



Virtuozzo Hybrid Infrastructure 4.7

Storage User Guide

12/6/2021

Table of contents

Supported storage types	3
Accessing S3 buckets	4
Managing buckets via the Virtuozzo Hybrid Infrastructure user panel	4
Logging in to the user panel	4
Adding, deleting, and listing S3 buckets	5
Creating, deleting, and listing folders	5
Uploading and downloading files	6
Obtaining and validating file certificates	6
Accessing S3 storage with CyberDuck	7
Managing S3 bucket versions	7
Mounting S3 storage with Mountain Duck	8
Creating S3 buckets on Mounted S3 Storage	9
S3 bucket and key naming policies	10
Accessing iSCSI targets	11
Accessing iSCSI targets from VMware ESXi	11
Accessing iSCSI targets from Linux	12
Accessing iSCSI targets from Microsoft Hyper-V	14
Accessing NFS shares	21
Mounting NFS exports on Linux	21
Mounting NFS exports on macOS	21

Supported storage types

Your service provider can configure Virtuozzo Hybrid Infrastructure to keep your data in three storage types:

- S3 object storage for storing an unlimited number of objects (files).
- iSCSI block storage for virtualization, databases, and other needs.
- NFS shares for storing an unlimited number of files via a distributed filesystem.

The following sections describe the ways to access data in Virtuozzo Hybrid Infrastructure in detail.

Accessing S3 buckets

To access S3 buckets, get the following information (credentials) from your system administrator:

- User panel IP address
- DNS name of the S3 endpoint
- Access key ID
- Secret access key

Virtuozzo Hybrid Infrastructure allows you to access your S3 data in several ways:

- Via the Virtuozzo Hybrid Infrastructure user panel
- Via a third-party S3 application like Cyberduck, Mountain Duck, etc.

Managing buckets via the Virtuozzo Hybrid Infrastructure user panel

This section describes how to manage buckets and their contents from the Virtuozzo Hybrid Infrastructure user panel.

Logging in to the user panel

To log in to the Virtuozzo Hybrid Infrastructure user panel, do the following:

1. On any computer with access to the web interface, in a web browser visit `http://<user_panel_IP_address>:8888/s3/`.

Note

If you use a self-signed certificate, add it to the browser's exceptions.

Log in

ENDPOINT

s3.example.com

Use secure transfer (SSL)

ACCESS KEY ID

d9fde6a530879f59HB8U

SECRET ACCESS KEY

.....

LOG IN

2. On the login screen, enter your credentials, and then click **Log in**.


Once you log in to the web interface, you will see the **Buckets** screen with the list of your buckets. From here, you can manage buckets, as well as folders and files stored inside the buckets.

To log out, click the user icon in the upper right corner of any screen, and then click **Log out**.

Adding, deleting, and listing S3 buckets

On the **Buckets** screen:

- To add a new bucket, click **Add bucket**, specify a name, and click **Add**.



The image shows a modal dialog box titled "Add bucket". It features a close button (an 'X' icon) in the top-left corner. The main content area contains a text input field with the label "Bucket name" and the text "bucket1" entered. At the bottom right of the dialog, there are two buttons: "Add" and "Cancel".

Use bucket names that comply with DNS naming conventions. For more information on bucket naming, refer to "S3 bucket and key naming policies" (p. 10).

- To delete a bucket, select it, and then click **Delete**.
- To list the bucket contents, click the bucket name on the list.

Listing S3 bucket contents in a browser

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, **s3.example.com/mybucket**.

Note

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

Creating, deleting, and listing folders

On the bucket contents screen:

- To create a folder, click **New folder**, specify the folder name in the **New folder** window, and then click **Add**.

✕ New folder

Folder name

Add Cancel

- To delete a folder, select it, and then click **Delete**.
- To list the folder contents, click the folder name.

Uploading and downloading files


On the bucket or folder contents screen:

- To upload files to S3, click **Upload**, and then choose files to upload.

Buckets > bucket1 👤

Type	Name	Size	Last modified	
				📁 New folder
📄	file1.test	5.3 MB	Dec 07 18:35	📄 Download file
📁	folder1/			📁 Upload file
				⚙️ Get certificate
				✅ Verify

Uploading file 1 of 1



file2.test 103.8 MB

- To download files, select them, and then click **Download**.

Obtaining and validating file certificates

Virtuozzo Hybrid Infrastructure offers integration with the Acronis Notary service to leverage blockchain notarization and ensure the immutability of data saved in S3 buckets.

