



Virtuozzo Storage

Administrator's Guide

September 29, 2017

Parallels International GmbH

Vordergasse 59

8200 Schaffhausen

Switzerland

Tel: + 41 52 632 0411

Fax: + 41 52 672 2010

<https://virtuozzo.com>

Copyright ©2016-2017 Parallels 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 Virtuoizzo Storage	1
2. Managing Virtuoizzo Storage	2
2.1 Configuring Node Network Interfaces	2
2.2 Creating Virtuoizzo Storage Clusters	6
2.2.1 Creating the Cluster on the First Node	7
2.2.2 Adding Nodes to Cluster	9
2.2.3 Assigning Disk Roles Manually	10
2.3 Releasing Nodes from Cluster	13
2.4 Removing Nodes from the Unassigned List	14
2.5 Managing Tier Encryption	15
2.6 Managing Virtuoizzo Storage Users	16
2.6.1 Creating User Accounts	16
2.6.2 Managing User Accounts	18
2.6.3 Adding LDAP or Active Directory Users	20
2.7 Managing Virtuoizzo Storage Updates	23
2.8 Allowing root Access to Cluster Nodes Over SSH	24
2.9 Backing Up and Restoring Management Database	25
2.9.1 Restoring Management Database from Backup	26
2.10 Managing Virtuoizzo Storage Licenses	27
3. Monitoring Virtuoizzo Storage Clusters	28
3.1 Monitoring Cluster Status	28
3.2 Monitoring Cluster Storage Space	28
3.2.1 Physical Space Chart	29

3.2.2	Logical Space Chart	29
3.2.2.1	Understanding Logical Space	30
3.2.3	Monitoring Chunk Status and Replication	30
3.2.4	Monitoring Cluster Services	31
3.2.5	Monitoring Cluster I/O Activity	32
3.3	Monitoring Virtuozzo Cluster Objects via SNMP	33
3.3.1	Enabling SNMP Access	34
3.3.2	Accessing Virtuozzo Objects via SNMP	35
3.3.2.1	Listening to SNMP Traps	35
3.3.3	Monitoring Clusters with Zabbix	36
3.3.4	Description of Cluster Objects and Traps	40
4.	Monitoring Virtuozzo Storage Nodes	43
4.1	Node Statuses	43
4.2	Monitoring Node Performance	43
4.2.1	Monitoring Node Disks	45
4.2.1.1	Monitoring the S.M.A.R.T. Status of Node Disks	46
4.3	Monitoring Node Network	46
5.	Viewing Virtuozzo Storage Alerts and Audit Log	47
5.1	Viewing Alerts	47
5.2	Viewing Audit Log	48
6.	Exporting Virtuozzo Storage Cluster Data	50
6.1	Exporting Data via iSCSI	50
6.1.1	Creating Virtuozzo Storage iSCSI Targets	51
6.1.1.1	Performance Tips	53
6.1.2	Listing, Stopping, and Deleting Virtuozzo Storage iSCSI Targets	53
6.1.3	Configuring Virtuozzo Storage iSCSI Targets	53
6.1.3.1	Listing LUNs	54
6.1.3.2	Adding LUNs	55
6.1.3.3	Configuring LUNs	57
6.1.3.4	Deleting LUNs	57
6.1.4	Managing iSCSI Users	57
6.1.4.1	Creating CHAP Accounts for Virtuozzo Storage iSCSI Targets	57
6.1.4.2	Creating Virtuozzo Storage iSCSI Targets Bound to CHAP Accounts	58
6.1.4.3	Changing CHAP Account Passwords	59

6.2	Exporting Data via S3	60
6.2.1	Object Storage Infrastructure Overview	61
6.2.2	Planning the S3 Cluster	62
6.2.3	Sample Object Storage	63
6.2.4	Creating the S3 Cluster	65
6.2.5	Managing Object Storage Users	70
6.2.5.1	Adding S3 users	70
6.2.5.2	Managing S3 Access Key Pairs	71
6.2.6	Managing Object Storage Buckets	73
6.2.6.1	Listing Bucket Contents	73
6.2.6.2	Managing Acronis Notary in Buckets	74
6.2.7	Best Practices for Using S3 in Virtuozzo Storage	75
6.2.7.1	Bucket and Key Naming Policies	75
6.2.7.2	Improving Performance of PUT Operations	76
6.2.8	Replicating Data Between Geographically Distributed Datacenters with S3 Clusters	76
6.2.9	Monitoring S3 Access Points	78
6.2.10	Releasing Nodes from S3 Clusters	78
6.2.11	Supported Amazon S3 Features	79
6.2.11.1	Supported Amazon S3 REST Operations	79
6.2.11.2	Supported Amazon Request Headers	81
6.2.11.3	Supported Amazon Response Headers	81
6.2.11.4	Supported Amazon Error Response Headers	82
6.2.11.5	Supported Authentication Scheme and Methods	83
6.3	Using Virtuozzo Storage with Acronis Backup Cloud	83
6.3.1	Understanding the Infrastructure	84
6.3.2	Configuring Acronis Backup Gateway Clusters	84
6.3.3	Releasing Nodes Running Acronis Backup Gateway Services	88
6.3.4	Connecting External NFS Shares through Acronis Backup Gateway	89
6.4	Managing Datastores	92
6.4.1	Creating Datastores	93
6.4.2	Placing Virtual Environments and Backups in Datastores	94

CHAPTER 1

Introduction

To support the growing demand for both high performance and high data availability, modern data centers need a fast, flexible storage solution. Existing solutions, however, are often difficult to manage and maintain, or not flexible enough (e.g., local RAID arrays), or too expensive (e.g., storage area networks).

Virtuozzo Storage is designed to solve these issues. It can run on commodity hardware, so no significant infrastructure investments are needed. It is also is easy to set up and grow on demand.

1.1 About Virtuozzo Storage

Virtuozzo Storage is a software-defined storage solution that allows you to quickly and easily transform low-cost commodity hardware and network equipment into protected enterprise-grade storage like storage area networks (SAN) or network-attached storage (NAS).

Virtuozzo Storage is optimized for storing large amounts of data and provides data redundancy (replication and erasure coding), high availability, self-healing, and storage sharing.

In Virtuozzo Storage, user data is stored on organized clusters of servers in the form of fixed-size chunks. These chunks are automatically replicated and distributed across available servers in the cluster to ensure high availability of user data.

Cluster storage space can be exported through access points like iSCSI, S3, or Acronis Backup Gateway.

CHAPTER 2

Managing Virtuozzo Storage

To start managing Virtuozzo Storage, log in to the management panel at its IP address on port 8888 as admin (or superadmin) and make sure that storage nodes are shown on the **NODES** screen.

The first step to perform, before you can create any clusters, is to create the internal and public networks required by Virtuozzo Storage. You can do that by configuring the network interfaces of all nodes. Having created the networks, you can proceed to creating Virtuozzo Storage clusters.

2.1 Configuring Node Network Interfaces

As described in **Planning Network** in the *Installation Guide*, Virtuozzo Storage requires one internal network for node traffic and one public network for exporting the storage space. You need to create these networks by assigning correct network roles to network interfaces on each node.

Important: To be able to create a cluster, you will need to assign a storage role to a node's network interface.

To assign a network role to a network interface, do the following:

1. On the **NODES** screen, click the node to configure the network interface(s) of.

2.1. Configuring Node Network Interfaces

Nodes ADD NODE 


Search Any status ▾ Any role ▾

UNASSIGNED 4 Hide unassigned nodes ^

172.16.24.126 	172.16.25.102 	172.16.25.85 	172.16.25.97 
--	--	---	---

2. On the node overview screen, click **NETWORK**.

Nodes > 172.16.24.126 HEALTHY



3. Select a network interface and click **Choose role**.

Nodes > 172.16.24.126 > Network

ADD NODE



DISKS NETWORK

☐	Name ↓	Status	IPs		Details
☐	ens37	UNPLUGGED			Performance
☒	br0	OK	172.16.24.126...	Management,...	Choose role

- On the **Choose roles** panel, select roles to assign to the network interface (for details, see **Network Interface Roles** in the *Installation Guide*).

2.1. Configuring Node Network Interfaces

× Choose roles

☒ Internal

- ☒ Management
- ☒ Storage
- ☒ S3 private
- ☒ ABGW private

☒ Public

- ☒ iSCSI
- ☒ S3 public
- ☒ ABGW public
- ☒ Web CP
- ☒ SSH

☐ Custom (configure)

Done

5. If you need to open specific ports on a network interface with public roles, do the following:

5.1. Click **Configure**.

✕ **Configure custom role**

Here you can manage allowed ports for custom role.

	Name	Port	
☒	FTP data	20	
☐	FTP cmd	21	

 Add  Remove

Done

5.2. On the **Configure custom role** panel, create custom roles: click **Add** and specify role names and ports. Custom roles can later be assigned to any network interface in a cluster.

To remove a custom role, make sure it is not assigned to any interface, select it, and click **Remove**.

5.3. Click **Done** to return to the **Choose roles** panel.

6. Select the required roles and click **Done** to assign them.

2.2 Creating Virtuozzo Storage Clusters


To create a Virtuozzo Storage cluster means to create a cluster on one (first) node, then populate it with more nodes.

Important: To be able to create a cluster, you will need to assign a storage role to a node's network interface.

2.2. Creating Virtuozzo Storage Clusters

2.2.1 Creating the Cluster on the First Node

1. Open the **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 dashes ("-").


The screenshot shows a modal dialog titled 'New cluster' with a close button (X). It contains two input fields: 'Cluster' with the text 'cluster1' and 'Storage interface' with a dropdown menu showing 'eth0 - 10.100.6.105'. Below these fields is a checkbox labeled 'Encryption' which is checked, followed by a gear icon. At the bottom, there are two buttons: 'New cluster' and 'Advanced configuration'.

4. Make sure a configured network interface with a storage role is selected from the **Storage interface** drop-down list.

Note: If the network was not previously configured, click the cogwheel icon and, on the **Network Configuration** screen, configure a storage role for a network interface.

5. If required, enable data encryption. To do this, check the **Encryption** box and proceed to create the cluster. Encryption will be enabled for all tiers by default.

To enable encryption for particular tiers, click the cogwheel icon to open the **Encryption Configuration** panel, select tiers to encrypt, and click **Done**.



Note: You can later disable encryption for new chunk services (CS) on the **SETTINGS > Advanced settings** panel.

6. Click **New cluster** to have Virtuoizzo Storage assign the roles to disks automatically. Alternatively, click **Advanced configuration** to assign the roles to each drive manually (see [Assigning Disk Roles Manually](#) on page 10).


You can monitor cluster creation progress in the **HEALTHY** list of the **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.

2.2. Creating Virtuozzo Storage Clusters

2.2.2 Adding Nodes to Cluster


To add an unassigned node to a cluster, do the following:

1. On the **NODES** screen, click an unassigned node.



2. On the node overview screen, click **Join cluster**.
3. Make sure a configured network interface with a storage role is selected from the **Storage interface** drop-down list.

