



Storage

Administrator's Guide

January 22, 2025

Virtuozzo International GmbH

Vordergasse 59

8200 Schaffhausen

Switzerland

Tel: + 41 52 632 0411

Fax: + 41 52 672 2010

<https://virtuozzo.com>

Copyright ©2016-2025 Virtuozzo International GmbH. All rights reserved.

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

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 Storage	1
1.2 Storage Cache Architecture	2
1.2.1 Reads and Writes Behavior	2
1.2.2 Caching Benefits	2
2. Managing Storage	4
2.1 Configuring Node Network Interfaces	4
2.2 Creating the Storage Cluster	7
2.2.1 Creating the Cluster on the First Node	8
2.2.2 Adding Nodes to the Cluster	10
2.2.3 Assigning Disk Roles Manually	12
2.3 Releasing Nodes from Cluster	15
2.4 Removing Nodes from the Unassigned List	16
2.5 Managing Tier Encryption	16
2.6 Managing Storage Users	17
2.6.1 Creating User Accounts	18
2.6.2 Managing User Accounts	20
2.6.3 Adding LDAP or Active Directory Users	22
2.7 Allowing root Access to Cluster Nodes Over SSH	25
2.8 Backing Up and Restoring the Management Database	26
2.8.1 Restoring the Management Database from Backup	27
2.9 Accessing the Management Panel via SSL	28
2.10 Managing Storage Licenses	28
3. Monitoring Storage Clusters	30
3.1 Monitoring Cluster Status	30

3.2	Monitoring Cluster Storage Space	30
3.2.1	Physical Space Chart	31
3.2.2	Logical Space Chart	31
3.2.2.1	Understanding Logical Space	32
3.2.3	Monitoring Chunk Status and Replication	32
3.2.4	Monitoring Cluster Services	33
3.2.5	Monitoring Cluster I/O Activity	34
3.3	Monitoring Virtuozzo Cluster Objects via SNMP	35
3.3.1	Enabling SNMP Access	36
3.3.2	Accessing Virtuozzo Objects via SNMP	37
3.3.2.1	Listening to SNMP Traps	37
3.3.3	Monitoring Clusters with Zabbix	38
3.3.4	Description of Cluster Objects and Traps	42
4.	Monitoring Storage Nodes	45
4.1	Node Statuses	45
4.2	Monitoring Node Performance	45
4.2.1	Monitoring Node Disks	47
4.2.1.1	Monitoring the S.M.A.R.T. Status of Node Disks	47
4.2.1.2	Slow Disk and Slow CS Analyzers	48
4.3	Monitoring Node Network	48
5.	Viewing Alerts, Audit Log, and Sending E-mail Notifications	49
5.1	Viewing Alerts	49
5.2	Viewing Audit Log	50
5.3	Sending E-mail Notifications	51
6.	Exporting Storage Cluster Data	54
6.1	Exporting Data via iSCSI	54
6.1.1	Creating iSCSI Targets	55
6.1.1.1	Performance Tips	57
6.1.2	Listing, Stopping, and Deleting iSCSI Targets	57
6.1.3	Configuring iSCSI Targets	57
6.1.3.1	Listing LUNs	58
6.1.3.2	Adding LUNs	59
6.1.3.3	Configuring LUNs	61
6.1.3.4	Deleting LUNs	61

6.1.4	Managing iSCSI Users	61
6.1.4.1	Creating CHAP Accounts for iSCSI Targets	61
6.1.4.2	Creating iSCSI Targets Bound to CHAP Accounts	62
6.1.4.3	Changing CHAP Account Passwords	63
6.2	Exporting Data via S3	64
6.2.1	Object Storage Infrastructure Overview	65
6.2.2	Planning the S3 Cluster	67
6.2.3	Sample Object Storage	67
6.2.4	Creating the S3 Cluster	69
6.2.5	Managing Object Storage Users	74
6.2.5.1	Adding S3 users	75
6.2.5.2	Managing S3 Access Key Pairs	76
6.2.6	Managing Object Storage Buckets	78
6.2.6.1	Listing Bucket Contents	78
6.2.6.2	Managing Notary in Buckets	79
6.2.7	Best Practices for Using S3 in Storage	80
6.2.7.1	Bucket and Key Naming Policies	80
6.2.7.2	Improving Performance of PUT Operations	81
6.2.8	Replicating Data Between Geographically Distributed Datacenters with S3 Clusters	81
6.2.9	Monitoring S3 Access Points	83
6.2.10	Releasing Nodes from S3 Clusters	84
6.2.11	Supported Amazon S3 Features	84
6.2.11.1	Supported Amazon S3 REST Operations	84
6.2.11.2	Supported Amazon Request Headers	86
6.2.11.3	Supported Amazon Response Headers	87
6.2.11.4	Supported Amazon Error Response Headers	88
6.2.11.5	Supported Authentication Scheme and Methods	88
6.3	Managing Datastores	89
6.3.1	Creating Datastores	89
6.3.2	Placing Virtual Environments and Backups in Datastores	91
6.3.3	Moving Virtual Environments Between Datastores	92

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).

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

1.1 About Storage

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).

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 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 or S3.

1.2 Storage Cache Architecture

The terms “cache” and “journal” are sometimes used interchangeably. In the storage cluster, however, cache refers to a fast hardware device (for example, SSD- or NVMe-based) that is used to store the chunk service journal. A journal, in its turn, is a buffer that is used by the chunk service and stored in a cache device. As multiple chunk services can share the same cache device, one cache can contain multiple journals.

As such, caching does not count as an additional storage tier in the cluster. Instead, each cache device can be associated with multiple chunk services that are assigned to different tiers, and used to store data journals.

By default, the chunk service stores its journal on the same device as its data. This configuration is called “inner cache.” In order to use a fast cache device, the chunk service must be configured to use an “external cache.”

1.2.1 Reads and Writes Behavior

In the storage cluster, cache is mainly used for writing data: when new data is ingested in the system, it is temporarily stored in the cache. As a cache device is faster than a capacity one, writing data on the cache device improves performance. For a certain amount of time, data only exists in the cache, with remote replicas on the other cluster nodes, if remote replication is configured. During this time, all read operations hit the cache and benefit from the performance boost, as well. When the cache is reclaimed and data is removed from it, all subsequent read operations are redirected to a capacity device.

The journal is used as a ring buffer: it stores data until there is a need to reclaim space and make room for new data. When this happens, data is offloaded to a capacity device in a first-in, first-out fashion (FIFO).

1.2.2 Caching Benefits

Caching helps to significantly improve write speed and write latency, with only a slight increase in the system cost. Systems with cache can benefit from high capacity of low-cost and low-performance hard disk devices (HDD), while providing fast write access by using flash devices, such as solid-state drives (SSD) or non-volatile memory devices (NVMe). Though the cost of cache devices is higher than that of capacity devices, only a few cache devices are needed, thereby making the overall system cost generally low. Moreover, the performance boost often justifies such an upgrade.

You can benefit from using cache in the following scenarios:

- “Hot” data storage
- Random writes
- Smaller block sizes or smaller files
- Databases and environments with several clients/threads

On the other hand, scenarios that usually have small advantage when using cache include:

- “Cold” data storage
- Constant throughput workloads, such as video surveillance recording
- Sequential writes of very large files
- Read-intensive workloads

In these cases, you might use an all-HDD solution, which will provide the same performance at a lower cost; or an all-flash solution, if your aim is to increase performance.

CHAPTER 2

Managing Storage

To start managing 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 the cluster, is to create the internal and public networks required by Storage. You can do that by configuring the network interfaces of all nodes. Having created the networks, you can proceed to creating Storage clusters.

2.1 Configuring Node Network Interfaces


As described in [Planning Network](#), 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.


Nodes


ADD NODE 


Search Any status ▾ Any role ▾

UNASSIGNED 4 Hide unassigned nodes ^

172.16.24.126





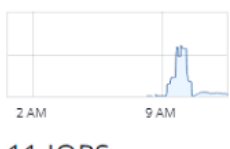
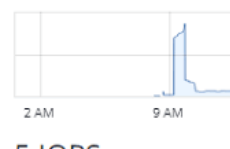


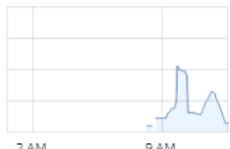
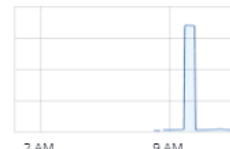
172.16.25.102


172.16.25.85


172.16.25.97


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

Nodes > 172.16.24.126 HEALTHY

CPU Cores 4 <hr/> RAM 967.83 MB	CPU Usage 4.37% 	Memory usage 70% 
DISKS >  3 hdd	Read 0.26 MB/s 	Write 0.1 MB/s 
NETWORK   172.16.24.126 1 GB	TX 0.51 KB/s 	RX 3.32 KB/s 

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

Nodes > 172.16.24.126 > Network

ADD NODE



DISKS NETWORK

<input type="checkbox"/>	Name ↓	Status	IPs		Details
<input type="checkbox"/>	ens37	UNPLUGGED			Performance
<input checked="" type="checkbox"/>	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](#).

Note: Assigning the network role **ISCSI** automatically enables high availability for virtual machines, containers, and iSCSI targets on this node. For detailed information on high availability, refer to [Managing High Availability Clusters](#).

✕ Choose roles


- Internal
 - Management
 - Storage
 - Object Storage private
- Public
 - iSCSI
 - S3 public
 - Web CP
 - SSH
- Custom (configure)
 - WebCockpit



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	
<input checked="" type="checkbox"/>	FTP data	20	
<input type="checkbox"/>	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.

