A shelved VM remains bootable and retains its configuration, including the IP addresses.

Virtual machines in other states can be shelved by clicking **Shut down** or **Power off** and selecting the checkbox **Shelve virtual machine** in the confirmation window.

- **Unshelve** spawns a shelved VM on a node with enough resources to host it.
- **Suspend** saves the current VM state to a file.

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.

- **Resume** restores a VM from suspended state.
- **Download console log** downloads the console log. Make sure logging is enabled inside the VM, otherwise the log will be empty (for more information, see *Enabling Logging inside Virtual Machines* (page 18)).

Examining console logs may be useful in troubleshooting failed virtual machines.

- **Reset state** resets the VM stuck in a failed or transitional state to its last stable state: active, shut down or shelved.
- **Delete** removes a VM from the compute cluster.

4.1.4 Enabling Logging inside Virtual Machines

VM's console log will contain log messages only if the TTY1 and TTYS0 logging levels are enabled inside the VM. For example, you can enable them as follows in Linux VMs:

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.

In Windows VMs, you can enable Emergency Management Services (EMS) console redirection for this purpose. Do the following:

1. Start **Windows PowerShell** with 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
```

You may also enable driver status logging to see the list of loaded drivers. This can be useful for troubleshooting a faulty driver or long boot process. You can do this as follows:

1. Start **System Configuration** with administrator privileges.
2. In the **System Configuration** windows, open the **Boot** tab, select the checkboxes **OS boot information** and **Make all boot settings permanent**.

3. Confirm the changes and restart the system.

4.1.5 Reconfiguring and Monitoring Virtual Machines

To monitor virtual machine's CPU, storage, and network usage, select the VM and open the **Monitoring** tab.

To reconfigure a VM, select it and, on the **Overview** tab, click the pencil icon next to a parameter you need to change. You cannot do the following:

- change, detach, or delete the boot volume;
- manage non-boot volumes except attaching and detaching;
- modify previously added network interfaces;
- attach and detach network interfaces to and from shelved VMs.

4.2 Managing Images

Acronis Cyber Infrastructure 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. In turn, a template is a ready volume in the QCOW2 format with an installed operating system and applications and a set minimum size. Many OS vendors offer templates of their operating systems under the name "cloud images". For a list of guest OSes supported in virtual machines, see *Supported Guest Operating Systems* (page 9).

To add an image, do the following:

1. On the **Images** screen, click **Add image**.
2. In the **Add image** window, do the following:
 1. Click **Browse** and select a template or ISO file.
 2. Specify an image name to be shown in the admin panel.
 3. Select a correct OS type from the drop-down list.

Important: OS type affects VM parameters like hypervisor settings. VMs created from an image with a wrong OS type may not work correctly, e.g., crash.

Add image ✕

Image file
Fedora-LXDE-Live-x86_64-27-1.6.iso Browse

Name
Fedora-LXDE-Live-x86_64-27-1.6.iso

Select OS distribution
Generic Linux ▼

Share between all projects

Cancel Add

3. Click **Done** to start uploading the image. Upload progress will be shown in the bottom right corner.

