



Virtuozzo Infrastructure Platform 3.0

Quick Start Guide

November 20, 2019

Virtuozzo International GmbH

Vordergasse 59

8200 Schaffhausen

Switzerland

Tel: + 41 52 632 0411

Fax: + 41 52 672 2010

<https://virtuozzo.com>

Copyright ©2016-2019 Virtuozzo International GmbH. All rights reserved.

This product is protected by United States and international copyright laws. The product's underlying technology, patents, and trademarks are listed at <https://virtuozzo.com>.

Microsoft, Windows, Windows Server, Windows NT, Windows Vista, and MS-DOS are registered trademarks of Microsoft Corporation.

Apple, Mac, the Mac logo, Mac OS, iPad, iPhone, iPod touch, FaceTime HD camera and iSight are trademarks of Apple Inc., registered in the US and other countries.

Linux is a registered trademark of Linus Torvalds. All other marks and names mentioned herein may be trademarks of their respective owners.

Contents

1. Introduction	1
1.1 About This Guide	1
1.2 Hardware Requirements	1
2. Installing Virtuozzo Infrastructure Platform	3
3. Creating the Storage Cluster	5
4. Enabling Management Node High Availability	7
5. Deploying Compute Cluster	10
6. Creating a Virtual Machine	14

CHAPTER 1

Introduction

1.1 About This Guide

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.

1.2 Hardware Requirements

A minimum Virtuozzo Infrastructure Platform 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 the “Hardware Requirements” section of the *Installation Guide*.

Table 1.2.1: Node hardware requirements

Type	Management and storage/compute node
CPU	64-bit x86 processors with AMD-V or Intel VT hardware virtualization extensions enabled. 16 logical CPUs*
RAM	32 GB**

Continued on next page

Table 1.2.1 – continued from previous page

Type	Management and storage/compute node
Storage	1 disk: system + metadata, 100+ GB SATA HDD 1 disk: storage, SATA HDD, size as required
Network	1 GbE for storage traffic 1 GbE (VLAN tagged) for other traffic

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

** Each chunk server (CS), e.g., storage disk, requires 1 GB of RAM (0.5 GB anonymous memory + 0.5 GB page cache). The total page cache limit is 12 GB. In addition, each metadata server (MDS) requires 0.2 GB of RAM + 0.1 GB per 100TB of physical storage space.

CHAPTER 2

Installing Virtuozzo Infrastructure Platform

To install Virtuozzo Infrastructure Platform, do the following:

1. Prepare 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).
2. Boot the server from the chosen media.
3. On the welcome screen, choose **Install Virtuozzo Infrastructure Platform**.
4. On step 1, please carefully read the End-User License Agreement. Accept it by ticking the **I accept the End-User License Agreement** checkbox and click **Next**.
5. On step 2, configure a static IP address for the NIC and provide a hostname: either a fully qualified domain name (hostname, domainname) or a short name (hostname).
6. 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.
7. 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.
 - 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. Panel's IP address is shown in the console after deploying the primary node. Use the default user name

shown on the login screen and the primary node's root password. In the admin panel, open **INFRASTRUCTURE > Nodes** and click **ADD NODE** to invoke a screen with the management node address and the token.

The node may appear on the **INFRASTRUCTURE > Nodes** screen in the **UNASSIGNED** list as soon as token is validated. However, you will be able to join it to the storage cluster only after the installation is complete.

8. 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.
9. On step 6, enter and confirm the password for the root account and 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.

CHAPTER 3

Creating the Storage Cluster

To create the storage cluster, do the following:



1. Open the **INFRASTRUCTURE** > **Nodes** screen and click a node in the **UNASSIGNED** list.
2. On the node overview screen, click **Create cluster**.
3. In the **Cluster** field, type a name for the cluster. The name may only contain Latin letters (a-z, A-Z), numbers (0-9), underscores ("_") and hyphens ("-").