To certify files stored in your buckets, ask your system administrator to enable the Acronis Notary service for the buckets.


After that, you will be able to do the following:

- To get a notarization certificate for a file, select it, and then click **Get Certificate**.
- To check the validity of a file's certificate, click **Verify**.

Accessing S3 storage with CyberDuck


To access Virtuozzo Hybrid Infrastructure with CyberDuck, do the following:

1. In CyberDuck, click **Open Connection**.
2. Specify your credentials:
 - The DNS name of the S3 endpoint.
 - The **Access Key ID** and the **Password**, the secret access key of an object storage user.



By default, the connection is established over HTTPS. To use CyberDuck over HTTP, you must install a special [S3 profile](#).

3. Once the connection is established, click **File > New Folder** to create a bucket.




4. Specify a name for the new bucket, and then click **Create**. Use bucket names that comply with DNS naming conventions. For more information on bucket naming, refer to "S3 bucket and key naming policies" (p. 10).

The new bucket will appear in CyberDuck. You can manage it and its contents.

Managing S3 bucket versions

Versioning is a way of keeping multiple variants of an object in the same bucket. You can use versioning to preserve, retrieve, and restore every version of every object stored in your S3 bucket. With versioning, you can easily recover from both unintended user actions and application failures. For more information about bucket versioning, refer to [the Amazon documentation](#).

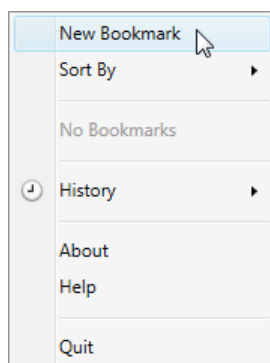
Bucket versioning is turned off by default. In CyberDuck, you can enable it in bucket properties. For example:



Mounting S3 storage with Mountain Duck

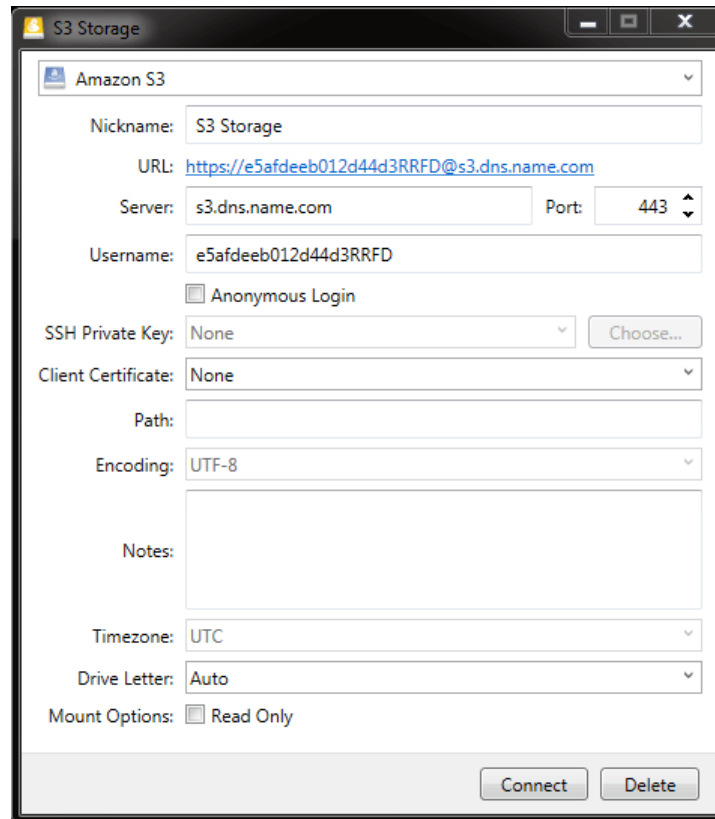
Mountain Duck enables you to mount and access Virtuozzo Hybrid Infrastructure S3 storage as a regular disk drive. Do the following:

1. If your service provider has provided you with an SSL certificate, install it.
2. In Mountain Duck, click **New Bookmark**.

