



Virtuozzo Hybrid Infrastructure 4.6

Evaluation Guide

8/30/2021

Table of contents

1 Introduction	4
2 Planning the infrastructure	5
2.1 Hardware requirements	5
2.2 Understanding storage policies	6
3 Managing the storage cluster	9
3.1 Installing Virtuoizzo Hybrid Infrastructure	9
3.2 Configuring networks	10
3.3 Creating the storage cluster	11
4 Managing the compute cluster	15
4.1 Creating the compute cluster	15
4.2 Allocating resources	16
4.2.1 Creating domains, projects, and users	16
4.2.2 Creating storage policies	18
4.2.3 Creating compute networks	18
4.3 Creating virtual machines	19
5 Exporting storage space	21
5.1 Exporting storage space via iSCSI	21
5.1.1 Creating target groups	21
5.1.2 Creating volumes	22
5.1.3 Attaching volumes to target groups	22
5.1.4 Accessing iSCSI targets from VMware vSphere	22
5.2 Exporting storage space via S3	23
5.2.1 Creating an S3 cluster	23
5.2.2 Managing S3 users and buckets	24
5.3 Exporting storage space via NFS	25
5.3.1 Creating an NFS cluster	25
5.3.2 Creating NFS shares	26
5.3.3 Creating and mounting NFS exports	26
6 Connecting Acronis Cyber Protect Cloud software to backup storage	27
6.1 Creating backup storage	27
6.2 Configuring Acronis Cyber Protect Cloud	28
7 Monitoring the storage cluster	30
8 Enabling high availability	31
8.1 High availability for the management node	31
8.2 High availability for the services	33

8.3 Testing high availability	34
-------------------------------------	----

1 Introduction

Virtuozzo Hybrid 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 gets you started on Virtuozzo Hybrid Infrastructure and outlines the following steps to evaluate its main features:

1. Install and configure Virtuozzo Hybrid Infrastructure.
2. Create a storage cluster.
3. Create a compute cluster and allocate its resources.
4. Create virtual machines.
5. Export storage space via iSCSI, S3, NFS, Backup Gateway.
6. Explore built-in monitoring tools.
7. Enable and test high availability.

There are many different scenarios possible, but in this guide we'll walk you through the most common ones. The procedures outlined in this guide are typical and simplified for evaluation purposes. However, Virtuozzo Hybrid Infrastructure is supported by a complete and detailed documentation suite, which you should refer to for further guidance. For more information, refer to the Quick Start Guide, Installation Guide, Administrator Guide, Self-Service Guide, and Storage User Guide.

2 Planning the infrastructure

2.1 Hardware requirements

There are many hardware configurations supported and described in "System requirements" in the Administrator Guide. However, for the evaluation purposes, we recommend deploying three nodes. This is to ensure that the cluster can survive the failure of one node without data loss. The following table lists the *minimum* hardware requirements for each of the three nodes. To enable high availability, you will need three nodes meeting the requirements for the management node with storage and compute.

Minimum node hardware requirements

Type	Management node with storage and compute	Secondary node with storage and compute	Management node with storage and Backup Gateway
CPU	64-bit x86 processors with AMD-V or Intel VT hardware virtualization extensions enabled		
	16 cores*	8 cores*	4 cores*
RAM	32 GB	16 GB	8 GB
Storage	1 disk: system + metadata, 100+ GB SATA HDD 1 disk: storage, SATA HDD, size as required	1 disk: system, 100 GB SATA HDD 1 disk: metadata, 100 GB SATA HDD (only on the first three nodes in the cluster) 1 disk: storage, SATA HDD, size as required	1 disk: system + metadata, 120 GB SATA HDD 1 disk: storage, SATA HDD, size as required
Network	10 GbE for the private network 1 GbE for the public network		

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

You will also need at least two IPs in the public network (for the highly available admin panel and for each NFS share), one IP in the private network (for the highly available management node), and two DNS names (for S3 and Backup Gateway).

Following is a sample network layout for our evaluation scenario. Your host names and addresses may differ depending on your settings.

RESERVED ADDRESSES

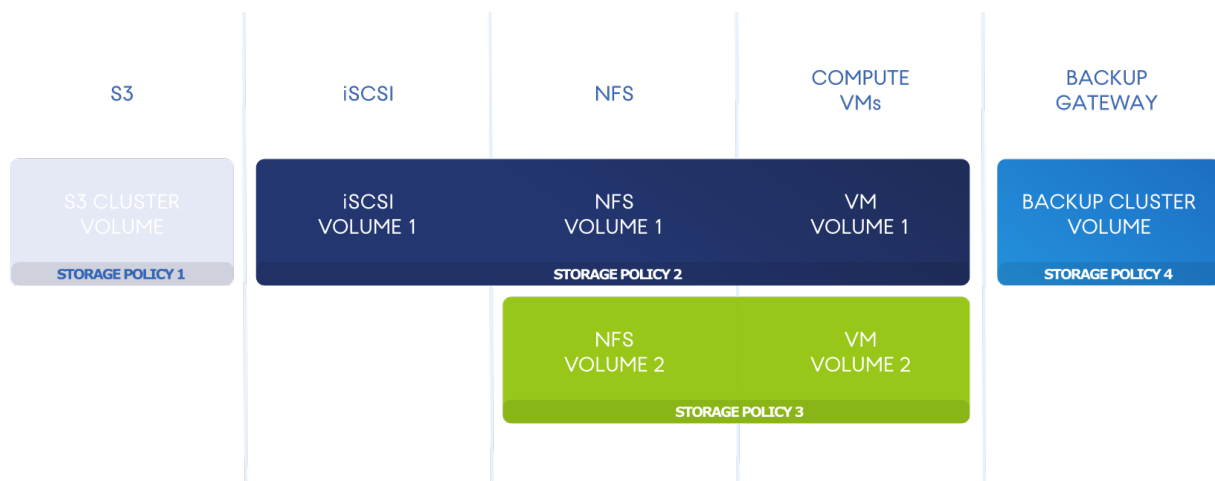


The DNS names are used by backup agents to send backups to and retrieve them from cloud storage. If your client machines and Virtuozzo Hybrid Infrastructure nodes are located in different datacenters, we strongly recommend configuring the DNS name in one of your preferred DNS services: Azure DNS for Azure, Amazon Route 53 for Amazon, or Google Cloud DNS for Google. If it is not possible for some reason, you have to manually add the DNS name to the `/etc/hosts` file on each node. On the other hand, if the client machines and Virtuozzo Hybrid Infrastructure nodes are located in the same datacenter, you can use a local DNS instead of a public one to speed up the backup. A local DNS can work only with a local IP address, so we recommend using NAT in this case.

We recommend using the Google Chrome browser for accessing Virtuozzo Hybrid Infrastructure and S3 storage in our evaluation scenario.

2.2 Understanding storage policies

Virtuozzo Hybrid Infrastructure can be used for the following scenarios: iSCSI block storage, NFS file storage, S3 object storage, Backup Storage (to store the backups created in Acronis Cyber Protect solutions). You can also use its built-in hypervisor to create compute virtual machines (VM). In all these scenarios, the common unit of data is a volume. For the compute service, a volume is a virtual drive that can be attached to a VM. For iSCSI, S3, Backup Gateway, and NFS, a volume is the data unit used for exporting space. In all these cases, when you create a volume, you need to define its *redundancy mode*, *tier*, and *failure domain*. These parameters make up a *storage policy* defining how redundant a volume must be and where it needs to be located.