To create a volume from an uploaded image, click the desired image then click **Create volume**. In the pop-up window that opens, specify volume name, size and choose a storage policy. Click **Create**.

Create volume ✕

Name
vol1

Size (GiB) Min. 1 GiB, Max. 512 TiB
10

Storage policy
default ▼

Image: centos7-minimal

Cancel Create

To download or remove an image, click the ellipsis button next to it and click the desired action.

For information on how to create Linux templates, see the *Administrator's Command Line Guide*.

4.3 Managing Volumes

A volume in Acronis Cyber Infrastructure is a virtual disk drive that can be attached to a VM. The integrity of data in volumes is protected by a redundancy mode specified in a storage policy.

Note: Additional virtual disks attached to VMs need to be initialized inside the guest OS by standard means before they can be used.

4.3.1 Creating, Editing, and Removing Volumes

To create a volume, do the following:

1. On the **Volumes** tab, click **Create volume**.

The screenshot shows a 'Create volume' dialog box with the following fields and values:

- Name:** vol1
- Size (GiB):** 1 (with a range of Min. 1 GiB, Max. 512 TiB)
- Storage policy:** default

Buttons at the bottom: Cancel, Create

2. In the **Create volume** window, specify volume name and size in gigabytes; select a storage policy from

the corresponding drop-down list; click **Add**.

To edit a volume, select it and click the pencil icon next to a parameter you need to change. Note the following restrictions:

- You cannot shrink volumes.
- To extend volumes that are in use, stop the VM first.
- You cannot change the volume redundancy type.

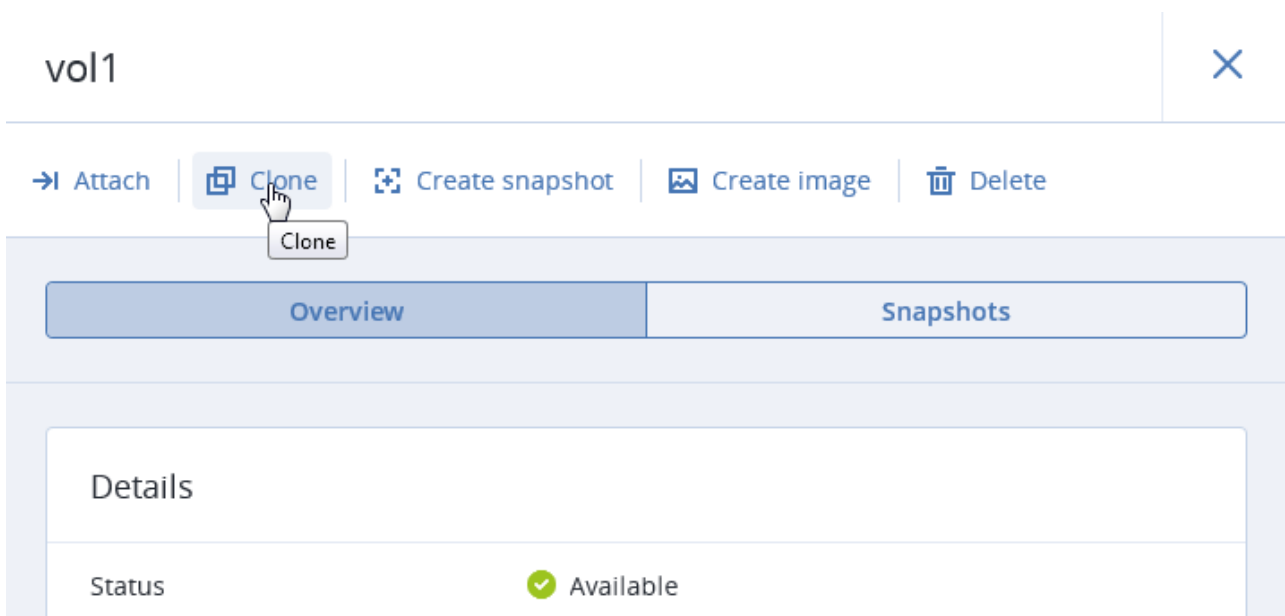
To remove a volume, click the corresponding ellipsis button then click **Delete**. To remove multiple volumes at once, select them and click **Delete**. To remove a volume that is in use, detach it first.

Note: A volume is removed along with all its snapshots.


4.3.2 Cloning Volumes

You can clone volumes that are not attached to VMs or attached to stopped VMs. To clone a volume, do the following:

1. On the **Volumes** screen, click a volume.
2. In volume details that opens, click **Clone**.



3. In the **Clone volume** window that opens, specify a volume name, size, and storage policy. Click **Clone**.



The screenshot shows a 'Clone volume' dialog box with the following fields and values:

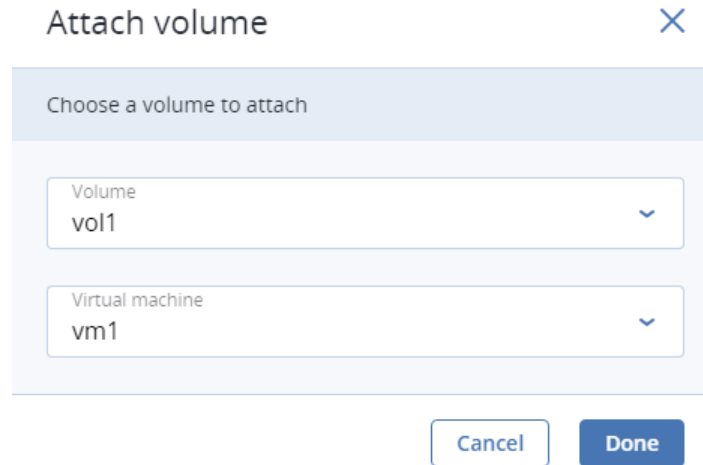
- Name:** Clone_vol1
- Size (GiB):** 1 (with a note: Min. 1 GiB, Max. 512 TiB)
- Storage policy:** default

Buttons at the bottom: Cancel, Clone

4.3.3 Attaching and Detaching Volumes

To add a writable virtual disk drive to a VM, attach a volume to it. To do this:

1. On the **Volumes** screen, click the ellipsis button next to an unused volume and click **Attach** in the context menu.
2. In the **Attach volume** window, select the VM from the drop-down list and click **Done**.



To detach a volume, do the following:

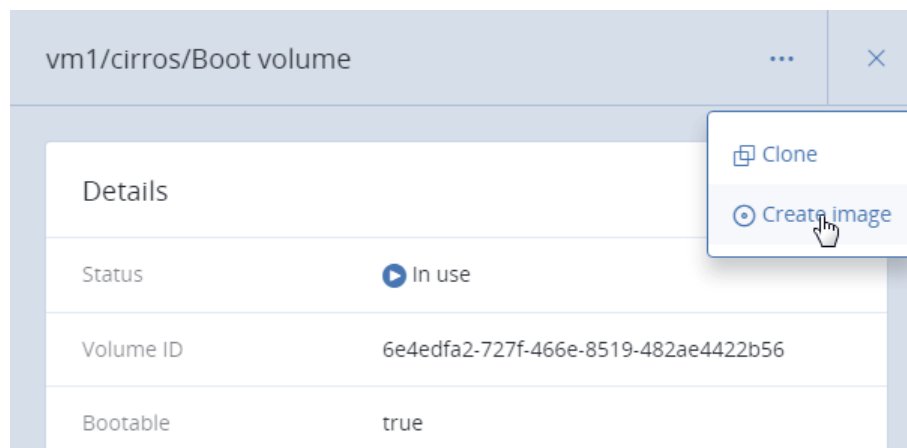
1. Click the ellipsis button next to the volume that is in use.
2. If the VM is not running, click **Detach**. If the VM is running, you can only click **Force detach** to immediately detach the volume with a risk of data loss.

4.3.4 Creating Images from Volumes


To create multiple VMs with the same boot volume, you can create an image from an existing boot volume and deploy VMs from it. Make sure to install cloud-init in the volume before creating the image.

Do the following:

1. Power off the VM that the original volume is attached to.
2. Switch to **Volumes**, click volume's ellipsis button and **Create image**.