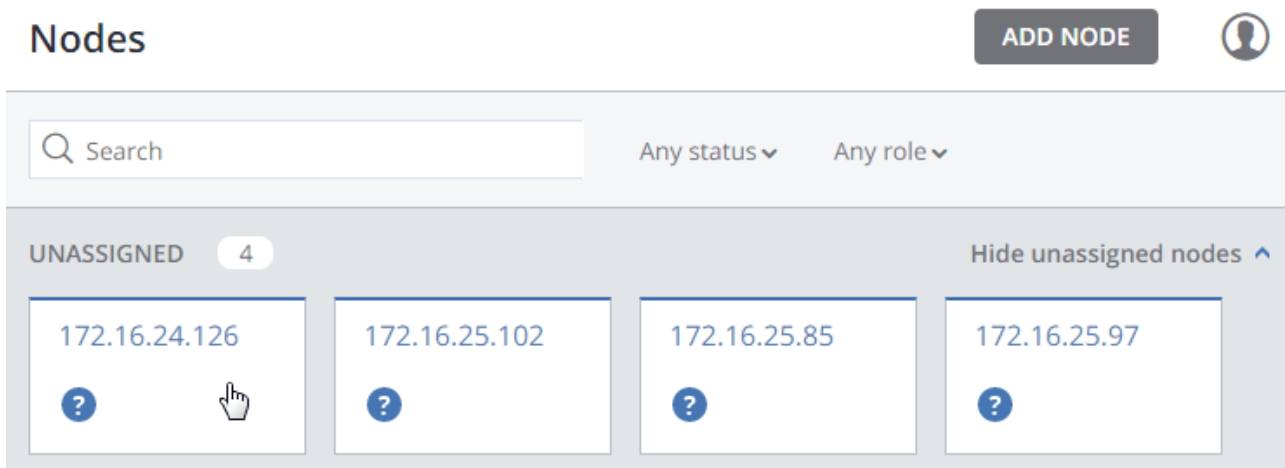
2.2 Creating the Storage Cluster

To create the Storage cluster means to create the 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.

2.2.1 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 ("-").



✕ New cluster

Cluster

cluster1

Storage interface

eth0 - 10.100.6.105  

 Encryption 

New cluster Advanced configuration

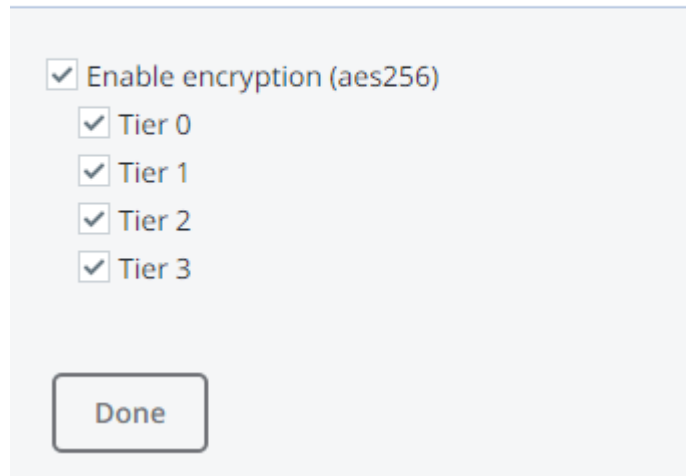
4. Make sure a configured network interface with a storage role is selected from the **Storage interface** drop-down list.

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**.

✕ Encryption Configuration



Enable encryption (aes256)

Tier 0

Tier 1

Tier 2

Tier 3

Done

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

6. Click **New cluster** to have 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 12).

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.

2.2.2 Adding Nodes to the 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.

- Click **Join cluster** to have 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 12).

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:

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

- **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, you can also select a tier from the **Tier** drop-down list. To make better use of data redundancy, do not assign all the disks on a node to the same tier. Instead, make sure that each tier is evenly distributed across the cluster with only one disk per node assigned to it. For more information, see the [Installation Guide](#).

Note: If the disk contains old data that was not placed there by Storage, the disk will not be considered suitable for use in 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 the two roles described above.
- **Unassigned.** Remove the roles from the disk.

Note the following:

- If a physical server has a system disk with the capacity greater than 100GB, that disk can be additionally assigned the **Metadata** or **Storage** role. In this case, a physical server can have at least 2 disks.
- It is recommended to assign the **System+Metadata** role to an SSD. Assigning both these roles to an HDD will result in mediocre performance suitable only for cold data (e.g., archiving).
- The **System** role cannot be combined with the **Cache** and **Metadata+Cache** roles. The reason is that is I/O generated by the operating system and applications would contend with I/O generated by journaling, negating its performance benefits.

3. Click **Done**.

4. Repeat steps 1 to 3 for every disk you want to be used in the 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	2	0	0
	Tier 1	1	0	0
	Tier 3	1	0	0

SSD CACHE

	Any tier			1
	Tier 0			1
	Tier 1			1

Unused disks 2

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, reconfigure DNS for the S3 access point to remove the node from DNS records. Next, release the node from S3 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 Storage.

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

Nodes completely removed from 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 obtained in the management panel.
- By reinstalling Storage on the node from scratch.

2.5 Managing Tier Encryption

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. Storage stores disk encryption keys in cluster's metadata (MDS).

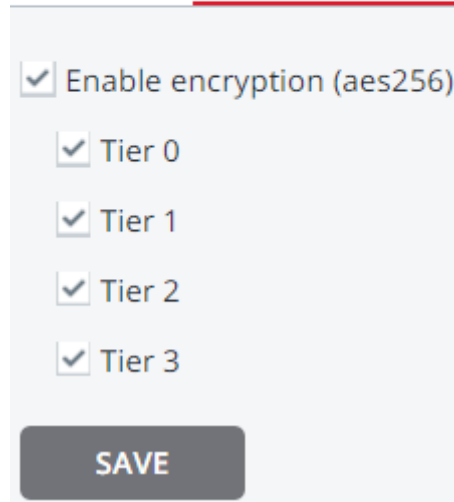
Encryption can be enabled or disabled only for the newly created chunk services (CS). Once tier encryption is enabled, you can decrypt disks (CSs) by manually releasing them from encrypted tiers. Correspondingly, simply enabling encryption on the disk's tier will not encrypt its data (CS). To encrypt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

Note the following:

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

Advanced

LDAP/AD ENCRYPTION



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 Storage Users

During the management panel installation on the first node, Storage creates the default unique administrator, superadmin. The user name for this account is root 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.

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

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

Note: All users can change their own passwords (see *Managing User Accounts* on page 20).

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**.
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

- SSH**
Can add and remove SSH keys for cluster nodes access.
- Cluster**
Can create cluster, join nodes to cluster, and manage (assign and release) disks.
- iSCSI**
Can create and manage iSCSI targets and LUNs.
- Administrator**
Can perform all management operations.
- S3**
Can create and manage S3 cluster.
- Network**
Can modify network settings and roles.

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

SELECT ROLES

- SSH**
Can add and remove SSH keys for cluster nodes access.
- Cluster**
Can create cluster, join nodes to cluster, and manage (assign and release) disks.
- iSCSI**
Can create and manage iSCSI targets and LUNs.
- Administrator**
Can perform all management operations.
- S3**
Can create and manage S3 cluster.
- Network**
Can modify network settings and roles.

2.6.3 Adding LDAP or Active Directory Users

You can add users and user groups to 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 Storage will be defined by the roles you assign in Storage (listed in *Managing Storage Users* on page 17).

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

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

The screenshot shows the 'Advanced' settings screen with the 'LDAP/AD' tab selected. The configuration fields are as follows:

Field	Value
Directory type	Microsoft Active Directory
Address	10.250.35.97
Port	389
Login	cn=admin, cn=Users, dc=ya, dc=skol, dc=
Password
Search Base DN	dc=ya, dc=skol, dc=ko

There is an unchecked checkbox labeled 'Advanced' and a 'SAVE' button at the bottom.

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;

- (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 Storage and click **Add**.

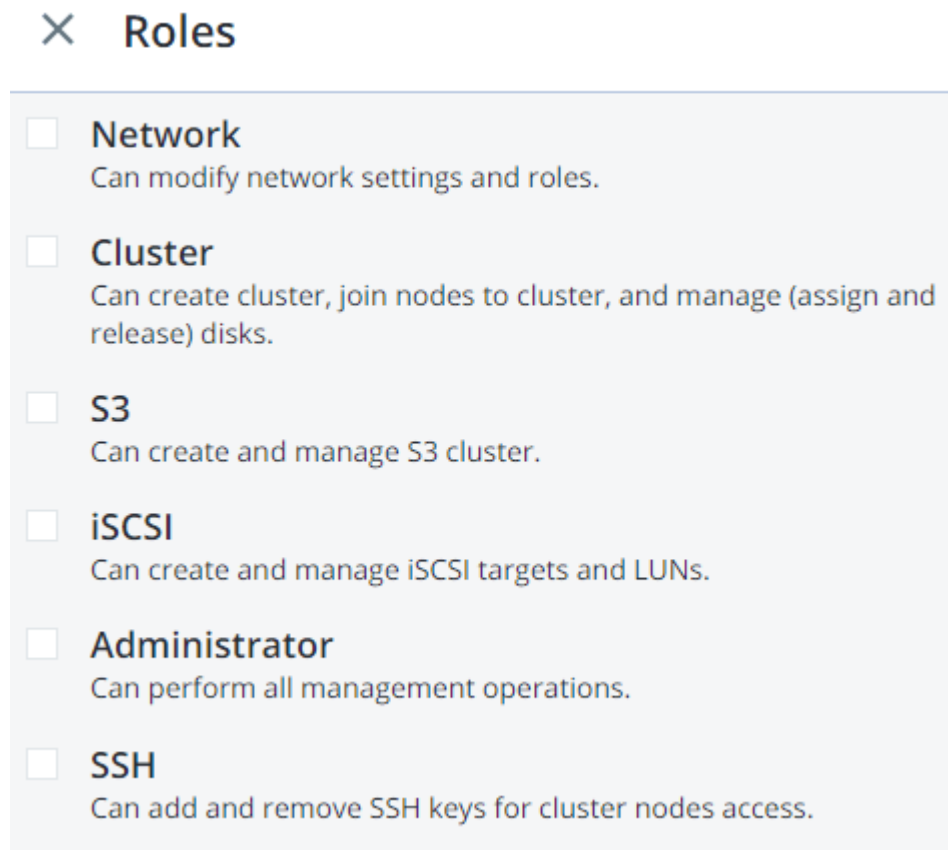
✕ Add LDAP users

<input type="checkbox"/>	Type	Name ↓	Description	
<input type="checkbox"/>		WinRMRemoteWMIUs...	Members of this gr...	
<input checked="" type="checkbox"/>		Administrators	Administrators hav...	
<input type="checkbox"/>		Users	Users are prevente...	
<input type="checkbox"/>		Guests	Guests have the sa...	
<input type="checkbox"/>		Print Operators	Members can admi...	
<input type="checkbox"/>		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.








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

Users ADD USER ADD LDAP USER 

 Configure

 Delete

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

2.7 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 > Security > 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/elFmb1zuUFijQJKprX88XypNDvjYN
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**.