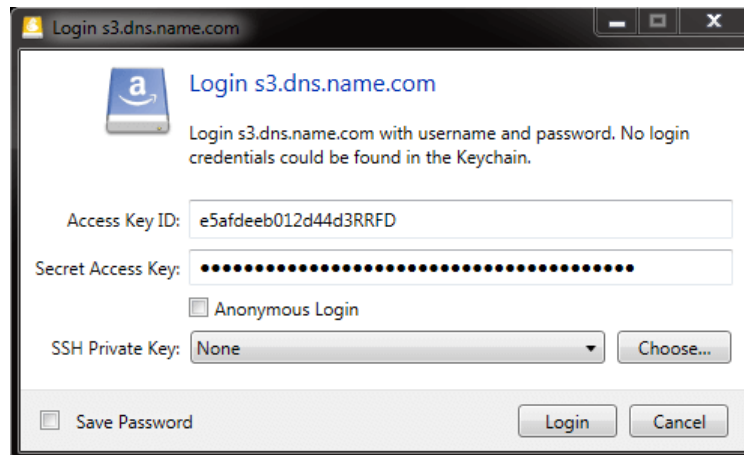


3. In the properties window, select **Amazon S3** profile from the first drop-down list and specify the following parameters:
 - Disk drive name in the **Nickname** field
 - Endpoint DNS name in the **Server** field

- Access key ID in the **Username** field
- Click **Connect**.



4. In the login window, specify **Secret Access Key** and click **Login**.



Mountain Duck will mount the S3 storage as a disk drive. On the disk, you can manage buckets and store files in them.

Creating S3 buckets on Mounted S3 Storage

Windows and macOS, operating systems supported by Mountain Duck, treat buckets as folders in case the S3 storage is mounted as a disk drive. In both operating systems, the default folder name contains spaces. This violates bucket naming conventions (refer to "S3 bucket and key naming policies" (p. 10)), therefore you cannot create a new bucket directly on the mounted S3 storage. To

create a bucket on a mounted S3 storage, create a folder with a name complying with DNS naming conventions elsewhere and copy it to the root of the mounted S3 storage.

S3 bucket and key naming policies

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

- Can be from 3 to 63 characters long
- Must start and end with a lowercase letter or number
- Can contain lowercase letters, numbers, periods (.), hyphens (-), and underscores (_)
- Can be a series of valid name parts (described previously) separated by periods

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

Accessing iSCSI targets

This section describes ways to attach iSCSI targets to operating systems and third-party virtualization solutions that support the explicit ALUA mode.

Accessing iSCSI targets from VMware ESXi

Before using Virtuozzo Hybrid Infrastructure volumes with VMware ESXi, you need to configure it to properly work with ALUA Active/Passive storage arrays. It is recommended to switch to the VMW_PSP_RR path selection policy (PSP) to avoid any issues. For example, on VMware ESXi 6.5:

- To set the default PSP for all devices, run:

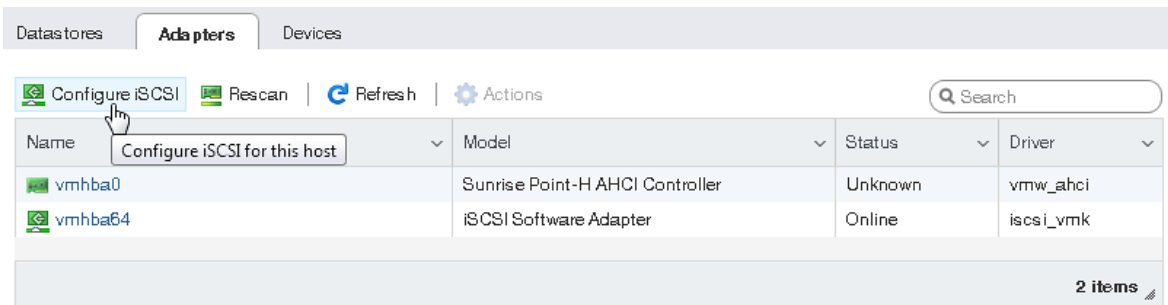
```
# esxcli storage nmp satp rule add --satp VMW_SATP_ALUA --vendor VSTORAGE \  
--model VSTOR-DISK --psp VMW_PSP_RR -c tpgs_on
```

- To set the PSP for a specific device, run:

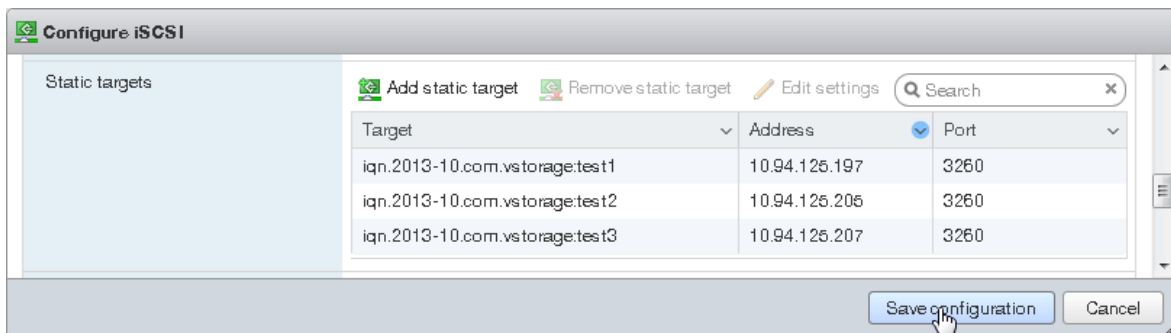
```
# esxcli storage core claimrule load
```

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

- In the Navigator, go to the **Storage > Adapters** tab and click **Configure iSCSI**.



- In the **Configure iSCSI** window, click **Add static target** in the **Static targets** section, fill out target IQNs, IP addresses, and ports. Click **Save configuration**.



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

Datstores Adapters Devices						
Name	Status	Type	Capacity	Queue ...	Vendor	
VSTORAGE iSCSI Disk (eui.6164383063623739)	Normal	Disk	10 GB	128	VSTORAGE	

1 items

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

Warning!

Partitioning the disk will erase all data from it.

The ready-to-use disk will appear in the list of datstores. You can now view its contents it with the datastore browser and provision it to VMs.

Datstores Adapters Devices							
Name	Drive Ty...	Capacity	Provisi...	Free	Type	Thin pr...	Access
datastore01	Non-SSD	9.75 GB	1.41 GB	8.34 GB	VMFS6	Supported	Single

1 items

Note

If your ESXi host loses connectivity to VMFS3 or VMFS5 datstores, follow the instructions in [KB article #2113956](#).

Accessing iSCSI targets from Linux

To connect a Linux-based iSCSI initiator to iSCSI targets of Virtuozzo Hybrid Infrastructure working in the ALUA mode, do the following:

- Make sure the required packages are installed.
 - On RPM-based systems (CentOS and other), run:

```
# yum install iscsi-initiator-utils device-mapper-multipath
```

- On DEB-based systems (Debian and Ubuntu), run:

```
# apt-get install open-iscsi multipath-tools
```

- Create and edit the configuration file `/etc/multipath.conf` as follows:

```
...
devices {
```

```

device {
    vendor "VSTORAGE"
    product "VSTOR-DISK"
    features "2 pg_init_retries 50"
    hardware_handler "1 alua"
    path_grouping_policy group_by_node_name
    path_selector "round-robin 0"
    no_path_retry queue
    user_friendly_names no
    flush_on_last_del yes
    failback followover
    path_checker tur
    detect_prio no
    prio alua
}
}
...

```

3. Load the kernel module and launch the multipathing service.

```

# modprobe dm-multipath
# systemctl start multipathd; systemctl enable multipathd

```

4. If necessary, enable CHAP parameters `node.session.auth.*` and `discovery.sendtargets.auth.*` in `/etc/iscsi/iscsid.conf`.

5. Launch the iSCSI services:

```

# systemctl start iscsi iscsid
# systemctl enable iscsi iscsid

```

6. Discover all targets by their IP addresses. For example:

```

# iscsiadm -m discovery -t st -p 10.94.91.49 10.94.91.49 3260,1 \
iqn.2014-06.com.vstorage:target1
# iscsiadm -m discovery -t st -p 10.94.91.54 10.94.91.54:3260,1 \
iqn.2014-06.com.vstorage:target2
# iscsiadm -m discovery -t st -p 10.94.91.55 10.94.91.55:3260,1 \
iqn.2014-06.com.vstorage:target3

```

7. Log in to the discovered targets. For example:

```

# iscsiadm -m node -T iqn.2014-06.com.vstorage:target1 -l
# iscsiadm -m node -T iqn.2014-06.com.vstorage:target2 -l
# iscsiadm -m node -T iqn.2014-06.com.vstorage:target3 -l