3. In the **Create image** window, enter an image name and click **Create**.



The screenshot shows a 'Create image' dialog box. The title bar contains the text 'Create image' and a close button 'X'. The main area of the dialog has a text input field with the label 'Name' and the value 'vol1-image'. Below the input field, the text 'Volume: vm1/cirros/Boot volume' is displayed. At the bottom of the dialog, there are two buttons: 'Cancel' and 'Create'.

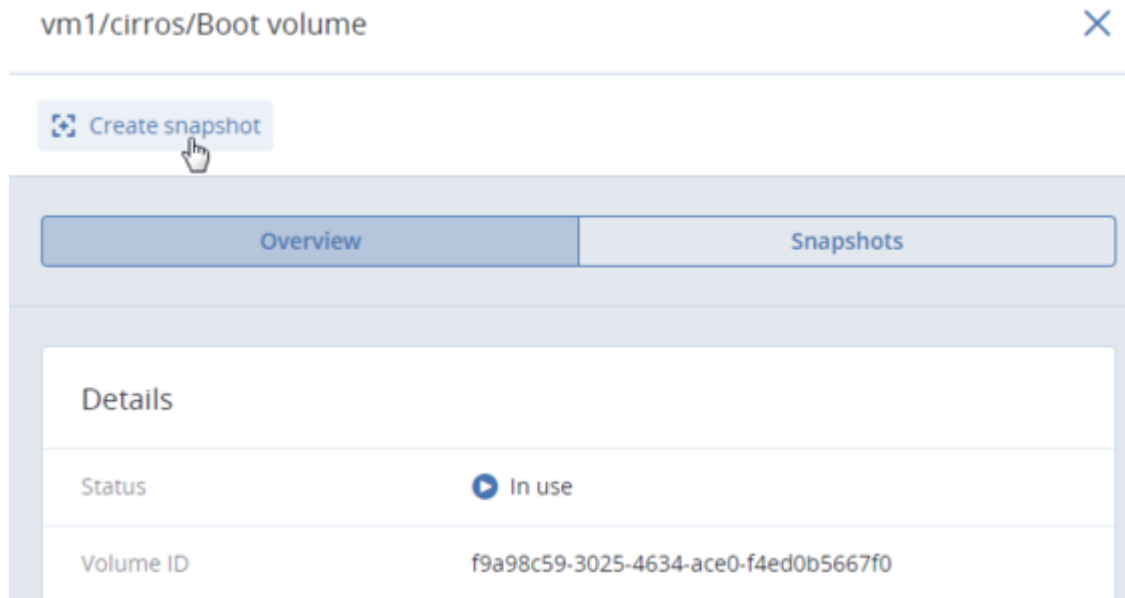
The new image will appear on the **IMAGES** tab.

4.3.5 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.


To create a snapshot of a volume, do the following:

1. On the **Volumes** screen, click a volume.
2. In the volume panel that opens, switch to **Snapshots** and click **Create snapshot**.



Note: To create a consistent snapshot of a running VM's volume, make sure the guest tools are installed in the VM. QEMU guest agent included in the guest tools image automatically quiesces the filesystem during snapshotting. For the instructions on installing the guest tools, see the section "Installing Guest Tools" in the *Administrator's Command Line Guide*.

Once the snapshot is created, you can see and manage it on the **Snapshots** tab on the volume panel.



To see the full list of available actions, click the ellipsis button next to a snapshot. Actions include:

- **Create volume** creates a new volume from the snapshot.
- **Create image** creates a template image from the snapshot.
- **Revert to snapshot** discards 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.

- **Edit** changes the snapshot name and description.
- **Reset** resets the snapshot stuck in the “Error” state or one of transitional states to the “Available” state.
- **Delete** removes the snapshot.

4.4 Managing Private Virtual Networks

To add a new virtual private network, do the following:

1. On the **Networks** screen, click **Create virtual network**.
2. In the **Network configuration** section, configure the network parameters:
 1. Enable or disable IP address management.

With IP address management enabled, Acronis Cyber Infrastructure will handle virtual machine IP addresses and provide the following features:

- **Allocation pools.** You can specify ranges of IP addresses that will be automatically assigned to VMs.
- **Built-in DHCP server.** Assigns IP addresses from allocation pools to virtual machines. With the DHCP server enabled, VM network interfaces will be automatically assigned IP addresses from allocation pools or, if no pools are specified, network’s entire IP range. With the DHCP server disabled, VM network interfaces will be allocated IP addresses, but you will have to manually assign them inside VMs.

- Custom DNS servers. You can specify DNS servers that will be used by VMs. These servers will be delivered to virtual machines via the built-in DHCP server.

With IP address management disabled:

- VMs connected to a network will be able to obtain IP addresses from DHCP servers in that network.
- Spoofing protection will be disabled for all VM network ports. Each VM network interface will accept all traffic, even frames addressed to other network interfaces.

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

2. Choose network type.
3. Specify a name. If IP address management is enabled, specify network's IPv4 address range in **Subnet CIDR**. Optionally specify a gateway. If you leave the **Gateway** field blank, the gateway will be omitted from network settings.

Click **Next**.

Create virtual network ✕

- Network configuration
- DHCP and DNS
- Summary

Configure network settings.

IP address management