✕ **New cluster**

Create cluster on node **node001**

Cluster

Storage interface

 Encryption 



NEW CLUSTER **ADVANCED CONFIGURATION**

4. Click **NEW CLUSTER**.
5. Click the next node in the **UNASSIGNED** list and click **JOIN CLUSTER**.

✕ Join cluster

Node to join: **node002**

Storage interface

eth1 - 10.37.130.127  

6. Repeat the previous step for the remaining unassigned nodes.

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

CHAPTER 4

Enabling Management Node High Availability

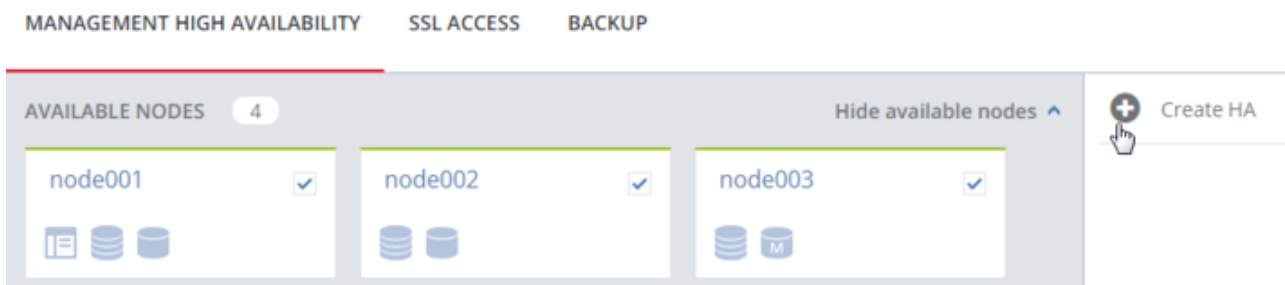
To make your infrastructure more resilient and redundant, you can create a high availability configuration of three nodes.

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

1. All nodes in the HA configuration will be added to the compute cluster.
2. 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. Make sure that each node is connected to a network with the **Admin panel** and **Internal management** traffic types.
2. On the **SETTINGS > Management node** screen, open the **MANAGEMENT HIGH AVAILABILITY** tab.



3. Select three nodes and click **Create HA**. The management node is automatically selected.
4. On **Configure network**, check that correct network interfaces are selected on each node. Otherwise, click the cogwheel icon for a node and assign networks with the **Internal management** and **Admin panel** traffic types to its network interfaces. Click **PROCEED**.

✕ **Configure network**

 node001 	
Management	Admin panel
eth1 - 10.37.130.109	br-eth0 - 10.94.16.12
<hr/>	
 node002 	
Management	Admin panel
eth1 - 10.37.130.127	br-eth0 - 10.94.16.34
<hr/>	
 node003 	
Management	Admin panel
eth1 - 10.37.130.148	br-eth0 - 10.94.16.47

5. Next, on **Configure network**, provide one or more unique static IP addresses for the highly available admin panel, compute API endpoint, and interservice messaging. Click **DONE**.

< Configure network

Assign unique dedicated virtual IP addresses to these services:

- Admin panel (public access to this web UI)
- Compute API (public access to compute APIs)
- Internal management (private interservice messaging)

In a high availability event, virtual IP addresses will automatically migrate to a healthy node in the high availability cluster to keep services accessible.

Virtual IP address for
Compute API, Admin panel

i The IP address must belong to the network **Public** (10.94.0.0/16)

Virtual IP address for
Internal management

i The IP address must belong to the network **Private** (10.37.130.0/24)

DONE

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).

CHAPTER 5

Deploying Compute Cluster

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

1. Network is set up according to recommendations in the “Managing Networks and Traffic Types” section of the *Administrator’s Guide*. The basic requirements are: (a) the traffic types **VM private**, **VM public**, and **Compute API** must be assigned to networks; (b) the nodes to be added to the compute cluster must be connected to these networks.
2. All nodes to be added to the compute cluster are connected to the same network with the **VM public** traffic type.
3. High availability for the management node is enabled (see the *Enabling Management Node High Availability* (page 7)).

Also take note of the following:

1. Creating the compute cluster prevents (and replaces) the use of the management node backup and restore feature.
2. If nodes to be added to the compute cluster have different CPU models, consult the section “Setting Virtual Machines CPU Model” in the *Administrator’s Command Line Guide*.

To create the compute cluster, open the **COMPUTE** screen, click **Create compute cluster** and do the following in the **Configure compute cluster** window:

1. In the **Nodes** section, select nodes to add to the compute cluster, make sure the network state of each selected node is **Configured**, and click **Next**.

Nodes in the management node high availability cluster are automatically selected to join the compute cluster.

Configure compute cluster ×

● Nodes
Select nodes to add to the compute cluster.

- Public network
- DHCP and DNS
- Summary

<input checked="" type="checkbox"/>	Name ↓	Node status	IP address	Network state
<input checked="" type="checkbox"/>	node001.vsto... ⓘ	Healthy	10.37.130.109	✔ Configured ⚙
<input checked="" type="checkbox"/>	node002.vstorage...	Healthy	10.37.130.127	✔ Configured ⚙
<input checked="" type="checkbox"/>	node003.vstorage...	Healthy	10.37.130.148	✔ Configured ⚙

Next

If node network interfaces are not configured, click the cogwheel icon, select networks as required, and click **Apply**.

- In the **Public network** section, enable IP address management if needed and provide the required details for the public network.

With IP address management enabled, Virtuozzo Infrastructure Platform 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 to virtual machines. 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 network's entire IP range. With the DHCP server disabled, VM network interfaces will still get 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.