```

8. Find out the multipath device ID. For example:

```

# multipath -ll
36000000000000000000000b50326ea44e3 dm-10 VSTORAGE,VSTOR-DISK
size=200G features='2 pg_init_retries 50' hwhandler='1 alua' wp=rw
|+- policy='round-robin 0' prio=50 status=active

```

```
| `- 6:0:0:1 sdf 8:80 active ready running
|+- policy='round-robin 0' prio=1 status=enabled
| `- 8:0:0:1 sdj 8:144 active ghost running
`+- policy='round-robin 0' prio=1 status=enabled
  `- 7:0:0:1 sdh 8:112 active ghost running
# fdisk -l | grep 360000000000000000000000b50326ea44e3
Disk /dev/mapper/360000000000000000000000b50326ea44e3: 10.7 GB, \
10737418240 bytes, 20971520 sectors
```

You can also find out the multipath device ID by adding 360000000000000000000000 to the last six bytes of the volume ID. In the example above, 360000000000000000000000b50326ea44e3 is the multipath device ID mapped from the volume ID 61c9d567-4666-4c16-8030-b50326ea44e3.

Now you can create partitions on the iSCSI device (/dev/mapper/360000000000000000000000b50326ea44e3 in this example), as well as format and mount it to your initiator node using standard Linux tools.

When you no longer need the external iSCSI device, you can remove it from the initiator node. Do the following:

1. Make sure the iSCSI device is not in use.
2. Disable multipathing to the device. For example:

```
# multipath -f /dev/mapper/360000000000000000000000b50326ea44e3
```

3. Log out of the iSCSI targets. For example:

```
# iscsiadm -m node -T iqn.2014-06.com.vstorage:target1 -p 10.94.91.49:3260 -u
# iscsiadm -m node -T iqn.2014-06.com.vstorage:target2 -p 10.94.91.54:3260 -u
# iscsiadm -m node -T iqn.2014-06.com.vstorage:target3 -p 10.94.91.55:3260 -u
```

4. Delete the iSCSI targets. For example:

```
# iscsiadm -m node -o delete -T iqn.2014-06.com.vstorage:target1 \
-p 10.94.91.49:3260
# iscsiadm -m node -o delete -T iqn.2014-06.com.vstorage:target2 \
-p 10.94.91.54:3260
# iscsiadm -m node -o delete -T iqn.2014-06.com.vstorage:target3 \
-p 10.94.91.55:3260
```

Accessing iSCSI targets from Microsoft Hyper-V

Before connecting an iSCSI initiator of Microsoft Hyper-V to iSCSI targets working in the ALUA mode, you need to install and configure Multipath I/O (MPIO). This feature can be used starting from Windows Server 2008 R2. To connect the initiator, for example, on Microsoft Hyper-V Server 2016, do the following:

1. Run Windows PowerShell with administrator privileges and install MPIO.

```
> Enable-WindowsOptionalFeature -Online -FeatureName MultiPathIO
```

Your server will automatically reboot to finalize the installation.

2. In the Windows PowerShell console, configure MPIO as follows:

a. Enable support for iSCSI disks:

```
> Enable-MSDSMAutomaticClaim -BusType iSCSI
```

b. Set the failover policy to Fail Over Only. The policy uses a single active path for sending all I/O, and all other paths are standby. If the active path fails, one of the standby paths is used. When the path recovers, it becomes active again.

```
> Set-MSDSMGlobalDefaultLoadBalancePolicy -Policy FOO
```