Name

Subnet CIDR

Gateway (optional)

3. If you enabled IP address management on the previous step, you will move on to the **DHCP and DNS** section. In it, enable or disable the built-in DHCP server and specify one or more allocation pools and DNS servers. Click **Next**.

Create virtual network ✕

- Network configuration
- DHCP and DNS
- Summary

Set DHCP and specify one or more allocation pools for the public virtual network.

Enable the built-in DHCP server.

Allocation pools + Add pool

192.168.0.2 — 192.168.0.128	127 addresses available	✎	🗑️
-----------------------------	-------------------------	---	----

DNS servers + Add server

10.10.10.10	✎	🗑️
-------------	---	----

Back
Next

4. In the **Summary** section, review the configuration and click **Add virtual network**.

Create virtual network ✕

- Network configuration
- DHCP and DNS
- Summary

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

Type	Private
Name	privnet1
Subnet CIDR	192.168.0.0/24
Gateway	192.168.0.1
DHCP	Enabled
Allocation pools	192.168.0.2 — 192.168.0.128 127 addresses available
DNS servers	10.10.10.10

Back
Create virtual network

To view and edit parameters of a virtual network, click it on the **Networks** screen. On the virtual network panel, you can change the virtual network name, gateway, DHCP settings, allocation pools, and DNS servers.

To do this, click the pencil icon, enter a new value, and click the check mark icon to confirm.

To delete a virtual network, click the ellipsis icon next to it and **Delete**. To remove multiple virtual networks at once, select them and click **Delete**. Before deleting virtual networks, make sure no VMs are connected to them.

4.5 Managing Virtual Routers

Virtual routers provide L3 services such as routing and Source Network Address Translation (SNAT) between private and public networks or different private networks:

- a virtual router between private and public networks provides access to public networks, such as the Internet, for VMs connected to this private network;
- a virtual router between different private networks provides network communication for VMs connected to these private networks.

A virtual router has two types of ports:

- an external gateway that is connected to a public network,
- an internal port that is connected to a private network.

Note: A router can only connect networks with enabled IP management.

To create a virtual router, do the following:

1. On the **COMPUTE > Networks** tab, make sure the virtual networks that are to be connected to a router have a specified gateway.
2. Navigate to the **COMPUTE > Routers** tab and click **Add router**.
3. In the **Add virtual router** window, perform these actions:
 1. Specify a router name.
 2. From the **Network** drop-down menu, select a public 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 public network.
 3. In the **Add internal interfaces** section, select one or more private networks to connect to a router

via internal interfaces. The new internal interfaces will attempt to use the gateway IP address of the selected private networks by default.

4. Optionally, select or deselect the **SNAT** checkbox to enable or disable SNAT, respectively, 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 public networks will be accessed.

Network
public: 10.94.0.0/16

SNAT ⓘ

Add internal interfaces + Add

private: 192.168.128.0/24 🗑️

Cancel Create

4. Click **Create**.

To edit a router name, click the ellipsis icon next to it and **Rename**.