2.8 Backing Up and Restoring the Management Database

In 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 > Management node > Backup** screen and click **BACKUP NOW**.

Backup OK

System stores node information, statistics and configuration in a database. Its backups are created automatically.

[How to recover?](#)

↻ Last backup
 Mar 09, 2017, 11:32 AM

🕒 Backup period
 Daily

📍 Backup location
 /mnt/vstorage/webcp/backup/

BACKUP NOW

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

2.8.1 Restoring the 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.

2.9 Accessing the Management Panel via SSL

When configuring various Storage features, you may need to enter sensitive information like credentials for user and e-mail accounts, S3 services, and such. To secure communication with the management panel, you can switch to the HTTPS protocol as follows:

1. On the **SETTINGS > Management node > SSL ACCESS** tab, click **UPLOAD**.
2. Upload an SSL certificate from a trusted certificate authority.
3. Click **SAVE**.

The uploaded certificate will be added to the configuration of the web server hosting the management panel and you will be able to access it over HTTPS.

You can also generate a self-signed certificate, although it will not be trusted and you will have to manually accept it in your browser.

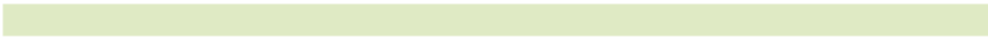
2.10 Managing Storage Licenses

To activate or upgrade a 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.

Licenses



Virtuozzo Storage	Expires on 08/06/2017
383 B of 9.54 PB licensed space used	
	
Register key	

Click **Activate**.

CHAPTER 3

Monitoring Storage Clusters

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

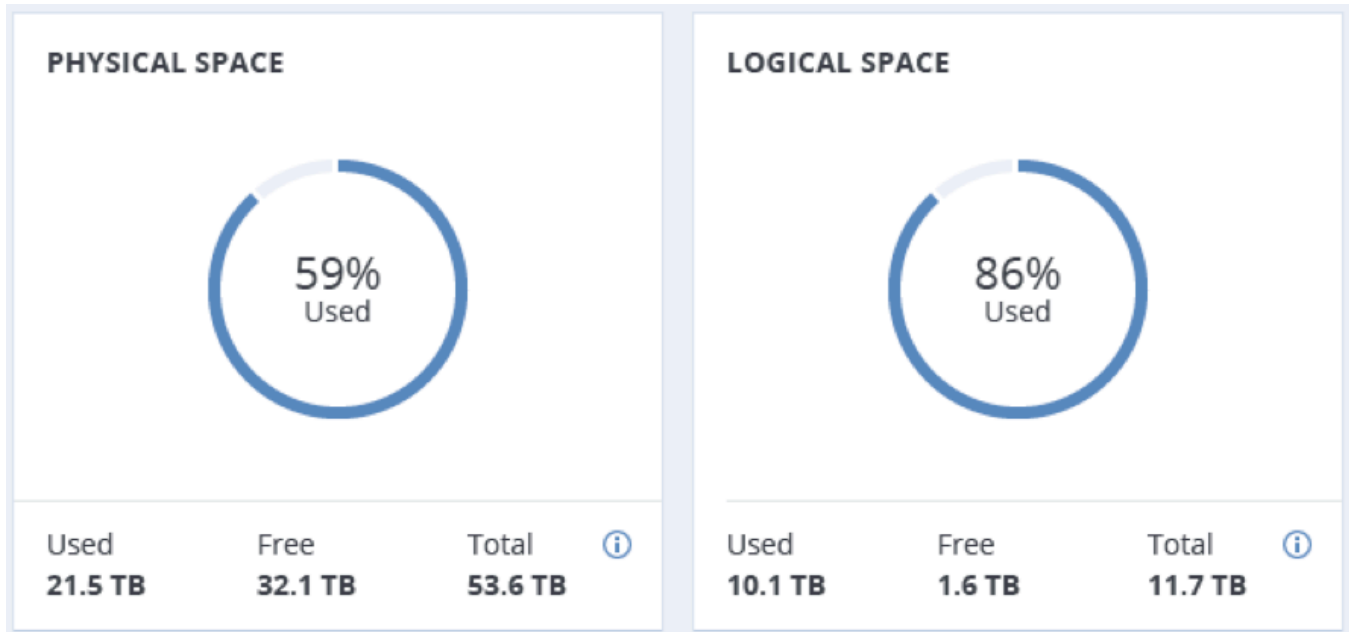
3.1 Monitoring Cluster Status

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:



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.

Continued on next page

Table 3.2.3.1 -- continued from previous page

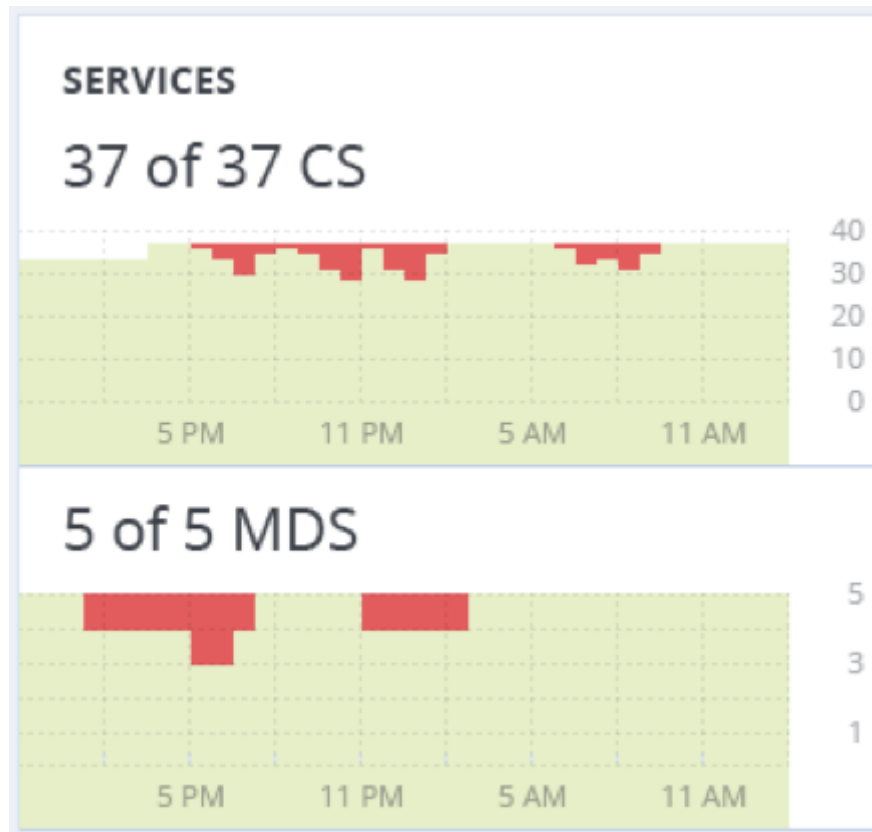
State	Description
offline	<p>Percentage of chunks all replicas of which are offline. Such chunks are completely inaccessible for the cluster and cannot be replicated, read from or written to. All requests to an offline chunk are frozen until a CS that stores that chunk's replica goes online.</p> <p>Get offline cluster nodes back online as soon as possible to avoid data loss.</p>
blocked	<p>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.</p> <p>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.</p>
degraded	<p>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.</p>

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.

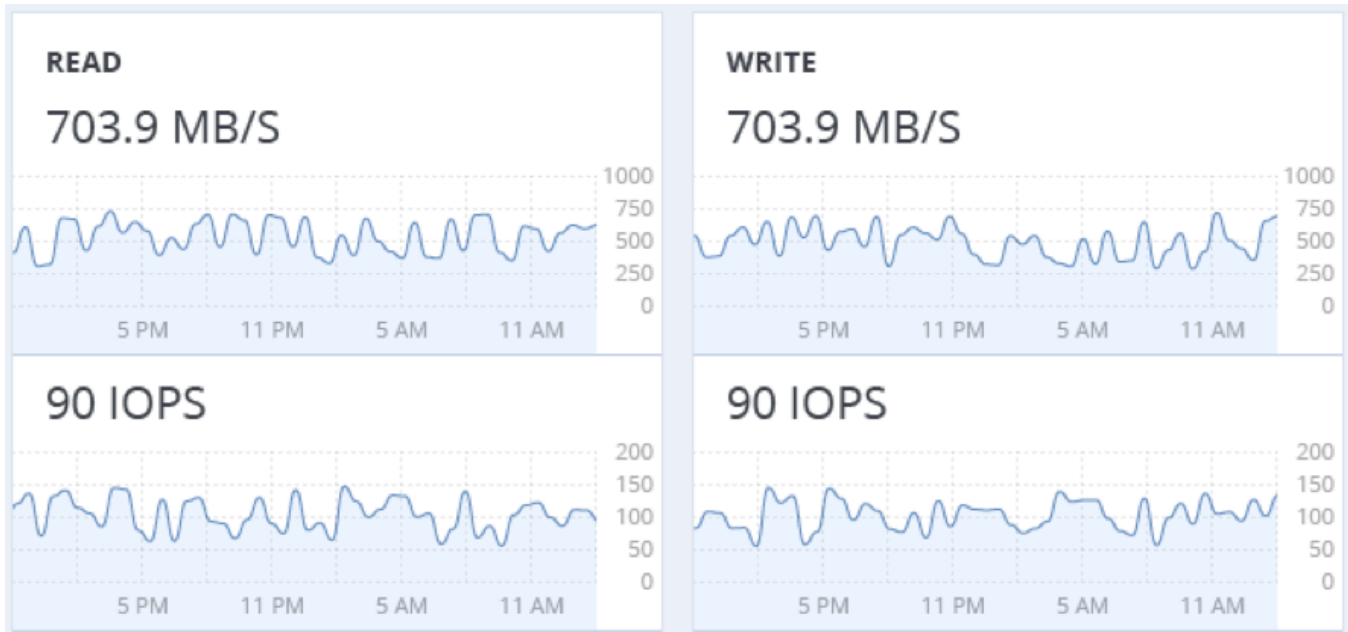
Typical statistics may look like this:



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:



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** EMAIL NOTIFICATIONS

Enable SNMP on the management node

You can access SNMP on the management node on port 161
MIB file can be [downloaded here](#)
Zabbix template can be [downloaded here](#)

Send SNMP traps to this network management system

IP address

Port

Community

SEND TEST TRAP

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

You can access Virtuozzo 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**.