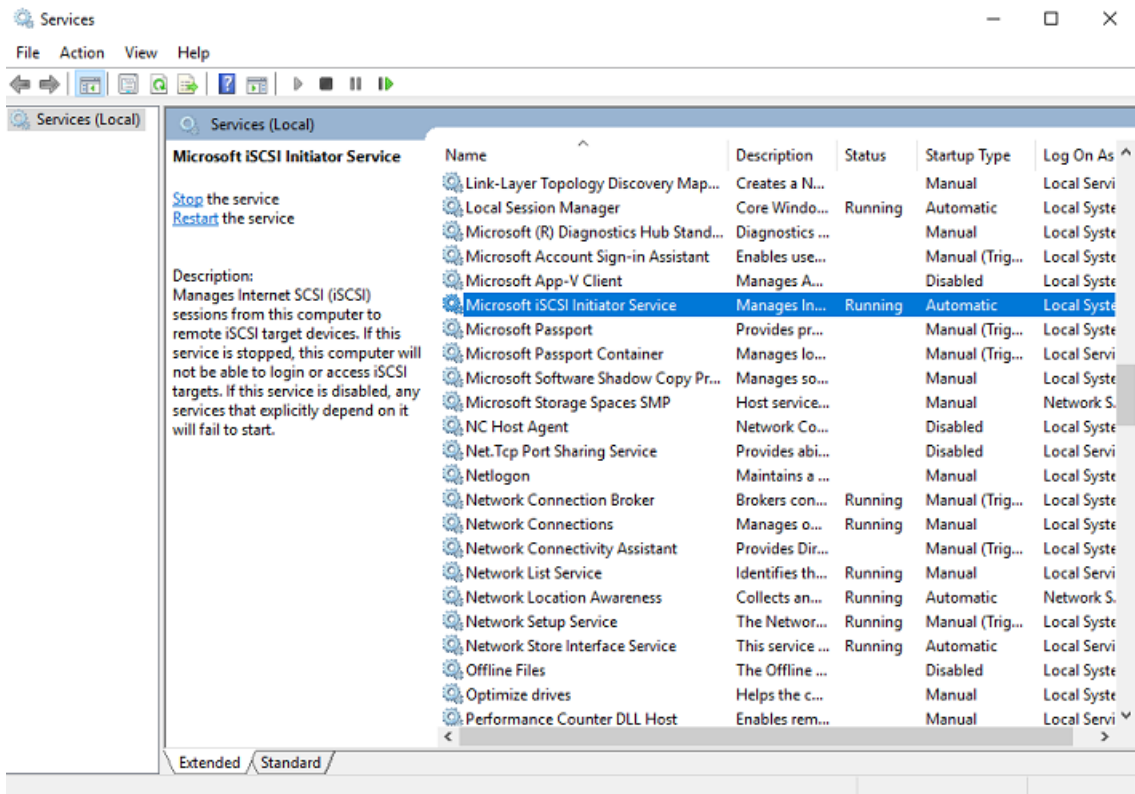
c. Enable path verification. By default, the initiator will verify each path every 30 seconds.

```
> Set-MPIOSetting -NewPathVerificationState Enabled
```

d. Reboot the server.

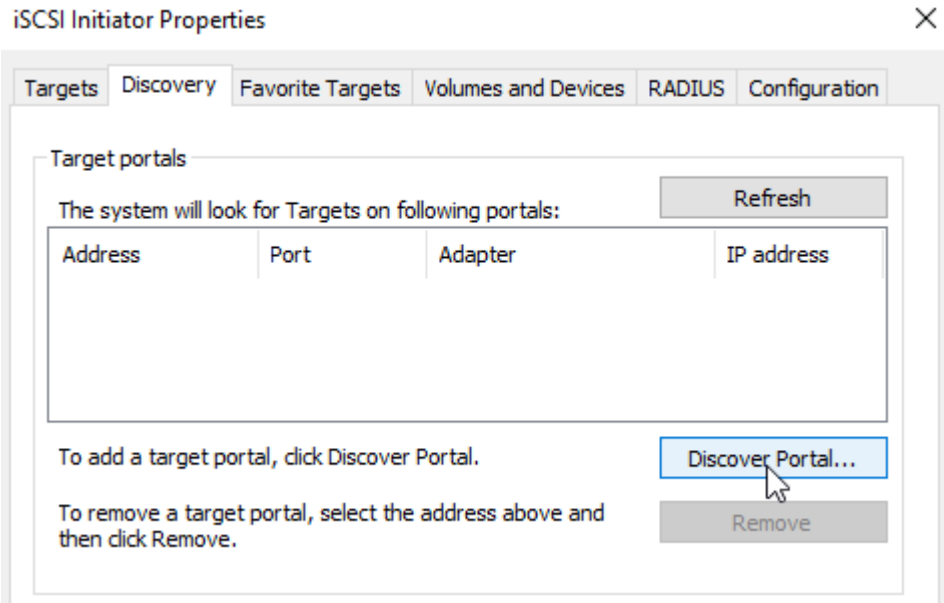
3. Connect your targets to the iSCSI initiator as follows:

a. In the **Control Panel > System and Security > Administrative Tools > Services** window, make sure that **Microsoft iSCSI Initiator Service** is running and its startup type is set to **Automatic**.




b. Launch **iSCSI Initiator**.


c. In the **iSCSI Initiator Properties** window, open the **Discovery** tab and click **Discover Portal**.