To remove a virtual router, click the ellipsis icon next to it and **Delete**. To remove multiple virtual networks at once, select them and click **Delete**. Before deleting a virtual router, make sure no floating IP addresses are associated with any network it is connected to.

4.5.1 Managing Router Interfaces

You can add an external router interface as follows:

1. On **Routers** screen, click the router name to open the list of its interfaces.
2. Delete the current external interface by clicking the ellipsis button and then **Delete**.
3. Click **Add**.
4. In the **Add interface** window, do the following:
 1. Choose **External gateway**.
 2. From the **Network** drop-down menu, select a public network to connect to the router. The new interface will pick an unused IP address from the selected public network. You can also provide a specific IP address from the selected public network to assign to the interface in the **IP address** field.
 3. Optionally, select or deselect the **SNAT** checkbox to enable or disable SNAT, respectively, 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: 10.94.0.0/16

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 public network. You can also provide a specific IP address from the selected public network to assign to the interface.

SNAT ⓘ

Cancel Add

5. Click **Add**.

To edit the external gateway parameters, click the ellipsis icon next to it and **Edit**. In the **Edit interface** window, you can change the external gateway IP address and enable or disable SNAT on it. To save your changes, click **Save**.

You can add an internal router interface as follows:

1. On **Routers** screen, click the router name to open the list of its interfaces.
2. Click **Add**.
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 private network by default. If it is in use, specify an unused IP address from the selected private network to assign to the interface in the **IP address** field.

Add interface
✕

Specify new interface parameters

Network
Select ▼

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 private network by default. If it is in use, specify an unused IP address from the selected private network to assign to the interface.

Cancel
Add

4. Click **Add**.

To remove a router interface, click the ellipsis icon next to it and **Delete**. To remove multiple interfaces at once, select them and click **Delete**.

4.5.2 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 private networks and want only one private network to be accessible from the other.

Consider the following example:

- the virtual machine `vm1` is connected to the private 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 private 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 public network via the external gateway with the IP address 10.94.129.73,
- the router `router2` connects the network `private2` to the public 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.

To create a static route for a router, do the following:

1. On the **Routers** screen, select router's checkbox and click **Manage static routes** above.

The screenshot shows the 'Routers' management interface. At the top, there are navigation links: 'Manage interfaces', 'Manage static routes', and 'Delete'. A status indicator shows '1 item selected'. Below this is a table with columns: 'Name', 'Status', 'External net...', 'SNAT', and a settings icon. The table contains one entry for 'router1' which is 'Active' and has 'public' as the external network and 'Enabled' for SNAT.

<input checked="" type="checkbox"/>	Name ↓	Status ↑	External net...	SNAT	⚙️
<input checked="" type="checkbox"/>	router1	Active	public	Enabled	⋮

2. On the next screen, click **Add static route**.
3. 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.