Redundancy means that the data is stored across different storage nodes and stays highly available even if some nodes fail. If a storage node is inaccessible, the data copies on it are replaced by new ones that are distributed among healthy storage nodes. When the storage node goes up after the downtime, out-of-date data on it is updated.

With replication, Virtuozzo Hybrid Infrastructure breaks a volume into fixed-size pieces (data chunks). Each chunk is replicated as many times as is set in the storage policy. The replicas are stored on different storage nodes if the failure domain is host, so that each node has only one replica of a given chunk.

With erasure coding (or just encoding), the incoming data stream is split into fragments of a certain size. Then, each fragment is not copied itself; instead, a certain number (M) of such fragments are grouped and a certain number (N) of parity pieces are created for redundancy. All pieces are distributed among M+N storage nodes (selected from all available nodes). The data can survive the failure of any N storage nodes without data loss. The values of M and N are indicated in the names of erasure coding redundancy modes. For example, in the 5+2 mode, the incoming data is split into 5 fragments, and 2 more parity pieces (same size) are added for redundancy. Refer to the Administrator Guide for detailed information on redundancy, data overhead, number of nodes, and raw space requirements.

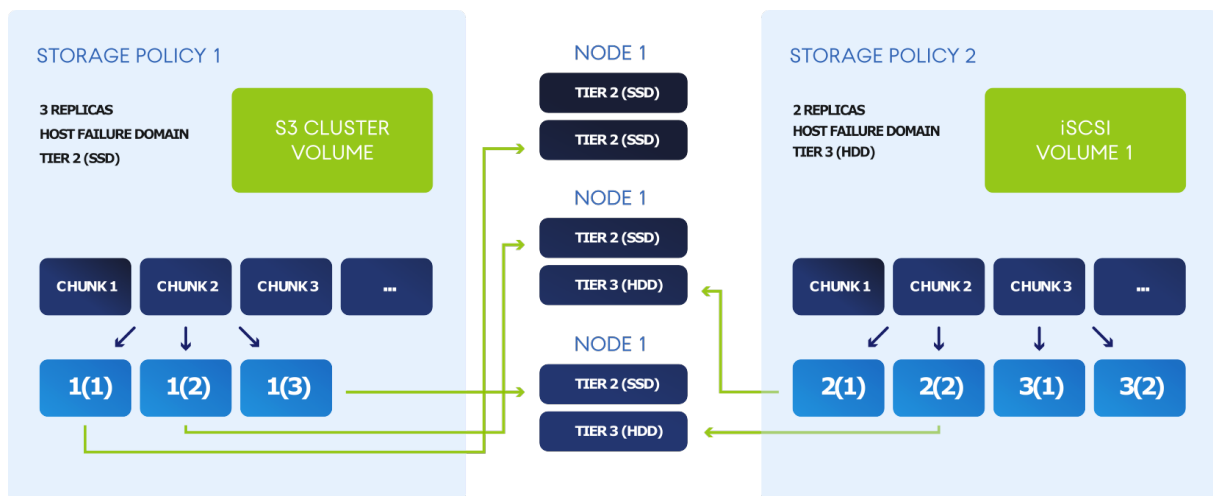
To better understand a storage policy, let's have a look at its main components (tiers, failure domains, and redundancy), for a sample scenario. For example, you have three nodes with a number of storage nodes: fast SSDs and high-capacity HDDs. Node 1 has only SSDs; nodes 2 and 3 have both SSDs and HDDs. You want to export storage space via iSCSI and S3, so you need to define a suitable storage policy for each workload.

- The first parameter, tier, defines a group of disks united by criteria (drive type, as a rule) tailored to a specific storage workload. For this sample scenario, you can group your SSD drives into tier 2, and HDD drives into tier 3. You can assign a disk to a tier when creating a storage cluster or adding nodes to it (refer to "Creating the storage cluster" (p. 11)). Note that only nodes 2 and 3 have HDDs and will be used for tier 3. The first node's SSDs cannot be used for tier 3.
- The second parameter, failure domain, defines a scope within which a set of storage services can fail in a correlated manner. The default failure domain is host. Each data chunk is copied to

different storage nodes, just one copy per node. If a node fails, the data is still accessible from the healthy nodes. A disk can also be a failure domain, though it is only relevant for one-node clusters. As you have three nodes in this scenario, we recommend choosing the host failure domain.

- The third parameter, redundancy, should be configured to fit the available disks and tiers. In our evaluation example, you have three nodes: all of them have SSDs on tier 2. So, if you select tier 2 in your storage policy, you can use the three nodes for 1, 2, or 3 replicas. But only two of your nodes have HDDs on tier 3. So, if you select tier 3 in your storage policy, you can only store 1 or 2 replicas on the two nodes. In both cases, you can also use encoding, but for our evaluation, let's stick to replication: 3 replicas for SSDs and 2 replicas for HDDs.

To sum it up, the resulting storage policies are:



3 Managing the storage cluster

This chapter outlines the steps to install Virtuozzo Hybrid Infrastructure and configure its initial settings for further deployment. First, you need to create a basic infrastructure on the management node, and then add secondary nodes to it in a similar manner. We recommend adding two secondary nodes for evaluation. Then, configure the networks and create a storage cluster on your nodes.

3.1 Installing Virtuozzo Hybrid Infrastructure

To install Virtuozzo Hybrid 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 Hybrid 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.

To get detailed information on the node, log in to the admin panel on port 8888, go to the **Infrastructure > Nodes** screen, and then click the node name. Go to the **Disks** tab to configure or view the node disks, or go to the **Network** tab to configure the node's network interfaces.

3.2 Configuring networks

Now that you have installed Virtuozzo Hybrid Infrastructure on the management and two secondary nodes, you need to set up networks and interfaces. Use separate networks for internal and public traffic. Doing so will prevent public traffic from affecting cluster I/O performance and also prevent possible denial-of-service attacks from the outside.

1. To configure networks, go to the **Infrastructure > Networks** screen on the admin panel. The advanced configuration is discussed in the Administrator guide, but for our simplified deployment, it is enough to customize the default **Public** and **Private** networks as follows:

Network	Traffic types
Public	Compute API, S3 public, iSCSI, NFS, Backup (ABGW) public, Admin panel, SSH, VM public, SNMP, Self-service panel
Private	Storage, Internal management, OSTOR private, Backup (ABGW) private, VM private, Admin panel, SSH

2. To configure interfaces, go to the **Infrastructure > Nodes** screen and click a node's name. Go to the **Network interfaces** tab. For the management node, both interfaces are already set up. You still need, however, to configure public network interfaces for every secondary node. Select an interface and click **Edit** on the right menu. In the **Network** field, select **Public**. You should now have one interface connected to the private network and the other assigned to the public network. Repeat these steps for every secondary node to connect them to the private and the public networks.
3. Ports that will be opened on cluster nodes depend on services that will run on the node and traffic types associated with them. For more information on the ports and services, refer to "Network ports" in the Administrator Guide.
4. Make sure your DNS settings are correct. To do that, go to the **Settings > Cluster DNS** screen. Check that the cluster DNS is configured properly and points to a DNS that can resolve external host names.