Import file vstorage.xml

Rules	Update existing	Create new	Delete missing
Groups		<input checked="" type="checkbox"/>	
Hosts	<input type="checkbox"/>	<input type="checkbox"/>	
Templates	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
Template screens	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Template linkage		<input checked="" type="checkbox"/>	
Applications		<input checked="" type="checkbox"/>	<input type="checkbox"/>
Items	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Discovery rules	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Triggers	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Graphs	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Web scenarios	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Screens	<input type="checkbox"/>	<input type="checkbox"/>	
Maps	<input type="checkbox"/>	<input type="checkbox"/>	
Images	<input type="checkbox"/>	<input type="checkbox"/>	
Value mappings	<input type="checkbox"/>	<input checked="" type="checkbox"/>	

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

Host name

Visible name

Groups

In groups

Other groups

- Discovered hosts
- Hypervisors
- Linux servers
- Templates
- Virtual machines
- Zabbix servers

New group

Agent interfaces

IP address	DNS name	Connect to	Port	Default
Add				

SNMP interfaces

IP address	DNS name	Port	Remove
<input type="text" value="10.250.14.15"/>	<input type="text"/>	<input type="text" value="161"/>	<input type="button" value="Remove"/>

Use bulk requests

[Add](#)

JMX interfaces [Add](#)

IPMI interfaces [Add](#)

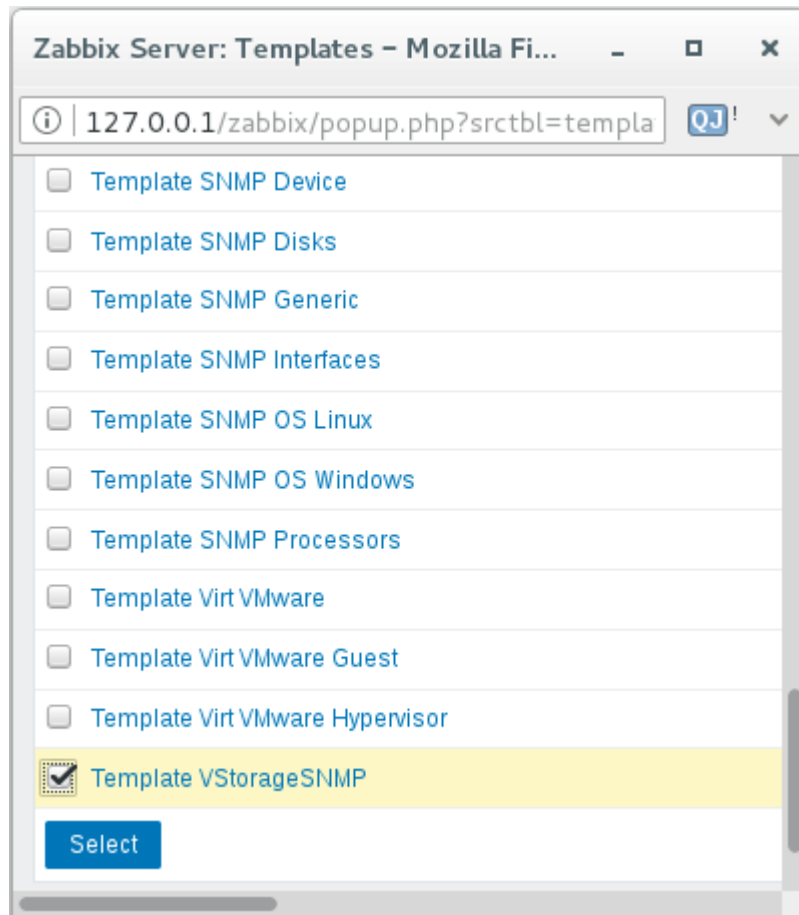
Description

Monitored by proxy

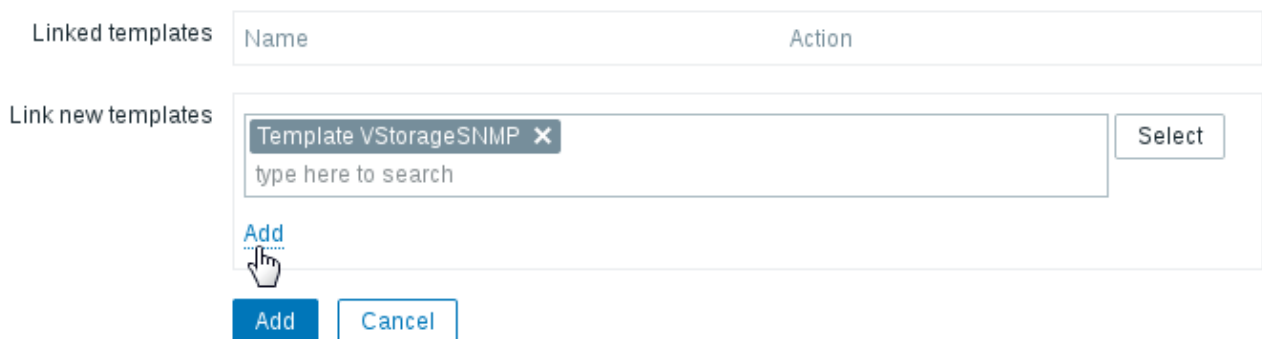
Enabled

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.

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

type here to search

[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.

Continued on next page

Table 3.3.4.1 -- continued from previous page

Object	Description
VSTORAGE-MIB::licenseStatus	License status.
VSTORAGE-MIB::licenseCapacity	The maximum disk space available as defined by license.
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 the time spent waiting for synchronization.

Continued on next page

Table 3.3.4.2 -- continued from previous page

Object	Description
VSTORAGE-MIB::csloReadS	Current read speed of a chunk server in bytes per second.
VSTORAGE-MIB::csloWriteS	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.
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 Storage Nodes

Nodes added to the Storage infrastructure are listed on the **NODES** screen, grouped by their statuses. If there are no clusters in 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:

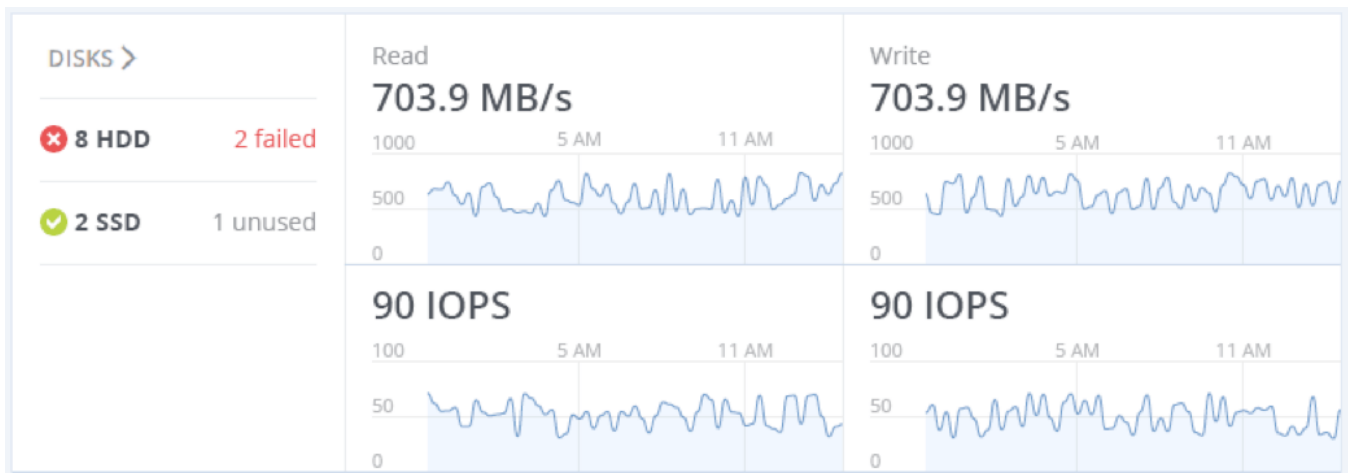
- The number of CPUs and the amount of RAM

- 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



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 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

4.2.1.2 Slow Disk and Slow CS Analyzers

Slow disk and CS analyzers calculate disk health according to the average I/O latency over time (15 minutes).

The following table shows the default thresholds:

Analyzer	OK latency, in seconds	Fatal latency, in seconds
Slow CS	0.1	60
Slow HDD Disk	0.1	30
Slow SSD Disk	0.01	10

If disk latency is less than **OK latency**, the disk health is considered to be 100%. If disk latency exceeds **FATAL latency**, the disk health is considered to be 0%. Disk latency that lies within these two thresholds will vary linearly from 100% to 0%.

When disk health becomes 0%, the service generates an alert and marks this CS as unresponsive. Such a CS does not trigger automatic replication but is no longer available for chunk allocation.

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 Alerts, Audit Log, and Sending E-mail Notifications

This chapter describes Storage alerts and audit log and e-mail notifications settings.

5.1 Viewing Alerts

The **ALERTS** tab lists all the alerts logged by 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.
- Licensed storage capacity is low or critically low.
- Cluster is about to or has run out of available space.
- CS journal is about to run out of available space.
- Metadata disk has run out of available space.
- Disk has run out of available space.

Alerts



<input type="text" value="Search"/>				
<input type="checkbox"/>	Type	Message	Date and time	
<input type="checkbox"/>		Node 172.16.25.85 is offline	Jul 19, 2017, 10:11 PM	
<input type="checkbox"/>		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.

Audit log 

Search  Show more details

Date and time ↓	User	Activity	Details	
Aug 19, 2022, 2:36 PM	root	User login	User "root" has logged in	
Aug 19, 2022, 2:19 PM	root	Register cluster license	Register cluster license	
Aug 19, 2022, 2:19 PM	root	User login	User "root" has logged in	
Jul 19, 2022, 10:41 AM	root	User login	User "root" has logged in	
Jul 14, 2022, 10:04 PM	root	Start backup process	Start management panel backu...	
Jul 14, 2022, 10:03 PM	root	Create datastore	Create datastore "dstore1replicas"	
Jul 14, 2022, 10:03 PM	root	Create datastore	Create datastore "dstore2replicas"	

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

5.3 Sending E-mail Notifications

Storage can send automatic e-mail notifications about errors, warnings, and alerts.