If you choose to enable IP address management, select a physical network to connect the public virtual

network to and optionally specify its gateway. The subnet IP range in the CIDR format will be filled in automatically. If you choose to leave IP address management disabled, select a physical network to connect the public virtual network to.

The screenshot shows a dialog box titled "Configure compute cluster" with a close button (X) in the top right corner. On the left side, there is a vertical navigation menu with four items: "Nodes", "Public network" (which is highlighted with a grey background), "DHCP and DNS", and "Summary". The main content area on the right is titled "Specify the subnet CIDR and gateway for the external virtual network." and contains the following elements:

- A toggle switch for "IP address management" which is currently turned on (green).
- A dropdown menu for "Physical network" with "Public" selected.
- A text input field for "Subnet CIDR" containing the value "10.94.0.0/16".
- A text input field for "Gateway (optional)" which is currently empty.

At the bottom right of the dialog, there are two buttons: "Back" and "Next".

The selected public network will appear in the list of virtual networks on compute cluster's **NETWORKS** tab.

Click **Next**.

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**.

Configure compute cluster ✕

- Nodes
- Public network
- 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

10.94.10.128 — 10.94.10.254	126 addresses available	✎ 🗑
-----------------------------	-------------------------	-----

DNS servers + Add server

10.94.0.10	✎ 🗑
------------	-----

Back
Next

4. In the **Summary** section, review the configuration and click **Create cluster**.

Configure compute cluster ✕

- Nodes
- Public network
- DHCP and DNS
- Summary

Review the compute cluster details and go back to change them if necessary.

Nodes	node001.vstoragedomain. (10.37.130.109) node003.vstoragedomain. (10.37.130.148) node002.vstoragedomain. (10.37.130.127)
Subnet CIDR	10.94.0.0/16
Physical network	Public
DHCP	Enabled
Allocation pools	10.94.10.128 — 10.94.10.254 126 addresses available
DNS servers	10.94.0.10

Back
Create cluster

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

CHAPTER 6

Creating a Virtual Machine

Note: For supported guest operating systems and other information, see the “Managing Virtual Machines” section in the *Administrator’s Guide*.

To create a VM, do the following:









1. On the **COMPUTE > Virtual machines > VIRTUAL MACHINES** tab, 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.

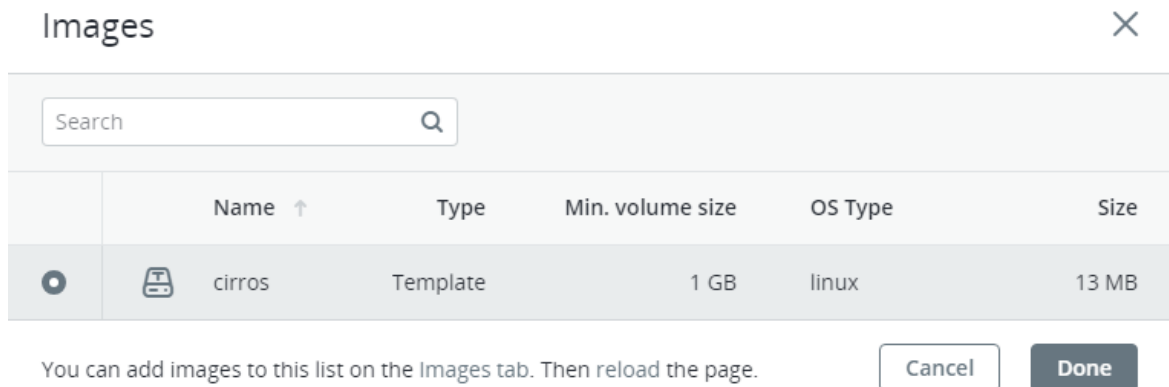
Name
vm1

Deploy from: Image Volume

 Image	Specify	
 Volumes	Specify	
 Flavor	Specify	
 Networks	Specify	

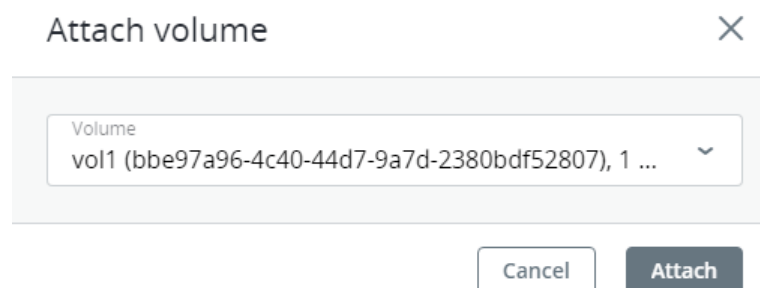
Deploy

2. Specify a name for the new VM.
3. In **Deploy from**, choose **Volume** if you have a boot volume or want to create one. Otherwise, choose **Image**.
4. 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**.