Note: If the network was not previously configured, click the cogwheel icon and, on the **Network Configuration** screen, configure a storage role for a network interface.



5. Click **Join cluster** to have Virtuozzo Storage assign the roles to disks automatically and add the node to the current cluster. Alternatively, click **Advanced configuration** to assign the roles to each drive manually (see [Assigning Disk Roles Manually](#) on page 10).


2.2.3 Assigning Disk Roles Manually

If you clicked **Advanced configuration** while creating a cluster or adding nodes to it, you will be taken to the list of drives on the node where you can manually assign roles to these drives.

Note: You can also access the list of drives on the **NODES > <node> > Disks** screen.

Do the following:

1. On the **Join cluster** or **New cluster** panel, select a drive or check multiple drives in the list and click **Configure**.
2. On the **Choose role** screen, select one of the following roles for the disk:



Choose role

☒ Storage

☐ Metadata

☐ Cache

☐ Metadata + Cache

Caching and checksumming

Enable checksumming

Tier

Tier 0 (Encrypted)

Done Cancel

- **Storage.** Use the disk to store chunks and run a chunk service on the node. From the **Caching and checksumming** drop-down list, select one of the following:

- Use SSD for caching and checksumming (available and recommended only for SSD).

Data caching improves cluster performance by placing the frequently accessed data on an SSD.

Data checksumming generates checksums each time some data in the cluster is modified.

When this data is then read, a new checksum is computed and compared with the old checksum. If the two are not identical, a read operation is performed again, thus providing better data reliability and integrity.

- Enable checksumming (better reliability). Checksumming is available for both SSD and HDD.
- Disable checksumming (better performance).

With this role, also select a tier from the **Tier** drop-down list (see **Understanding Storage Tiers** in the *Installation Guide*).

Note: If the disk contains old data that was not placed there by Virtuozzo Storage, the disk will not be considered suitable for use in Virtuozzo Storage.

- **Metadata.** Use the disk to store metadata and run a metadata service on the node.
- **Cache.** Use the disk to store write cache. This role is only for SSDs. To cache a specific storage tier, select it from the drop-down list. Otherwise, all tiers will be cached.

- **Metadata+Cache.** A combination of two roles described above.
- **Unassigned.** Remove the roles from the disk.

Note: Both the System and Metadata roles can be assigned to the same disk with capacity of 100GB or higher—preferably to an SSD. Assigning both these roles to an HDD will result in mediocre performance suitable only for cold data (e.g., archiving).

3. Click **Done**.
4. Repeat steps 1 to 3 for every disk you want to be used in the VirtuoZZo Storage cluster.
5. Click **NEW CLUSTER** or **JOIN CLUSTER**. On the **Configuration summary** screen, check the number of disks per each configuration category.

✕ Configuration summary

METADATA SERVICE				
	Metadata			1
STORAGE SERVICE				
		SSD cache	Same disk cache	No cache
	Tier 0	1	0	0
	Tier 1	3	0	0
	Tier 2	0	1	0
SSD CACHE				
	Tier 0			1
	Tier 1			1
PROCEED		CANCEL		

6. Click **PROCEED**. You can monitor disk configuration progress in the **HEALTHY** list of the **NODES** screen.

2.3 Releasing Nodes from Cluster

To release a node means to remove it from the cluster (e.g., for maintenance). As the node may be running services needed by the cluster, do the following prior to releasing it to avoid cluster degradation:

1. If the node runs one of the five required metadata services, add a metadata role to another node. You need to make sure that the cluster has at least five metadata services running at any time.

2. If the node has any access points, make sure that the same access points are configured on other nodes in the cluster as well.
3. If the node has iSCSI targets, move them to a different node.
4. If the node has an S3 gateway or ABGW, reconfigure DNS for S3 and ABGW access points to remove the node from DNS records. Next, release the node from S3 and ABGW in the corresponded sections of the **SERVICES** screen.
5. Make sure the cluster has enough storage space to accommodate the data from the released node.

Once you initiate the release, the cluster will start replicating data chunks that were stored on the released node and distributing them among other storage nodes in the cluster. Depending on the amount of data to replicate, the process may take as much as several hours.

If necessary, you can also release a node forcibly, that is, without replication.

Warning: Releasing nodes forcibly may result in data loss.

To release a node from a cluster, do the following:

1. On the **NODES** screen, click the node to release.
2. On the node overview screen, click **Release**.
3. If necessary, in the **Release** node window, check force to release the node forcibly (highly not recommended).
4. Click **Yes**. The released node will return to the **UNASSIGNED** list on the **NODES** screen.

2.4 Removing Nodes from the Unassigned List

Nodes in the **UNASSIGNED** list can be completely removed from Virtuozzo Storage.

Do the following: on the **NODES** screen, select the node in the **UNASSIGNED** list and click **Remove (forget)**.

Nodes completely removed from Virtuozzo Storage can be re-added to the **UNASSIGNED** list in two ways:

- By logging in to the node via SSH and running
`/usr/libexec/vstorage-ui-agent/bin/register-storage-node.sh -m <MP_addr> -t <token>` in the node's console, where <MP_addr> is the IP address of the management panel container and <token> is the token

2.5. Managing Tier Encryption

obtained in the management panel.

- By reinstalling Virtuozzo Storage on the node from scratch.

2.5 Managing Tier Encryption

Virtuozzo Storage can encrypt data stored on disks with the AES-256 standard, so if a disk gets lost or stolen the data will be safe. Virtuozzo Storage stores disk encryption keys in cluster's metadata (MDS).

Encryption can be enabled or disabled only for the newly created chunk services (CS). Once tier encryption is enabled, you can decrypt disks (CSs) by manually releasing them from encrypted tiers. Correspondingly, simply enabling encryption on the disk's tier will not encrypt its data (CS). To encrypt a disk, you must assign it to an encrypted tier.

Note:

1. Virtuozzo Storage does not encrypt data transmitted over the internal network.
2. Enabled encryption slightly decreases performance.

Advanced

LDAP/AD

ENCRIPTION

☒ Enable encryption (aes256)

☒ Tier 0☒ Tier 1☒ Tier 2☒ Tier 3

SAVE

To enable or disable tier encryption, on the **SETTINGS > Advanced settings** panel, select or deselect tiers

and click **SAVE**.

2.6 Managing Virtuoizzo Storage Users

During the management panel installation on the first node, Virtuoizzo Storage creates the default unique administrator, superadmin. The user name for this account is **admin** and the password is specified during installation. This account cannot be deleted and its privileges cannot be changed. Other than that, superadmin does not differ from a user account assigned the **Administrator** role (i.e. an admin).

An admin can create user accounts and assign to them one or more roles listed below:

- **Administrator**. Can fully manage cluster and users.
- **Network**. Can modify network settings and roles.
- **Cluster**. Can create cluster, join nodes to cluster, and manage (assign and release) disks.
- **iSCSI**. Can create and manage iSCSI targets and LUNs.
- **S3**. Can create and manage S3 cluster.
- **Acronis Backup Gateway**. Can create and manage Acronis Backup Gateway.
- **SSH**. Can add and remove SSH keys for cluster nodes access.
- **Updates**. Can install updates.

User accounts to which no roles are assigned are guest accounts. Guests can monitor Virtuoizzo Storage performance and parameters but cannot change any settings.

Note: All users can change their own passwords (see [Managing User Accounts](#) on page 18).

2.6.1 Creating User Accounts

To create a user account in the web-based user interface, do the following:

1. Log in to the management panel as admin.
2. Open the **SETTINGS > Users** screen and click **ADD USER**.

2.6. Managing Virtuozzo Storage Users

3. On the **Add user** panel, specify the user name, password, and, if required, a user description in the corresponding fields.

× **Add user**

Name

Enabled

☒

Password

Description

SELECT ROLES

☐ **S3**
Can create and manage S3 cluster.

☐ **Administrator**
Can perform all management operations.

☐ **ABGW**
Can create and manage Acronis Backup Gateway.

☒ **SSH**
Can add and remove SSH keys for cluster nodes access.

☐ **iSCSI**
Can create and manage iSCSI targets and LUNs.

☒ **Cluster**
Can create cluster, join nodes to cluster, and manage (assign and release) disks.

☒ **Updates**
Can install updates.

☒ **Network**
Can modify network settings and roles.

ADD

4. Check the roles to assign to the account and click **Done**.

2.6.2 Managing User Accounts

Any user can change their account password by clicking the user icon in the top right corner of the management panel and then clicking **Change password**.

An admin can create/delete other users' accounts, add/remove roles from them, change their descriptions and passwords (although superadmin's password can only be changed by superadmin), as well as enable/disable user accounts (i.e. allow/prohibit user login). To manage a user account, login as an admin, open the **Settings** -> **Users** screen, select a user from the list, and click **Configure** or **Delete** depending on what you need to do.

× **Configure user**

Name

user1

Enabled

Password

change

Description

ROLES

☐ **Administrator**

Can perform all management operations.

☐ **Network**

Can modify network settings and roles.

☐ **Updates**

Can install updates.

☐ **iSCSI**

Can create and manage iSCSI targets and LUNs.

☒ **SSH**

Can add and remove SSH keys for cluster nodes access.

☐ **ABGW**

Can create and manage Acronis Backup Gateway.

☐ **S3**

Can create and manage S3 cluster.

☒ **Cluster**

Can create cluster, join nodes to cluster, and manage (assign and release) disks.

DONE

2.6.3 Adding LDAP or Active Directory Users

You can add users and user groups to Virtuoizzo Storage from an external LDAP-compliant database or Microsoft Active Directory. These users will be able to log in using their respective user names and passwords. The set of actions these users will be able to perform in Virtuoizzo Storage will be defined by the roles you assign in Storage (listed in *Managing Virtuoizzo Storage Users* on page 16).

To add an LDAP (or AD) user or group to Virtuoizzo Storage, do the following:

1. On the **SETTINGS > Advanced settings** screen, open the **LDAP/AD** tab.

Advanced

LDAP/AD ENCRYPTION SNMP MANAGEMENT NODE HA

Directory type

Microsoft Active Directory

▼

Address

10.250.35.97

Port

389

Login

cn=administrator, cn=Users,dc=ya,dc=skol,dc=

Password

.....

Search Base DN

dc=ya,dc=skol,dc=ko

☐ Advanced

SAVE

2. Select LDAP or Microsoft Active Directory from the **Type** drop-down list.

3. Specify the following parameters:

- IP **Address** of an LDAP server or AD domain controller;

2.6. Managing Virtuozzo Storage Users

- (optional) LDAP **Port**;
 - **Bind DN** (a distinguished name of an LDAP authentication database user) or **Login** (AD);
 - **Bind Password** (LDAP) or **Password** (AD);
 - **Search Base DN**, a distinguished name of a search starting point;
 - (optional) **Advanced** LDAP or AD parameters.
4. Click **Save** to authenticate in Active Directory or LDAP server.
 5. On the **SETTINGS > Users** screen, click **ADD LDAP USER**.
 6. On the **Add LDAP users** panel, select users or user groups to add to Virtuozzo Storage and click **Add**.