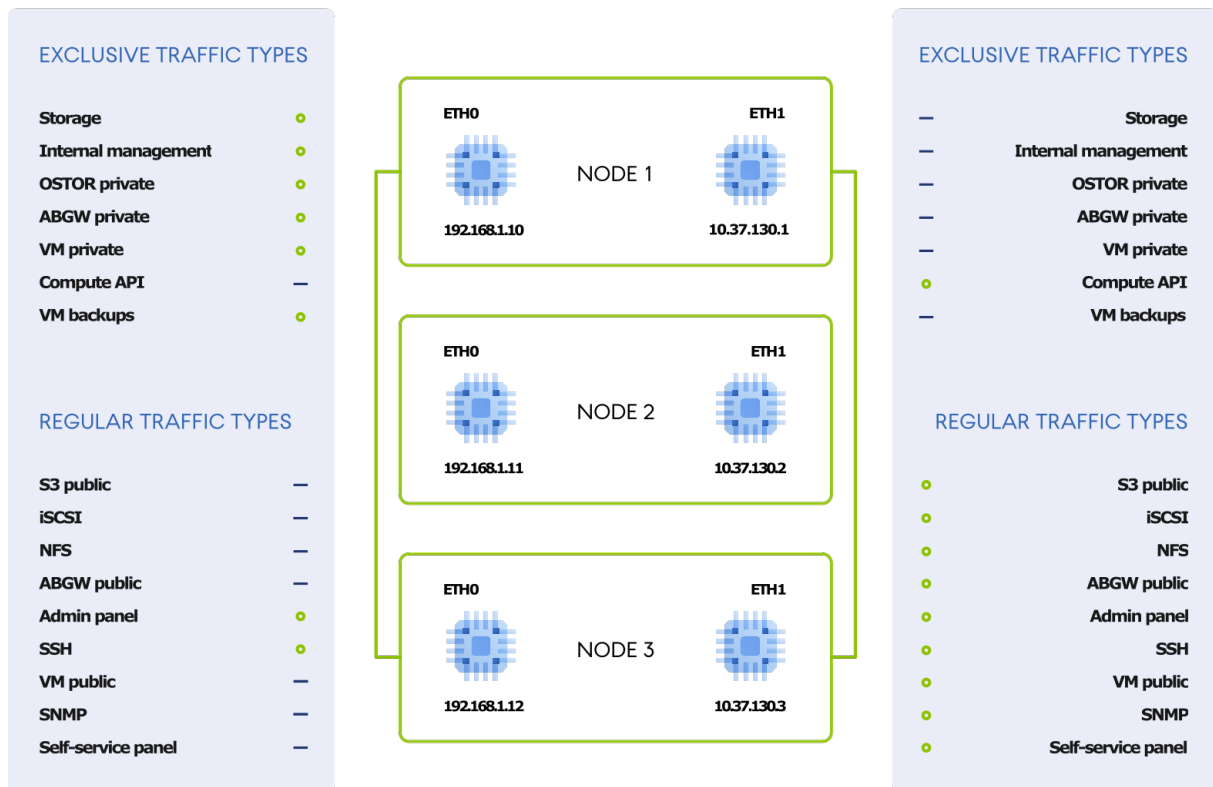
The figure below shows the sample network infrastructure we are going to build for our evaluation scenario:

PRIVATE NETWORK

192.168.1.0/24

PUBLIC NETWORK

10.37.130.0/24



Note

If you only have one network, do not connect one node to it via two interfaces. In case of one network, work with one public interface of the node.

3.3 Creating the storage cluster

Now that you have created the management and secondary nodes and configured networking, you can proceed to create the storage cluster.

1. Open the **Infrastructure > Nodes** screen, and then click **Create storage cluster**.
2. [Optional] To configure the disk roles or node location, click the cogwheel icon.
 - To configure disks
 - a. Select **Disks**.
 - b. Select the required disk, and then click **Configure**.
 - c. In the **Choose role** window, select a role:
 - To store data
 - i. Select the **Storage** role.
 - ii. Select a storage tier where to store your data. To make better use of data redundancy, do not assign all of the disks on a node to the same tier. Instead, make sure that each tier is evenly distributed across the cluster with only one disk per node assigned to it.
 - iii. Enable data caching and checksumming:

- **Enable SSD caching and checksumming.** Available and recommended only for nodes with SSDs.
 - **Enable checksumming** (default). Recommended for nodes with HDDs as it provides better reliability.
 - **Disable checksumming.** Not recommended for production. For an evaluation or testing environment, you can disable checksumming for nodes with HDDs, to provide better performance.
- To store cluster metadata
Select the **Metadata** role.

Note

It is recommended to have only one metadata service per node and maximum five metadata services for a cluster.

- [Only for SSD drives] To store write cache
 - i. Select the **Cache** role.
 - ii. Select a storage tier that you want to cache.
- [Only for SSD drives] To store both metadata and write cache
 - i. Select the **Metadata+Cache** role.
 - ii. Select a storage tier that you want to cache.

✕ Choose role

- d. Repeat the previous steps for every disk you want to be used in the storage cluster, and click **Done**.
- e. On the **Configuration summary** screen, check the number of disks per each configuration category, and then click **Proceed**.
- To configure node location
By default, the nodes are added to the **Default rack** in the **Default row** in the **Default room**.

- a. Select **Location**.
 - b. To create a room, click **Create room** in the **Move nodes** window. Specify the room name, and then click **Create**.
 - c. To add a new row, click the created room. Click **Create row**, enter a name for it, and then click **Create**.
 - d. To add a new rack, click the created row. Click **Create rack**, enter a name for it, and then click **Create**.
 - e. Click the created rack, and then click **Move**.
3. Enter a name for the cluster. It may only contain Latin letters (a-z, A-Z), numbers (0-9), and hyphens ("-").
 4. Enable encryption, if required.
 5. Click **Create**.

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 a secondary node, do the following:

1. On the **Infrastructure > Nodes** screen, click an unassigned node.
2. On the node right pane, click **Join to cluster**.
3. Click **Join** to assign the roles to disks automatically and add the node to the default location. Alternatively, click the cogwheel icon to configure the disk roles or node location.

For the evaluation scenario, you need to assign a node's storage disks to various tiers (refer to "Understanding storage policies" (p. 6)). For each node, assign SSDs to tier 2, and HDDs to tier 3.

Assign role



Select the role to assign to the disk "sdc"

☒ **Storage**

Use the disk to store data.

☐ **Cache**

Use the disk to store write cache. This disk does not add capacity to the cluster but improves its performance.

☐ **Metadata**

Use the disk to store cluster metadata.

☐ **Metadata + Cache**

Use the disk to store both cluster metadata and write cache.

Storage tier

Tier 0



Caching and checksumming

Enable checksumming



Cancel

Assign

4 Managing the compute cluster

Virtuozzo Hybrid Infrastructure offers high-performance virtualization, and its fundamental component is a compute cluster. It allows admins to create multiple tenants, virtual machines, and software-defined networks, as well as easily deploy container orchestration solutions like Kubernetes.

In this section, we will have a look at two common scenarios for the compute cluster:

- A service provider (SP) offers virtualization services to end customers. In this case, the SP can benefit from self-service panels with white-labeling, multitenancy, and easy resource management and allocation.
- An enterprise plans to adopt new software across its infrastructure. It can deploy virtual machines and install the software in them, for all employees to access and work with.