Then, in the **Volumes** window, make sure the default boot volume is large enough to accommodate the guest OS and click **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**.



- Otherwise, click **Add**. In the **Create volume** window, specify a name, size in GB, and select a storage policy from the drop-down list. Click **Add**.

Create volume ✕

Name
vol1

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

Storage policy
default ▼






Delete on termination

CancelAdd

The top volume in the list is considered bootable. So the first created or attached volume becomes the boot volume by default.

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

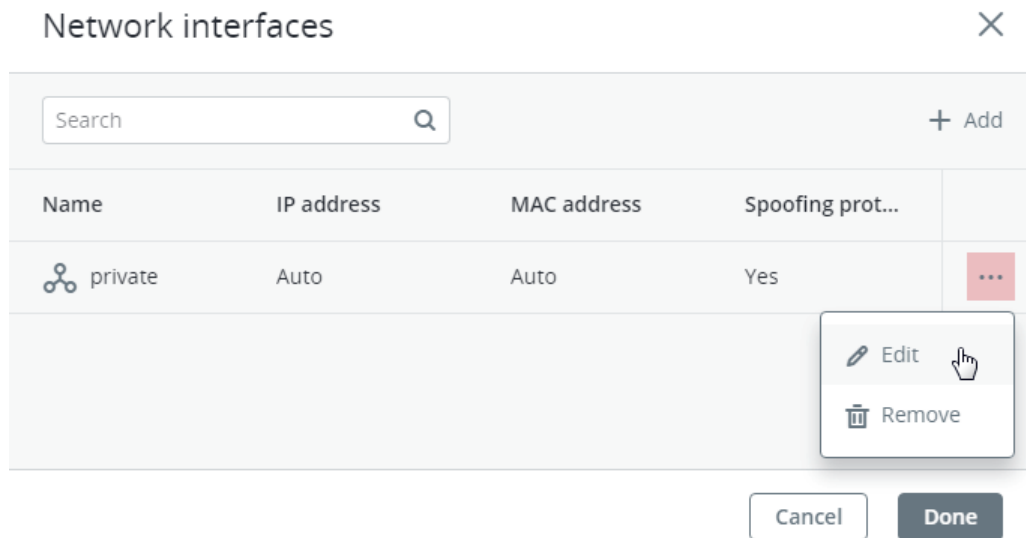
You can add flavors to this list on the Flavors tab. Then reload the page.

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

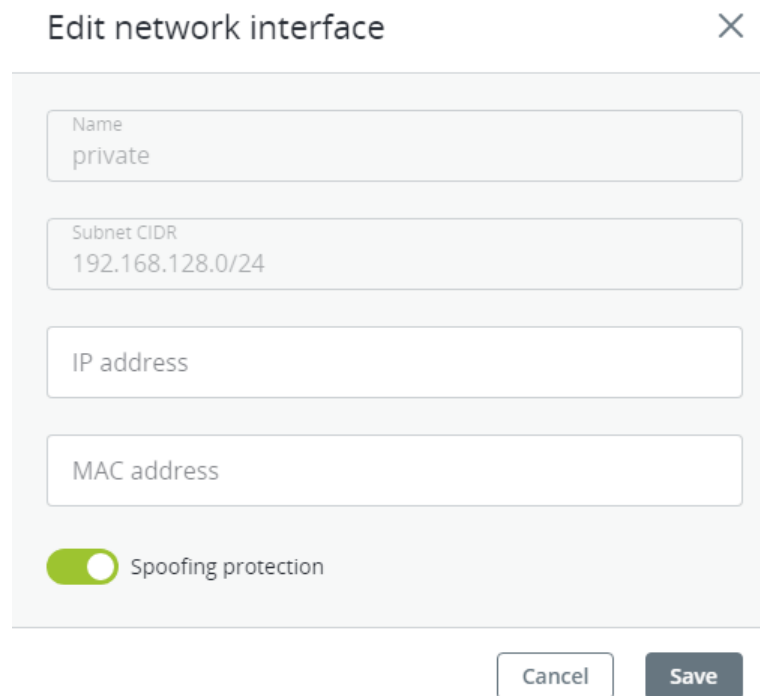
Add network interface ✕

Network
public ▼

You can edit additional parameters of newly added network interfaces, like IP and MAC addresses and spoofing protection. To do this, click interface's ellipsis icon, then **Edit**, and set parameters in the **Edit network interface** window.



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 11, 2019 11:34 AM	⋮

ⓘ 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 “Creating SSH-Enabled Templates” section 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
```