× **Add LDAP users**

☐	Type	Name ↓	Description	
☐		WinRMRemoteWMIUs...	Members of this gr...	
☒		Administrators	Administrators hav...	
☐		Users	Users are prevente...	
☐		Guests	Guests have the sa...	
☐		Print Operators	Members can admi...	
☐		Backup Operators	Backup Operators ...	

ADD

7. On the **Roles** panel, select the roles to assign to selected users or user groups.

Note: If a role is assigned to a group, every user in it is granted the corresponding privileges.

× **Roles**

☐ **S3**
Can create and manage S3 cluster.

☒ **Administrator**
Can perform all management operations.

☐ **ABGW**
Can create and manage Acronis Backup Gateway.

☐ **SSH**
Can add and remove SSH keys for cluster nodes access.

☐ **iSCSI**
Can create and manage iSCSI targets and LUNs.

☐ **Cluster**
Can create cluster, join nodes to cluster, and manage (assign and release) disks.

☐ **Updates**
Can install updates.

ADD

8. Click **Add** to add users to Virtuoizzo Storage.

2.7. Managing Virtuozzo Storage Updates

Users

ADD USERADD LDAP USER

Type	Name ↓	Role	
	user1	Administrator	
	admin	Superadmin	
	Administrators	Administrator	
	ACuser1@vstorage.local	Updates, Cluster, Guest, Network	

 Configure

 Delete

2.7 Managing Virtuozzo Storage Updates

Note: For instructions on installing major Virtuozzo updates, refer to the [Virtuozzo 7 User's Guide](#).

You can update your Virtuozzo Storage infrastructure using the web-based user interface.

Important: To check for and download updates, the cluster must be healthy and each node in the infrastructure must be able to open outgoing Internet connections.

To update Virtuozzo Storage, do the following:

1. Open the **SETTINGS > Updates** screen and click **CHECK FOR UPDATES**. The script will run `yum update` on each node. If updates are available for a node, said node's status will change to `Update available`.

Updates



Update available

UPDATE NOW

Node	Status	
172.16.24.126	Update available	

2. To apply all available updates, click **UPDATE NOW**.

While updates are being applied, some of the Virtuozzo Storage services might be unavailable for a short period of time.

2.8 Allowing root Access to Cluster Nodes Over SSH

In certain situations, you or the technical support team may need root access to cluster nodes via SSH. To allow root access to all nodes in the cluster, do the following:

1. Obtain an SSH public key from the technical support team.
2. Open the **SETTINGS > SSH** screen, click **ADD KEY**, paste the key, and click **Add key**.

× **Add Public Key**

Key

```
ssh-rsa
AAAAB3NzaC1yc2EARRABIwAAAEAKIOUpkDHrfHY17SbrmTIpNLTG
K9Tjom/BWDSU
GPI+nafzIHDTYW7hdI4yZ5ew18JH4JW9jbhUFrviQzM7xIELEVf4h9IFX5
QVkbPppSwg0cda3
Pbv7kOdJ/MTyRRWXFCR+HAo3FXRitBqxiX1nYUXpHAZsMciLq8V6Rjs
NAQwdsdMFvSIVK/7XA
t3FaoJoAsncM1Q9x5+3V0Ww68/eIFmb1zuUFJjQJKprX88XypNDvjYN
by6vw/Pb0rwert/En
mZ+AW4OZPnTPI89ZPmVMLuayrD2cE86Z/il8b+gw3r3+1nKatmlkjn2
so1d01QraTIMqVSsbx
NrRFi9wrf+M7Q== user@localdomain.local
```


Add key

To delete the key after root access is no longer required, select the key and click **Delete**.

2.9 Backing Up and Restoring Management Database

In Virtuozzo Storage, node information, statistics, and configurations are stored in a database inside the container with the management panel. Database backups are created automatically every day.

To back up the database manually, open the **SETTINGS > Backup** screen and click **BACKUP NOW**.



Once backup is completed, the **Last backup** date will be refreshed.

2.9.1 Restoring Management Database from Backup

You can restore a management database from backup to a new container on any node your cluster is mounted to. In the new container, a new management panel will be deployed, and its database will be restored from the specified backup file.

To restore the database, run the following script as shown below:

```
# restore-storage-ui-ct -p|--rootpw <CT_root_passwd> -i|--ip <CT_IP_addr> \
[-t|--tarball <backup_tarball>] [-n|--nameserver <NS1>[,<NS2>]]
```


where <CT_root_passwd> and <CT_IP_addr> are the root password and new IP address of the new management panel container, <backup_tarball> is the database backup file (e.g., /mnt/vstorage/webcp/backup/backup-20170529122721.tar), and <NS1>, <NS2> are the nameserver IPs.

If you omit the `--tarball`, a container with a clean management panel will be created (without restoring a database backup). If you omit `--nameserver`, you will not be able to update the management panel in the future.

2.10 Managing Virtuozzo Storage Licenses

To activate or upgrade a Virtuozzo Storage license, do the following:

1. On the **LICENSES** screen, click **Register key**.
2. On the **Register key** panel, enter a license key and click **Done**.
3. Back on the **Licenses** screen, the expiration date and storage space limit will change according to the licenses entered.



Click **Activate**.

CHAPTER 3

Monitoring Virtuozzo Storage Clusters

From the management panel, you can monitor the performance of both the whole cluster and its parts.

3.1 Monitoring Cluster Status


The overall cluster statistics are available on the cluster **OVERVIEW** screen. Pay attention to the cluster status that can be one of the following:

- **HEALTHY.** All cluster components are active and operate normally.
- **UNKNOWN.** Not enough information about the cluster state (e.g., because the cluster is inaccessible).
- **DEGRADED.** Some of the cluster components are inactive or inaccessible. The cluster is trying to heal itself, data replication is scheduled or in progress.
- **FAILURE.** The cluster has too many inactive services, automatic replication is disabled. If the cluster enters this state, troubleshoot the nodes or contact the support team.

3.2 Monitoring Cluster Storage Space

You can monitor cluster storage space on the cluster **OVERVIEW** screen. Typical statistics may look like this:

3.2. Monitoring Cluster Storage Space



The two charts that provide information on how storage space is used are **PHYSICAL SPACE** and **LOGICAL SPACE**. They are described in the following sections in more detail.

3.2.1 Physical Space Chart

The **PHYSICAL SPACE** chart shows the combined space of all disks available to the cluster. The following statistics are available:

- **Used space.** The space occupied by all data chunks and their replicas plus the space occupied by any other data stored on cluster nodes' disks.
- **Free space.** The unused space on all cluster nodes' disks.
- **Total space.** The total space on all cluster nodes' disks.

3.2.2 Logical Space Chart

The **LOGICAL SPACE** chart represents all the space that can be allocated and used by the cluster for storing user data. This space includes the following:

- **Total space.** The maximum disk space available as defined by license.
- **Used space.** The space occupied exclusively by user data. Replicas and erasure coding metadata are

not taken into account.

- **Free space.** The difference between the two above.

3.2.2.1 Understanding Logical Space

When monitoring disk space information in the cluster, keep in mind that logical space is the amount of free disk space that can be used for storing user data in the form of data chunks and all their replicas. Once this space runs out, no data can be written to the cluster.

To better understand how logical disk space is calculated, consider the following example:

- The cluster has three disks with the storage role. The first disk has 200 GB of space, the second one has 500 GB, and the third one has 1 TB.
- If the redundancy mode is set to three replicas, each data chunk must be stored as three replicas on three different disks with the storage role.

In this example, the available logical disk space will be 200 GB, that is, equal to the capacity of the smallest disk with the storage role. The reason is that each replica must be stored on a different disk. So once the space on the smallest disk (i.e. 200 GB) runs out, no new chunk replicas can be created unless a new disk with the storage role is added or the redundancy mode is changed to two replicas.

With the two replicas redundancy mode, the available logical disk space would be 700 GB, because the two smallest disks combined can hold 700 GB of data.

3.2.3 Monitoring Chunk Status and Replication

You can monitor the state of all chunks in the cluster in the **CHUNKS** section of the cluster **OVERVIEW** screen.

The table below lists all possible states a chunk can have.

State	Description
healthy	Percentage of chunks that have enough active replicas. The normal state of chunks.
offline	Percentage of chunks all replicas of which are offline. Such chunks are completely inaccessible for the cluster and cannot be replicated, read from or written to. All requests to an offline chunk are frozen until a CS that stores that chunk's replica goes online. Get offline cluster nodes back online as soon as possible to avoid data loss.

3.2. Monitoring Cluster Storage Space


State	Description
blocked	Percentage of chunks which have fewer active replicas than the set minimal amount. Write requests to a blocked chunk are frozen until it has at least the set minimum amount of replicas. Read requests to blocked chunks are allowed, however, as they still have some active replicas left. Blocked chunks have higher replication priority than degraded chunks. Having blocked chunks in the cluster increases the risk of losing data, so postpone any maintenance on working cluster nodes and get offline chunk servers back online as fast as possible.
degraded	Percentage of chunks with the number of active replicas lower than normal but equal to or higher than the set minimum. Such chunks can be read from and written to.

3.2.4 Monitoring Cluster Services

You can monitor two types of services in the **SERVICES** section on the cluster **OVERVIEW** screen:

- MDS, metadata services. Ensure that five are running at all times.
- CS, chunk services. With this chart, you can also keep track of all disks with the storage role.

Typical statistics may look like this:



If some of the services were not in the healthy state for some time, these time periods will be highlighted in red on the charts.

3.2.5 Monitoring Cluster I/O Activity

You can monitor the history of the cluster I/O activity on the **READ** and **WRITE** charts on the cluster **OVERVIEW** screen. Typical statistics may look like this:

3.3. Monitoring Virtuozzo Cluster Objects via SNMP



The current cluster I/O activity averaged for the last 10 seconds is shown as:

- the speed of read and write I/O operations, in megabytes per second (MB/s).
- the number of read and write I/O operations per second (IOPS).

3.3 Monitoring Virtuozzo Cluster Objects via SNMP

You can monitor cluster objects via the Simple Network Management Protocol (SNMP). The implementation conforms to the same Structure of Management Information (SMI) rules as the data in the standard SNMP context: all objects are organized in a tree; each object identifier (OID) is a series of integers corresponding to tree nodes and separated by dots.

General information:

- The OID of the root subtree with all the objects you can monitor is 1.3.6.1.4.1.8072.161.1.
- The VSTORAGE-MIB.txt information base file is required to monitor the objects. You can download the file at http://<management_panel_IP>:8888/api/v2/snmp/mibs/VSTORAGE-MIB.txt.

The following subsections describe ways to enable and use SNMP to monitor cluster objects.