To set up e-mail notifications, do the following:

1. On the **SETTINGS > Advanced settings > EMAIL NOTIFICATIONS** tab, specify the following information:
 - 1.1. In the **From** and **Sender name** fields, the notification sender's e-mail and name.
 - 1.2. In the **To** field, one or more notification recipient e-mails, one per line.
 - 1.3. In the **User account** fields, the credentials of the notification sender registered on the SMTP server.
 - 1.4. In the **Outgoing SMTP server** field, the DNS name of the SMTP server, either public (e.g., smtp.gmail.com) or the one in your organization.

Note: The management node must be able to access the SMTP server.

1.5. If required, a custom **SMTP port** the server uses.

1.6. In the **Security** field, the security protocol of the SMTP server.

Advanced

LDAP/AD ENCRYPTION SNMP **EMAIL NOTIFICATIONS**

From

Sender name

User account

Outgoing SMTP server

SMTP port

Security

To

Send notifications about

- Errors
- Warnings
- Information

2. Tick the checkboxes for alerts you want to get notified about.

3. Click **SAVE**.

To send a test e-mail, specify your e-mail registered on the SMTP server in both the **From** and **To** fields and click **TEST**.

CHAPTER 6

Exporting Storage Cluster Data

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 datastore for Virtuozzo virtual machines, containers, and backups.

6.1 Exporting Data via iSCSI

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 Storage, you can create and run multiple iSCSI targets per 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 Storage disk space over iSCSI.



In this example there are three hardware nodes working in a 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 iSCSI Targets

Prerequisites:

- Each iSCSI target must be assigned at least one unique IP address from DC network's static pool.
- The name of each iSCSI target must be unique in the Storage cluster.
- 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

Node

 ▼

Enable CHAP

 ▼

Target portal IP address

172.16.24.125

+ Add − Remove

Enable limits

IOPS: − + **Bandwidth (MB/s):** − +

3. In the node drop-down list, select a node on which the target will be located. The node must have the **iSCSI** role assigned to one of its network interfaces to appear in the list.

Note: Assigning the network role **iSCSI** automatically enables high availability for virtual machines, containers, and iSCSI targets on this node. For detailed information on high availability, refer to

Managing High Availability Clusters.

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 61).
5. Click **Add** to specify one or more IP addresses for the target.
6. If necessary, enable and specify IOPS and bandwidth limits for the target. If both limits are set, the first one that is hit is applied. Setting a limit value to zero disables the limit.
7. 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 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 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

Node

Enable CHAP

3. Click **Done**.

6.1.3.1 Listing LUNs

In 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** 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			
Q 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

1 ▼

LUN size

- 20 + GB Tier 0 ▼ ✓

Data redundancy:

Erasure coding

Replication

Failure domain:

Disk ▼

Encoding 1+0	0% overhead	
Encoding 1+2	200% 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

- 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](#).
- From the **Failure domain** drop-down list, choose a placement policy for replicas. For more details, see [Understanding Failure Domains](#).

6. Choose a data redundancy mode. For more details, see [Understanding Data Redundancy](#).
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 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

Password

Description

Enabled

Done

3. Click **Done**.

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

6.1.4.2 Creating iSCSI Targets Bound to CHAP Accounts

To create a 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

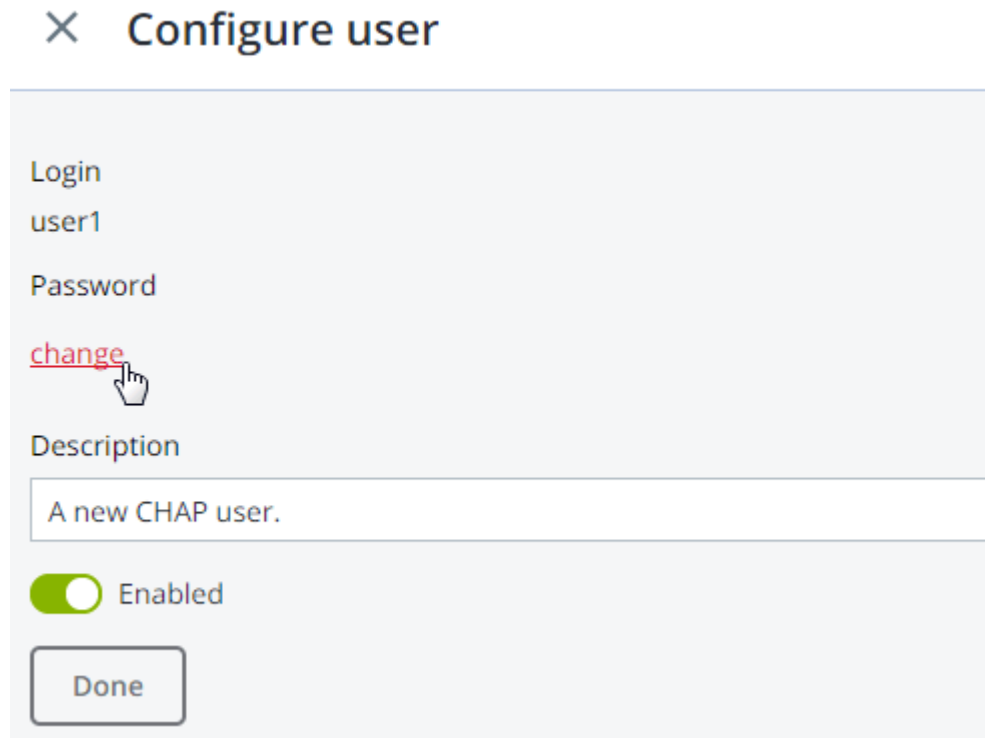
Done

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

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

Storage is implemented as an Amazon S3-like API, which is one of the most common object storage APIs. End users can work with 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 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, 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.

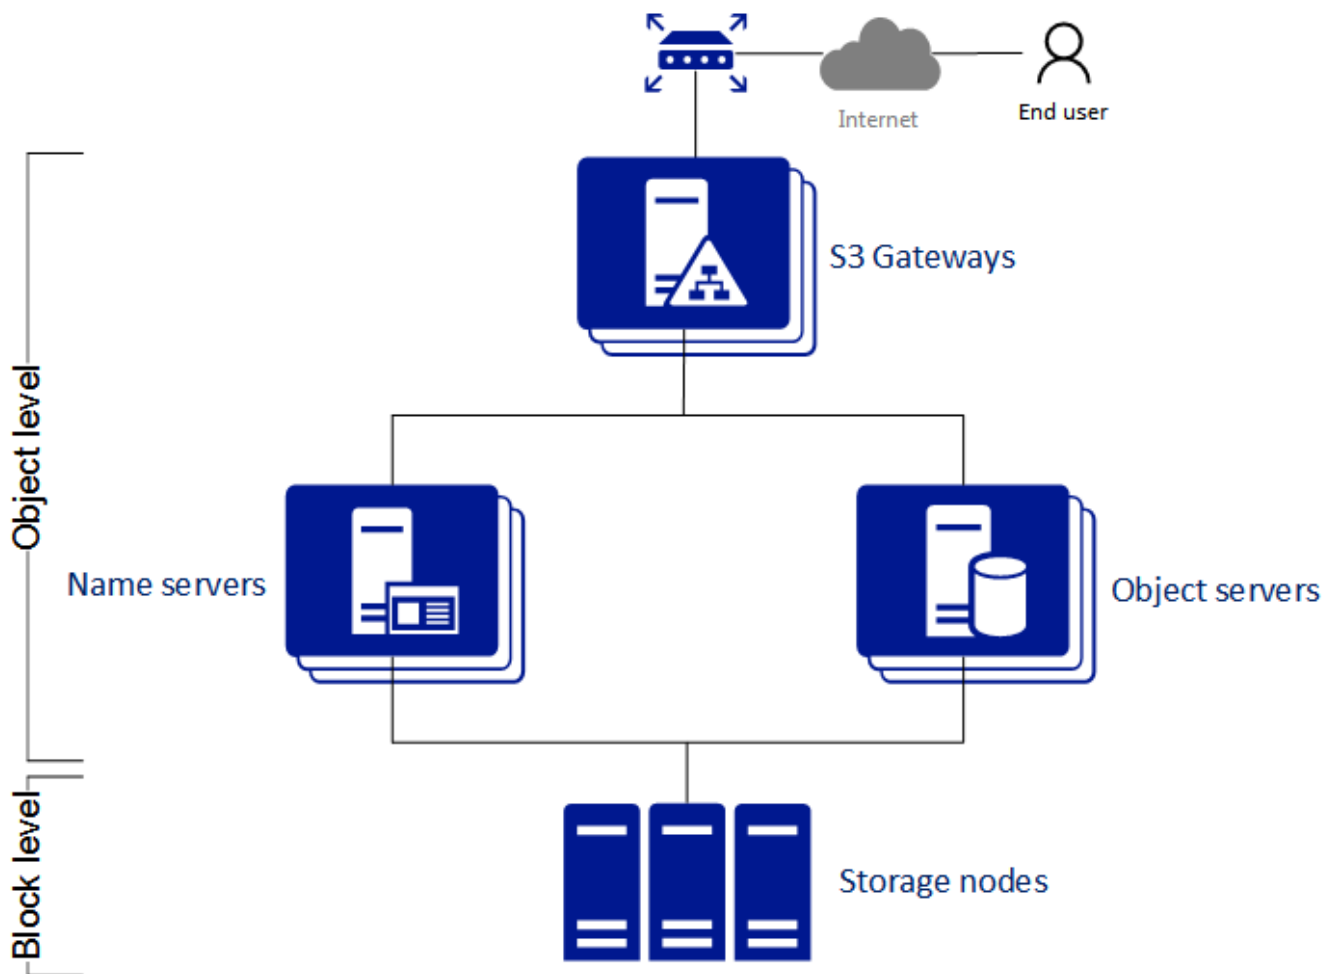
Compared to other types of storage, the key difference of object storage is that parts of an object cannot be modified, so if the object changes a new version of it is spawned instead. This approach is extremely important for maintaining data availability and consistency. First of all, changing an object 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 Storage nodes. Each service should be deployed on multiple 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:

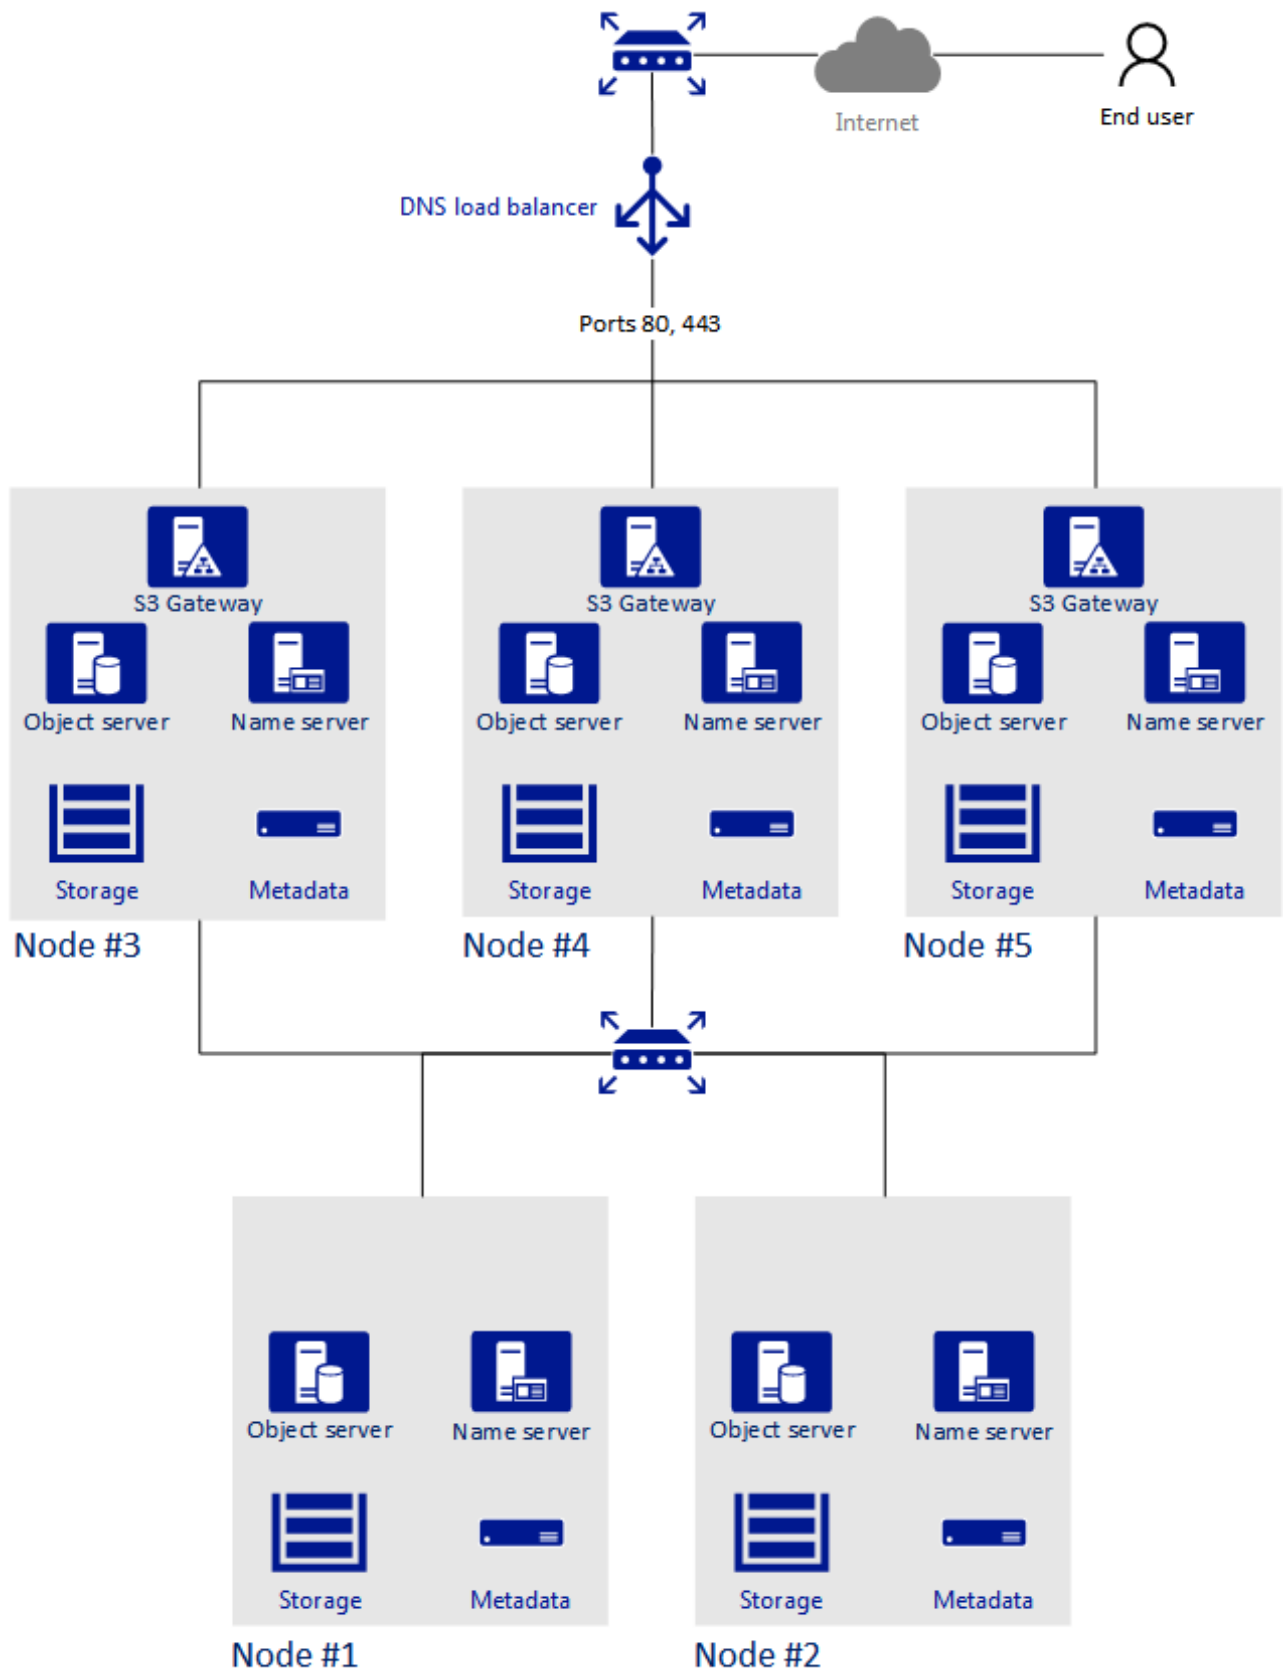
1. Define which nodes of the Storage cluster will run the S3 storage access point services. It is recommended to have all nodes available in Storage run these services.
2. Configure the network so that the following is achieved:
 - All components of the S3 cluster communicate with each other via the S3 private network. All nodes of an S3 cluster must be connected to the S3 private network. 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 a 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

Note: Joining a node to an S3 cluster automatically enables high availability for virtual machines, containers, and S3 resources on this node. For detailed information on high availability, refer to [Managing High Availability Clusters](#).

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 > S3 > Nodes** screen, check the box of each cluster node where object storage services will run.

S3 Nodes 

OVERVIEW **NODES** USERS BUCKETS Create S3 cluster

Search Any status







AVAILABLE NODES 4 Hide available nodes

172.16.24.126 <input checked="" type="checkbox"/>	172.16.25.102 <input checked="" type="checkbox"/>
172.16.25.85 <input checked="" type="checkbox"/>	172.16.25.97 <input checked="" type="checkbox"/>

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	S3 public
<input type="text" value="eno16777736 - 172.16.24.1:"/> 	
<hr/>	
 172.16.25.102 	
S3 private	S3 public
<input type="text" value="eno16777736 - 172.16.25.102"/> 	<input type="text" value="eno16777736 - 172.16.25.102"/>
<hr/>	
<input type="button" value="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	i
Encoding 7+2	29% overhead	i
Encoding 17+3	18% overhead	i

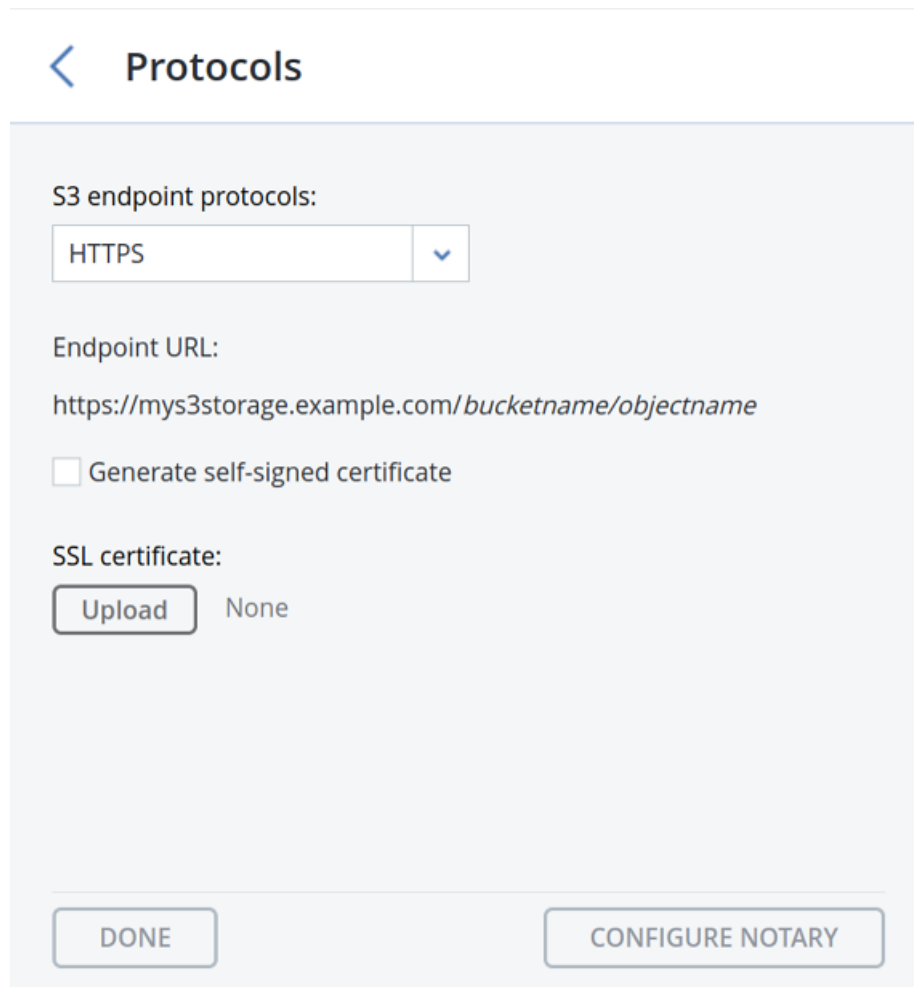
Proceed

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.