Although those two scenarios differ from each other significantly, the procedure within Virtuozzo Hybrid Infrastructure is similar:

1. Create the compute cluster.
2. Allocate resources to domains, projects (tenants), and users.
3. Create virtual machines for the end users.

In this chapter, we will take a look at each of these steps.

4.1 Creating the compute cluster

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. The network state of each selected node should be **Configured**, if you followed the instructions in "Configuring networks" (p. 10). Click **Next**.

Configure compute cluster

- Nodes
- Public network
- Add-on services
- Summary

Select nodes to add to the compute cluster.

Search

<input type="checkbox"/>	Name ↑	Node status	IP address	Network state	
<input checked="" type="checkbox"/>	host ⓘ	Healthy	192.168.1.10	✔ Configured	⚙
<input type="checkbox"/>	host3	Healthy	192.168.1.11	✔ Configured	⚙
<input type="checkbox"/>	node2	Healthy	192.168.1.12	✔ Configured	⚙

Next

2. In the **Physical network** section, you can enable IP address management if you have a pool of IP addresses that you can use. It is also necessary for virtual routers, floating public IP addresses,

and network load balancers. Otherwise, you can leave IP address management disabled. Select an infrastructure network to connect the physical network to and its type: select **VLAN** and specify a VLAN ID to create a VLAN-based network, or select **Untagged** to create a flat physical network.

3. In the **Add-on services** section, you can install additional services if you want to evaluate them as well. Or you can do it later, as described in the Administrator Guide.
4. On the **Summary** step, review the configuration, and then click **Create cluster**.

Configure compute cluster

Nodes

Public network

Add-on services

Summary

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

Nodes

node001 (192.168.1.10)
node002 (192.168.1.11)
node003 (192.168.1.12)

Physical network

Public

Add-on services

Kubernetes service
Load balancer service
Billing metering service

Back

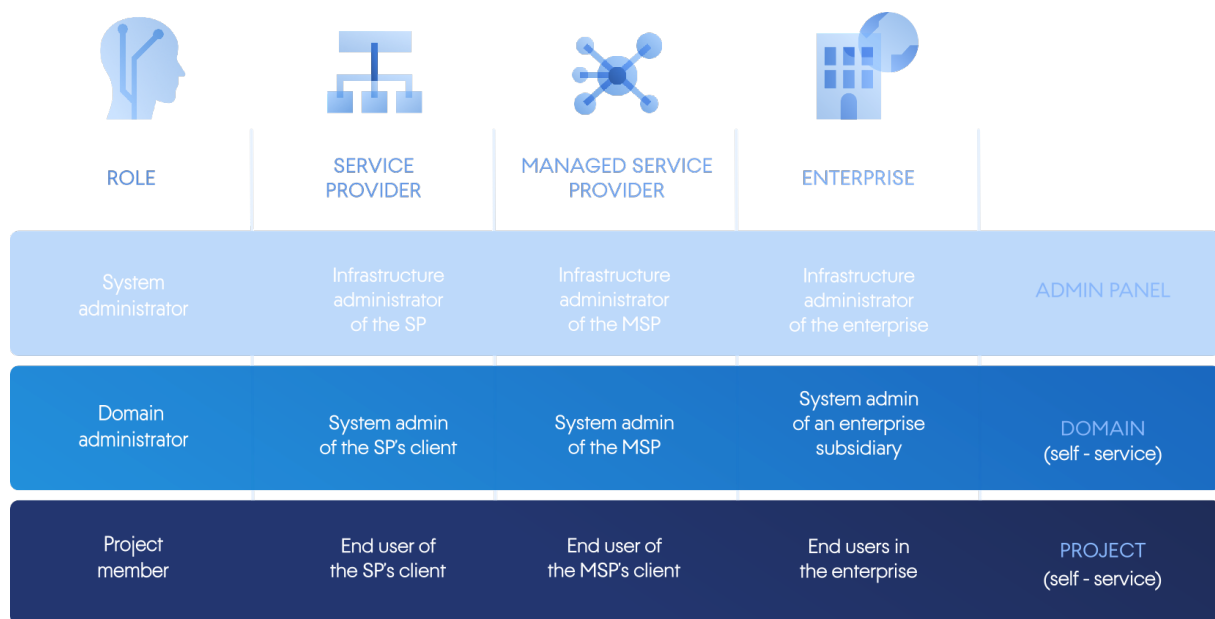
Create cluster

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

4.2 Allocating resources

4.2.1 Creating domains, projects, and users

There are three user roles in Virtuozzo Hybrid Infrastructure: a system administrator, a domain administrator, and a project member. The following chart shows typical users with these roles working at service providers and enterprises, along with their workspaces: admin or self-service panels.



- System administrators have full control over Virtuozzo Hybrid Infrastructure and can access the admin panel. This is the role you get by default when installing Virtuozzo Hybrid Infrastructure. These are usually infrastructure administrators of an MSP or the main IT department of an enterprise, depending on your business case.
- Domain administrators are in charge of their domains. A domain is a collection of virtualization projects (tenants) and users (end customers). Domain administrators have access to the self-service panel. They can create users, as well as use project resources within allowed quotas: deploy and manage virtual machines, images, volumes, networks, routers, floating IPs, and SSH keys.
- Project members can manage resources within their projects by using the self-service panel: deploy and manage virtual machines, images, volumes, networks, routers, floating IPs, and SSH keys. A project is a set of compute and storage resources defined by quotas and accessible by assigned users.

Both the domain administrator and project member roles have certain limitations. For instance, they cannot migrate virtual machines between nodes, as nodes are not present at that level of abstraction.

In our evaluation scenario, you are the system administrator. Once you have created the compute cluster, you need to create a domain, a project, some end users, and assign them to the project. Then, create a storage policy for VM volumes and define their redundancy parameters. Next, configure virtual networking. After that, domain users will have access to their domains and projects via the self-service panel. There, they will be able to create their own virtual machines, volumes, networks, etc.

Note

The self-service panel IP address is shown on the **Settings > Self-service** screen in the admin panel.

The actions to perform in the self-service panel are described in the Self-Service Guide. For our evaluation, we shall stick to the operations with the compute cluster done from the admin panel.

1. Create a domain. To do this, log in to the admin panel and open the **Settings > Projects and users** screen. Click **Create domain** in the upper-right corner. Specify a name and a description for the new domain. Click **Create**.
2. Create an admin account for the new domain. To do this, select the newly created domain and click **Create user**. Specify a login and a password, and then select the **Domain administrator** role. Check the **Image uploading** box to allow the new administrator to upload images for deploying virtual machines. Click **Create**.
3. Create a project. To do this, navigate to the domain's **Projects** tab and click **Create project**. Set the quotas and click **Create**. Make sure that you have enough CPU, RAM, storage, and network resources to deploy virtual machines (and add-on services, should you choose to enable them).
4. Create a project member. To do this, open the **Domain users** tab and click **Create user**. Specify a login and a password, and then select the **Project member** role. Choose the project to assign the new member to, and then click **Create**.
5. Optionally, customize the self-service panel's look on the **Settings > Self-service** screen by adding logos and selecting a color scheme. This way, managed service providers, for example, can offer branded virtualization services to end customers.