3.3.1 Enabling SNMP Access

To monitor cluster objects, enable the SNMP access on the node. Do the following in the management panel:

1. On the **SETTINGS > Advanced settings > SNMP** tab, check **Enable SNMP on management node**.
Doing so lets your network management system (SNMP monitor) access the cluster via the SNMP protocol on the management node's port 161.

Advanced

LDAP/AD

ENCRYPTION

SNMP

MANAGEMENT NODE HA

☒ Enable SNMP on management node

You can access SNMP at management node at port 161

MIB file can be [downloaded here](#)

Zabbix template can be [downloaded here](#)

☒ Send SNMP traps to the following Network Management System

IP 10.250.14.17

Port 162

Community public

SEND TEST TRAP

SAVE

2. Click the corresponding link to download the MIB file and set it up in your SNMP monitor.
3. If required, have Virtuoizzo Storage send SNMP traps to your SNMP monitor. Do the following:

3.3. Monitoring Virtuoizzo Cluster Objects via SNMP

- 3.1. Check **Send SNMP traps to Network Management System**.
- 3.2. Specify the **IP** of the system, and, if required, change the default **Port** and **Community**.
- 3.3. If required, click **SEND TEST TRAP** to test the service.
4. Click **SAVE** to apply changes.

3.3.2 Accessing Virtuoizzo Objects via SNMP

You can access Virtuoizzo objects with SNMP tools of your choice, e.g., the free Net-SNMP suite for Linux.

To display cluster information on the node with management panel, place the MIB file to `/usr/share/snmp/mibs` and run the `snmpwalk` command. For example:

```
# snmpwalk -M /usr/share/snmp/mibs -m VSTORAGE-MIB -v 2c -c public \
localhost:161 VSTORAGE-MIB:cluster
```

Typical output may be the following:

```
VSTORAGE-MIB::clusterName.0 = STRING: "cluster1"
VSTORAGE-MIB::healthStatus.0 = STRING: "healthy"
VSTORAGE-MIB::usedSpace.0 = Counter64: 173732322
VSTORAGE-MIB::totalSpace.0 = Counter64: 1337665179648
VSTORAGE-MIB::freeSpace.0 = Counter64: 1318963253248
VSTORAGE-MIB::licenseStatus.0 = STRING: "unknown"
VSTORAGE-MIB::licenseCapacity.0 = Counter64: 1099511627776
VSTORAGE-MIB::licenseExpirationStatus.0 = STRING: "None"
VSTORAGE-MIB::ioReadOpS.0 = Counter64: 0
VSTORAGE-MIB::ioWriteOpS.0 = Counter64: 0
VSTORAGE-MIB::ioReads.0 = Counter64: 0
VSTORAGE-MIB::ioWrites.0 = Counter64: 0
VSTORAGE-MIB::csActive.0 = Counter64: 11
VSTORAGE-MIB::csTotal.0 = Counter64: 11
VSTORAGE-MIB::mdsAvail.0 = Counter64: 4
VSTORAGE-MIB::mdsTotal.0 = Counter64: 4
<...>
```

3.3.2.1 Listening to SNMP Traps

To start listening to SNMP traps, do the following:

1. Configure the `snmptrapd` daemon to log SNMP traps, allow them to trigger executable actions, and resend data to the network. To do this, add the following `public` community string to the `/etc/snmp/snmptrapd.conf` file:

```
authCommunity log,execute,net public
```

2. Start the daemon and specify the MIB file:

```
# snmptrapd -M /usr/share/snmp/mibs -m VSTORAGE-MIB -n -f -Lf /tmp/traps.log
```

3. Send a test trap from the **SETTINGS > Advanced settings > SNMP** tab in the management panel.
4. View the log file:

```
# tail -f /tmp/traps.log
2017-04-23 02:48:18 UDP: [127.0.0.1]:58266->[127.0.0.1]:162 [UDP: \
[127.0.0.1]:58266->[127.0.0.1]:162]:
SNMPv2-SMI::mib-2.1.3.0 = Timeticks: (1687405) 4:41:14.05      \
SNMPv2-SMI::snmpModules.1.1.4.1.0 = OID: VSTORAGE-MIB::generalAlert \
VSTORAGE-MIB::trapType = STRING: Test Case      VSTORAGE-MIB::trapMsg = \
STRING: This Is Text Message to end-user      \
VSTORAGE-MIB::trapPriority = Counter64: 1
```

The test trap is considered a `generalAlert`.

3.3.3 Monitoring Clusters with Zabbix

To configure cluster monitoring in Zabbix, do the following:

1. On the **SETTINGS > Advanced settings > SNMP** tab, click the corresponding link to download a template for Zabbix.

Note: The template is compatible with Zabbix 3.x.

2. In Zabbix, click **Configuration > Templates > Import** and **Browse**.

3.3. Monitoring Virtuozzo Cluster Objects via SNMP

Import file vstorage.xml

Rules	Update existing	Create new	Delete missing
Groups		☒	
Hosts	☐	☐	
Templates	☒	☒	
Template screens	☒	☒	☐
Template linkage		☒	
Applications		☒	☐
Items	☒	☒	☐
Discovery rules	☒	☒	☐
Triggers	☒	☒	☐
Graphs	☒	☒	☐
Web scenarios	☒	☒	☐
Screens	☐	☐	
Maps	☐	☐	
Images	☐	☐	
Value mappings	☐	☒	

3. Navigate to the template, select it, and click **Import**.
4. Click **Configuration > Hosts > Create host**.

The screenshot shows the Zabbix configuration form for a new host. The form includes the following sections:

- Host name:** management.node.hostname
- Visible name:** Storage Cluster
- Groups:** In groups (empty), Other groups (Discovered hosts, Hypervisors, Linux servers, Templates, Virtual machines, Zabbix servers)
- New group:** vstorage
- Agent interfaces:** IP address DNS name Connect to Port Default. Add
- SNMP interfaces:** 10.250.14.15 IP DNS 161. Use bulk requests. Add
- JMX interfaces:** Add
- IPMI interfaces:** Add
- Description:** (empty text area)
- Monitored by proxy:** (no proxy)
- Enabled:** ☒
- Buttons:** Add, Cancel

5. On the **Host** tab, do the following:

5.1. Specify the **Host name** of the management node and its **Visible name** in Zabbix.

5.2. Specify vstorage in the **New group** field.


5.3. **Remove** the **Agent Interfaces** section.

5.4. **Add** an **SNMP interfaces** section and specify the IP of the management node in the corresponding field.


6. On the **Templates** tab, click **Select** next to the **Link new templates** field.

3.3. Monitoring Virtuozzo Cluster Objects via SNMP

7. In the **Zabbix Server: Templates** window, check the Template VStorageSNMP template and click **Select**.



8. Back on the **Templates** tab, click the **Add** link in the **Link new templates** section. The VStorageSNMP template will appear in the **Linked templates** group.



9. Having configured the host and added its template, click the **Add** button.

Linked templates

Name	Action
Template VStorageSNMP	Unlink

Link new templates

[Add](#)

In a few minutes, the cluster's **SNMP** label in the **Availability** column on the **Configuration > Hosts** screen will turn green.

☐ Name ▲	Applications	Items	Triggers	Graphs	Discovery	Web	Interface	Templates	Status	Availability	Agent encryption	Info		
☐ Cluster	Applications 2	Items 32	Triggers 7	Graphs 3	Discovery 2	Web	10.250.14.15: 161	VStorageSNMP	Enabled	ZBX	SNMP	JMX	IPMI	NONE

To monitor cluster's parameters, open the **Monitoring > Latest data** screen, set the filter's **Host groups** to **vstorage** and click **Apply**.

You can create performance charts on the **Configuration > Hosts > <cluster> > Graphs** tab and a workplace for them on the **Monitoring > Screens** tab.

3.3.4 Description of Cluster Objects and Traps

The table below describes cluster-related objects you can monitor:

Object	Description
VSTORAGE-MIB:cluster	General cluster information.
VSTORAGE-MIB:csStatTable	Chunk server statistics table.
VSTORAGE-MIB:mdsStatTable	Metadata server statistics table.
VSTORAGE-MIB::clusterName	Cluster name.
VSTORAGE-MIB::healthStatus	Cluster health status.
VSTORAGE-MIB::usedSpace	The space occupied by all data chunks and their replicas plus the space occupied by any other data stored on cluster nodes' disks.
VSTORAGE-MIB::totalSpace	The total space on all cluster nodes' disks.
VSTORAGE-MIB::freeSpace	The unused space on all cluster nodes' disks.
VSTORAGE-MIB::licenseStatus	License status.
VSTORAGE-MIB::licenseCapacity	The maximum disk space available as defined by license.

3.3. Monitoring Virtuozzo Cluster Objects via SNMP

Object	Description
VSTORAGE-MIB::licenseExpirationStatus	License expiration status.
VSTORAGE-MIB::ioReadOpS	Current read speed in operations per second.
VSTORAGE-MIB::ioWriteOpS	Current write speed in operations per second.
VSTORAGE-MIB::ioReads	Current read speed in bytes per second.
VSTORAGE-MIB::ioWrites	Current read write in bytes per second.
VSTORAGE-MIB::csActive	The number of active chunk servers.
VSTORAGE-MIB::csTotal	The total number of chunk servers.
VSTORAGE-MIB::mdsAvail	The number of running metadata servers.
VSTORAGE-MIB::mdsTotal	The total number of metadata servers.
VSTORAGE-MIB::s3OsAvail	The number of running S3 object servers.
VSTORAGE-MIB::s3OsTotal	The total number of S3 object servers.
VSTORAGE-MIB::s3NsAvail	The number of running S3 name servers.
VSTORAGE-MIB::s3NsTotal	The total number of S3 name servers.
VSTORAGE-MIB::s3GwAvail	The number of running S3 gateways.
VSTORAGE-MIB::s3GwTotal	The total number of S3 gateways.

The table below describes the CS-related objects you can monitor:

Object	Description
VSTORAGE-MIB::csId	Chunk server identifier.
VSTORAGE-MIB::csStatus	Current chunk server status.
VSTORAGE-MIB::csIoReadOpS	Current read speed of a chunk server in operations per second.
VSTORAGE-MIB::csIoWriteOpS	Current write speed of a chunk server in operations per second.
VSTORAGE-MIB::csIoWait	The percentage of time spent waiting for I/O operations. Includes time spent waiting for synchronization.
VSTORAGE-MIB::csIoReadS	Current read speed of a chunk server in bytes per second.
VSTORAGE-MIB::csIoWriteS	Current write speed of a chunk server in bytes per second.

The table below describes MDS-related objects you can monitor:

Object	Description
VSTORAGE-MIB::mdsId	Metadata server identifier.
VSTORAGE-MIB::mdsStatus	Current metadata server status.

Object	Description
VSTORAGE-MIB::mdsMemUsage	The amount of memory used by a metadata server.
VSTORAGE-MIB::mdsCpuUsage	The percentage of the CPU's capacity used by a metadata server.
VSTORAGE-MIB::mdsUpTime	Time since the startup of a metadata server.