The screenshot shows a configuration panel titled "Protocols" with a back arrow on the left. The panel contains the following elements:

- S3 endpoint protocols:** A dropdown menu with "HTTPS" selected and a downward arrow on the right.
- Endpoint URL:** A text input field containing the URL `https://mys3storage.example.com/bucketname/objectname`.
- Generate self-signed certificate:** An unchecked checkbox.
- SSL certificate:** A section with a rounded button labeled "Upload" and the text "None" next to it.
- Bottom navigation:** Two rounded buttons: "DONE" on the left and "CONFIGURE NOTARY" on the right.

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 purposes.

Note the following:

- S3 geo-replication requires a certificate from a trusted authority. It does not work with

self-signed certificates.

- To access the data in the S3 cluster via a browser, add the self-signed certificate to browser's exceptions.
- Acquire a key and a trusted wildcard SSL certificate for endpoint's bottom-level domain. For example, the endpoint `s3.storage.example.com` would need a wildcard certificate for `*.s3.storage.example.com` with the subject alternative name `s3.storage.example.com`.

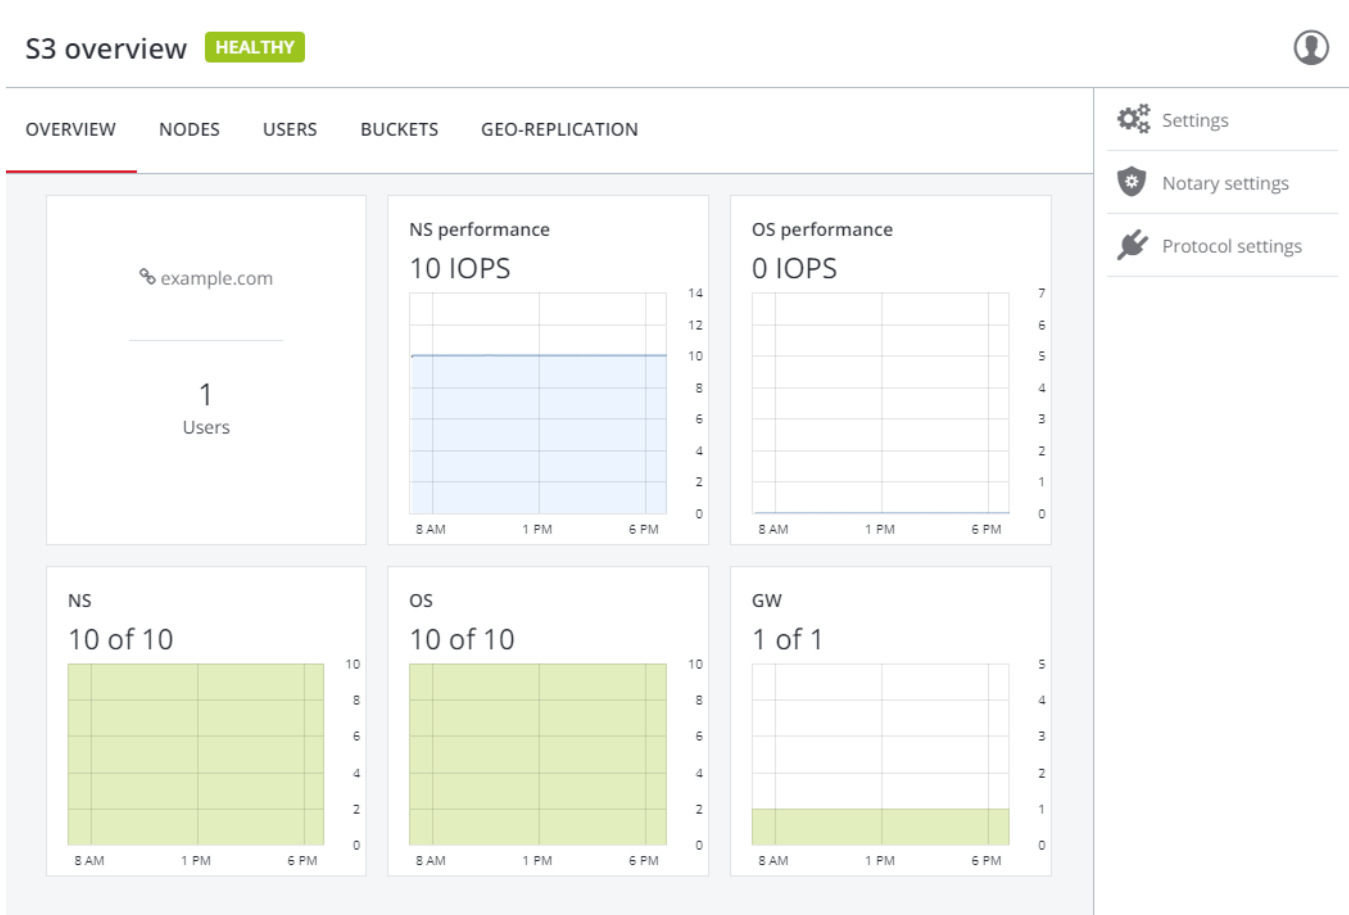
Upload the certificate, and, depending on the certificate type, do one of the following:

- If the certificate is contained in a PKCS#12 file, specify the passphrase.
- Upload the SSL key.

12. If required, click **Configure Notary** and specify **Notary DNS name** and **Notary user key**. For more information on Notary, see **Managing 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 Storage comply with the Amazon S3 user model and access policies.

User management scenarios in 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**.



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

✕ Add user

Login

user1@example.com

Enabled

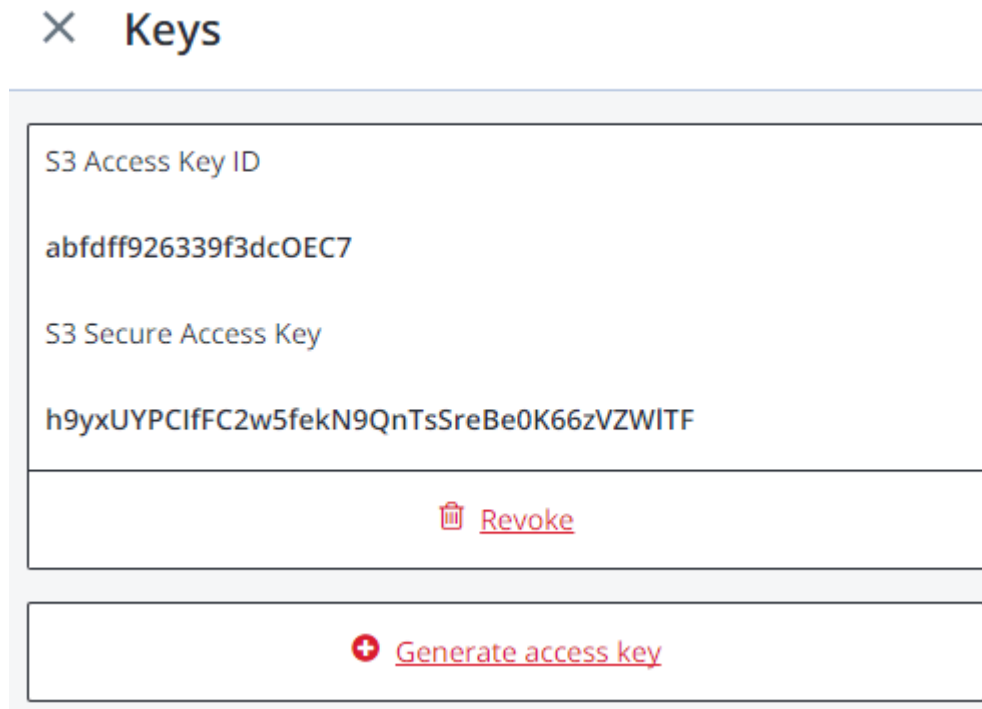
Done

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 the [Amazon documentation](#).) The access keys are generated and stored locally in the 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**.

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 Storage, you can enable and disable 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 Storage management panel. However, you can do it via the 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.