4.2.2 Creating storage policies

In order to create a new storage policy, go to the **Compute > Storage** screen, open the **Storage policies** tab, and then click **Create storage policy**. Specify a name, a tier, a failure domain, and a redundancy scheme. For the evaluation scenario, select the **2 replicas** mode and **Host** as the **Failure domain**.

Now that you have created a storage policy, you can select it for volumes when creating virtual machines (refer to "Creating virtual machines" (p. 19)). You can also apply it when creating volumes directly on the **Volumes** tab.

Storage policies can be used in project quotas. A policy created before a project will be enabled in its quotas. A policy created after a project will not be enabled for that project automatically. You will need to edit that project's quotas and select the policy manually.

4.2.3 Creating compute networks

Before deploying virtual machines, you need to configure networking on the **Compute > Network** screen. A virtual network lets virtual machines connected to it communicate with each other. A physical network connects your virtual machines to an existing infrastructure network, so they can access the Internet, for example.

In order to create a new network, click **Create network** and specify its type. Provide a name for the new network, a subnet CIDR, for example, **192.168.0.1/24**, and a gateway. Click **Next** to proceed. If you have a pool of IP addresses for your virtual machines, you can enable the built-in DHCP server to have these IP addresses automatically assigned to virtual machines. Click **Create network** to complete the process.

Along with compute networks, you can create floating IPs. A floating IP is a public IP address that you can manually assign to a private IP address of a virtual machine. It will let you access the virtual machine from the public network, even though it only has a private IP address. To create a floating IP, you will first need to link a physical and a virtual network with a virtual router. For more details, refer to "Managing floating IP addresses" in the Administrator Guide.

4.3 Creating virtual machines

Now that you have created a compute cluster, a domain and a project, a storage policy, and networks, you can proceed to create virtual machines.

In our evaluation scenario, you will create a virtual machine from an image. You can either upload one or use the Cirros image shipped by default. To upload an image, open the **Compute > Virtual Machines > Images** tab. You can use ISO images and templates (ready-to-use volumes in the QCOW2 cloud image format with the OS and apps installed). Click **Add Image** and select an ISO image from your local machine. Specify a name for the new image and select a compatible operating system from the drop-down list. Check the **Share between all projects** box, if you want to use this image as a template for future VM deployments. Click **Add**.

Note

Virtuozzo Hybrid Infrastructure supports a wide range of Windows and Linux guest operating systems that you can deploy virtual machines from (refer to "Supported guest operating systems" in the Administrator Guide). Moreover, it uses a number of patented innovations to optimize the performance of deployed VMs. For example, Windows-based VMs should perform as if they were deployed on Hyper-V.

1. On the **Compute > Virtual Machines** screen, click **Create virtual machine** and specify a name.
2. To create a VM from an image, select the image uploaded earlier. In the section **Volumes**, you will see the bootable volume automatically added, based on the image data. You can change the volume's storage policy. To do that, click the pencil sign in the **Volumes** section, click the ellipsis icon in the **Volumes** window, select **Edit**, and then change the policy in the **Edit volume** window. You can also add new volumes to the VM here.
3. In the section **Flavor**, select a flavor. This is a preset defining how many virtual CPUs and how much memory the virtual machine will have.
4. In the section **Networks**, add interfaces to the virtual networks that your virtual machine should be connected to.
5. Click **Deploy** to start virtual machine creation. Watch the status of the new virtual machine. Once it becomes **Active**, the VM is ready.

In order to access the newly created virtual machine, click its name, and then click **Console** on the right pane. On the **Monitoring** tab of the VM, you can see how much resources it uses.

Note

You can also migrate virtual machines from VMware vCenter by using the `virt-v2v` tool, as described in "Migrating virtual machines from VMware vCenter" in the Administrator Command Line Guide.

Once your virtual machine is ready, you can perform a wide range of operations on it: stop and start, suspend and resume, reboot, migrate, and more. For more details, refer to "Managing virtual machines" in the Administrator Guide.

5 Exporting storage space

Virtuozzo Hybrid Infrastructure allows you to export storage space:

- As block storage via iSCSI for virtualization, databases, and other needs. You can export cluster disk space to external physical or virtual hosts, in the form of LUN block devices over iSCSI and in a SAN-like manner.
- As object storage for storing an unlimited number of files via an S3-compatible protocol. The S3-like object storage can store data like media files, backups, and Open Xchange files, with the access provided via Dropbox-like applications. End users can run the applications for S3 after the data migration from Amazon S3 to Virtuozzo Hybrid Infrastructure. They can also build their own object storage services which are compatible with Amazon S3.
- Via NFS. You can organize nodes into a highly available NFS cluster where you can create NFS shares. In each share, you can create multiple NFS exports that are actual exported directories for user data. Each export can be mounted by using standard commands. On the technical side, NFS volumes are based on object storage. An NFS cluster makes a perfect cold and warm file storage, but is not recommended for hot workloads with high performance requirements. For the best integration with VMware vSphere, it's recommended to use iSCSI protocol.

5.1 Exporting storage space via iSCSI

Block storage enables managing data as blocks, as opposed to files in file systems or objects in S3 storage. Those blocks can be stored in different operating systems in a SAN-like manner.

Virtuozzo Hybrid Infrastructure allows you to create groups of redundant targets running on different storage nodes. To each target group, you can attach multiple storage volumes with their own redundancy provided by the storage layer. Targets export these volumes as LUNs.

You can create multiple target groups on same nodes. A volume, however, may only be attached to one target group at any moment in time.

Each node in a target group can host a single target for that group. If one of the nodes in a target group fails along with its targets, healthy targets from the same group continue to provide access to the LUNs previously serviced by the failed targets.

5.1.1 Creating target groups

To create a target group, open **Storage services > Block storage > Target groups** and click **Create target group**. A wizard will open. Do the following:

1. For **Name and type**, enter a target group name and select iSCSI. Next, select at least two nodes to add to the target group for high availability.
2. For **Targets**, select iSCSI interfaces to add to the target group. For **Volumes**, select volumes to attach to target group LUNs, or you can add them later. For the evaluation scenario, skip the **Access control** settings.
3. For **Summary**, review the target group details. Click **Create**.

The newly created target group will appear on the **Target groups** tab. Its targets will start automatically. Click the group name to view its details. On the **Target** tab, you can add more nodes for new targets. You can also view or add LUNs on the **LUNS** tab.

5.1.2 Creating volumes

1. Open **Storage services > Block storage > Volumes** and click **Create volume**. A wizard will open.
2. For **Name and size**, enter a volume name and specify its size. Note that volumes can be extended later, but not shrunk.
3. For **Storage policy**, select a redundancy mode, a storage tier, and a failure domain.
4. For **Summary**, review the volume details. Click **Create**.

5.1.3 Attaching volumes to target groups

1. Open **Storage services > Block storage > Target groups**, click the ellipsis icon of the desired target group, and click **Add LUNs**.
2. In the **Attach** window that opens, select volumes to attach to the target group or create them. Click **Apply**.
3. On the **Target groups** tab, click the required target group name and go to the **LUNS** tab. Here, you can see all the available LUNs.

5.1.4 Accessing iSCSI targets from VMware vSphere

You can access iSCSI targets from Linux, Microsoft Hyper-V, and VMware vSphere. Refer to "Accessing iSCSI targets" in the Storage User Guide for more details on access from Linux and Microsoft Hyper-V. The following section describes the VMware vSphere scenario for evaluation.