The table below describes SNMP traps triggered by the specified alerts:

Trap	Alert
licenseExpired	The license has expired.
tooFewClusterFreeLogicalSpace	Too few free space is left.
tooFewClusterFreePhysicalSpace	Too few physical space is left.
tooFewNodes	Too few nodes are left.
tooFewMdses	Too few MDSs are left.
generalAlert	Other.

CHAPTER 4

Monitoring Virtuozzo Storage Nodes

Nodes added to the Virtuozzo Storage infrastructure are listed on the **NODES** screen, grouped by their statuses. If there are no clusters in Virtuozzo Storage, you will only see a list of the **UNASSIGNED** nodes. If there are clusters, you can select one in the drop-down list on the left. The cluster nodes will be listed on the screen alongside the unassigned nodes.

4.1 Node Statuses

A node can have one of the following statuses:

- **HEALTHY.** All the storage services on the node are running.
- **OFFLINE.** The node cannot be reached from the management panel, although it may still be up and its services may be running.
- **FAILED.** One or more storage services on the node have failed.
- **UNASSIGNED.** The node is not assigned to a cluster.

4.2 Monitoring Node Performance

To monitor the performance of a cluster node, open the NODES screen and click the node. On the node overview screen, you will see performance statistics described below.


The overall statistics include:

- the number of CPUs and the amount of RAM,
- CPU usage, in percent over time,
- RAM usage, in percent over time.



The **DISKS** section shows:

- the number of HDD and SSD drives and their statuses,
- node I/O activity over time on the read and write charts.



The **NETWORK** section shows:

- the list of network interfaces and their statuses,
- the amount of transmitted (TX) and received (RX) traffic over time.

4.2. Monitoring Node Performance



The following sections provide more information on disk and network usage.

4.2.1 Monitoring Node Disks

To monitor the usage and status of node disks, click the **DISKS** link on the node overview screen. You will see a list of all disks on the node and their status icons.

A disk status icon shows the combined status of S.M.A.R.T. and the service corresponding to the disk role. It can be one of the following:

- **Ok.** The disk and service are healthy.
- **Failed.** The service has failed or S.M.A.R.T. reported an error.
- **Releasing.** The service is being released. When the process finishes, the disk status will change to **Ok**.

On this screen, you can:

- monitor the details and performance of each disk,
- manage disk roles,
- have the disk blink its activity LED. Works only for LSI and PERC controllers.

To monitor performance of a particular disk, select it and click **Performance**. The **Drive performance** panel will display the I/O activity of the disk.

To view information about the disk, including its S.M.A.R.T. status, click **Details**.

To have the disk blink its activity LED, select the disk, and click **Blink**. To have the disk stop blinking, click **Unblink**.

4.2.1.1 Monitoring the S.M.A.R.T. Status of Node Disks

The S.M.A.R.T. status of all disks is monitored by a tool installed along with Virtuozzo Storage. Run every 10 minutes, the tool polls all disks attached to nodes, including journaling SSDs and system disks, and reports the results to the management node.

Note: For the tool to work, make sure the S.M.A.R.T. functionality is enabled in node's BIOS.

If a S.M.A.R.T. warning message is shown in the node status, one of that node's disks is in pre-failure condition and should be replaced. If you continue using the disk, keep in mind that it may fail or cause performance issues.

Pre-failure condition means that at least one of these S.M.A.R.T. counters is not zero:

- Reallocated Sector Count
- Reallocated Event Count
- Current Pending Sector Count
- Offline Uncorrectable

4.3 Monitoring Node Network

To monitor the node's network usage, click **NETWORK** on the node overview screen.

To display the performance charts of a specific network interface, select it in the list and click **Performance**.

When monitoring network performance, keep in mind that if the **TX DROPS** and/or **RX DROPS** charts are not empty, the network is experiencing issues and requires attention.

To display the details of a network interface, click **Details**. The **Network details** panel shows the interface state, bandwidth, MTU, MAC address, and all IP addresses.

CHAPTER 5

Viewing Virtuozone Storage Alerts and Audit Log

This chapter describes Virtuozone Storage alerts and audit log.

5.1 Viewing Alerts

The **ALERTS** tab lists all the alerts logged by Virtuozone Storage. An alert is generated and logged each time one of the following conditions is met or events happen:

- a critical issue has happened with a cluster, its components (CS, MDS), disks, nodes, or services;
- cluster requires configuration or more resources to build or restore its health;
- network requires configuration or is experiencing issues that may affect performance;
- license is about to expire or has expired;
- cluster is about to or has run out of available space.

Alerts



Search				
☐	Type	Message	Date and time	
☐		Node 172.16.25.85 is offline	Jul 19, 2017, 10:11 PM	
☐		Cluster "cluster1" has fewer than 3 Metad...	Jul 19, 2017, 10:12 PM	

To view an alert details, select an alert on the **ALERTS** tab and click **Details**.

Alerts can be ignored (deleted from the alerts list) or postponed for several hours. Postponed alerts reappear in the list after some time.

To ignore or postpone an alert, select it and click the corresponding button.

5.2 Viewing Audit Log

The **AUDIT LOG** tab lists all management operations performed by users and their activity events.

5.2. Viewing Audit Log

Audit log



 Show extended details

Date and time ↓	User	Activity	
Jul 19, 2017, 4:45 PM	admin	Create S3 cluster	Create S3 cluster
Jul 19, 2017, 3:06 PM	admin	Join node	Join node to clus...
Jul 19, 2017, 3:06 PM	admin	Join node	Join node to clus...
Jul 19, 2017, 3:05 PM	admin	Create cluster	Create cluster "c...
Jul 19, 2017, 3:04 PM	admin	Configure network	Configure netwo...
Jul 19, 2017, 3:04 PM	admin	Configure network	Configure netwo...
Jul 19, 2017, 3:03 PM	admin	Configure network	Configure netwo...
Jul 19, 2017, 1:15 PM	admin	Login user	User "admin" lo...

To view detailed information on a log entry, select it and click **Show extended details**.

CHAPTER 6

Exporting Virtuoizzo Storage Cluster Data

Virtuoizzo Storage allows you to export storage space as:

- Block storage via iSCSI for virtualization, databases and other needs.
- Object storage for storing unlimited number of files via an Amazon S3 compatible protocol. You can store data like media files, backups, dropbox-like applications, and Open Xchange files. You can build your own Amazon S3 compatible object storage services as a part of your cloud offering or for internal needs.
- A back-end for Acronis Backup Cloud backups.
- A datastore for Virtuoizzo virtual machines, containers, and backups.


6.1 Exporting Data via iSCSI

Virtuoizzo Storage allows you to export cluster disk space to external operating systems in the form of LUN block devices over iSCSI in a SAN-like manner.

In Virtuoizzo Storage, you can create and run multiple iSCSI targets per Virtuoizzo Storage cluster node. In turn, each iSCSI target can have multiple LUNs (virtual disks). At any given moment, each iSCSI target runs on a single node. If a node fails, iSCSI targets hosted on it are moved to and re-launched on a healthy node.

The figure below shows a typical network configured for exporting Virtuoizzo Storage disk space over iSCSI.

6.1. Exporting Data via iSCSI



In this example there are three hardware nodes working in an Virtuozzo Storage cluster. Two nodes host one iSCSI target each, while the third hosts two iSCSI targets. Each node has two network connections: one internal for cluster communication and one external for iSCSI exporting. Each iSCSI target has its own static IP address assigned from the datacenter network. The iSCSI network must be properly configured to run iSCSI targets. Each iSCSI target has a unique IP address from the iSCSI subnetwork different from the host IP address. So make sure you have a range of unused IP addresses for iSCSI targets.

6.1.1 Creating Virtuozzo Storage iSCSI Targets

Note:

1. Each iSCSI target must be assigned at least one unique IP address from DC network's static pool.
2. The name of each iSCSI target must be unique in the Virtuozzo Storage cluster.
3. Virtuozzo Storage iSCSI targets support persistent reservations to allow iSCSI initiators obtain exclusive access to the specified target's LUNs.

To create and start a target, do the following:

1. On the **SERVICES > iSCSI > Targets** screen, click **ADD TARGET**.
2. On the **Add target** panel, type a name for the new target in the **Name** field.

× **Add target**

Name

target1

Node

172.16.24.126

▼

☐ **Enable CHAP**

▼

Target portal IP address

☐ 172.16.24.125

+ Add

− Remove

Done

3. In the node drop-down list, select a node on which the target will be located. The node should have an iSCSI role assigned to one of its network interfaces to appear in the list.
4. If necessary, check the **Enable CHAP** box and select an iSCSI user in the corresponding drop-down list (For more information on CHAP users, see [Managing iSCSI Users](#) on page 57).

6.1. Exporting Data via iSCSI

5. Click **Add** to specify one or more IP addresses for the target.
6. Click **Done** to create the target.

The iSCSI target will be automatically started after creation and the initiators will be able to access the target via the specified IP address.

6.1.1.1 Performance Tips

- Spread iSCSI targets evenly across nodes in the cluster. For example, ten nodes with one iSCSI target per each will perform better than a single node with ten iSCSI targets on it.
- Fewer LUNs per more iSCSI targets will perform better than more LUNs per fewer iSCSI targets.

6.1.2 Listing, Stopping, and Deleting Virtuozone Storage iSCSI Targets

On the iSCSI targets screen, you can list and manage all iSCSI targets and their LUNs, and display detailed information about specific iSCSI targets registered on a node.

To stop or delete an iSCSI target, select it on the iSCSI targets screen and click **Stop** or **Delete**, respectively. Doing so will disconnect the iSCSI initiator from the target. However, breaking the connection in such a way may result in I/O errors on the iSCSI initiator's side.

6.1.3 Configuring Virtuozone Storage iSCSI Targets

To configure an iSCSI target, do the following:

1. On the **SERVICES > iSCSI > Targets** screen, select the necessary target and click **Configure**.
2. On the **Configure** target screen, specify the necessary parameters.

× Configure target

IQN

iqn.2014.06.com.vstorage:target7

Node

172.16.24.126

▼

☐ Enable CHAP

▼

☐ 172.16.24.125

+

 Add

−

 Remove

Done

3. Click **Done**.

6.1.3.1 Listing LUNs

In Virtuozzo Storage, each iSCSI target can have multiple LUNs (virtual disks or volumes). You can list the LUNs of a target and iSCSI initiators that are currently connected.

To list the LUNs, open the **SERVICES > iSCSI > Targets** screen, select a target and click a link in the **LUNs**

6.1. Exporting Data via iSCSI

column.

To list the initiators that are currently connected to iSCSI targets, open the **INITIATORS** tab on the same screen.