- d. In the **Discover Target Portal** window, enter the target IP address and click **OK**. Repeat this step for each target from the target group.



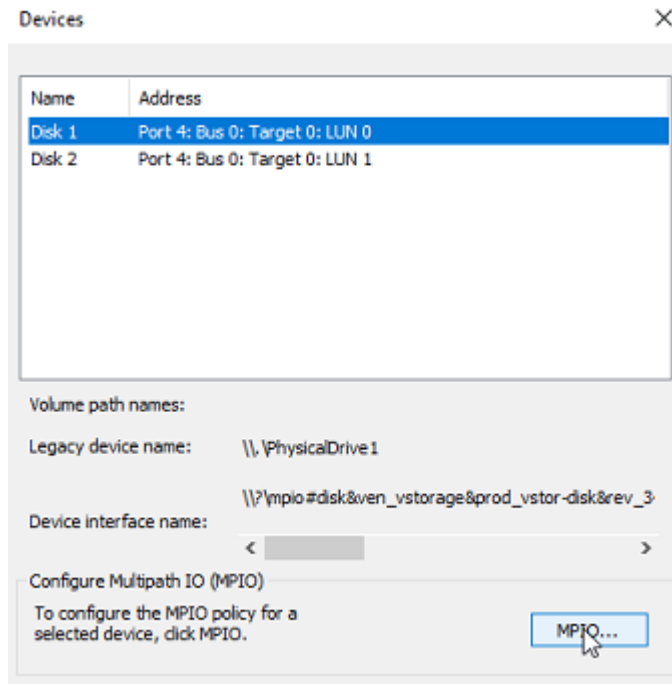
- e. On the **Targets** tab, click **Refresh** to discover the added targets.



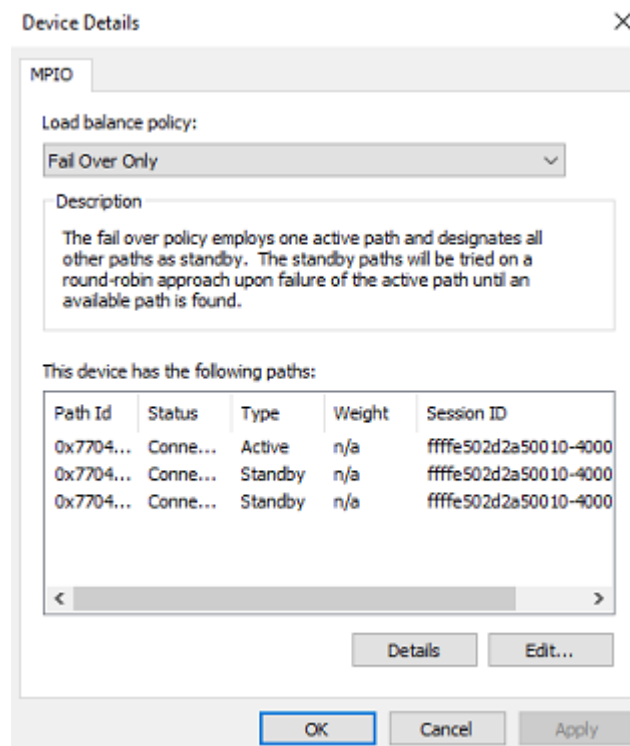
- f. Click **Connect** for each target to connect it to the initiator. In the **Connect To Target** window, select the **Enable multi-path** checkbox and click **OK**.



- g. On the **Targets** tab, click **Devices..**, select the connected LUN, and click **MPIO...**




h. Make sure the connected LUN has several paths.




You can now initialize the newly added disk for use in Microsoft Hyper-V. Do the following:

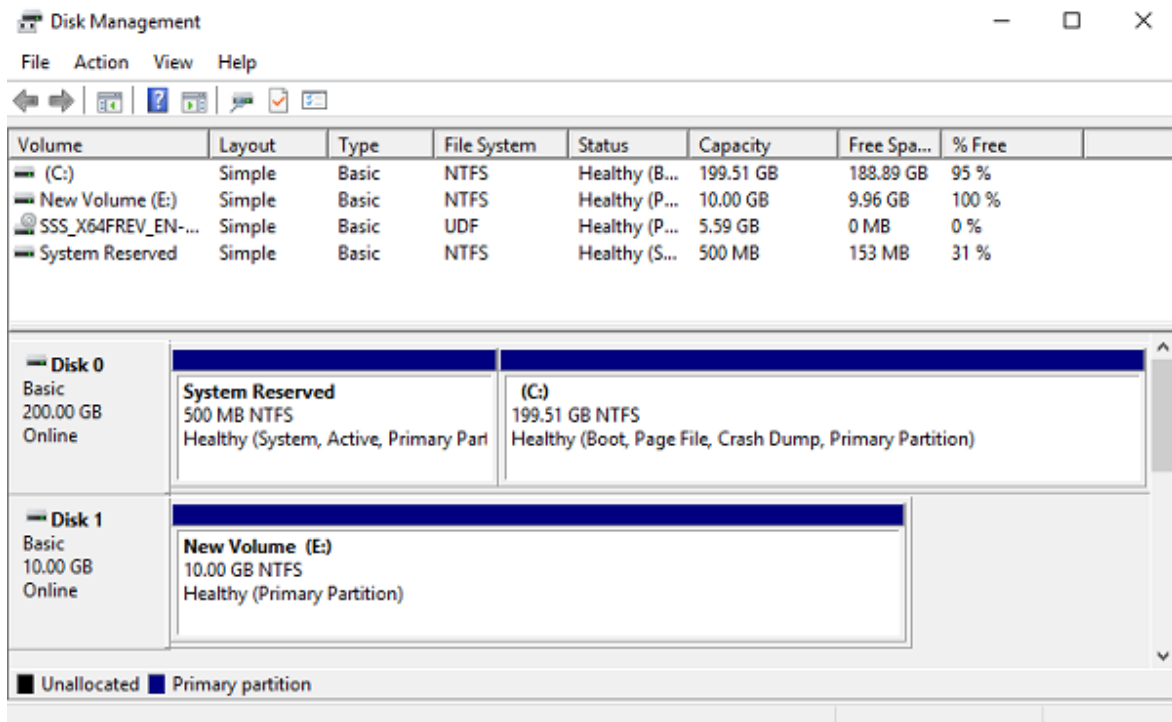
1. Open **Disk Management**, right-click the added disk, and choose **Properties** from the drop-down menu.



2. Check the settings on the **MPIO** tab. The first connected target becomes **Active/Optimized** and the preferred path.



3. Partition and format the disk as usual.



Accessing NFS shares

This section describes ways to mount Virtuozzo Hybrid Infrastructure NFS shares on Linux and macOS.

Note

Virtuozzo Hybrid Infrastructure currently does not support the Windows built-in NFS client.

Mounting NFS exports on Linux

You can mount an NFS export created in Virtuozzo Hybrid Infrastructure like any other directory exported via NFS. You will need the share IP address (or hostname) and the volume identifier.

In console, run the following commands:

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

where:

- `-o vers=4.0` is the NFS version to use.
Virtuozzo Hybrid Infrastructure supports NFS versions 4.0 and 4.1.
- `<share_IP>` is the share IP address. You can also use the share hostname.
- `/<share_name>/` is the root export path, like `share1`. For user exports, specify their full path, for example: `/<share_name>/export1`.
- `/mnt/nfs` is an existing local directory to mount the export to.

Mounting NFS exports on macOS

You can mount an NFS export created in Virtuozzo Hybrid Infrastructure like any other directory exported via NFS. You will need the share IP address (or hostname) and the volume identifier.

You can use the command-line prompt or Finder:

- In console, run the following commands:

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

where:

- `-o vers=4.0` is the NFS version to use.
Virtuozzo Hybrid Infrastructure supports NFS versions 4.0 and 4.1.
- `<share_IP>` is the share IP address. You can also use the share hostname.
- `/<share_name>/` is the root export path, like `share1`. For user exports, specify their full path, for example: `/<share_name>/export1`.
- `/mnt/nfs` is an existing local directory to mount the export to.

- In Finder, do the following:
 1. Set the NFS version to 4.0. To do this, add the `nfs.client.mount.options = vers=4.0` line to the `/etc/nfs.conf` file.
 2. In the **Finder > Go > Connect to server** window, specify `nfs://192.168.0.51:/<share_name>/` where:
 - `192.168.0.51` is the share IP address. You can also use the share hostname.
 - `/<share_name>/` is the root export path. For user exports, specify their full path, for example: `/<share_name>/export1`.
 3. Click **Connect**.The Finder will mount the export to `/Volumes/<share_name>/`.