Before using Virtuozzo Hybrid Infrastructure volumes with VMware vSphere, you need to configure it properly to work with ALUA Active/Passive storage arrays. It is recommended that you change the default path policy to RR policy with the command:

```
# esxcli storage nmp satp set -s VMW_SATP_ALUA -P VMW_PSP_RR
```

Now you can reboot the host and create datastores from Virtuozzo Hybrid Infrastructure volumes exported via iSCSI. Log in to the VMware ESXi web panel and do the following:

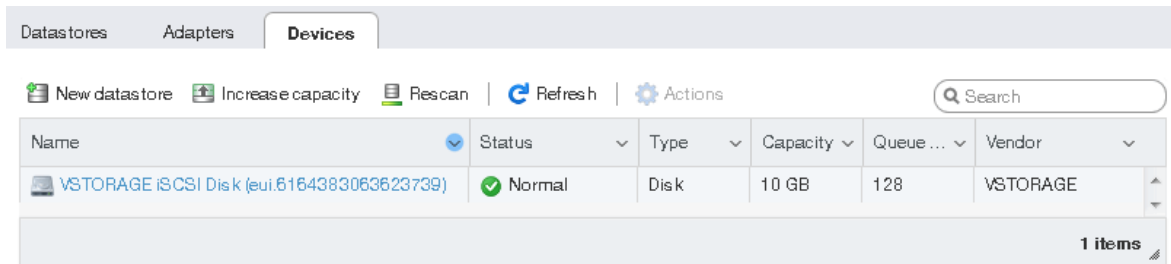
1. In the Navigator, go to the **Storage > Adapters** tab and click **Software iSCSI**.
2. In the **Configure iSCSI** window, select **Enabled**. In the **Dynamic targets** section, click **Add dynamic target** and enter the IP addresses of your nodes.

Note

You can see the IP addresses in the admin panel. On the **Infrastructure > Nodes** screen, click the required node name. Then, go to the **Network** tab and copy the public network IP.

3. Click **Save configuration**.

4. Proceed to the **Devices** tab and click **Refresh**. The newly added disk will appear in the list of devices.



Name	Status	Type	Capacity	Queue...	Vendor
VSTORAGE iSCSI Disk (lun:6164383063623739)	Normal	Disk	10 GB	128	VSTORAGE

5. Select the disk and click **New datastore**. In the wizard that appears, enter a name for the datastore and select partitioning options. Click **Finish** to partition the disk.

Warning!

Partitioning the disk will erase all data from it.

6. The ready-to-use disk will appear in the list of datastores. Select it and click **Database browser** to view its contents and upload files, so that you could proceed to check its accessibility.

5.2 Exporting storage space via S3

Virtuozzo Hybrid Infrastructure allows you to export cluster disk space to customers, in the form of an S3-compatible object storage.

Object storage is optimized for storing billions of objects, in particular for application storage, static web content hosting, online storage services, big data, and backups. Compared to other types of storage, the key difference is that parts of an object cannot be modified; so if the object changes, a new version of it is created instead. This approach eliminates the issue of conflicts.

Virtuozzo Hybrid Infrastructure can store replicas of S3 cluster data and keep them up to date in multiple geographically distributed datacenters. Geo-replication reduces the response time for local S3 users accessing the data in a remote S3 cluster, or remote S3 users accessing the data in a local S3 cluster, as they do not need an Internet connection.

Geo-replication schedules the update of the replicas as soon as any data is modified. Its performance depends on the Internet connection speed, the redundancy mode, and cluster performance.

If you have multiple datacenters with enough free space, it is recommended to set up geo-replication between S3 clusters residing in these datacenters, as described in "Replicating S3 data between datacenters" in the Administrator Guide.

Before creating an S3 cluster, make sure you have a DNS name for the S3 gateway.

5.2.1 Creating an S3 cluster

To create an S3 cluster, do the following:

1. In the left menu, click **Storage services** > **S3**. Select three nodes for our evaluation scenario, and click **Create S3 cluster** in the right menu.

2. Next, select a storage policy.
3. Specify the external (publicly resolvable) DNS name for the S3 endpoint that will be used by the end users to access the object storage. For example, **s3.example.com**. Click **Proceed**.
4. From the drop-down list, select an S3 endpoint protocol: HTTP or HTTPS, or both. For our simple evaluation scenario, we recommend selecting HTTPS and selecting the **Generate self-signed certificate** check box. Click **Done** to create an S3 cluster.

After the S3 cluster is created, open the **S3 Overview** screen to view cluster status, hostname, used disk capacity, the number of users, I/O activity, and the state of S3 services.

To check if the S3 cluster is successfully deployed and can be accessed by users, visit `https://<S3_DNS_name>` in your browser. You should receive the following XML response:

```
<Error>
<Code>AccessDenied</Code>
<Message/>
</Error>
```

To start using the S3 storage, you will also need to create at least one S3 user.

5.2.2 Managing S3 users and buckets

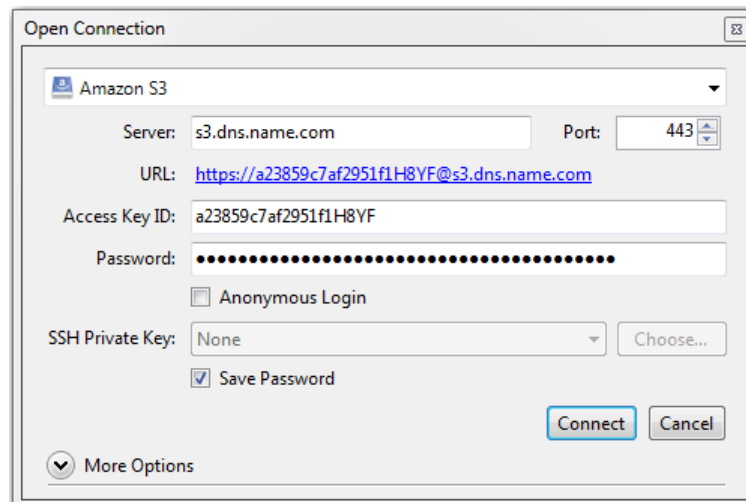
To add an S3 user, do the following:

1. On the **Storage services > S3 > Users** screen, click **Add user**.
2. Specify an email address as the user login, and then click **Add**.

To log in to the S3 portal automatically with user credentials using the generated keys, go to the admin panel, select a user, and then click **Browse**. In this workspace, you can create new buckets and monitor the contents of existing ones.

You can also log in to S3 storage by using a third-party application, like CyberDuck, MountainDuck, or Backup Exec. For our evaluation scenario, connect to your S3 storage via CyberDuck by following these steps:

1. In CyberDuck, click **Open Connection**.
2. Obtain your credentials from the Virtuozzo Hybrid Infrastructure admin panel:
 - Get the DNS name of the S3 endpoint on the **Storage services > S3 > Overview** tab.
 - Get the **Access Key ID** and the **Password**, on the **Storage services > S3 > Users** tab. Select the required user and click **Keys** on the right. This will display the access key ID and the secure access key.
3. Specify your credentials in CyberDuck:



4. Once the connection is established, you can see the existing buckets and create new ones. Click **File > New Folder** to create a bucket. Specify a name for the new bucket, and then click **Create**. Use bucket names that comply with DNS naming conventions.