The 'Add static route' dialog box is shown. It has a title bar with a close button (X). The main content area is titled 'Specify static route parameters' and contains two input fields. The first field is labeled 'Destination subnet and mask' and contains the text '192.168.30.0/24'. The second field is labeled 'Next hop' and contains the text '10.94.129.74'. Below the input fields, there is a note: 'The next hop's IP address must belong to one of the networks that the router is connected to.' At the bottom of the dialog, there are two buttons: 'Cancel' and 'Add'.

4. Click **Add**.

To edit a static route, click the ellipsis icon next to it and **Edit**. In the **Edit static route** window, change the desired parameters and click **Save**.

To remove a static route, click the ellipsis icon next to it and **Delete**. To remove multiple routes at once, select them and click **Delete**.

4.6 Managing Floating IP Addresses


A virtual machine connected to a virtual private 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 public network and mapped to 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 private 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.

Note the following prerequisites:

1. A VM must have a fixed private IP address.
2. A virtual router must connect the public network from which a floating IP will be picked with VM's private network.

You can create a floating IP address and assign it to a VM as follows:

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



The screenshot shows a dialog box 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 shows the value "public: 10.94.0.0/16". The second dropdown is labeled "Virtual machine" and shows the value "myvm — private: 192.168.128.5". At the bottom of the dialog, there are two buttons: "Cancel" and "Add".

3. Click **Add**.

A floating IP address can be re-assigned to another virtual machine. Do the following:

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 choose **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, unassign it from a VM as described above, then click the ellipsis icon again and choose **Delete**.

4.7 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 Acronis Cyber Infrastructure 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.

Note: You can specify an SSH key only if you deploy a VM from a template or boot volume (not an ISO image).

Before using the SSH keys feature, make sure the following requirements are met:

- The `cloud-init` utility is installed in a VM template or boot volume.
- OpenSSH Server is installed in a Windows template or boot volume.

For the instructions on preparing templates or boot volumes, see the section “Creating SSH-Enabled Templates” in the *Administrator’s Command Line Guide*.

To add a public key, do the following:

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

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

2. On the **SSH Keys** screen, click **Add key**.
3. 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.

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
9MANMUTVzgDu/xFh0Nm2HKNV4GWGVAGGbGNqBfkjDBOq/wfj
OrrwXQXghgmvD+FCeGIeh3YCxeVIMS6/PgnbZefOG9o4QianAGs8
kMrrF8zL6svL8qOvIWUxsGoJT+3WmXT+fF5OExm01XDau0vhmhT
6VI6KDON2Y14YthzBQxGheUEhjUC45xvklQXi0oYxa0eGi1Ed3s3bX
ICWbDQsJSvaluRviqMKE7x6M+iWSgm9wuzBwM1+SKHtiaKsDKyQ
zPqpmGVkl4tj7X9gWRhM2trKqd0CkKkd2lgezDReTgQOerJ5+YTPg
qIKnbNPAMSn root@node001.vstoragedomain

Cancel Add

To delete one or more keys, select them and click **Delete**.

Note: If a key has been injected into one or more VMs, it will remain inside those VMs even if you delete it from the admin panel.