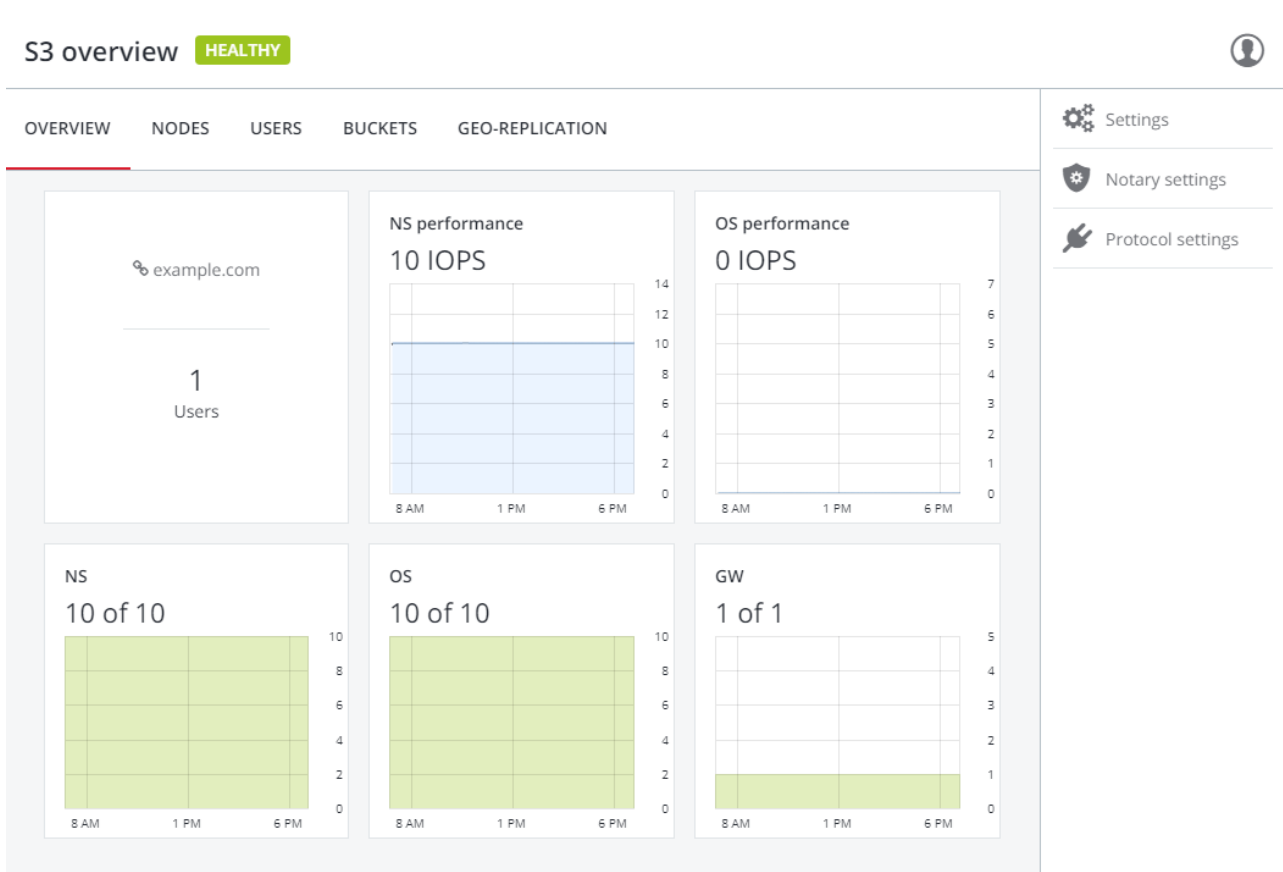
6.2.6.2 Managing Notary in Buckets

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.

6.2.6.2.1 Setting Up Notary

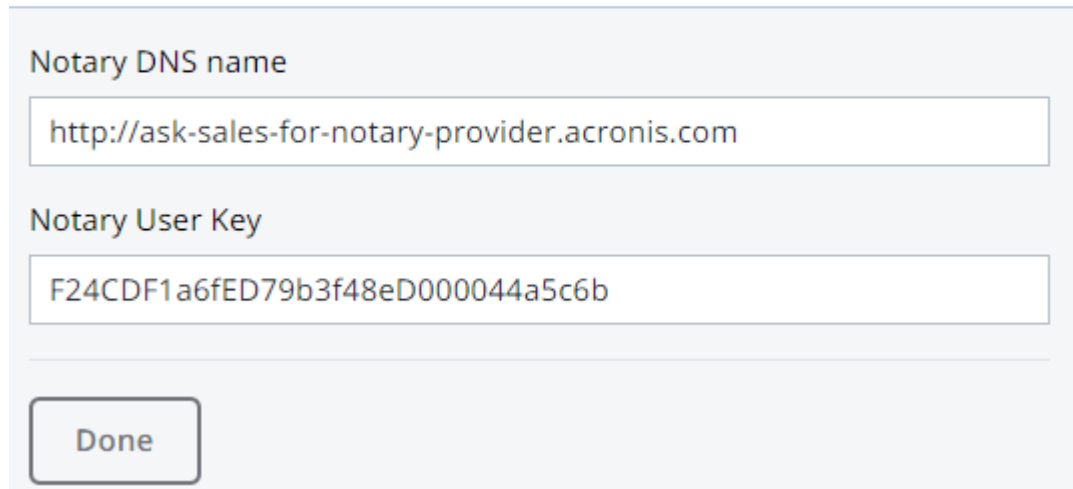
To set up Acronis Notary, do the following:

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

Notary User Key

Done

6.2.6.2.2 Enabling and Disabling 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 Storage

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

6.2.7.1 Bucket and Key Naming Policies

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

- 3 to 63 characters long.

- Start and end with a lowercase letter or number.
- Contain only 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

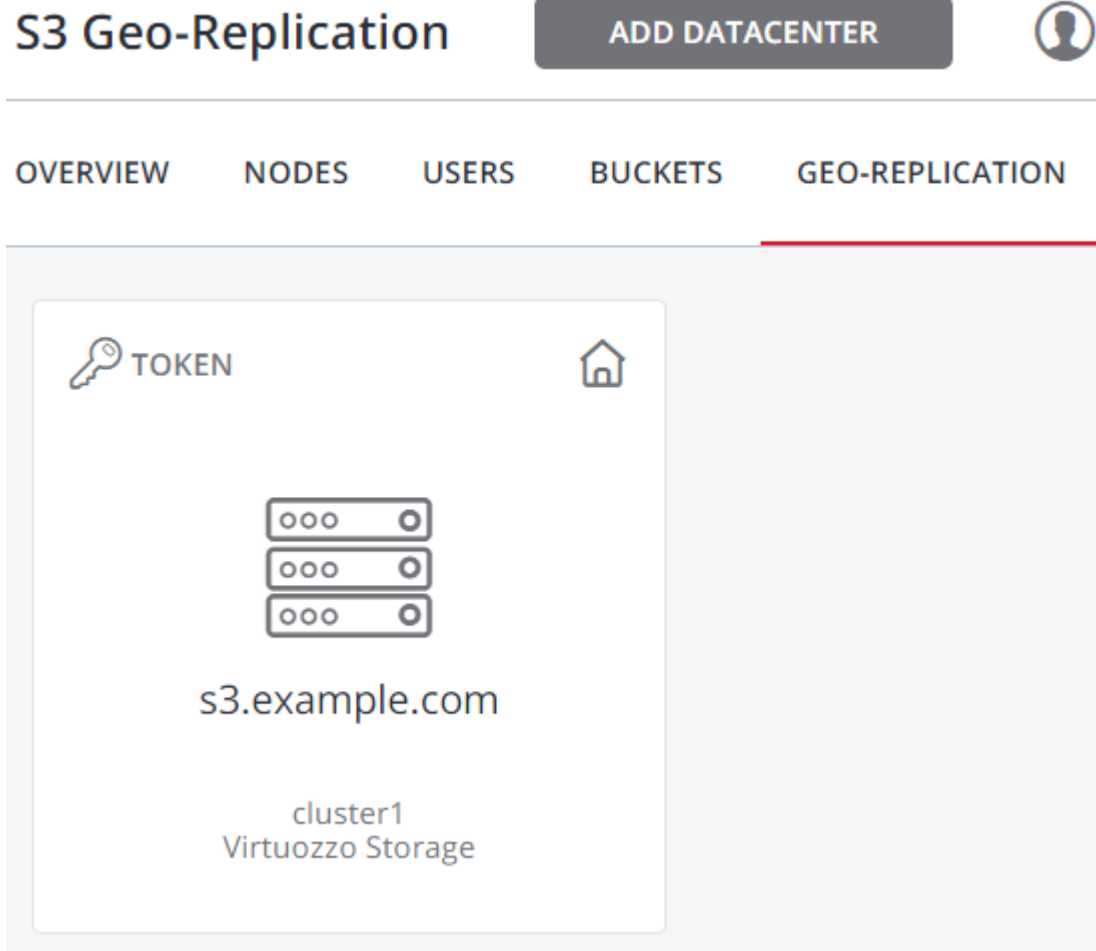
Storage can store replicas of S3 cluster data and keep them up-to-date in multiple geographically distributed datacenters with S3 clusters based on Storage. Geo-replication reduces the response time for local S3 users accessing the data in a remote S3 cluster or remote S3 users accessing the data in 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.



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
FkYWJsZV9uYW1lIjogImNsdXN0ZXIxlwglVzZXJfc2VjcmV0X2tleSI6ICJk
cDI1eUwxWEpLc3BGcEN2U3BCd2IDS1BydkpxVINtRncwakpPWW45li
wglVybCI6ICJodHRwczovL3MzMmV4YW1wbGUuY29tOjQ0MyIsICJpc1
9zZWxmljogdHJ1ZSwglVpZCI6IClxOWI4N2M0YTRkOGJkOTIhIn0=
```

ADD

4. Enter the copied token and click **Done**.
5. Configure the remote 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.

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 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 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

Continued on next page

Table 6.2.11.1.1 -- continued from previous page

Operation	Supported
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
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

Continued on next page

Table 6.2.11.1.2 -- continued from previous page

Operation	Supported
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

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 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 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 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

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 Storage implementation of the Amazon S3 protocol:

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

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

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

6.3 Managing Datastores

To keep your virtual machines and containers in Storage, you can place them in datastores. A datastore is a directory in a 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!

6.3.1 Creating Datastores

To create a datastore, do the following:

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

✕ Add Datastore

Name

Tier:
 ▼

Data redundancy: Erasure coding Replication

Failure-domain:
 ▼

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.

2. On the **Add Datastore** panel, specify a name.
3. In **Tier**, select the storage tier that will be used for the object storage (explained in [Understanding Storage Tiers](#)).
4. In **Failure domain**, choose a placement policy for replicas (explained in [Understanding Failure Domains](#)).
5. In **Data redundancy**, select the redundancy mode that the datastore will use (explained in

Understanding Data Redundancy).

The general rule is as follows:

- choose replication for highly loaded VMs, Windows VMs, and other workloads that generate lots of IOPS;
- choose erasure coding for lightly loaded Linux VMs and backups.

6. Click **Done**.

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

6.3.2 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 [Configuring Virtual Environments Host Settings](#)) 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:

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

2.3. Backup files:

```
# 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
```

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

6.3.3 Moving Virtual Environments Between Datastores

In Storage, you can move containers and virtual machines, online or offline, between datastores using the `prlctl move` command. For example, to move the virtual machine `MyVM` to the datastore `vms_new`, run:

```
# prlctl move MyVM --dst=/mnt/vstorage/vols/datastores/vms_new
```

Take note of the following:

- For large virtual machines the process may take considerable time as it implies copying all files of a virtual environment.
- After moving a virtual machine to another datastore, you cannot revert to its old snapshots.
- Destination datastore's redundancy mode is not applied automatically to the moved container. For example, a VE moved from a datastore with 3:2 replication to one with a different replication or encoding scheme retains its original 3:2 replication setting.

If the destination datastore does not exist, it will be automatically created with the default replication parameters. You can then configure new datastore's redundancy mode, tier, and failure domain in the management panel.

To check that the virtual environment files have been successfully moved, use the `prlctl list <CT|VM_name> -i | grep "Home"` command. For example:

```
# prlctl list MyVM -i | grep "Home"
Home: /mnt/vstorage/vols/datastores/vms_new/d35d28e5-11f7-4b3f-9065-8fef6178bc5b/
```