To manage files in buckets, you have to log in to the S3 portal as a user. For more information, refer to "Accessing S3 buckets" in the Storage User Guide.

5.3 Exporting storage space via NFS

File storage is a storage architecture that uses the Network File System (NFS) protocol to manage data as files. Virtuozzo Hybrid Infrastructure allows you to organize nodes into a highly available NFS cluster in which you can create NFS shares. An NFS share is an access point for a volume and it can be assigned an IP address or a DNS name. The volume, in turn, can be assigned a redundancy scheme, a tier, and a failure domain. In each share, you can create multiple NFS exports that are actual exported directories for user data. Each export has, among other properties, a path that, combined with share's IP address, uniquely identifies the export on the network and allows you to mount it using standard tools.

On the technical side, NFS volumes are based on object storage. Aside from offering high availability and scalability, object storage eliminates the limit on the amount of files and the size of data you can keep in the NFS cluster. Each share is perfect for keeping billions of files of any size. However, such scalability implies I/O overhead that is wasted on file size changes and rewrites. For this reason, an NFS cluster makes a perfect cold and warm file storage, but is not recommended for hot and high performance, and data that is often rewritten (like running virtual machines). Integration of Virtuozzo Hybrid Infrastructure with solutions from VMware, for example, is best done via iSCSI to achieve better performance.

5.3.1 Creating an NFS cluster

1. In the left menu, click **Storage services > NFS**.
2. Select node(s) and click **Create NFS cluster** in the right menu. For the evaluation scenario, we

recommend selecting three nodes.

3. Click **Create**.

After the NFS cluster has been created, you can proceed to create NFS shares.

5.3.2 Creating NFS shares

1. On the **Storage services > NFS > Shares** screen, click **Add NFS share**.
2. On the **Add NFS Share** panel, specify a name (for example, **share1**) and a unique resolvable static IP address from the public network.
3. In **Share size**, specify the size. For users accessing exports, this value will be the filesystem size.
4. Select the desired tier, failure domain, and data redundancy type. Click **Done**.

After the share has been created, you can proceed to create NFS exports.

5.3.3 Creating and mounting NFS exports

1. On the **Storage services > NFS > Shares** screen, click the number in the **Exports** column in the row of the desired share. This will open the share screen.
2. On the share screen, click **Add export**, specify **root** as the export name and **/** as the path, and select the **Read and write** access mode. This will create a directory with a default path that designates the export location inside the share and is used (alongside share's IP address) to mount the export. The root export will be shown in the export list.
3. After creating the root export, you can mount it on Linux or macOS, as described in the Storage User Guide. For our evaluation scenario, mount it on Linux with the following commands:

```
# mkdir /mnt/nfs
# mount -t nfs -o vers=4.0 <share_IP>:<share_name>/ /mnt/nfs
```

where:

- `-o vers=4.0` is the NFS version to use.
- `<share_IP>` is the share IP address. You can also use the share hostname.
- `/<share_name>/` is the root export path, like `share1`.
- `/mnt/nfs` is an existing local directory to mount the export to.

To check the mounted storage, you can run `df -h`.

6 Connecting Acronis Cyber Protect Cloud software to backup storage

Backup storage uses Backup Gateway as a storage access point. It is intended for service providers who use Acronis Cyber Protect and/or Acronis Cyber Protect Cloud and want to store their clients' backed-up data in the local cluster, in the cloud (like Google Cloud, Microsoft Azure, and AWS S3), or on NAS (via the NFS protocol).

Backup storage enables a service provider to easily configure storage for the proprietary deduplication-friendly data format used by Acronis. In addition, the backup storage data can be geo-replicated.

Backup storage supports the following backup destinations:

- Virtuozzo Hybrid Infrastructure storage clusters with erasure coding providing for data redundancy
- NFS shares
- Public clouds, including a number of S3 solutions, as well as Microsoft Azure, OpenStack Swift, and Google Cloud Platform

In this section, we will show how to deploy backup storage in the Virtuozzo Hybrid Infrastructure, then create a new customer in the Acronis Cyber Cloud, and then set up a storage for backups in the Acronis Cyber Protect Cloud. The corresponding steps for Acronis Cyber Protect are similar.

Note

We assume that you have already created a partner account for the Acronis Cyber Cloud and have all the required credentials for it. If not, you can go to the [product page](#) and submit a request.

6.1 Creating backup storage

Before creating backup storage, make sure that the DNS configuration meets the requirements outlined in "Hardware requirements" (p. 5). In addition, port 44445 should be open for inbound/outbound connections for network interface with the **Backup (ABGW) public** role (this is the public network in our evaluation scenario).

1. Configure a new storage for storing and managing your customers' backups by using the Virtuozzo Hybrid Infrastructure admin panel. To do this, log in to the Virtuozzo Hybrid Infrastructure and navigate to **Storage services**, and then to **Backup storage**.
2. Click **Create backup storage**.
3. On the **Backup destination** step, select **Virtuozzo Hybrid Infrastructure cluster**. In our evaluation scenario, the customers' data will be stored and managed on the storage cluster nodes.
4. On the **Nodes** step, select nodes to add to the backup storage cluster, and then click **Next**.
5. On the **Storage policy** step, select the desired tier, failure domain, and data redundancy mode that will be applied to your customers' backups. Then, click **Next**.

6. On the **DNS** step, specify a DNS name that will be associated with the selected cluster and used to register that cluster within Acronis Cyber Protect Cloud (like **backup.example.com**). The new DNS name is associated with each node's IP address in the selected cluster. A specific node for backup operations is selected automatically by the backup agent. It depends on a number of factors, such as node availability and load. Click **Next**.
7. On the **Acronis account** step, specify the URL of your Acronis Cyber Protect Cloud instance. By default, it is <https://cloud.acronis.com>. If you use Acronis Cyber Protect, at this stage you should use the IP address of the corresponding machine to access the Backup Management Console. For the evaluation scenario, provide a login and password for Acronis Cyber Cloud administrator account. Click **Next**.
8. [Optional] Test the DNS setup locally before going public. To do this, you can add the DNS name to the `/etc/hosts` file on machines you will use to access the backup storage. For example, **192.168.1.10 backup.example.com**.
9. On the **Summary** step, review the configuration, and then click **Create**.

The deployment will start immediately. As soon as it is over, you will see three tabs: **Overview**, **Nodes**, **Geo-replication**, and **Settings**. On the **Overview** tab, for example, you can see the information about the registered gateways and their performance.

Note

If the current storage does not have a public IP address and a DNS name, the Web Restore tool for Acronis Cyber Cloud cannot work properly.

6.2 Configuring Acronis Cyber Protect Cloud

To create a new customer and assign the new backup destination in Acronis Cyber Cloud

1. Log in to the Acronis Cyber Cloud management console.
2. Navigate to **Settings > Locations**. Ensure the system created a new backup destination with the corresponding name derived from the DNS name.
3. Set up the backup agents, as described in the [Acronis Cyber Protect User Guide](#).
4. Create a new customer account:
 - a. Click **New** in the upper-right corner and select **Customer**.
 - b. Provide the customer general information: name, mode and language. Then specify customer's email, language, first and last names for an administrator account.
 - c. Select services that you would like to provide to the new customer.
 - d. Specify the customer's devices and workloads, such as servers and workstations.
 - e. In the section **Location**, click the current location name to display all the available options. Select the required storage.
 - f. Click **Done** to complete the whole process.
5. To confirm your account, check your email and follow the steps in the activation request.

