

# Virtuozzo

## Virtuozzo Infrastructure 7.3

### Quick Start Guide

3/31/2026

# Table of contents

- Introduction ..... 3**
- Hardware requirements ..... 4**
- Installing Virtuozzo Infrastructure ..... 5**
- Creating the storage cluster ..... 7**
- Enabling management node high availability ..... 8**
- Deploying the compute cluster ..... 9**
- Creating a virtual machine ..... 13**

# Introduction

Virtuozzo Infrastructure represents a new generation of hyperconverged infrastructures targeted at both service providers and end customers. It is a scale-out, cost-efficient, and multi-purpose solution that combines universal storage and high-performance virtualization.

This guide describes how to set up a full-fledged storage cluster on three nodes, deploy a compute cluster on top of it, and create a virtual machine.

# Hardware requirements

A minimum Virtuozzo Infrastructure installation recommended for production consists of three nodes for storage and compute services with enabled high availability for the management node. This is to ensure that the cluster can survive failure of one node without data loss. The following table lists the *minimal* hardware requirements for all the three nodes. The recommended configurations are provided in "System requirements" in the Administrator Guide.

Type	Management node with storage and compute
CPU	64-bit x86 processors with AMD-V or Intel VT hardware virtualization extensions enabled. 16 cores*
RAM	32 GB
Storage	1 disk: system + metadata, 100+ GB SATA HDD 1 disk: storage, SATA HDD, size as required
Network	10 GbE for storage traffic 1 GbE for other traffic

\* A CPU core here is a physical core in a multicore processor (hyperthreading is not taken into account).

# Installing Virtuozzo Infrastructure

---

## Important

The time needs to be synchronized via NTP on all nodes in the same cluster. Make sure that the nodes can access the NTP server.

---

To install Virtuozzo Infrastructure, do the following:

1. Obtain the distribution ISO image. To do that, visit the [product page](#) and submit a request for the trial version.
2. Prepare the bootable media using the distribution ISO image (mount it to an IPMI virtual drive, create a bootable USB drive, or set up a PXE server).
3. Boot the server from the chosen media.
4. On the Welcome screen, choose **Install Virtuozzo Infrastructure**.
5. On step 1, carefully read the End-User License Agreement. Accept it by selecting the **I accept the End-User License Agreement** check box, and then click **Next**.
6. On step 2, configure a static IP address for the network interface and provide a host name: either a fully qualified domain name (**<hostname>.<domainname>**) or a short name (**<hostname>**). A dynamic IP is not recommended as it might cause issues with reaching the nodes. Check that the network settings are correct.
7. On step 3, choose your time zone. Date and time will be set via NTP. You will need an Internet connection for synchronization to complete.
8. On step 4, specify what type of node you are installing. First, deploy one primary node. Then, deploy as many secondary nodes as you need.
  - If you chose to deploy the primary node, select two network interfaces: for internal management and configuration and for access to the admin panel. Also create and confirm a password for the superadmin account of the admin panel. This node will be the management node.
  - If you chose to deploy a secondary node, provide the IP address of the management node and the token. Both are obtained from the admin panel. Log in to the admin panel on port 8888. The panel's IP address is shown in the console after deploying the primary node. Enter the default username **admin** and the superadmin account password. In the admin panel, open **Infrastructure > Nodes**, and then click **Connect node**, to invoke a screen with the management node address and the token.

The node may appear on the **Infrastructure > Nodes** screen with the **Unassigned** status as soon as the token is validated. However, you will be able to join it to the storage cluster only after the installation is complete.
9. On step 5, choose a disk for the operating system. This disk will have the supplementary role **System**, although you will still be able to set it up for data storage in the admin panel. You can also create software RAID1 for the system disk, to ensure its high performance and availability.
10. On step 6, enter and confirm the password for the root account, and then click **Start installation**.

Once the installation is complete, the node will reboot automatically. The admin panel IP address will be shown in the welcome prompt.