6.1.3.2 Adding LUNs

To add a LUN to an iSCSI target, do the following:

1. On the **SERVICES > iSCSI > Targets** screen, select the necessary target and click a link in the **LUNs** column.

iSCSI targets

TARGETS USERS					
Search					
IQN	State	IP/Address	LUNs	Initiators	
iqn.2014.06.com.vstorage:target1	running	172.16.24.125	0	0	

2. To add a new LUN to the list, click **ADD LUN**.

× Add LUN

LUN

2 ▼

LUN size

- 20 + GB
 Tier 0 ▼
 ☒ Thin provisioning

Data redundancy:

☒ Erasure coding
 ☐ Replication

Failure-domain:

Disk ▼

Encoding 1+0	0% overhead	
Encoding 3+2	67% overhead	
Encoding 5+2	40% overhead	i
Encoding 7+2	29% overhead	i
Encoding 17+3	18% overhead	i

Done

- On the **Add LUN** screen, select the LUN's number from the drop-down list.
- In the **LUN Size** field, specify the size of the LUN in GB; select a tier from the drop-down list to the right. For more information on tiers, see **Understanding Storage Tiers** in the *Installation Guide*.
- From the **Failure domain** drop-down list, choose a placement policy for replicas. For more details, see **Understanding Failure Domains** in the *Installation Guide*.

6.1. Exporting Data via iSCSI

6. Choose a data redundancy mode. For more details, see **Understanding Data Redundancy** in the *Installation Guide*.
7. Click **Done**.

6.1.3.3 Configuring LUNs

To configure a LUN of an iSCSI target, do the following:

1. On the **SERVICES > iSCSI > Targets** screen, stop the target to which the LUN belongs.
2. Click a link in the target's **LUNs** column.
3. On the **Configure LUN** screen, specify the LUN size in the corresponding field.
4. Click **Done**.

6.1.3.4 Deleting LUNs

To delete a LUN, do the following:

1. On the **SERVICES > iSCSI > Targets** screen, select the necessary target and click a link in the **LUNs** column.
2. Select the necessary LUN in the list and click **Delete**.

6.1.4 Managing iSCSI Users

6.1.4.1 Creating CHAP Accounts for Virtuoizzo Storage iSCSI Targets

To create a CHAP account, do the following:

1. On the **SERVICES > iSCSI Users** screen, click **Add user**.
2. Specify login, password, and, if necessary, a description for the account. The password should be 12 to 16 characters long for Windows clients to be able to establish connections.

× **Add user**

Login

user1

Password

.....

Description

A new CHAP user.

☒ Enabled

Done

3. Click **Done**.

The newly created CHAP user account will be listed on the **iSCSI Users** screen.

6.1.4.2 Creating Virtuoizzo Storage iSCSI Targets Bound to CHAP Accounts

To create an Virtuoizzo Storage iSCSI target bound to a CHAP account, do the following:

1. On the **SERVICES > iSCSI > Targets** screen, select an iSCSI target and click **Configure**.
2. On the **Configure** target screen, check **Enable CHAP** and/or **Enable mutual CHAP** and select users in the corresponding drop-down lists. If you enable CHAP, the target will authenticate the initiator. If you enable mutual CHAP, the initiator will authenticate the target. These options can be enabled in any combination.

× **Configure target**

IQN

Node

☒ **Enable CHAP**

▼

user1

user2

+ Add

− Remove


Done

3. Click **Done**.

6.1.4.3 Changing CHAP Account Passwords

To change the password of a CHAP account, do the following:

1. On the **SERVICES > iSCSI Users** screen, select a user and click **Configure**.
2. In the **Password** section on the **Configure** user screen, click change.



×

Configure user

Login
user1

Password
[change](#)

Description
A new CHAP user.

☒ Enabled

Done

3. Type a new password in the corresponding field and click **Done**. The password should be 12 to 16 characters long for Windows clients to be able to establish connection.

The new password will become active after target reboot.

6.2 Exporting Data via S3

Virtuozzo Storage allows you to export cluster disk space to customers in the form of an S3-like object-based storage.

Virtuozzo Storage is implemented as an Amazon S3-like API, which is one of the most common object storage APIs. End users can work with Virtuozzo Storage as they work with Amazon S3. You can use the usual applications for S3 and continue working with it after the data migration from Amazon S3 to Virtuozzo Storage.

Object storage is a storage architecture that enables managing data as objects (like in a key-value storage) as opposed to files on file systems or blocks in a block storage. Except for the data, each object has metadata that describes it as well as a unique identifier that allows finding the object in the storage. Object storage is optimized for storing billions of objects, in particular for application storage, static web content hosting,

6.2. Exporting Data via S3

online storage services, big data, and backups. All of these uses are enabled by object storage thanks to a combination of very high scalability and data availability and consistency.


Compared to other types of storage, the key difference of object storage is that parts of an object cannot be modified, so if the object changes a new version of it is spawned instead. This approach is extremely important for maintaining data availability and consistency. First of all, changing an object as a whole eliminates the issue of conflicts. That is, the object with the latest timestamp is considered to be the current version and that is it. As a result, objects are always consistent, i.e. their state is relevant and appropriate.

Another feature of object storage is eventual consistency. Eventual consistency does not guarantee that reads are to return the new state after the write has been completed. Readers can observe the old state for an undefined period of time until the write is propagated to all the replicas (copies). This is very important for storage availability as geographically distant data centers may not be able to perform data update synchronously (e.g., due to network issues) and the update itself may also be slow as awaiting acknowledges from all the data replicas over long distances can take hundreds of milliseconds. So eventual consistency helps hide communication latencies on writes at the cost of the probable old state observed by readers. However, many use cases can easily tolerate it.

6.2.1 Object Storage Infrastructure Overview

The infrastructure of the object storage consists of the following entities: object servers (OS), name servers (NS), and the S3 gateways (GW).

These entities run as services on the Virtuozzo Storage nodes. Each service should be deployed on multiple Virtuozzo Storage nodes for high availability.



- An object server stores actual object data received from S3 gateway, packed into special containers to achieve high performance. The containers are redundant, you can specify the redundancy mode while configuring object storage.
- A name server stores information about objects (metadata) received from the S3 gateway. Metadata includes object name, size, ACL, location, owner, and such.
- S3 gateway is a data proxy between object servers and users. It receives and handles Amazon S3 protocol requests and uses nginx web server for external connections. S3 gateway handles S3 user authentication and ACL checks. It has no data of its own (i.e. is stateless).

6.2.2 Planning the S3 Cluster

Before creating an S3 cluster, do the following:

6.2. Exporting Data via S3


1. Define which nodes of the Virtuozzo Storage cluster will run the S3 storage access point services. It is recommended to have all nodes available in Virtuozzo Storage run these services.
2. Configure the network so that the following is achieved:
 - All components of the S3 cluster communicate with each other via the S3 private network. All nodes of an S3 cluster must be connected to the S3 private network. Virtuozzo Storage internal network can be used for this purpose.
 - The nodes running S3 gateways must have access to the public network.
 - The public network for the S3 gateways must be balanced by an external DNS load balancer.

For more details on network configuration, refer to **Planning Network** in the *Installation Guide*.

3. All components of the S3 cluster should run on multiple nodes for high-availability. Name server and object server components in the S3 cluster are automatically balanced and migrated between S3 nodes. S3 gateways are not automatically migrated; their high availability is based on DNS records. You should maintain the DNS records manually when adding or removing the S3 gateways.

6.2.3 Sample Object Storage

This section shows a sample object storage deployed on top of an Virtuozzo Storage cluster of five nodes that run various services. The final setup is shown on the figure below.



6.2.4 Creating the S3 Cluster

To set up object storage services on a cluster node, do the following:

1. Make sure that S3 private network is configured on each node that will run object storage services.
2. On the **SERVICES > Nodes** screen, check the box of each cluster node where object storage services will run.

S3 Nodes



OVERVIEW

NODES

USERS

BUCKETS

Create S3 cluster

Q Search

Any status ▾

AVAILABLE NODES

4

Hide available nodes ^

172.16.24.126 ✓


172.16.25.102 ✓


172.16.25.85 ✓


172.16.25.97 ✓


3. Click **Create S3 cluster**.
4. Make sure a network interface with an **S3 (private)** role is selected in the drop-down list. The corresponding interfaces with S3 public roles will be selected automatically.

Note: If necessary, click the cogwheel icon and, on the **Network Configuration** screen, configure S3 roles.

× **Create S3 cluster**

 172.16.24.126 

S3 private

eno16777736 - 172.16.24.1: 

S3 public

 172.16.25.102 

S3 private

eno16777736 - 172.16.25.102: 

S3 public

eno16777736 - 172.16.25.102

Proceed

5. Click **Proceed**.
6. In **Tier**, select the storage tier that will be used for the object storage. For information about storage tiers, consult **Understanding Storage Tiers** in the *Installation Guide*.
7. In **Failure domain**, choose a placement policy for replicas. For more details, see **Understanding Failure Domains** in the *Installation Guide*.
8. In **Data redundancy**, select the redundancy mode that the object storage will use. For more details, see **Understanding Data Redundancy** in the *Installation Guide*.

< Volume parameters

Tier:

Tier 0

Data redundancy:

☒ Erasure coding ☐ Replication

Failure-domain:

Disk

Encoding 1+0	0% overhead	
Encoding 3+2	67% overhead	
Encoding 5+2	40% overhead	
Encoding 7+2	29% overhead	
Encoding 17+3	18% overhead	

Proceed

Note: You can later change the redundancy mode on the **S3 > Settings** panel.

9. Click **Proceed**.
10. 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, `mys3storage.example.com`. Click **Proceed**.

Important: Configure your DNS server according to the example suggested in the management panel.

11. From the drop-down list, select an S3 endpoint protocol: HTTP, HTTPS or both.

 **Protocols**

S3 endpoint protocols:



Endpoint URL:

`https://mys3storage.example.com/bucketname/objectname`

☐ Generate self-signed certificate

SSL certificate:

Valid

SSL key:

Valid

Note: It is recommended to use only HTTPS for production deployments.

If you have selected HTTPS, do one of the following:

- Check **Generate self-signed certificate** to get a self-signed certificate for HTTPS evaluation.

6.2. Exporting Data via S3


Note: To access the data in S3 cluster via a browser, add the self-signed certificate to the browser's exceptions.

- Acquire an SSL certificate and key from a third party, upload it, and, depending on the certificate type, do one of the following:
 - in case the certificate is contained in a PKCS#12 file, specify the passphrase;
 - otherwise, upload the SSL key.

12. If required, click **Configure Acronis Notary** and specify **Notary DNS name** and **Notary user key**. For more information on Acronis Notary, see **Managing Acronis Notary in Buckets**.

13. Click **Done** to create an S3 cluster.

After the cluster is created, on the **S3 Overview** screen, you can 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> or http://<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.

6.2.5 Managing Object Storage Users

The concept of S3 user is one of the base concepts of object storage along with those of object and bucket (container for storing objects). The Amazon S3 protocol uses a permission model based on access control lists (ACLs) where each bucket and each object is assigned an ACL that lists all users with access to the given resource and the type of this access (read, write, read ACL, write ACL). The list of users includes the entity owner assigned to every object and bucket at creation. The entity owner has extra rights compared to other users. For example, the bucket owner is the only one who can delete that bucket.