To set up backup storage in Acronis Cyber Protect Cloud or Acronis Cyber Protect

1. Log in to Acronis Cyber Cloud as the administrator.
2. Open the **Clients** screen. Click the created customer, and then click **Manage service** on the **Overview** screen. The customer's Cyber Backup Management Console will open.
3. On the **Devices** screen, click **Add** on the toolbar. Select the device you want to add. The backup agent installer will be downloaded.
4. On the **Devices** screen, click **Add** on the toolbar. Select the device you want to add. For our evaluation scenario, select a workstation with the operating system currently in use. The backup agent installer will be downloaded.
5. In the backup agent installer:
 - a. Click **Install**.
 - b. On the **Almost done...** screen, click **Register the machine**.
 - c. Enter the device registration info and confirm it.
 - d. Ensure you are using the customer's account you've created: check the user in the upper-right corner.

When the registration is complete, the added device will be displayed on the **Devices > All devices** screen of the customer's Backup Management Console.

To create a backup from the customer's machine, do the following:

1. Click the device and select **Protect** in the right menu.
2. Click **Add plan** and specify its details. For our evaluation scenario, enable only the **Backup** feature:
 - a. In **What to back up**, select **Files/folders**.
 - b. In **Items to back up**, select the required file or folder.
 - c. In **Where to back up**, select the destination cloud storage.
 - d. In **Schedule**, set **None** by turning it off.
3. Click **Create** and the backup plan will be listed on the left.
4. Click **Run now** to start the backup process.

When the process is completed, you can view the backed-up files on the **Backup storage > Locations** screen. Click the required customer to see the files you uploaded earlier. Double-click the backup name to display its details on the right. You can click **Recover files/folders** to navigate to the uploaded files and download them if necessary.

7 Monitoring the storage cluster

Virtuozzo Hybrid Infrastructure provides built-in monitoring tools, including a preintegrated Prometheus engine and preconfigured Grafana dashboards, which show service state, availability, and performance, as well as network bandwidth, replication backlog, memory, and CPU usage. Integration with third-party systems is possible via 100%-compatible OpenStack APIs. You can ensure that systems are up and running smoothly, and troubleshoot problems before they impact end users or third-party systems.

1. To monitor the storage cluster, go to the **Monitoring > Dashboard** screen. Here you can get general information about the selected storage cluster for the last 30 minutes; 1, 6, and 12 hours; and for the last seven days. The information displayed includes the read and write operations, chunk services' health, and the physical and logical space usage. For more details, refer to "Monitoring the storage cluster" in the Administrator Guide.
2. For advanced monitoring, go to the **Monitoring > Dashboard** screen and click **Grafana dashboard**. A separate browser tab will open with preconfigured Grafana dashboards of the storage cluster, hardware nodes, exports, etc. Two dashboards are dedicated to backup storage. To see a detailed description for each chart, refer to "Monitoring backup storage" in the Administrator Guide.
3. You can also monitor the backup storage on the **Storage services > Backup storage** screen. Here, you can see the information about the deployed backup storage cluster and its performance. Moreover, you can get the information about geo-replication of the selected backup storage clusters, as well as the information about the storage usage in public clouds such as Amazon S3, Microsoft Azure, Google Cloud, or Alibaba Cloud.

8 Enabling high availability


High availability keeps Virtuozzo Hybrid Infrastructure services operational even if the node they are located on fails. In such cases, services from a failed node are relocated to healthy nodes.

You have previously built a cluster of three nodes, and can now make it more resilient and redundant. To do that, enable high availability for the management node, the admin panel, and services.

8.1 High availability for the management node



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



1. On the **Settings > Management node** screen, open the **Management high availability** tab.





2. Select three nodes, and then click **Create HA**. The management node is automatically selected.
3. On **Configure network**, verify that the 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

 node002	
Management	Admin panel
eth1 - 10.37.130.127	br-eth0 - 10.94.16.34

 node003	
Management	Admin panel
eth1 - 10.37.130.148	br-eth0 - 10.94.16.47

PROCEED

4. 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

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

Virtual IP address for
Internal management

❗ 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).

8.2 High availability for the services

Virtuozzo Hybrid Infrastructure provides additional high availability for the following services:

- Admin panel. If the management node fails or becomes unreachable over the network, an admin panel instance on another node takes over the panel's service so that it remains accessible at the same dedicated IP address. The relocation of the service can take several minutes. Admin panel HA is enabled manually along with management node HA.
- Virtual machines. If a compute node fails or becomes unreachable over the network, virtual machines hosted on it are evacuated to other healthy compute nodes based on their free resources. The compute cluster can survive the failure of only one node. By default, high availability for virtual machines is enabled automatically after creating the compute cluster and can be disabled manually, if required.
- iSCSI service. If the active path to volumes exported via iSCSI fails (for example, a storage node with active iSCSI targets fails or becomes unreachable over the network), the active path is rerouted via targets located on healthy nodes. Volumes exported via iSCSI remain accessible as long as there is at least one path to them.
- S3 service. If an S3 node fails or becomes unreachable over the network, the name server and object server components hosted on it are automatically balanced and migrated between other S3 nodes. S3 gateways are not automatically migrated; their high availability is based on DNS records. You need to maintain the DNS records manually when adding or removing S3 gateways. High availability for the S3 service is enabled automatically after enabling management node HA and creating an S3 cluster from three or more nodes. The S3 cluster of three nodes may lose one node and remain operational.
- Backup Gateway service. If a node included in the Backup Gateway cluster fails or becomes unreachable over the network, other nodes in the Backup Gateway cluster continue to provide access to the chosen storage backend. Backup gateways are not automatically migrated; their high availability is based on the DNS records. You need to maintain the DNS records manually when adding or removing backup gateways. High availability for the backup gateway is enabled automatically after creating a Backup Gateway cluster from two or more nodes. Access to the storage backend remains until at least one node in the Backup Gateway cluster is healthy.
- NFS shares. If a storage node fails or becomes unreachable over the network, the NFS volumes located on it are migrated between other NFS nodes. High availability for NFS volumes on a storage node is enabled automatically after creating an NFS cluster.

8.3 Testing high availability

This section simulates an event in which the management node has failed:

1. Forcibly power off your Virtuozzo Hybrid Infrastructure management node.

Note

High availability (HA) keeps services operational if the node they are located on fails due to kernel crash, power outage, or becomes unreachable over the network. Graceful shutdown is not considered a failure event. To test HA, you should forcibly power off the node or disconnect the network cable from it.

2. Open the **Infrastructure > Nodes** screen. The failed node has the **Unhealthy** status and is highlighted in red.
3. Even though one node has failed and is now unavailable, you can still access the following services:
 - Admin panel.
 - Virtual machines.
 - iSCSI: in the VMware vSphere, you can still access the volumes exported via iSCSI.
 - S3: you can still access your buckets via CyberDuck.
 - NFS: in the mounted root export, you can still access the data you uploaded.
 - Backup Gateway: in the Backup Management Console, you can still navigate to the backup you created earlier (it is accessible if you have properly configured the DNS name).

You have just run a demonstration where a node was forcibly powered off, causing the services and the data to be evacuated to healthy nodes and remain available without downtime.