# Creating the storage cluster

To create the storage cluster, do the following:

1. Open the **Infrastructure > Nodes** screen, and then click **Create storage cluster**.
2. In the **Create storage cluster** window, enter a name for the cluster. The cluster name may only contain Latin letters (a-z, A-Z), numbers (0-9), and hyphens ("-"). It must start with a letter and end with a letter or number.
3. [Optional] Enable disk encryption for tiers. You can also enable it later.
4. Select one node to create the storage cluster from, and then click **Next**.
5. In the next window, check the default disk configuration. If it is correct, proceed to create the storage cluster.

Also, you can assign roles to your disks manually or use **Disk actions** to work with the disks.

6. Once you finish configuring the disks, click **Create**, to create the storage cluster.

You can monitor cluster creation on the **Infrastructure > Nodes** screen. The creation might take some time, depending on the number of disks to be configured. Once the configuration is complete, the cluster is created.

To add more nodes to the storage cluster, do the following:

1. On the **Infrastructure > Nodes** screen, click an unassigned node.
2. On the node right pane, click **Join to cluster**.
3. In the **Join node to storage cluster** window, check the default disk configuration. If it is correct, proceed to join the node to the storage cluster.

Also, you can assign roles to your disks manually or use **Disk actions** to work with the disks.

Alternatively, you can copy the disk configuration from another node by clicking **Copy configuration from** and selecting the desired node.

4. Once you finish configuring the disks, click **Join**, to add the node to the storage cluster.

# Enabling management node high availability

To make your infrastructure more resilient and redundant, you can create a high availability (HA) configuration of three nodes.

Management node HA and compute cluster are tightly coupled, so changing nodes in one usually affects the other. Take note of the following:

- All nodes in the HA configuration will be added to the compute cluster.
- Single nodes cannot be removed from the compute cluster as they are included in the HA configuration. In such a case, the compute cluster can be destroyed completely, but the HA configuration will remain. This is also true vice versa, the HA configuration can be deleted, but the compute cluster will continue working.

To enable high availability for the management node and admin panel, do the following:

1. Go to **Settings > System settings > Management node high availability**, and then click **Create HA configuration**.
2. In the **Create HA configuration** window, select three nodes, and then click **Next**.
3. Depending on your network configuration, specify one or multiple unique static IP addresses for the highly available admin and self-service panels, compute API endpoint, and interservice messaging.
4. Click **Create**.

Once the high availability of the management node is enabled, you can log in to the admin panel at the specified static IP address (on the same port 8888).

# Deploying the compute cluster

Before creating a compute cluster, make sure the following requirements are met:

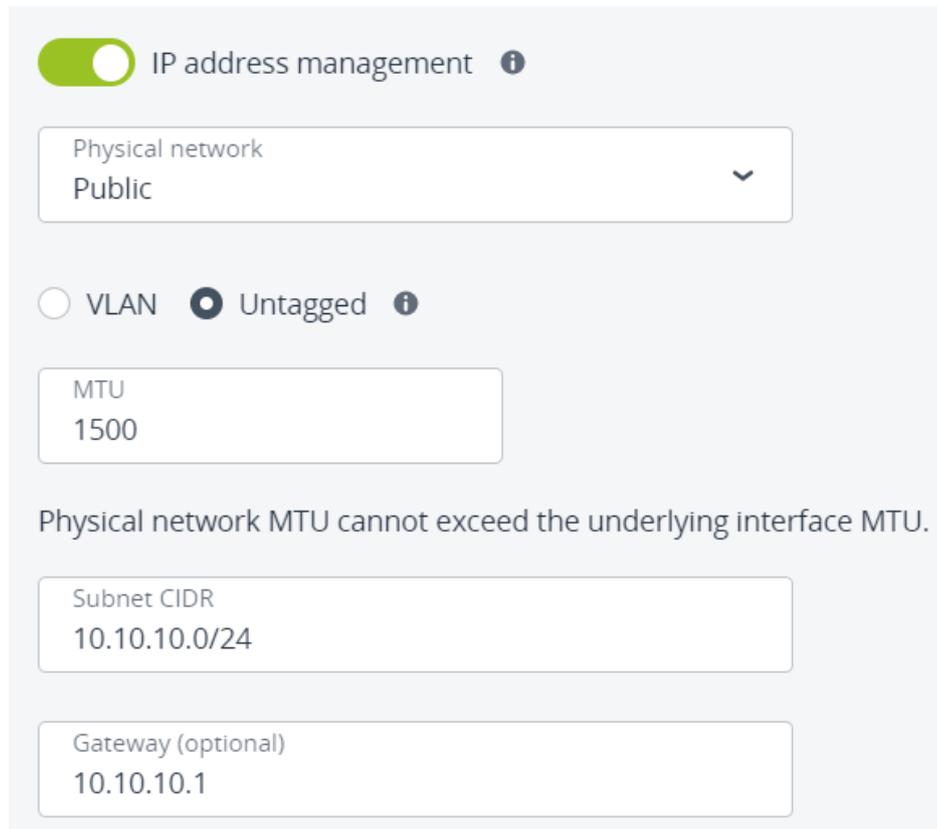
- The traffic types **VM private**, **VM public**, **Compute API**, and **VM backups** are assigned to networks. The full recommended network configuration is described in "Setting up networks for the compute cluster" in the Administrator Guide.
- The nodes to be added to the compute cluster are connected to these networks and to the same network with the **VM public** traffic type.
- The nodes to be added to the compute cluster have the same CPU model (refer to "Setting virtual machine CPU model" in the Administrator Guide).
- (Recommended) High availability for the management node is enabled (refer to "Enabling management node high availability" (p. 8)).

To create the compute cluster, do the following:

1. On the **Infrastructure > Networks** screen, make sure that these traffic types are added to the networks you intend to use: **VM private**, **VM public**, **Compute API**, **VM backups**.
2. Open the **Compute** screen, and then click **Create compute cluster**.
3. On the **Nodes** step, select the nodes to add to the compute cluster. You can only select nodes with the **Configured** network state. Nodes in the management node high availability cluster are automatically selected to join the compute cluster. Then, click **Next**.
4. On the **Physical network** 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. Provide the required details for the physical network:
    - i. Select an infrastructure network to connect the physical network to.
    - ii. Select the physical network type: select **VLAN** and specify a VLAN ID to create a VLAN-based network, or select **Untagged** to create a flat physical network.
    - iii. The network MTU is set to 1500 by default. If required, you can adjust this value according to the MTU of the underlying network interface.

- iv. If you enabled IP address management, the subnet IP range in the CIDR format will be filled in automatically. Optionally, specify a gateway. If you leave the **Gateway** field blank, the gateway will be omitted from network settings.
- c. Click **Next**.



The screenshot shows a configuration panel for IP address management. At the top, there is a green toggle switch labeled "IP address management" with an information icon. Below this is a dropdown menu for "Physical network" currently set to "Public". There are two radio buttons: "VLAN" (unselected) and "Untagged" (selected) with an information icon. Below the radio buttons is a text input field for "MTU" containing the value "1500". A warning message states: "Physical network MTU cannot exceed the underlying interface MTU." Below this is another text input field for "Subnet CIDR" containing "10.10.10.0/24". At the bottom is a text input field for "Gateway (optional)" containing "10.10.10.1".

The selected physical network will appear in the list of compute networks on compute cluster's **Network** tab. By default, it will be shared between all future projects. You can disable the network access on the network right pane later.

5. On the **MAC address prefix** step, specify a unique prefix to be used for all virtual network ports requiring MAC addresses, such as virtual machines, routers, and DHCP servers. This prefix defines the first octets of the MAC addresses assigned to these interfaces. Click **Generate** to create a prefix automatically or enter two or three octets manually, for example, fa:16 or fa:16:33. Then, click **Next**.