User model and access policies implemented in Virtuozzo Storage comply with the Amazon S3 user model and access policies.


User management scenarios in Virtuozzo Storage are largely based on the Amazon Web Services user management and include the following operations: create, query, and delete users as well as generate and revoke user access key pairs.

6.2.5.1 Adding S3 users

To add an S3 user, do the following:

1. On the **SERVICES > S3 Users** screen, click **Add user**.

6.2. Exporting Data via S3



2. Specify a valid email address as login for the user and click **Done**.

× Add user

Login

Enabled


Done

6.2.5.2 Managing S3 Access Key Pairs

Each S3 user has one or two key pairs (access key and secret key) for accessing the S3 cloud. You can think of the access key as login and the secret key as password. (For more information about S3 key pairs, refer to <http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSGettingStartedGuide/AWSCredentials.html>.) The access keys are generated and stored locally in the Virtuozzo Storage cluster on S3 name servers. Each user can have up to two key pairs. It is recommended to periodically revoke old and generate new access key pairs.

To view, add, or revoke the S3 access key pairs for an S3 user, do the following:

1. Select a user in the list and click **Keys**.



2. The existing keys will be shown on the **Keys** panel.

- To revoke a key, click **Revoke**.
- To add a new key, click **Generate access key**.

To access a bucket, a user will need the following information:

- management panel IP address,
- DNS name of the S3 cluster specified during configuration,
- S3 access key ID,
- S3 secret access key,
- SSL certificate if the HTTPS protocol was chosen during configuration.

Note: The certificate file can be found in the `/etc/nginx/ssl/` directory on any node hosting the S3 gateway service.

To automatically log in to S3 with user credentials using the generated keys, select a user and click **Browse**.

6.2. Exporting Data via S3

Note: To **Browse** using an SSL certificate, make sure it is valid or, in case of a self-signed one, add it to browser's exceptions.

6.2.6 Managing Object Storage Buckets

All objects in Amazon S3-like storage are stored in containers called buckets. Buckets are addressed by names that are unique in the given object storage, so an S3 user of that object storage cannot create a bucket that has the same name as a different bucket in the same object storage. Buckets are used to:

- group and isolate objects from those in other buckets,
- provide ACL management mechanisms for objects in them,
- set per-bucket access policies, for example, versioning in the bucket.

In the current version of Virtuozzo Storage, you can enable and disable Acronis Notary for object storage buckets and monitor the space used by them on the **SERVICES > S3 > Buckets** screen. You cannot create and manage object storage buckets from Virtuozzo Storage management panel. However, you can do it via the Virtuozzo Storage user panel or by using a third-party application. For example, the applications listed below allow you to perform the following actions:

- CyberDuck: create and manage buckets and their contents.
- MountainDuck: mount object storage as a disk drive and manage buckets and their contents.
- Backup Exec: store backups in the object storage.

6.2.6.1 Listing Bucket Contents

You can list bucket contents with a web browser. To do this, visit the URL that consists of the external DNS name for the S3 endpoint that you specified when creating the S3 cluster and the bucket name. For example, `mys3storage.example.com/mybucket`.

Note: You can also copy the link to bucket contents by right-clicking it in CyberDuck, and then selecting Copy URL.


6.2.6.2 Managing Acronis Notary in Buckets

Virtuozzo Storage offers integration with the Acronis Notary service to leverage blockchain notarization and ensure the immutability of data saved in object storage clusters. To use Acronis Notary in user buckets, you need to set it up in the S3 cluster and enable it for said buckets.

Setting Up Acronis Notary

To set up Acronis Notary, do the following:

1. Get the DNS name and the user key for the notary service from your Acronis sales contact.
2. On the **SERVICES > S3** screen, click **Notary settings**.



3. On the **Notary Settings** screen, specify the DNS name and user key in the respective fields and click **Done**.

< Notary Settings

Notary DNS name

http://ask-sales-for-notary-provider.acronis.com

Notary User Key

F24CDF1a6fED79b3f48eD000044a5c6b

Done

Enabling and Disabling Acronis Notary

To enable or disable blockchain notarization for a bucket, select a bucket on the **SERVICES > S3 > Buckets** screen and click **Enable Notary** or **Disable Notary**, respectively.

Notarization is disabled for new buckets by default.

Note: Once you enable notarization for a bucket, certificates are created automatically only for the newly uploaded files. The previously uploaded files are left unnotarized. Once a file was notarized, it will remain notarized even if you disable notarization later.

6.2.7 Best Practices for Using S3 in Virtuozzo Storage

This section offers recommendations on how to best use the S3 feature of Virtuozzo Storage.

6.2.7.1 Bucket and Key Naming Policies

It is recommended to use bucket names that comply with DNS naming conventions:

- can be from 3 to 63 characters long,

- must start and end with a lowercase letter or number,
- can contain lowercase letters, numbers, periods (.), hyphens (-), and underscores (_),
- can be a series of valid name parts (described previously) separated by periods.

An object key can be a string of any UTF-8 encoded characters up to 1024 bytes long.

6.2.7.2 Improving Performance of PUT Operations

Object storage supports uploading of objects as large as 5 GB in size with a single PUT request, or 5 TB in size with multipart upload. Upload performance can be improved, however, by splitting large objects into pieces and uploading them concurrently with multipart upload API. This approach will divide the load between multiple OS services.

It is recommended to use multipart uploads for objects larger than 5 MB.

6.2.8 Replicating Data Between Geographically Distributed Datacenters with S3 Clusters

Virtuozzo Storage can store replicas of S3 cluster data and keep them up-to-date in multiple geographically distributed datacenters with S3 clusters based on Virtuozzo Storage. 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 to have an Internet connection.

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

To set up geo-replication between S3 clusters, exchange tokens between datacenters as follows:

1. In the management panel of a remote datacenter, open the **SERVICES > S3 > GEO-REPLICATION** screen.

S3 Geo-Replication

[ADD DATACENTER](#)[OVERVIEW](#)[NODES](#)[USERS](#)[BUCKETS](#)[GEO-REPLICATION](#)**TOKEN**

s3.example.com

cluster1
Virtuozzo Storage

2. In the section of the home S3 cluster, click **TOKEN** and, on the **Get token** panel, copy the token.
3. In the management panel of the local datacenter, open the **SERVICES > S3 > GEO-REPLICATION** screen and click **ADD DATACENTER**.

✕ Add datacenter

To replicate data from another datacenter, insert the token obtained from its Web Control panel.

Token

```
eyJ1c2VyX2tleV9pZCI6ICI5ZGY5MmlyYTgxZGI0MGFhQTcyWCIsICJyZW
FkYWJsZV9uYW1lIjogImNsdXN0ZXIwglVzZXJfc2VjcmV0X2tleSI6ICJk
cDI1eUwxWEpLc3BGcEN2U3BCd2IDS1BydkpxVINtRncwakpPWW45li
wglVybCI6ICJodHRwczovL3MzLmV4YW1wbGUuY29tOjQ0MyIsICJpc1
9zZWxmIjogdHJ1ZSwglVpZCI6ICxOWI4N2M0YTRkOGJkOTIhIn0=
```

ADD

4. Enter the copied token and click **Done**.
5. Configure the remote Virtuozzo Storage S3 cluster the same way.

6.2.9 Monitoring S3 Access Points

The S3 monitoring screen enables you to inspect the availability of each S3 component as well as the performance of NS and OS services (which are highly available).

If you see that some of the NS or OS services are offline, it means that the S3 access point does not function properly, and you should contact support or consult the CLI guide for low-level troubleshooting. S3 gateways are not highly available, but DNS load balancing should be enough to avoid downtime if the gateway fails.

The performance charts represent the number of operations that the OS/NS services are performing.

6.2.10 Releasing Nodes from S3 Clusters

Before releasing a node, make sure that the cluster has enough nodes running the Name Server, Object Server, and S3 Gateway services left.

6.2. Exporting Data via S3

Warning: When the last node in the S3 cluster is removed, the cluster is destroyed, and all the data is deleted.

To release a node from an S3 cluster, do the following:

1. On the **SERVICES > S3 Nodes** screen, check the box of the node to release.
2. Click **Release**.

6.2.11 Supported Amazon S3 Features

This section lists Amazon S3 operations, headers, and authentication schemes supported by the Virtuozzo Storage implementation of the Amazon S3 protocol.

6.2.11.1 Supported Amazon S3 REST Operations

The following Amazon S3 REST operations are currently supported by the Virtuozzo Storage implementation of the Amazon S3 protocol:

Supported service operations:

- GET Service

Bucket operations:

Operation	Supported
DELETE/HEAD/PUT Bucket	Yes
GET Bucket (List Objects)	Yes (only version 1)
GET/PUT Bucket acl	Yes
GET Bucket location	Yes (returns US East)
GET Bucket Object versions	Yes
GET/PUT Bucket versioning	Yes
List Multipart Uploads	Yes
DELETE/GET/PUT Bucket analytics	No
DELETE/GET/PUT Bucket cors	No
DELETE/GET/PUT Bucket inventory	No
DELETE/GET/PUT Bucket lifecycle	No

Operation	Supported
DELETE/GET/PUT Bucket metrics	No
DELETE/GET/PUT Bucket policy	No
DELETE/GET/PUT Bucket replication	No
DELETE/GET/PUT Bucket tagging	No
DELETE/GET/PUT Bucket website	No
GET/PUT Bucket accelerate	No
GET/PUT Bucket logging	No
GET/PUT Bucket notification	No
GET/PUT Bucket requestPayment	No
List Bucket Analytics Configurations	No
List Bucket Inventory Configurations	No
List Bucket Metrics Configurations	No

Object operations:

Operation	Supported
DELETE/GET/HEAD/POST/PUT Object	Yes
Delete Multiple Objects	Yes
PUT Object - Copy	Yes
GET/PUT Object acl	Yes
Delete Multiple Objects	Yes
Abort Multipart Upload	Yes
Complete Multipart Upload	Yes
Initiate Multipart Upload	Yes
List Parts	Yes
Upload Part	Yes
Upload Part - Copy	No
DELETE/GET/PUT Object tagging	No
GET Object torrent	No
OPTIONS Object	No
POST Object restore	No

6.2. Exporting Data via S3

Note: For more information on Amazon S3 REST operations, see [Amazon S3 REST API documentation](#).

6.2.11.2 Supported Amazon Request Headers

The following Amazon S3 REST request headers are currently supported by the Virtuozzo Storage implementation of the Amazon S3 protocol:

- Authorization
- Content-Length
- Content-Type
- Content-MD5
- Date
- Host
- x-amz-content-sha256
- x-amz-date
- x-amz-security-token

The following Amazon S3 REST request headers are ignored:

- Expect
- x-amz-security-token

Note: For more information on Amazon S3 REST request headers, see [Amazon S3 REST API Common Request Headers](#).