Keeping the default prefix fa:16:3e is not recommended because it may cause duplicate MAC addresses on the physical network, but you can change the prefix later if needed.

6. If you enabled IP address management, you will move on to the **DHCP and DNS** step, where you can configure the network settings for IP address management:
  - a. 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**.

- On the **Add-on services** step, enable the additional services that will be installed during the compute cluster deployment. You can also install these services later. Then, click **Next**.

---

### Note

Installing Kubernetes automatically installs the load balancer service as well.

---

- On the **Storage policy** step, select a redundancy mode, storage tier, and failure domain for the default policy, which will be applied to uploaded images and base volumes created from these images. You can also use the default parameters, which include the first available storage tier, the host failure domain, and the best replication scheme allowed by the number of nodes in the storage cluster:
  - The **3 replicas** mode is used if the storage cluster has three or more nodes.
  - The **2 replicas** mode is used if the storage cluster has two nodes.
  - The **No redundancy** mode is used for a single-node deployment.

To discard your changes to the storage policy parameters and reset them to their defaults, click **Reset to default parameters**.

Then, click **Next**.

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

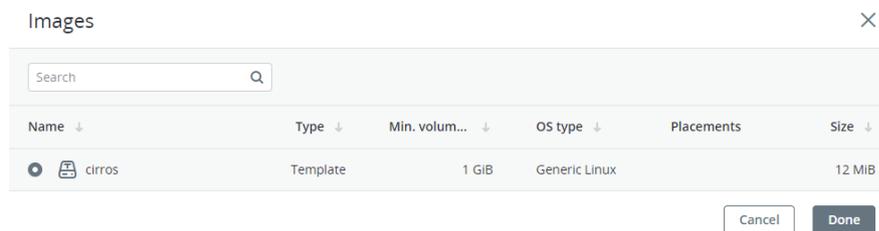
You can monitor compute cluster deployment on the **Compute** screen.

# Creating a virtual machine

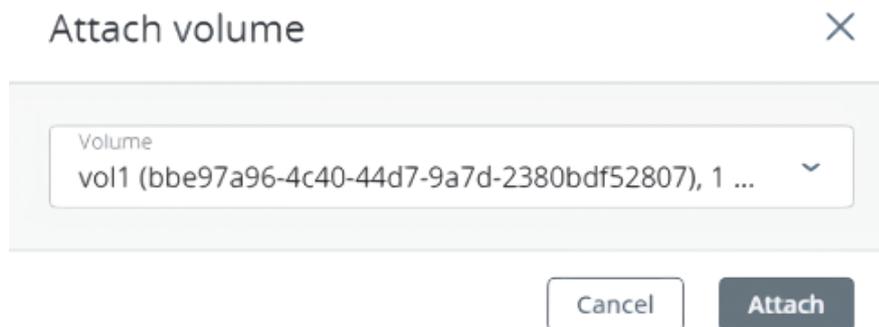
## Note

For supported guest operating systems and other information, refer to "Managing virtual machines" in the Administrator Guide.

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.

---

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 ⓘ + Add

IPv4:

Assign automatically

 Assign automatically 🗑

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.
- c. [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.
  - d. 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:

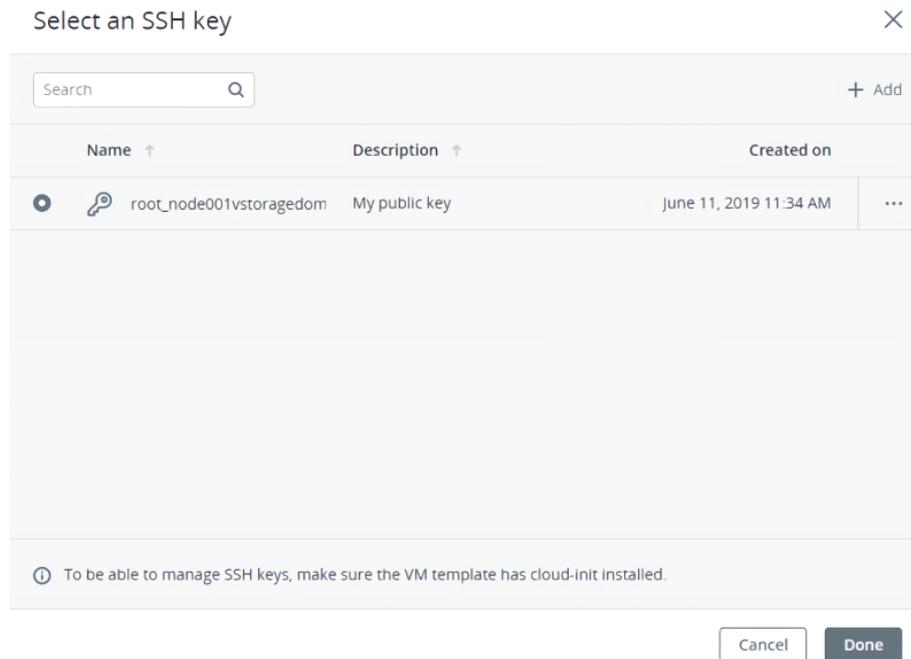
---

### Important

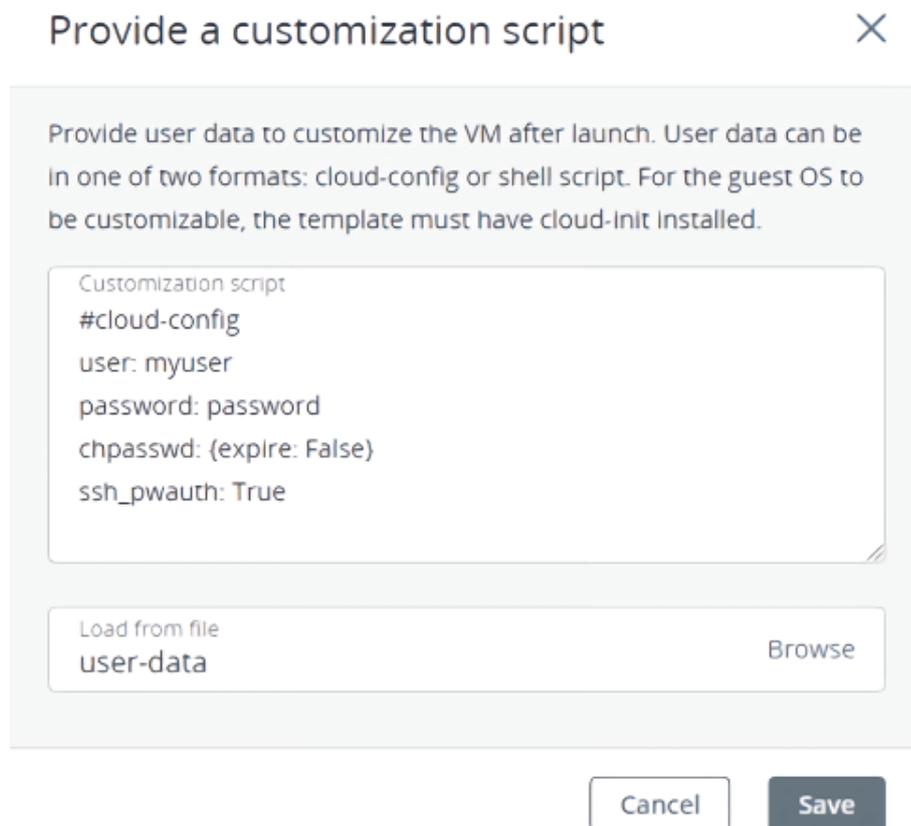
As cloud images have no default password, you can access VMs deployed from them only by using the key authentication method with SSH.

---

- 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**.



- 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.



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.