6.2.11.3 Supported Amazon Response Headers

The following Amazon S3 REST response headers are currently supported by the Virtuozzo Storage implementation of the Amazon S3 protocol:

- Content-Length
- Content-Type
- Connection
- Date
- ETag
- x-amz-delete-marker
- x-amz-request-id
- x-amz-version-id

The following Amazon S3 REST response headers are not used:

- Server
- x-amz-id-2

Note: For more information on Amazon S3 REST response headers, see [Amazon S3 REST API Common Response Headers](#).

6.2.11.4 Supported Amazon Error Response Headers

The following Amazon S3 REST error response headers are currently supported by the Virtuozzo Storage implementation of the Amazon S3 protocol:

- Code
- Error
- Message
- RequestId
- Resource

The following Amazon S3 REST error response headers are not supported:

- RequestId (not used)
- Resource

6.3. Using Virtuoizzo Storage with Acronis Backup Cloud

Note: For more information on Amazon S3 REST response headers, see [Amazon S3 REST API Error Response Headers](#).

6.2.11.5 Supported Authentication Scheme and Methods

The following authentication scheme is supported by the Virtuoizzo Storage implementation of the Amazon S3 protocol:

- [Signature Version 2](#).
- [Signature Version 4](#).

The following authentication methods is supported by the Virtuoizzo Storage implementation of the Amazon S3 protocol:

- [HTTP Authorization header](#).
- [Query string parameters](#).

6.3 Using Virtuoizzo Storage with Acronis Backup Cloud

The Acronis Backup Gateway storage access point is intended for service providers who use Acronis Backup Cloud and want to organize a storage on their premises to store their clients' backed-up data.

The Acronis Backup Gateway enables a service provider to easily configure the Virtuoizzo Storage cluster for the proprietary data format used by Acronis.

Acronis Backup Gateway supports the following storage back-ends:

- Virtuoizzo Storage clusters with software redundancy by erasure coding,
- NFS shares mounted on nodes in parallel with clusters.


In the latter case, you can use hardware redundancy, e.g. EMC NFS. Acronis Backup Gateway does not provide data redundancy or perform data deduplication itself.

When configuring the Acronis Backup Gateway, you must provide the credentials of your administrator

account in Acronis Backup Cloud.

6.3.1 Understanding the Infrastructure

The Acronis Backup Gateway storage access point runs as services on the Virtuozzo Storage nodes. It should be deployed on multiple Virtuozzo Storage nodes for high availability.




6.3.2 Configuring Acronis Backup Gateway Clusters

To set up an Acronis Backup Gateway cluster on top of an Virtuozzo Storage cluster, do the following:

1. Make sure that the Acronis Backup Gateway network is configured on each node that will run the gateway service.
2. On the **SERVICES > Acronis Backup Nodes** screen, select a check box next to each cluster node where the gateway services will run.

6.3. Using VirtuoZZo Storage with Acronis Backup Cloud



3. Click **Create Gateway** and **This VirtuoZZo Storage Cluster**.
4. For each node, select the network interface to which the Acronis Backup Gateway network role is assigned. The gateway service will listen on the IP address assigned to this interface.

× **Configure Network**

 172.16.14.11 

ABGW private

eth0 - 172.16.11.11 

ABGW public

eth0 - 172.16.11.11

 vip-2a 

ABGW private

eth0 - 172.16.11.12 

ABGW public

eth0 - 172.16.11.12

Proceed

5. Click **Proceed**.
6. In **Tier**, select the storage tier that will be used for the object storage. For information about storage tiers, consult **Understanding Storage Tiers** in the *Installation Guide*.
7. In **Failure domain**, choose a placement policy for replicas. For more details, see **Understanding Failure Domains** in the *Installation Guide*.
8. In **Data redundancy**, select an erasure coding mode. For more details, see **Understanding Data Redundancy** in the *Installation Guide*.

< Volume parameters

Tier:

Tier 0

Data redundancy:

☒ Erasure coding

Failure-domain:

Disk

Encoding 1+0	0% overhead	
Encoding 3+2	67% overhead	
Encoding 5+2	40% overhead	
Encoding 7+2	29% overhead	
Encoding 17+3	18% overhead	

Proceed

Note:

1. Redundancy by replication is not supported for Acronis Backup Gateway.
2. You can later change the erasure coding mode on the **Acronis Backup Gateway > Parameters** panel.

9. Click **Proceed**.

10. Specify the external DNS name for this gateway, e.g, backupgateway.example.com. Make sure that the required port is open for outgoing Internet connections and incoming connections from Acronis Backup Cloud on each node running the gateway service. Backup agents will use this address and port to upload the backed-up data.

Important:

1. Configure your DNS server according to the example suggested in the management panel.
2. Each time you add or remove a node to or from the Acronis Backup Gateway cluster, adjust the DNS settings accordingly.

11. Click **Proceed**.
12. In **Account Server Name**, specify the Acronis Backup Cloud management server address, which matches the account management console address (<https://cloud.acronis.com/>).
13. In **Acronis Account**, specify the credentials of your Acronis Backup Cloud administrator account.
14. Click **Done**.

6.3.3 Releasing Nodes Running Acronis Backup Gateway Services

Before releasing a node, make sure that the cluster has enough nodes running the Acronis Backup Gateway service left.

Warning: If the last node running the Acronis Backup Gateway is removed, the Acronis Backup Cloud customers will not be able to use this cluster as a backup storage backend.

To release a node running an Acronis Backup Gateway service from a cluster, do the following:

1. On the **SERVICES > Acronis Backup Gateway Nodes** screen, select the checkbox next to the necessary node.
2. Click **Release**.

6.3.4 Connecting External NFS Shares through Acronis Backup Gateway


Note:

1. Virtuozzo Storage does not provide data redundancy on top of NFS volumes. Depending on the implementation, NFS shares may use their own hardware or software redundancy.
2. In the current version of Virtuozzo Storage, only one cluster node may store backups of an NFS volume.

To set up an external NFS share on a cluster node through Acronis Backup Gateway, do the following:

1. Make sure that the Acronis Backup Gateway network is configured on the node that will run the gateway service.
2. On the **SERVICES > Acronis Backup Nodes** screen, select a check box next to the node where the gateway service will run.
3. Click **Create Gateway** and **Network File System**.

< Storage type



4. Make sure the network interface with the assigned Acronis Backup Gateway network role is selected. The gateway service will listen on the IP address assigned to this interface. Click **Proceed**.

✕ Configure Network



172.16.25.102

ABGW private

ens33 - 172.16.25.102


ABGW public

ens33 - 172.16.25.102

PROCEED

- On the **Volume parameters** screen, specify the **NFS hostname or IP** of the NFS share as well as the name of its volume in the **Export name** field.

< Volume parameters



NFS hostname or IP

hostname.example.com

Export name

/path/to/volume

PROCEED

- Specify the external DNS name for this gateway, e.g, backupgateway.example.com. Make sure that the required port is open for outgoing Internet connections and incoming connections from Acronis Backup Cloud on each node running the gateway service. Backup agents will use this address and port to upload backups.

< DNS Configuration

DNS name

backupgateway.example.com

This would probably require to change the configuration of the DNS server. The DNS configuration may look as follows:

```
$TTL 1h

@   IN  SOA  ns1.myhoster.com.
      root.backupgateway.example.com. (
        2017072013    ; serial
        1h    ; refresh
        30m    ; retry
        7d    ; expiration
        1h )    ; minimum

; primary name server
NS ns1.myhoster.com.

; secondary name server
NS ns2.myhoster.com.
```

PROCEED

Important:

1. Configure your DNS server according to the example suggested in the management panel.
2. Each time you add or remove a node to or from the Acronis Backup Gateway cluster, adjust the DNS settings accordingly.

7. Click **Proceed**.

< Registration in Acronis Backup Gateway

To register this cluster in Acronis Backup Cloud, provide the credentials of your administrator account in Acronis Backup Cloud.

Account Server Name

Administrator Account

DONE

8. In **Account Server Name**, specify the Acronis Backup Cloud management server address, which matches the account management console address (<https://cloud.acronis.com/>).
9. In **Acronis Account**, specify the credentials of your Acronis Backup Cloud administrator account and click **Done**.

6.4 Managing Datastores

To keep your virtual machines and containers in Virtuozzo Storage, you can place them in datastores. A datastore is a directory in a Virtuozzo Storage cluster with a number of parameters set for it, namely tier, redundancy, and failure domain.

Warning: If you delete a datastore, all VMs, containers, or backups in it will be deleted!

6.4.1 Creating Datastores

To create a datastore, do the following:

1. On the **SERVICES > COMPUTE** screen, click **ADD DATASTORES**.

× **Add Datastore**

Name

datastore1

Tier:

Tier 0

▼

Data redundancy:

Failure-domain:

☒ Erasure coding ☐ Replication

Host

▼

Encoding 1+0	0% overhead	
Encoding 3+2	67% overhead	i
Encoding 5+2	40% overhead	i
Encoding 7+2	29% overhead	i
Encoding 17+3	18% overhead	i

Warning: Datastore without redundancy should be configured only if underlying drives storage has its own built-in redundancy and data protection, provided by RAID controller or SAN/NAS devices.

Done

2. On the **Add Datastore** panel, specify a name.
3. In **Tier**, select the [storage tier](#) that will be used for the object storage.

4. In **Failure domain**, choose a [placement policy](#) for replicas.
5. In **Data redundancy**, select the [redundancy mode](#) that the object storage will use.
6. Click **Done**.

To change the datastore settings later, select it and click **Configure**.

6.4.2 Placing Virtual Environments and Backups in Datastores

After creating a datastore, you can set it to store your virtual environments and their backups

- via Virtuozzo Automator (for instructions, refer to the [Virtuozzo Automator Administrator's Guide](#)) or
- by creating symbolic links from the virtual environments and backups default location to this datastore as follows:

1. Log in to the server as root.

2. Check the default location of:

- 2.1. container private area in the `/etc/vz/vz.conf` file:

```
VE_PRIVATE=/vz/private/$VEID
```

- 2.2. virtual machine files by running:

```
# prlsrvctl info | grep "VM home"
VM home: /vz/vmprivate
```

- 2.3. backup files by running:

```
# prlsrvctl info | grep "Backup path"
Backup path: /vz/vmprivate/backups
```

3. If the `/vz/private`, `/vz/vmprivate`, and `/vz/vmprivate/backups` directories exist on the server, remove them.
4. Make symbolic links from the containers, virtual machines and backups default locations to the desired datastores. For example, to store your virtual machines in the `vmprivate` datastore, containers in the `private` datastore, and their backups in the `backups` datastore, execute the following commands:

```
# ln -s /mnt/vstorage/vols/datastores/vmprivate /vz/vmprivate
# ln -s /mnt/vstorage/vols/datastores/private /vz/private
# ln -s /mnt/vstorage/vols/datastores/backups /vz/vmprivate/backups
```

6.4. Managing Datastores

The newly created virtual environments and backups will be placed and stored in their respective datastores.